

Essays in Empirical Macroeconomics

DISSERTATION
zur Erlangung des Doktorgrades
der Wirtschafts- und Sozialwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen

vorgelegt von
M.SC. KNUT NIEMANN

Tübingen
2023

- 1. Betreuer: Prof. Dr. Gernot J. Müller
- 2. Betreuer: Prof. Dr. Joachim Grammig

Tag der mündlichen Prüfung: 06.12.2023

Dekan: Prof. Dr. Ansgar Thiel

- 1. Gutachter: Prof. Dr. Gernot J. Müller
- 2. Gutachter: Prof. Dr. Joachim Grammig

Acknowledgments

First and foremost, I thank my advisor Gernot Müller. Despite studying in Tübingen for ten years, I only first met him at the interview to be his Ph.D. student in 2019. I am thankful for his trust in me and for admitting me as his Ph.D. student, although my macroeconomic education was limited. He is co-author of the first two chapters of this dissertation, and I am grateful for the invaluable guidance and profound expertise he provided during my doctoral journey.

I am also thankful to my second advisor Joachim Grammig. He shaped my understanding of statistics and econometrics during my Bachelor's and Master's degrees. His education gave me the confidence to start a Ph.D. in a field where I had limited prior experience.

Furthermore, I thank my other co-authors: Benjamin Born, Zeno Enders, and Manuel Menkhoff. Their dedication inspired me, and working together was a great joy.

During my Ph.D. studies, I had the opportunity to join the Bundesbank for a five-month internship. Almira Enders and Felix Strobel, all other members of the team, and the other interns made it a great experience, and I learned a lot.

At a less scientific level, I am also thankful to members of and other Ph.D. students in the department. Jonas Adolph, Alexander Dietrich, Peter Eppinger, and Susanne Wellmann were great company and helped me navigate the expected and unexpected challenges that occur during a Ph.D.

I have known Constantin Hanenberg and Martin Kipp since I started my Bachelor's. Since then, we have stayed in Tübingen, obtained our M.Sc., and started PhDs at different chairs at the University of Tübingen. I am thankful for their company and discussions, both on academic and personal topics.

On a personal level, I am thankful to my family for their support during my Ph.D. The distance between Tübingen and my parents' home is 860km, so our time together was limited, but we made the best of it. I very much enjoyed our holidays in the Alps during the summer break, short weekends in Heidelberg, and the weeks at home after Christmas. Our time together helped me not to get lost in an ivory tower.

Lastly, I thank my girlfriend, Francesca Izzi. I am grateful for your support, perspective, and understanding over these last 3.5 years. The Ph.D. was the most significant endeavor of my life. But, thanks to you and the hours of Instagram reels you showed me, I was able to push through!

Contents

List of Figures	iii
List of Tables	iv
Introduction	1
1 Firm expectations about production and prices: Facts, determinants, and effects	3
1.1 Introduction	3
1.2 Surveying firm expectations	5
1.2.1 Background	5
1.2.2 Example: The ifo Business Expectations Panel	9
1.3 Stylized facts	10
1.4 Expectation formation	17
1.4.1 Determinants of expectations	17
1.4.2 Over- and underreaction to news	21
1.5 Firm expectations and firm decisions	24
1.5.1 The effect of firm expectations	24
1.5.2 Firm-level uncertainty and firm decisions	27
1.6 Conclusion	27
1.A Appendices	29
1.A.1 Expectation errors	29
1.A.2 Additional figures and tables	30
2 Firm Expectations and News: Micro v Macro	37
2.1 Introduction	37
2.2 Measuring forecast errors and news	41
2.2.1 The ifo survey	41
2.2.2 Forecast errors	42
2.2.3 Macro news	42
2.2.4 Micro news	44
2.3 How firm expectations respond to news	45
2.3.1 Empirical framework	46
2.3.2 Results	47
2.3.3 Measurement error and robustness	51
2.3.4 Accounting for heterogeneity	53

2.3.5	Reaction to news and firm performance	56
2.3.6	Further evidence for Italian firms	58
2.4	A model of island illusion	60
2.4.1	Setup and timing	60
2.4.2	Households	61
2.4.3	Firms	62
2.4.4	Island illusion	63
2.4.5	Monetary policy and market clearing	65
2.4.6	Accounting for over- and underreaction	65
2.5	Conclusion	67
2.A	Appendices	68
2.A.1	Additional figures and tables	68
2.A.2	SIGE Data	78
2.A.3	Model solution	81
2.A.4	Proofs	87
3	Political Distance and International Trade	95
3.1	Introduction	95
3.2	Gravity and political distance	100
3.2.1	A simple gravity model	100
3.2.2	Trade flows	101
3.2.3	Trade costs	104
3.2.4	Descriptive statistics	106
3.3	Bilateral trade and political distance	107
3.3.1	Estimation framework	107
3.3.2	Results	108
3.3.3	Robustness and alternative specifications	109
3.3.4	Heterogeneity	111
3.3.5	Identification	114
3.4	Counterfactuals	115
3.5	Conclusion	117
3.A	Appendices	119
3.A.1	Additional figures and tables	119
3.A.2	Constructing a sectoral trade panel	130
	Conclusion	133
	References	135

List of Figures

1.1	What matters for firm decisions?	4
1.2	BEP observations across both panel dimensions	9
1.3	Performance of firm expectations relative to benchmark models	12
1.4	Response of forecast error to forecast revision	23
1.A.1	Point estimates for constant and slope	31
2.1	The ifo survey, forecast errors, and news	43
2.1	Distribution of firm-level responses to news	50
2.2	Response to concurrent and lagged news	52
2.3	Response to news over time	56
2.A.1	Average forecast revisions and production growth	68
2.A.2	Relation between macro and micro coefficients at the firm-level	69
2.A.3	Firm-level regressions – univariate distribution of news coefficients	80
3.1	Fifty years of trade in TradeProd	103
3.2	Resolutions at the UN GA, votes, and political distance	105
3.2	Elasticity of political distance across time	114
3.A.1	Standardized dispersion of political distance to the US	129
3.A.2	Distribution of political distances: counterfactual vs. 2018	129

List of Tables

1.1	Surveys with firm expectations about firm-specific developments	7
1.2	Average unconditional expectation errors	11
1.3	Experience and expectation errors	14
1.4	Dispersion and volatility measures	16
1.5	Stickiness of firm expectations	18
1.6	Determinants of production and price expectations	20
1.7	Effects of increased and decreased production expectations	26
1.A.1	Definitions of qualitative expectation errors	29
1.A.2	Relevant questions from the ifo Survey	30
1.A.3	Summary statistics on firm-level average forecast errors	31
1.A.4	Summary statistics on firm-level average squared forecast errors	32
1.A.5	Definition of variable blocks	32
1.A.6	Overreaction to firm-specific news	33
1.A.7	Summary statistics firm-level constant estimates	34
1.A.8	Predictability of expectation errors	35
2.1	Macro news and forecast revisions	46
2.1	Over- and underreaction to news	48
2.2	Alternative specifications	53
2.3	Heterogeneity	54
2.4	Over- and underreaction to news and real activity	57
2.5	Over- and underreaction to news—Italian firms	59
2.A.1	Relevant questions from ifo survey	70
2.A.2	Alternative specifications	71
2.A.3	Relevant questions from SIGE	78
2.A.4	Additional regression results from the SIGE	79
3.1	Two panels of trade flows	102
3.2	Descriptive statistics for the combined TradeProd panel	106
3.1	Political distance and bilateral trade	108
3.2	Alternative specifications	110
3.3	Elasticities differ meaningfully across countries and sectors	112
3.4	Causality	115
3.1	Counterfactual v. actual exports	117
3.A.1	List of countries in the sectoral panel	119
3.A.2	List of countries in TradeProd	120
3.A.3	Details on alternative specifications	121

3.A.4	Full results for heterogeneity	125
3.A.5	An illustration of an $m : n$ conversion	132

Introduction

This dissertation features three chapters that each make relevant contributions to economic research. The first two chapters focus on expectations that take center stage in modern macroeconomics. We focus on firm expectations about their own variables. For the empirical analysis, we use the German ifo survey and, additionally, in the second chapter, an Italian firm survey run by the Banca d'Italia. The third chapter changes focus on the current debate about decoupling or de-risking in international trade. It introduces a novel type of trade costs into a gravity model and presents a new panel on sectoral trade.

The first chapter consolidates what we know about firm expectations about their own variables. We illustrate the findings gathered from different surveys using the German ifo survey of firms. It distills six stylized effects facts about firm expectations about their own variables. Then it discusses the expectation formation of firms and concludes by considering the causal effect of firm expectations.

The second chapter zooms in on the expectation formation of firms about their own prices and production. We find that how firms react to news crucially depends on the type of news. We distinguish micro news and macro news. Micro news is about firms' own developments, and macro news is about the aggregate economy. Based on the ifo survey and a survey of Italian firms, we show that firms overreact to micro news and underreact to macro news. We propose a general-equilibrium model with "island illusion" to explain these patterns in the data. This way, we contribute to efforts to flesh out the expectation-formation process in greater detail and, eventually, converge to a new paradigm for rational expectations.

The third chapter is single-authored. I started working on it during my internship at the Bundesbank in the summer of 2022. It introduces a novel trade cost to gravity models of international trade. I find that an increase in political distance predicts a significant decrease in bilateral trade, controlling for time-constant pair characteristics, tariffs, and economic integration agreements. I use this insight to construct a counterfactual decoupling scenario for 2018 that mimics political distances during the Cold War. There is a substantial reshuffling of trade in this scenario.

More specifically, Chapter 1 is based on a joint research project with Benjamin Born, Zeno Enders, and Gernot Müller. This chapter revisits survey evidence about firm expectations, focusing on firms' production and prices. We aim to synthesize the evidence established based on various firm surveys from different countries. We complement our discussion of existing work with new evidence based on the ifo Survey of German firms. This allows us, first, to put together five stylized facts regarding firm expectations and expectation errors. In addition, we present new evidence regarding the stickiness of firm expectations. Second, we use the same data set to revisit key results regarding the formation of firm expectations. Firm expectations react strongly to firm-specific developments, whereas aggregate variables are less important. Third, we summarize the evidence on how firm expectations drive firm decisions.

Chapter 2 is based on a joint research project with Benjamin Born, Zeno Enders, Gernot Müller, and Manuel Menkhoff. Using firm-level data, we study how firm expectations adjust to news while accounting for a) the heterogeneity of news and b) the heterogeneity of firms. We classify news as either micro or macro, that is, information about firm-specific developments or information about the aggregate economy. Survey data for German and Italian firms allows us to reject rational expectations: Both types of news predict forecast errors *at the firm level*. Yet, while firm expectations overreact to micro news, they underreact to macro news. We propose a general-equilibrium model where firms suffer from “island illusion” to explain these patterns in the data.

Chapter 3 is my single-authored project. Geopolitical tensions are increasing worldwide, and some foresee a decoupling of international trade from politically distant countries. Still, we know little about the role of political distance in international trade. I introduce a novel trade cost to the gravity model: political distance computed from countries’ voting behavior at the United Nations General Assembly. On average and over the last 50 years, I find an increase in political distance within a country pair by one standard deviation to predict a significant decrease in trade by 4 percent. To zoom in on different sectors, I construct a novel trade panel. I find the predicted decrease in trade to be more than twice as large for trade involving the US, the EU, or the UK and trade in strategic sectors. In a counterfactual decoupling scenario comparable to a New Cold War in 2018, the median absolute change of a trade flow is 56 percent of its actual value in 2018, suggesting substantial trade diversion.

In sum, this dissertation features two distinct contributions that advance our understanding of macroeconomics from a theoretical and empirical perspective. First, it contributes to our theoretical and empirical understanding of how firms form expectations about their own variables, and it contributes to efforts to converge to a new paradigm for rational expectations. Second, it informs the discussion about decoupling international trade by quantifying how much trade reshuffling would occur in such a decoupling scenario.

Firm expectations about production and prices: Facts, determinants, and effects

Joint with Benjamin Born, Zeno Enders, and Gernot J. Müller

1.1 Introduction

In this chapter, we review recent work which uses survey data to analyze firm expectations—with a particular focus on firms’ production and price expectations. These matter a great deal for actual firm decisions. To see this, consider the responses to a brief survey among German firms about their production and pricing decisions. As illustrated by Figure 1.1, firm-specific developments are as important for these decisions as the developments of the aggregate economy and a firm’s market segment (see also Freuding et al. 2021). At the same time, forecasting their own variables is potentially hard for firms and perhaps even harder than forecasting the aggregate economy (Bloom et al. 2021).¹


We revisit the evidence based on various surveys from different countries. Because the existing literature on the issue is still in a somewhat early stage, we complement our discussion of existing work with new evidence based on the ifo Survey of German firms. The ifo Survey is one of the oldest and largest surveys of firms currently available. It is based on a firm survey which has been conducted since 1949 and whose design has since then been adopted by other surveys as well (Becker and Wohlrabe 2008). We provide details about this survey and introduce basic concepts in Section 1.2.²

In Section 1.3, we use the ifo Survey to establish—on the basis of a common data set—five stylized facts which emerge robustly across various studies and surveys. First, firms’ expectation errors are unconditionally unbiased, that is, mostly not significantly different from zero. Second, survey responses are informative in that they outperform static and adaptive expectations in terms of forecasting firm-specific developments. Third, larger and older firms tend to do even better in terms of forecasting. Fourth, we find that firms make predictable forecast errors. Past information about firms’ own variables, in particular, predict expectation errors. Fifth, the dispersion and volatility of expectations and expectation errors is countercyclical, in line with the notion that uncertainty increases during recessions. In addition to those stylized facts, we present a sixth observation which has not been made in the survey literature so far: firm expectations are sticky, that is, they are adjusted only infrequently.

¹Giustinelli (2023) and Baumeister (2023) consider inflation expectations of households and firms, respectively.

²The ifo Survey is also one of the surveys discussed in greater detail in Carstensen and Bachmann (2023).

Figure 1.1: What matters for firm decisions?



Notes: responses to special question in the October 2020 wave of the ifo survey of German firms. “How important are the following domains for your production and/or pricing decisions?”, with answer scale 1 to 5. Categories: recent developments in the aggregate economy, the firm’s market segment, and within the firm. No. of responses: 1,666. Left bars show results for all firms, the other blocks show results for specific sectors.

In the second part of the chapter, we seek to shed light on both, expectation formation (Section 1.4) and the effects of expectations on firm actions (Section 1.5). We stick to our strategy and revisit for our sample results established in earlier work. As we do so, we focus on the main results in the literature but also offer some additional findings. A first important result concerning the expectation-formation process is that firm-specific variables account for almost all the variation in firm expectations regarding their own output and prices. Next, we consider the responsiveness of firm expectations to news. Here we discuss some recent results which pertain mostly to professional forecasters (Coibion and Gorodnichenko 2015; Bordalo et al. 2020, and Clements et al. (2023)). As a noteworthy exception, Born et al. (2022) study the response of firms’ forecast errors about their own variables to forecast revisions (news): firms tend to overreact to firm-specific news, but underreact to news about the aggregate economy.

Eventually, we care about firm expectations to the extent that they matter for actual outcomes—an issue we revisit last, following earlier work by Enders (2020). Here two results are key. First, firm expectations about future production significantly impact current production and pricing decisions. Second, this also holds for expectations that turn out to be incorrect from an ex-post point of view. This suggests that expectations not only operate as a transmission channel of news but also as a genuine source of shocks. There is also evidence that expectations are key for firms’ investment decisions.

Before getting started, we note that rather than relying on surveys, one may measure expectations or, relatedly, confidence through proxies extracted from observable behavior (e.g., Malmendier and Tate 2005b,a; Hirshleifer et al. 2012). Also, in our analysis, we treat firms and firm expectations as the primitives and abstract from within-firm dynamics and management practices and personality traits of CEOs (e.g., Bloom and Reenen 2007; Kaplan et al. 2012).

1.2 Surveying firm expectations

By now there is a sizeable number of firm surveys which collect direct evidence on firm expectations about their own variables, such as production and prices. In what follows we provide an overview. We then zoom in on the ifo Business Expectations Panel (BEP), which we will use throughout the chapter to replicate the most important findings in the literature and to generate some new results based on a single data set.

1.2.1 Background

Several surveys were initiated in the 1950s–1970s in order to provide early and additional information about the current state of the (national or regional) economy when official statistics were incomplete and available with a considerable lag only (INSEE 2007; Nerb and Sauer 2020; Bank of Japan 2020; Trebing and Fenske 2018).³ In these surveys, firms are typically asked only qualitative questions. They may respond that they expect, say, prices or production to increase, stay the same, or decrease, likewise for their business situation or related variables.⁴

Questions regarding realized values are typically structured analogously to those about expectations. For instance, firms report if production had risen, fallen, or stayed the same. Nerb and Sauer (2020) document that this format was adopted in order to increase the return rate of the survey. Moreover, the format is considered adequate because the surveys feature several questions which require subjective evaluations. Responding qualitatively to questions about, say, the current business situation or the adequateness of inventories, allows firms to weigh different aspects depending on current circumstances in a flexible manner. These types of questions also constitute the so-called ‘Judgement’ part of the Tankan Survey (Bank of Japan 2020).⁵ Rosewell (1987) adds, referring to the CBI Industrial Trends Survey, that the qualitative format increases chances that senior management answers the questionnaires (which is confirmed in Glynn 1969) and that questions about actual outcomes and expectations can be easily asked in the same context. By aggregating answers regarding current and expected firm-specific variables (most often by forming balances of positive and negative answers), the surveys turn out to have a high predictive value for sector-wide or even national economic developments, see Abberger and Wohlrabe (2006), Henzel and Rast (2013), and Lehmann (2023) for the ifo Survey, Trebing and Fenske (2018) for the Manufacturing Business Outlook Survey of the Philadelphia Fed, and Glynn (1969) for capital expenditure elicited in the CBI Industrial Trends Survey. Note that this result lends credibility to the choice of aggregating qualitative answers by calculating balances of positive and negative answers.

The large potential of business surveys for rigorous empirical analysis became more apparent over time (see, e.g., Nerb 1987; Seiler and Wohlrabe 2013, for the ifo Survey).⁶ To

³See also Carstensen and Bachmann (2023) for further details on individual firm surveys.

⁴See Table A.2 in the online appendix for examples of qualitative questions from the ifo Survey. Note that throughout this chapter, material in the online appendix will be marked with an “A.” prefix.

⁵The predecessor of the Tankan started in 1951, following the methodology of the ifo Survey (Bank of Japan 2022).

⁶This is not necessarily true for the underlying micro data, that is, the individual responses. They were often, after aggregation, not kept for later use.

increase the scope further still, quantitative questions have been added in several surveys.⁷ In this case, respondents are asked to provide a specific number or to choose from predefined ranges when responding to questions about, say, expected sales growth. Providing predefined ranges to elicit point estimates involves potential pitfalls, as the provision of ranges may have a bearing on the elicited answers (Schwarz et al. 1985). Even more recently, following Bloom (2009) and others, business-cycle research highlighted the role of uncertainty for economic developments and, as a consequence, several firm surveys now ask for probability distributions in addition to point forecasts to measure uncertainty.⁸ Specifically, survey participants are asked to assign probabilities to either several bins that cover predefined ranges for the future realizations of the variable of interest (e.g., Business Inflation Expectations Survey) or to freely selected bins (Survey of Business Uncertainty, SBU).⁹ However, in order to evaluate the answers to these questions additional assumptions need to be made regarding, for instance, probability-mass distribution inside the bins or the underlying models (formal or not) used by survey participants (Krüger and Pavlova 2020; Glas and Hartmann 2021).¹⁰

We provide an overview of existing firm surveys in Table 1.1, Panels (a) and (b). Here we focus on those surveys that are available for economic research on firm expectations about firms' own variables.¹¹

⁷For instance, the ifo Survey and the CBI Industrial Trends Survey introduced quantitative questions in 2005 and 2008, respectively. There is some evidence that using qualitative (elicited via visual analog scales) and quantitative expectation data yields similar results (Enders 2020). Similarly, we stress that the facts established in Section 1.3 hold for qualitative and quantitative data. Nevertheless, a systematic investigation of differences induced by choosing qualitative or quantitative answer possibilities, e.g., by randomizing this choice, seems fruitful.

⁸See Bruine de Bruin et al. (2023) for the use of probabilistic questions in household surveys.

⁹Bloom et al. (2020) analyze business expectations that are surveyed as part of the Census Bureau's Management and Organizational Practices Survey. For selected years, it elicits point estimates for current-year outcomes and five-point probability distributions for the next. Bloom et al. (2020) find that 85% of respondents provide logically sensible responses to the five-point distribution questions, suggesting that most managers can form and express detailed subjective probability distributions.

¹⁰See also Clements et al. (2023), for issues relating to constructing measures of disagreement and uncertainty in the context of surveys of professionals.

¹¹We only consider those surveys that include questions about firm expectations about their own variables and whose firm-level answers are generally provided to researchers. These criteria eliminate a moderate number of firm surveys.

Table 1.1: Surveys with firm expectations about firm-specific developments

(a) General information

Name	Country	Expectation Variables	From	Freq.	Format	Maintained by
ifo Business Climate Surv.	Germany	output, prices, employment, business situation	1949	m	ql, qt 2005+ d 2013+	ifo
Tankan Surv.	Japan	sales, exports, profits, investment	1951	q	ql, qt	METI
CBI Industrial Trends Surv.	UK	wages, sales prices, employment, unit costs, , new orders	1958	q	ql, qt 2008+	Confederation of British Industry
Monthly Outlook Surv. in Industry	France	sales, prices, employment	1962	m	ql, qt	INSEE
Surv. of Industrial Trends	Australia	output, employment, prices, stocks, overtime	1966	m	ql	Australian Chamber of Commerce
Surv. of Production Forecasts	Japan	production	1971	m	qt	METI
Surv. on Industrial and Service Firms	Italy	investment, production, turnover, prices, costs	1972	a	qt	Banca d'Italia
ifo Investment Surv.	Germany	investment	1973	s	qt	ifo
Basic Surv. on Overseas Business Activities	Japan	sales	1995	a	qt	METI
CFO Surv.	US	revenue, wages, unit costs, employment	1996	q	qt	FRB Richmond and FRB Atlanta
Surv. on Inflation and Growth Expectations	Italy	economic situation, prices, demand, investment, empl.	1999	q	ql, qt	Banca d'Italia
Business Outlook Surv.	Japan	sales, operating profits	2004	q	qt	Ministry of Finance of Japan
Monitoraggio Economia e Territorio Surv.	Italy	sales, prices	2008	a	ql, qt	MET Research Center
Management and Organizational Practices Surv.	US	production, capital expenditures, employment, costs	2010	5a	ql	U.S. Census Bureau
Business Inflation Expectations Surv.	US	unit costs	2011	m	qt, d	FRB Atlanta
Surv. of Business Uncertainty	US	employment, sales, capital expenditures (investment rate)	2014	m	d	FRB Atlanta
Bundesbank Online Panel - Firms	Germany	employment, sales, inputs, finances, inventories	2020	i	ql, qt, d	Bundesbank

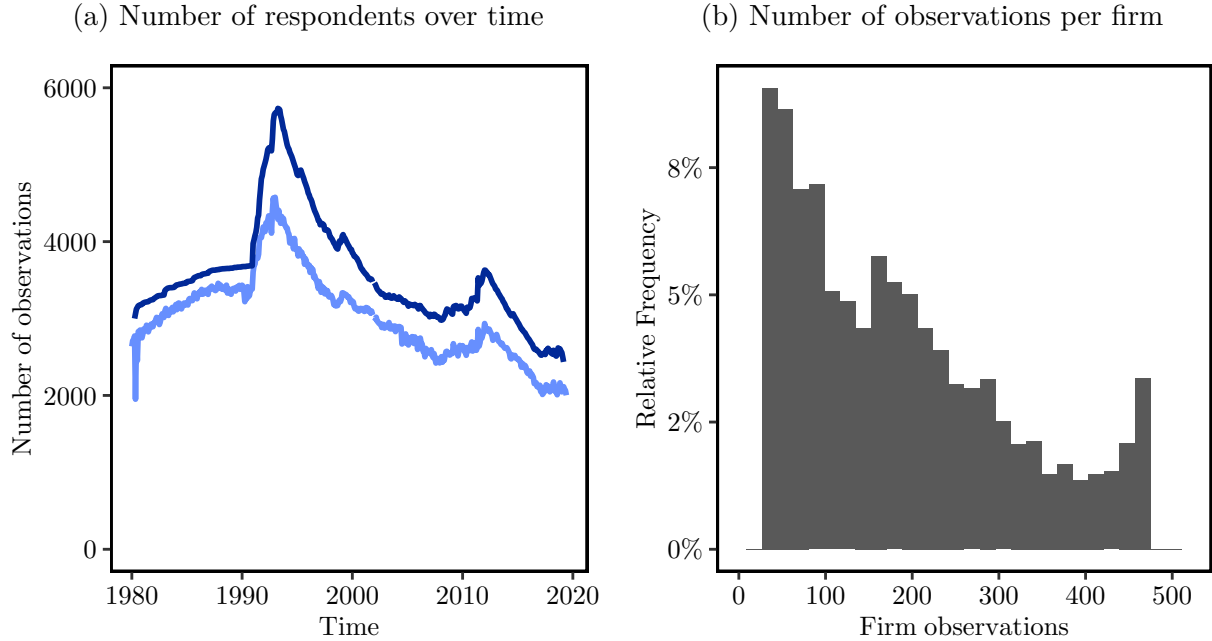
Notes: Frequencies (Freq.) are monthly (m), quarterly (q), semi-annually (s), annually (a), every 5 years (5a), and irregular (i). Formats are qualitative (ql), quantitative (qt), and distributional (d). METI is the Ministry of Economy, Trade, and Industry in Japan. Surveys ordered by their inception date, although the quality and scope of the initial waves may be much reduced (if they are available at all) relative to subsequent waves, e.g., data from the ifo Business Climate Survey is available for research since 1980. Only those surveys are listed whose firm-level data about firms' expectations about their own variables are generally provided to researchers. For this reason, the surveys of some central banks and regional Federal Reserve Banks (mostly <250 participants/month) are not included, e.g., the Business Outlook Surveys run by the Bank of Canada and the FRB Philadelphia. Similarly, the Joint Harmonised EU Programme of Business and Consumer Surveys consists of several national surveys but does not provide firm-level data.

(b) Additional information

Name	Selected Literature	Sectors	Resp.*	Firm Size	Documentation
ifo Business Climate Surv.	Nerlove (1983), Kawasaki and Zimmermann (1986), Bachmann et al. (2013), Bachmann and Elstner (2015), Massenot and Pettinicchi (2018), Enders et al. (2019, 2022), and Born et al. (2022)	man	2,000	nr	bit.ly/doc-ifo
Tankan Surv.	Morikawa (2016)	nr	11,000	20m.+ yen	bit.ly/doc-tankan
CBI Industrial Trends Surv.	Bennett (1984), McIntosh et al. (1989), Thomas (1995), Lui et al. (2010), and Boneva et al. (2020)	man	500	nr	bit.ly/doc-cbi
Monthly Outlook Surv. in Industry	König et al. (1981), Nerlove (1983), and Andrade et al. (2022)	man, extr	1,600	20+ empl	bit.ly/doc-mos-ind
Surv. of Industrial Trends	Smith and McAleer (1995)	man	250	nr	bit.ly/doc-sit
Surv. of Production Forecasts	Morikawa (2019)	man			bit.ly/doc-spf
Surv. on Industrial and Service Firms	Guiso and Parigi (1999) and Ma et al. (2020)	man, con, serv	5,000	20+ empl	bit.ly/doc-sisf
ifo Investment Surv.	Bachmann et al. (2017)	man, trade	2,000	nr	bit.ly/doc-ifo
∞ Basic Surv. on Overseas Business Activities	Chen et al. (2021)	nr	8,700	mult.nat.	bit.ly/doc-bsoba
CFO Surv.	Gennaioli et al. (2015)	nr	1,000	nr	bit.ly/doc-cfos
Surv. on Inflation and Growth Expectations	Coibion et al. (2020)	ind, serv	1,000	50+ empl	bit.ly/doc-sige
Business Outlook Surv.	Chen et al. (2021)	nr	11,500	nr	bit.ly/doc-bos
Monitoraggio Economia e Territorio Surv.	Balduzzi et al. (2020)	man	25,000	nr	bit.ly/doc-met
Management and Organizational Practices Surv.	Bloom et al. (2020)	man	37,000	nr	bit.ly/doc-mops
Business Inflation Expectations Surv.	Meyer et al. (2021a)	nr	300	nr	bit.ly/doc-bies
Surv. of Business Uncertainty	Altig et al. (2022) and Barrero (2021)	nr	1,300	nr	bit.ly/doc-sbu
Bundesbank Online Panel - Firms	Balleer et al. (2020)	nr	10,000	nr	bit.ly/doc-bopf

Notes: *Resp. refers to current respondents per wave. The ifo Business Climate Survey was initially launched for the manufacturing sector. Similar surveys were later added for the construction, trade, services, and insurance sectors. Sector refers to sectoral coverage: not restricted (nr), manufacturing (man), extraction (extr), construction (con), non-financials private services (serv), industry (ind), and trade. Firm size gives restrictions on target firms: not restricted (nr), minimum number of employees (empl), mult.nat. (multinationals). The Tankan Survey targets firms with capital of at least 20 million Yen (Bank of Japan 2020).

Figure 1.2: BEP observations across both panel dimensions



Notes: observations of the ifo Business Expectations Panel (BEP) across time and firms. Left panel: number of actual (light blue) and target observations (dark blue). The number of actual observations is the number of firms that respond in a given month. Target observations equal the number of firms that are in the survey during a given month. Due to the harmonization of survey periods introduced by the European Union, no survey was conducted in December 2001. We set the value to missing in this plot.

1.2.2 Example: The ifo Business Expectations Panel

Below we survey the existing literature on firm expectations and, in doing so, we replicate the most important findings on the basis of a single data set. Because of its large coverage in terms of firms, firm-specific variables, and its time dimension, we choose the Business Expectations Panel of the LMU-ifo Economics and Business Data Center (BEP or ifo Survey from now on). It is based on the ifo Business Climate Survey, one of the oldest firm surveys in existence. Specifically, the BEP combines survey data from the Business Climate Survey and balance sheet data from the Amadeus and Hoppenstedt databases (EBDC-BEP 2019). Because the wording of the questions and possible answers differs somewhat across sectors, we focus on firms in the manufacturing sector for our analysis, the sector with the largest number of firms and the longest time dimension. Since the BEP combines annual balance-sheet data with the monthly survey data, we use the most recent balance-sheet data at a given point in time to avoid using information that is not yet available when firms report expectations. The BEP starts in January 1980; the last observation available to us is for June 2019. The survey questions (regarding prices, production, etc.) refer to a specific product.¹²

In the following, we produce a set of descriptive statistics for the BEP sample. Panel (a) of

¹²Some firms, hence, respond to several questionnaires each month. In our sample, however, this is the case for less than 10% of firms. In our analysis below, we refer to the individual observation as a “firm” in order to ease the exposition.

Figure 1.2 displays the actual number of responses per month (light blue line) and the target observations (dark blue line), i.e., the number of firms that are in principle in the survey during a given month but did not return the questionnaire, over time. The difference between the two is usually small, that is, the average monthly response rate of 85% is quite high.¹³ Furthermore, the median firm responds in 92% of the months they are in the panel. The ifo institute enlarged the panel significantly at various points in time, for example, after the German reunification in 1990. The right panel of Figure 1.2 shows the number of responses per firm. While there are many firms that participate only a few times in the survey, there is still a relatively high number of firms that answer the survey more than 100 and up to almost 500 times.

1.3 Stylized facts

The literature has established a number of facts about firm expectations—they emerge consistently across surveys and for both qualitative and quantitative measures. In this section, we offer a synthesis of these facts with a focus on firms’ expectations (and expectation errors) about their own production and prices. We consolidate five facts that we illustrate using one consistent, mostly qualitative data set: the ifo Business Expectations Panel (BEP), introduced in the previous section. Afterwards, we present a new, sixth fact that—to the best of our knowledge—has not been documented in the literature so far.

Given that we not only look at firm expectations but also at expectation errors, we first have to define expectation errors. There are different ways to do this for qualitative business surveys. However, Table A.1 and the discussion in Section A.1 show that these yield very similar outcomes for the ifo Survey. In what follows, we employ the widely-used definition of Bachmann et al. (2013). It is based on firms’ reported realized monthly changes $x_{t+j,1}^i$ of production or prices over a 3-month period, $x_{t,3}^i = \sum_{j=1}^h x_{t+j,1}^i$, and their 3-months ahead expectations, $x_{t,3|t}^i$.¹⁴ The expectation error is then defined as

$$e_{t,3}^i = \begin{cases} 0 & \text{if } \text{sgn}(x_{t,3}^i) = \text{sgn}(x_{t,3|t}^i) \\ \frac{1}{3}(x_{t,3}^i - x_{t,3|t}^i) & \text{else} \end{cases} \quad (1.1)$$

When the sign of the summed-up realizations is equal to the expectation, no error is assigned. In all other cases, the error is equal to the sum of the realizations minus the expectation, standardized by the forecasting horizon $h = 3$.

¹³Firms do not receive any compensation for participating in the survey, except the aggregate and sectoral results of the survey itself. Andrade et al. (2022) report a response rate of 60% for the quarterly INSEE survey. Banca d’Italia (2019) indicate a response rate of 40% to 50% for its Survey of Inflation and Growth Expectations, similar to the monthly response rate of 45% for the SBU (FRB Atlanta 2021). Note, however, that our reported response rate refers to firms which have already answered at least once. Out of all firms that were contacted in mid 2021 for the first time, around 2/3 returned at least two surveys. For the SBU, around 1/3 of firms responded at least once after the initial contact (FRB Atlanta 2021).

¹⁴See Table A.2 for the exact wording in the ifo Survey.

Table 1.2: Average unconditional expectation errors

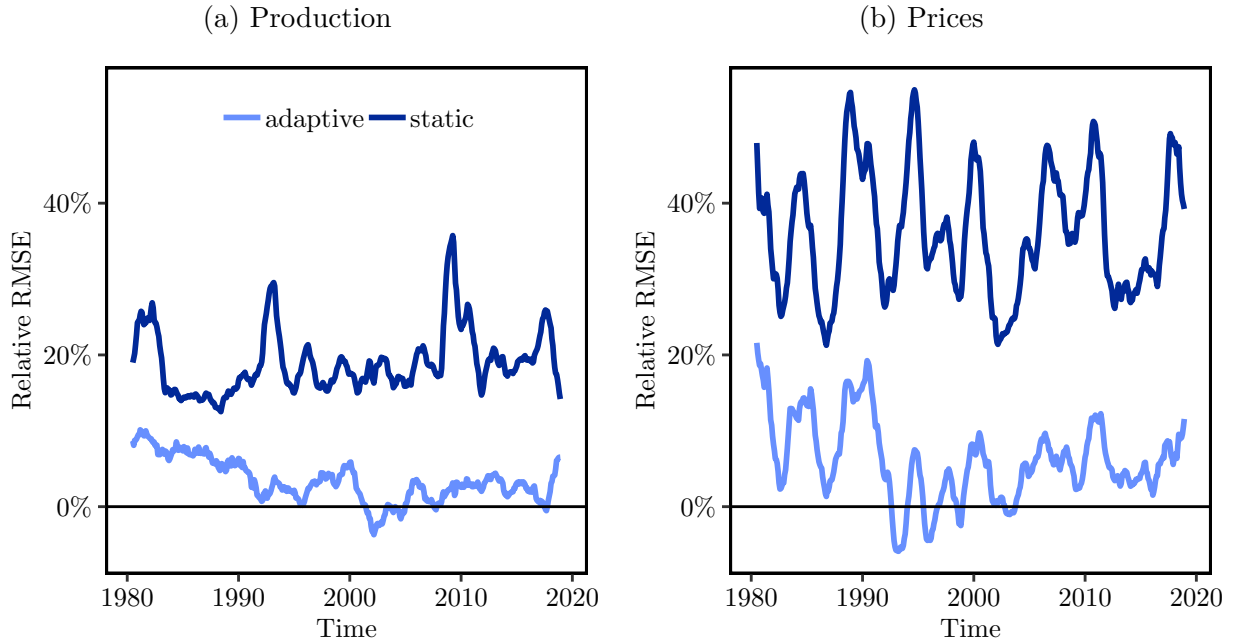
Grouped by	Group	Production			Prices		
		N	Median	% insig.	N	Median	% insig.
Overall		5122	-0.0183	77.59	5074	-0.0097	79.96
Number of Employees	Fewer than 50	801	-0.0128	76.40	779	-0.0056	81.51
	50-199	881	-0.0143	76.73	865	-0.0078	81.73
	200-499	410	-0.0097	81.22	410	-0.0048	84.88
	500-999	131	-0.0324	78.63	129	-0.0013	77.52
	More than 1000	95	-0.0041	77.89	93	-0.0051	75.27
Employees (Quartile)	First Quartile	566	-0.0115	77.56	548	-0.0048	81.02
	Second Quartile	588	-0.0172	76.19	578	-0.0085	82.87
	Third Quartile	582	-0.0154	77.15	569	-0.0076	81.20
	Fourth Quartile	582	-0.0097	79.38	581	-0.0039	81.76
Sales (Quartile)	First Quartile	566	-0.0191	74.56	546	-0.0046	82.97
	Second Quartile	576	-0.0147	77.08	557	-0.0071	81.33
	Third Quartile	562	-0.0169	80.25	564	-0.0058	82.27
	Fourth Quartile	571	-0.0159	78.98	574	-0.0063	79.27
Total Assets (Quartile)	First Quartile	672	-0.0159	75.60	652	-0.0070	82.82
	Second Quartile	673	-0.0113	77.86	655	-0.0065	81.07
	Third Quartile	666	-0.0193	78.53	668	-0.0079	83.98
	Fourth Quartile	676	-0.0153	79.29	677	-0.0056	79.03
Location	Eastern Germany	527	-0.0215	79.70	497	-0.0040	89.13
	Western Germany	1050	-0.0123	79.81	1052	-0.0041	82.60

Notes: firm-level average expectation errors (computed by regressing a firm’s expectation error on a constant); table entries provide the number of firms in each subgroup (N), the median of their average expectation errors (Median) and share of insignificant average expectation errors (% insig.), based on Newey-West standard errors. When grouping by location, we only consider firms that joined the ifo Survey after the German reunification.

Fact 1 - Unbiasedness. *Unconditionally, firms’ expectation errors are small and almost always insignificant.*

This fact emerges robustly from a number of studies. Evaluating a quantitative supplement to the ifo Business Climate Survey, Bachmann and Elstner (2015) find that more than two-thirds of firms in their sample of German manufacturing firms do not systematically over- or underpredict their production growth one quarter ahead. Using qualitative and quantitative questions from the same survey, Massenot and Pettinicchi (2018) also find that, on average, firms do not make unconditional expectations errors about their business situation. Altig et al. (2022) and Barrero (2021) again find little evidence of an unconditional bias in expected firm-level sales growth rates, using qualitative and quantitative data from the Survey of Business Uncertainty. Chen et al. (2020) document for a panel of Japanese firms small quantitative forecast errors on average. Andrade et al. (2022), in turn, show in a quantitative French firm survey that there is a strong positive relationship between firms’ anticipated and ex-post price changes. To illustrate Fact 1 further, Table 1.2 reports average expectation errors of individual firms for production, Panel (a), and prices, Panel (b), based on the

Figure 1.3: Performance of firm expectations relative to benchmark models



Notes: relative RMSE for production, Panel (a), and price expectations, Panel (b), both for adaptive (light blue line) and static expectations (dark blue line). Values above (below) zero mean that the respective benchmark model does not (does) beat the actual survey-based expectations. All series are plotted as moving averages over the previous and the next six months. All values expressed in percent.

BEP. For the full sample and across various classification schemes, we find robustly that the median forecast error is close to zero and the share of insignificant expectations errors is consistently above 75 percent. Table A.3 provides additional sectoral evidence in support of Fact 1.

Fact 2 - Information content. *Firm expectations outperform static and adaptive expectations.*

Firm expectations have significant information content because they help predicting future developments. To see this formally, we compute the root mean squared expectation error (RMSE), based on the actual expectations reported in the BEP, and compare it to two alternative models of expectation formation. The first assumes *adaptive* expectations: here, we simply carry forward as expectation the most recent realization (increase, no change, decrease) of either production or prices. The second model assumes *static* expectations: here we simply assume that no further change for either production or prices is expected. Figure 1.3 compares the RMSE of the benchmark models to reported production and price expectations. It shows that for almost all months, the benchmark models are less precise, that is, have larger RMSEs, than the reported expectations.

This observation is consistent with earlier work. Kawasaki and Zimmermann (1986) also find that ifo Survey-based qualitative price expectations beat adaptive expectations. Using the Confederation of Australian Industries (CAI)/Westpac Survey of Industrial Trends, Smith and McAleer (1995) also document the high information content of qualitative survey

expectations about firms’ output, prices, employment, stocks, and overtime relative to static expectations, and relative to a number of univariate/multivariate time-series models. Using quantitative survey questions, Chen et al. (2020) show for Japanese firms that a large majority of firms do not just use their realized sales to forecast the next period’s sales.

Fact 3 - Experience. *Larger and older firms are better at forecasting their own variables.*

While firm expectations generally reflect meaningful information (Facts 1 and 2), this is even more the case as firms get older and/or larger: experience, according to Fact 3, matters for the accuracy of firm expectations about their own variables. Massenet and Pettinicchi (2018), for instance, show, based on qualitative and quantitative questions in the ifo Survey, that older and larger firms make smaller expectation errors. Bachmann and Elstner (2015) for German firms in the ifo Business Climate Survey and Morikawa (2019) for Japanese firms in the Survey of Production Forecast document that larger firms make smaller quantitative expectation errors, presumably because they are able to spend more resources on forecasting than smaller firms. Experience also matters: Triebs and Tumlinson (2013) find that firms located in eastern Germany did worse, relative to their western peers, in predicting business conditions early after German reunification, but improved their forecasting performance over time. Similarly, Chen et al. (2020) show for a panel of Japanese firms that forecast precision increases with age. Related, there is also evidence that better-managed firms make smaller forecasting errors (Bloom et al. 2021).

We complement the existing work with new evidence based on the BEP and present it in Table 1.3. Panel (a) shows that mean squared expectation errors (MSEs) tend to be smaller for older firms and consistently so across decades. One exception are the 2000s: here older firms did worse. This result may be caused by the global financial crisis and deserves some future research. Panel (b) of Table 1.3 reports firm-level mean and median SEs for different firm sizes. In line with the literature, we observe that larger firms tend to make smaller MSEs.

Fact 4 - Predictability. *Firms make predictable expectation errors.*

Under rational expectations (RE), expectation errors should not be predictable on the basis of information that is available at the time when expectations are formed. The RE hypothesis can be framed in a regression setup as

$$e_{t,h}^i = x_t^i \beta + v_t^i, \quad (1.2)$$

where the forecast error $e_{t,h}^i$, at horizon $h = 3$ in our case, is the dependent variable and x_t^i contains candidate predictors. The β -coefficients should not be different from zero under the null of RE.¹⁵ We estimate the equation using the observations for the BEP and report results in Table A.8. While macroeconomic variables turn out to be mostly insignificant as predictors, many firm-specific variables—such as the order backlog, changes in demand, or past expectations—help in predicting expectation errors for production and prices. Overall, about 17 percent of the variance in expectations errors can be explained in our regressions.

¹⁵An alternative test for rationality is based on the regression $x_{t,h}^i = \beta_0 + \beta_1 x_{t,h|t}^i + v_t^i$, where $\beta_0 = 0$ and $\beta_1 = 1$ under the null of RE. This test is discussed in Clements et al. (2023).

Table 1.3: Experience and expectation errors

(a) Experience by age

Decade	Production				Prices			
	MSE _{old}	MSE _{young}	Difference	p-value	MSE _{old}	MSE _{young}	Difference	p-value
1980-89	0.1058	0.1121	-0.0064	0.00	0.0447	0.0498	-0.0051	0.00
1990-99	0.1185	0.1343	-0.0158	0.00	0.0533	0.0556	-0.0022	0.01
2000-09	0.1415	0.1405	0.0010	0.53	0.0674	0.0637	0.0037	0.00
2010-19	0.1303	0.1414	-0.0110	0.00	0.0607	0.0658	-0.0051	0.01

(b) Experience by size

Grouped by	Group	Production			Prices		
		N	Mean	Median	N	Mean	Median
Overall		5122	0.1278	0.1170	5074	0.0594	0.0372
Number of Employees	Fewer than 50	801	0.1319	0.1197	779	0.0617	0.0363
	50-199	881	0.1299	0.1217	865	0.0615	0.0386
	200-499	410	0.1233	0.1184	410	0.0556	0.0358
	500-999	131	0.1209	0.1052	129	0.0500	0.0372
	More than 1000	95	0.1088	0.0988	93	0.0615	0.0422
Employees (Quartile)	First Quartile	566	0.1312	0.1165	548	0.0622	0.0370
	Second Quartile	588	0.1323	0.1262	578	0.0579	0.0359
	Third Quartile	582	0.1302	0.1216	569	0.0645	0.0406
	Fourth Quartile	582	0.1187	0.1078	581	0.0549	0.0363
Sales (Quartile)	First Quartile	566	0.1348	0.1220	546	0.0587	0.0360
	Second Quartile	576	0.1326	0.1248	557	0.0655	0.0391
	Third Quartile	562	0.1240	0.1147	564	0.0558	0.0375
	Fourth Quartile	571	0.1199	0.1074	574	0.0615	0.0355
Total Assets (Quartile)	First Quartile	672	0.1310	0.1197	652	0.0611	0.0375
	Second Quartile	673	0.1326	0.1209	655	0.0624	0.0375
	Third Quartile	666	0.1284	0.1187	668	0.0589	0.0370
	Fourth Quartile	676	0.1188	0.1082	677	0.0586	0.0361

Notes: Panel (a) shows the difference of mean squared expectation errors (MSE) between young and old firms. At the time of being surveyed, a firm is considered young when it was founded at most 10 years ago. For each decade, we pool observations by age and estimate the difference in the MSE between old and young firms. Panel (b) shows the firm-level mean and median squared expectation errors; table entries provide summary statistics for different firm sizes. We measure size in terms of the absolute number of employees, as well as firms' location in the distributions of employees, sales, and total assets. N denotes the number of firms in each group.

Consistent with our results, Massenot and Pettinicchi (2018) find that firms extrapolate from past experience too much and end up making predictable expectation errors. Similarly, Barrero (2021), using distributional questions from the Survey of Business Uncertainty (SBU), documents that firm managers over-extrapolate: their forecasts are too optimistic after positive shocks and too pessimistic after negative shocks. Ma et al. (2020) analyze expectation errors of Italian firms about their sales and detect significant auto-correlation. Boneva et al. (2020) show that UK firms tend to have rational expectations of quantity variables, such as their own employment and new orders, but deviate from rational expectations when it comes to prices, wages, and unit costs. Hence, Fact 4.

At first sight, this fact is hard to reconcile with Fact 1. Note, however, that while Fact 1 is about the unconditional accuracy of expectations, Fact 4 shows that forecast errors are predictable conditional on specific information. As such, the two facts are not contradictory but raise challenges that need to be addressed in future research. At an empirical level, a more systematic investigation into the two facts seems warranted. At a conceptual level, one may explore models of learning and/or limited attention which can rationalize the patterns in the data.

Fact 5 - Countercyclical second moments. *The dispersion and volatility of expectations and expectation errors are countercyclical.*

This fact has been observed for a variety of survey-based measures (e.g., Bachmann et al. 2013, 2017; Bachmann et al. 2019; Enders et al. 2019; Morikawa 2016, 2019), based both on qualitative and quantitative survey questions. As before we corroborate these findings. While Panel (a) of Table 1.4 lists dispersion and volatility measures, Panel (b) reports their time-series properties based on BEP data. The first subpanel shows correlation coefficients between the measures for production (left) and prices (right). The correlation is generally quite high, in particular for the error-based measures.

The countercyclicity of the dispersion and volatility measures can be read off the second subpanel where we report correlation coefficients vis-à-vis monthly measures of economic activity: the growth rates of industrial production, hours worked, and employment. Across the board, the signs of the correlation coefficients are negative and mostly significantly so. We also regress the measures on recession dummies—as dated by the German Council of Economic Experts—and again find a significant increase in dispersion and volatility in economic downturns. Especially so in the Great Recession of 2008/09, where our measures increase by between 8.3 and 25 percent.

Fact 6 - Stickiness. *Firm expectations are updated infrequently; updates for production and prices often happen at the same time and in the same direction.*

This fact has not been documented in the literature. This is surprising in light of influential work which models firms' sticky information, that is, infrequent updating as key friction for business cycle dynamics (Mankiw and Reis 2002). As a first pass towards assessing the stickiness of expectations in the BEP, we compute mean and median spells of expectations, that is, the number of consecutive months for which expectations remain unchanged. Panel (a) of Table 1.5 shows results, both for production (left) and prices (right). For the whole sample, expectations are quite sticky: we observe, for instance, that production expectations

Table 1.4: Dispersion and volatility measures

(a) Definitions

Domain	Measure	Definition
firm & time	Absolute forecast error	$absfe_{i,t} = abs(e_{t,h}^i)$
	Rolling window standard deviation	$stdef_{i,t} = \sqrt{\frac{1}{3} \sum_{k \in \{-3,0,3\}} (e_{t+k,h}^i - \bar{e}_{t,h}^i)^2}$
time	Forecast dispersion	$fdisp_t = \sqrt{\text{frac}_t^+ + \text{frac}_t^- - (\text{frac}_t^+ - \text{frac}_t^-)^2}$
	Forecast error dispersion	$fedisp_t = \sqrt{Var(e_{t,h,i t})}$
	Mean absolute forecast error	$mae_t = \frac{1}{n_t} \sum_i absfe_{i,t}$
	Avg. rolling window standard deviation	$stdfe_t = \frac{1}{n_t} \sum_i stdef_{i,t}$

(b) Business cycle properties

Variable	Production				Prices			
	fdisp	fedisp	mae	stdfe	fdisp	fedisp	mae	stdfe
Correlation within measures								
fdisp	1.00	0.69***	0.56***	0.58***	1.00	0.40***	0.60***	0.46***
fedisp		1.00	0.93***	0.73***		1.00	0.94***	0.88***
mae			1.00	0.82***			1.00	0.87***
stdfe				1.00				1.00
Correlation with aggregates								
$\Delta \log$ Production	-0.12***	-0.04	-0.12***	-0.15***	0.06	-0.07	-0.07	-0.03
$\Delta \log$ Hours	-0.02	-0.08*	-0.18***	-0.14***	-0.01	-0.03	-0.04	-0.03
$\Delta \log$ Employment	-0.20***	-0.30***	-0.44***	-0.44***	-0.04	-0.21***	-0.22***	-0.20***
Recession Dummies								
Recession	0.019***	0.016**	0.043***	0.029***	0.024*	0.061***	0.114***	0.094***
Recession 2008/09	0.083***	0.084***	0.128***	0.140***	0.088***	0.154***	0.246***	0.243***

Notes: Panel (a): $e_{t,h}^i$ is the forecast error of Bachmann et al. (2013) defined in equation 1.1 and $\bar{e}_{t,h}^i$ is the average forecast error of the current value, its third lag, and its third lead. $\text{frac}_t^+ = \sum_i \mathbf{1}(x_{t,h|t}^i = +1)/n_t$ and $\text{frac}_t^- = \sum_i \mathbf{1}(x_{t,h|t}^i = -1)/n_t$ are the shares of expected increases and decreases at time t . $fdisp_t$, $fedisp_t$, and mae_t based on Bachmann et al. (2013); $stdfe_t$ on Bachmann et al. (2019). Panel (b) shows Spearman rank correlation among dispersion measures first, Spearman rank correlation with aggregate business cycle measures second, and regression results using recession dummies third. After standardizing each time series by its non-recession mean, we report coefficients for a general recession dummy and a dummy for the 2008/09 recession. One, two, and three stars (*) correspond to significance at the 10, 5, and 1 percent significance levels.

are not adjusted for more than 3 months on average. The panel also offers a breakdown into the stickiness of the three different response categories. Here, we observe the largest degree of stickiness for the “no change” category. Overall, price expectations tend to be more sticky than production expectations. Panel (b) of Table 1.5 shows that firms in the BEP tend to update expectations across variables at the same time. Specifically, observing an update in price expectations increases the probability of observing an update (upwards or downwards) in production expectations by 10 percentage points or 39 percent. A production expectation update increases the probability of observing a price expectation update by 9 percentage points or 46 percent. This is consistent with the findings for firms’ macroeconomic expectations discussed in Baumeister (2023). Calibrating sticky information models to capture the evidence put forward in Table 1.5 seems a promising venue for future research. Moreover, Panel (c) of Table 1.5 shows, that for the majority of cases, price and production expectations change in the same direction. In particular, if we observe a change in either production or price expectations, we find that the other variable is updated in the same direction at least twice as often as in the opposite direction. This pattern in the data suggests an important role for demand shocks for firm expectations and calls for further investigation.

1.4 Expectation formation

In this section, we turn to the expectation formation process of firms with a focus on recent survey evidence. This evidence often points to departures from the full information rational expectations (FIRE) benchmark. For instance, Fact 4 shows that firms make predictable forecast errors. At this point, however, there is no consensus about an alternative to FIRE. At a very basic level, there is a long tradition of noisy information models. Here, information processing is rational but information is incomplete. In the classic contributions by Lucas (1973), Woodford (2002), Sims (2003), or Maćkowiak and Wiederholt (2009), economic actors—and notably firms—process information and update expectations in a rational way. This goes some way to account for the evidence presented above. Likewise, more recent contributions emphasize that a (rational) focus on certain sectors/media distorts the information formation process (Chahrour et al. 2021; Kohlhas and Walther 2021). Other models, by contrast, allow for behavioral aspects in the expectation formation process (for instance, Shiller 2017; Bordalo et al. 2019), where, under certain conditions, behavioral models and incomplete information models give rise to equivalent equilibrium effects (Angeletos and Huo 2021).

In what follows, we seek to inform this discussion by first surveying the evidence on the determinants of expectations. In the second part of this section, we zoom in on the expectation formation process as we discuss recent evidence regarding the response of firms to news, both at the firm level and the aggregate level. As in the previous section, we revisit key findings on the basis of the BEP.

1.4.1 Determinants of expectations

We aim to provide a simple empirical characterization of the determinants of firm expectations. We first focus on the mean forecast (first moment). Afterwards, we also consider briefly the determinants of firm uncertainty (second moment).

Table 1.5: Stickiness of firm expectations

(a) Spell lengths

Spell type	Production			Prices		
	Share in %	Mean	Median	Share in %	Mean	Median
overall		3.38	2		4.85	2
decrease	24.73	2.17	1	18.25	2.21	1
no change	48.36	4.67	2	51.00	7.23	4
increase	26.91	2.15	1	30.74	2.45	2

(b) Conditional updating frequencies

Frequencies	Production		Prices	
	Updating freq. conditional on	Value	Updating freq. conditional on	Value
	Update in price exp.: yes	36.58%	Update in prod. exp.: yes	24.74%
	Update in price exp.: no	26.32%	Update in prod. exp.: no	16.91%
Difference				
	in percentage points	10.26pp		8.83pp
	in percent	38.98%		46.30%

(c) Conditional distribution of expectation updates

P(Y=y X=x)	Production			Prices		
	Y = Prod. update X= Price update			Y = Price update X= Prod. updates		
y=	downwards	no update	upwards	downwards	no update	upwards
x= downwards	25.63	63.64	10.73	17.17	75.47	7.37
no update	13.35	73.68	12.97	8.51	83.09	8.40
upwards	11.05	63.19	25.76	7.36	75.05	17.58

Notes: Panel (a) shows summary statistics for spell length of qualitative expectations for prices and production. Given qualitative expectations (increase, no change, decrease) we calculate the lengths of sequences with identical expectations (spells). We compute their average and median length in months both across spell types (overall) and for each spell type separately. Panel (b) shows relative frequencies of expectation updates (changes in the reported qualitative expectations) for production (prices) conditional on whether a firm reported an update for price (production) expectations. Observations are pooled across time and firms. Panel (c) shows the distribution of expectation updates for production conditional on price-expectation updates (left) and vice versa (right). Entries in the table are conditional probabilities of observing an update, as in the column labels, conditional on observing an update of the other variable, as in the row labels. Each row for production and prices sums to 100. Computation based on full ifo sample (manufacturing, 2002–2019).

Firm expectations

In terms of expectations, we focus, as before, on firm expectations about production and prices. To set the stage, we perform an analysis based on the ifo Survey which builds on earlier work by Enders (2020). Because firm answers regarding production and price expectations are qualitative in the ifo Survey, we estimate an ordered probit model. Specifically, using $j = \{-1, 0, 1\}$ to index the reported expectations $x_{t,h|t}^i$ about firms' prices or production, we estimate

$$\begin{aligned} Pr(x_{t,h|t}^i = j) &= Pr(a_{j-1} < x_{t,h|t}^{i*} \leq a_j) \\ &= \Phi(\alpha_j - X'_{it}\beta) - \Phi(\alpha_{j-1} - X'_{it}\beta) , \end{aligned} \tag{1.3}$$

where X_{it} contains the variables which may influence firm expectations, $x_{t,h|t}^{i*}$ is the latent variable, and α_{j-1} and α_j are threshold parameters. Since the set of potential variables is large, we consider different groups of variables and summarize their impact by focusing on the model fit, namely on the pseudo R^2 as defined by McFadden (1974).¹⁶ In terms of explanatory variables X_{it} , we distinguish three sets of variables. The first set contains variables that describe a firm's own condition as reported in the survey, such as, for instance, the current state of business, orders, and capacity utilization. In addition, it includes lags of expected production and prices. It also contains interaction terms that we include on the basis of a log-likelihood test. The second set consists of firm fundamentals as reported in the most recent balance sheet, such as, for instance, the debt share. Here our selection of variables follows Enders (2020). A third set of variables contains macro variables as observable by firms in real time, notably the unemployment rate in the previous month as well as industrial production. Table A.5 provides a full list of variables for each of the three sets. In addition, we always include sector fixed effects and the average reported state of business, both on a two-digit level.

We estimate model (1.3) using all combinations of the three sets of variables and show results in Table 1.6. Results are clear cut. The survey responses account for a fairly large share of the variation in firm expectations, with a pseudo R^2 of 25 and 32 percent for production and prices, respectively. The contributions of balance-sheet fundamentals and macro variables, on the other hand, appear negligible. We should stress, however, that balance sheet data ("fundamentals") is available only at an annual frequency and may therefore not matter much for changes in the short-term outlook of firms over the next three months. In addition to using the R^2 to judge the contribution of each group of variables, we also checked by how much the share of correctly predicted expectations increases when we include each group one-by-one. We find that the first set of variables helps to increase the performance of the model most strongly also in this case.

The result that firm-specific information, as reflected in survey responses, is a key determinant of firm expectations echoes early work based on the ifo Survey in the 1950s. Pioneering work by Anderson et al. (1956a), Anderson et al. (1956b), and somewhat later by Anderson and Strigel (1960) showed that unexpected changes in demand lead to changes in firms' production and pricing plans. This early work already established that production

¹⁶Formally, we consider: $R_{mf}^2 = 1 - \ln L_M / \ln L_0$, where R_{mf}^2 is the pseudo R^2 , L_M is the likelihood of the model and L_0 is the likelihood of a constant-only model.

Table 1.6: Determinants of production and price expectations

Variables	Production		Prices	
	Observations	Pseudo- R^2	Observations	Pseudo- R^2
Survey	181,329	0.2523	181,276	0.3204
Fundamentals	271,498	0.0002	277,890	0.0001
Macro	337,028	0.0057	345,828	0.0074
Survey + Fundamentals	180,686	0.2524	180,633	0.3204
Survey + Macro	172,428	0.2524	172,374	0.3244
Fundamentals + Macro	254,624	0.0064	260,988	0.0075
Survey + Fundamentals + Macro	172,327	0.2525	171,731	0.3244

Notes: summary statistics for ordered probit models using expectations about a firm’s own production and price as dependent variables. Explanatory variables are combinations of variables from the survey (business situation, orders, etc. with up to three lags and interaction terms), firm fundamentals from their balance sheet (debt share, financing coefficient), and macro variables (monthly growth rates of PPI, CPI, and IP and the unemployment rate, each with their publication lag). See Table A.5 for more details on the variables.

plans are more responsive to surprise demand changes than price plans. For the latter, cost changes are important. More recently, Carlsson and Skans (2012) document an influence of both current and expected future marginal cost on firms’ price-setting behavior, while Meyer et al. (2021a) find that firms’ year-ahead unit-cost expectations covary strongly with year-ahead price expectations.¹⁷ Massenot and Pettinicchi (2018), in turn, find for the ifo Survey that business expectations are responsive to past business developments. Similarly, Boneva et al. (2020) show for UK firms that past orders are important when it comes to accounting for price and wage expectations. Financial factors, too, matter for expectations: Balduzzi et al. (2020) study Italian firms during the Corona crisis and find that financially constrained firms expect to charge higher prices relative to their unconstrained counterparts.

Our results above suggest that firm-specific developments are considerably more important than macroeconomic developments when it comes to accounting for firm expectations. But there is also evidence that firm expectations are responsive to macroeconomic developments. Enders et al. (2019), for instance, find that firm expectations respond to monetary policy shocks. Similarly, Eminidou and Zachariadis (2022) document the effects of monetary policy shocks on firm expectations for a panel of euro area countries. For this purpose, they rely on the Joint Harmonised EU Programme of Business and Consumer Surveys (BCS). Strasser (2013) uses the ifo Survey and investigates to what extent firms’ export expectations respond to exchange-rate movements.

Several studies use survey data to explore the impact of the Covid-19 pandemic on firm expectations. Meyer et al. (2021b) rely on the Business Inflation Expectations Survey run by the Federal Reserve Bank of Atlanta. Balleer et al. (2020) and Bundesbank (2021) look at German firms, using ifo data and the Bundesbank Online Panel - Firms, respectively. These studies find consistently that firms’ price expectations have decreased in the early phase of the pandemic. In addition, there is evidence that lockdown measures matter for

¹⁷The former use Swedish firm-level data and the latter the Atlanta Fed’s Business Inflation Expectations Survey. Meyer et al. (2021a) also demonstrate that information treatments about aggregate inflation and policymakers’ forecasts have a negligible effect on firms’ unit-cost expectations.

firm expectations. Buchheim et al. (2022), using ifo data for Germany, show that the announcement of nationwide school closures on March 13, 2020 in response to the first wave of Corona infections was followed by the largest change in business perceptions by far.

Finally, there is evidence that the developments of the sectors or regions in which firms operate influence their expectations. Andrade et al. (2022) stress the importance of industry-level shocks, as distinct from aggregate and firm-specific shocks, for both firm actions and expectations. Their analysis is based on a survey of French firms. Kukuvec and Oberhofer (2020) use input-output tables and establish on the basis of the BCS that firms' business expectations are also influenced by expectations of other firms, in particular of those located upstream. Dovern et al. (2020) find for the ifo Survey that firms extrapolate from local economic conditions to aggregate growth expectations.

Firm uncertainty

So far, we have focused on the determinants of the first moment of firm expectations, that is, the mean forecast. But firm surveys also shed light on the determinants of the second moment of firm expectations, that is, into firm-level uncertainty. Altig et al. (2022) survey business executives about firm outcomes with a particular focus on business uncertainty. They find, among other things, that subjective uncertainty is higher when firms' have grown faster and when they have revised their growth expectations. Similarly, Bachmann et al. (2021), using data for German firms, show that firms' subjective uncertainty of future sales growth increases in the aftermath of unusual, in particular negative, growth experiences. In the cross-section of firms, large and fast-growing firms display, for a given shock volatility, lower subjective uncertainty than unsuccessful ones.

Dovern et al. (2020) document a negative relationship between firms' uncertainty about their own business outlook and expectations about GDP growth. There is also survey evidence that specific events raise uncertainty at the firm level, notably in the context of Brexit and Covid-19 (Bloom et al. 2019; Altig et al. 2020). Finally, we note that measuring firm uncertainty remains challenging from a methodological point of view. Bachmann et al. (2020), for instance, find that a majority of firms use an interval of probabilities instead of a single number at least once in their sample period. The authors interpret this behavior as reflecting Knightian uncertainty.

1.4.2 Over- and underreaction to news

How do firms form expectations? In an influential study, Coibion and Gorodnichenko (2015) propose a simple diagnostic in order to shed light on the expectation-formation process. Specifically, using the Survey of Professional Forecasters (SPF), they regress the upcoming forecast error on the current forecast revision. It turns out that forecast revisions predict forecast errors in the same direction. An upward revision, say, is followed by an underprediction of the same variable—forecasters seem to underreact to news, as reflected in the revision. This finding is in line with rational expectations models featuring noisy information. Yet, it has given rise to an intensive debate about the expectation-formation process and motivated new explorations, both empirically and in terms of theory.

In their original contribution, Coibion and Gorodnichenko (2015) study the response of the average forecast error in the SPF to the average forecast revision in the SPF. Against this background, Bordalo et al. (2020) stress that results change—from underreaction to overreaction—once one studies the relationship between forecast errors and forecast revisions *at the level of individual forecasters*. Other work, some of which we discuss below, establishes that whether there is over- or underreaction depends on the nature of the news which forecasters receive. Most of the evidence to date, however, is based on the SPF.

In what follows, we broaden the discussion and follow Born et al. (2022) in turning to firms’ forecasts and their expectation formation process. We estimate a simplified version of their empirical model on our BEP sample:¹⁸

$$e_{t,h}^i = \beta_0^i + \beta_1^i FR_{t,h}^i + v_{t+h}^i, \quad (1.4)$$

where index i denotes a specific firm, $e_{t,h}^i$ is the forecast error (as defined in equation (1.1)), $FR_{t,h}^i$ is the forecast revision defined as $\text{sgn}(x_{t+h|t}^i - x_{t-1+h|t-1}^i) \in \{+1, 0, -1\}$, and v_{t+h}^i is a zero-mean error. A positive β_1^i -coefficient implies underreaction to the news that is reflected in the forecast revision. We estimate this equation separately for each firm, for both price and production expectations.¹⁹

Figure 1.4 shows the distribution of the estimates for β_1^i across firms for production and price expectations. The mass of firms is characterized by negative betas, of which 32 percent are significant for production and 41 percent for prices. The overall mean estimate for production is -0.112 and -0.107 for prices. The overall result is in line with Born et al. (2022) and clear cut: firms tend to overreact to news.²⁰ This is particularly noteworthy because, in our analysis, news and forecast errors pertain to firms’ expectations about their own production and prices rather than the aggregate economy and rational expectations models with noisy information have a hard time rationalizing overreactions. A number of behavioral models have been put forward to account for overreaction in other contexts. Azeredo da Silveira and Woodford (2019), for instance, show that if memory is noisy, current realizations are extrapolated into the future disproportionately. Bordalo et al. (2020), instead, rely on diagnostic expectations to rationalize overreaction. Here, forecasters overweigh the probability of certain states in the light of recent signals.

Table A.6 shows that the coefficients are robustly below zero across different measures of firm size and location. The same holds if we consider distinct sectors. We conclude that overreaction of firm expectations to news is a robust and pervasive feature of the data, not driven by a particular group of firms.

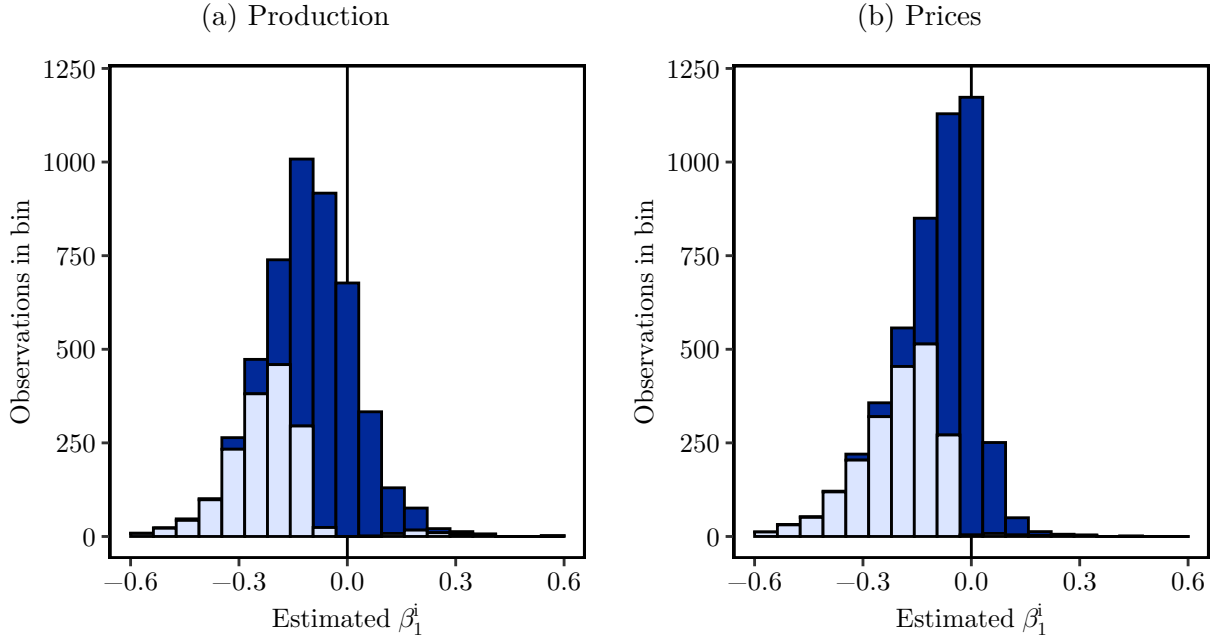
Born et al. (2022) also estimate equation (1.4) on pooled data while allowing for firm and time-fixed effects. For this specification, the estimate of β_1^i is significantly negative as well. They further distinguish the response to “macro news” (measured by unexpected changes in

¹⁸In the context of the qualitative ifo Survey data, there are a number of noteworthy conceptual issues and limitations that are discussed in Born et al. (2022).

¹⁹For a firm to be considered in the estimation we require it to provide us with at least 30 observations and a non-zero variance of forecast errors and forecast revisions, that is, a firm must have revised its expectation at least once.

²⁰Figure A.1 shows that estimates for the intercept in equation (1.4) are generally well-behaved in the sense that they are scattered evenly around zero. Moreover, there is no systematic pattern which would suggest a specific relationship between the estimate for the slope and the intercept.

Figure 1.4: Response of forecast error to forecast revision



Notes: histograms of estimated β_1^i -coefficients in firm-level regressions for production and price expectations, see equation (1.4); sample restricted to firms that initially report no expected change. Coefficients outside of the 1 and 99 percent quantiles (pooled over all subfigures) are dropped. Dark blue is for estimates that are insignificant at the 5%-level, and light blue is for significant estimates.

the aggregate ifo index or manufacturing orders) from the response to firm-specific micro news (as reflected in the revision of a firms' own production expectation net of time-fixed effects) and still find that firms overreact to micro news, but also that they underreact to macro news.²¹

Born et al. (2022) rationalize their findings in a general equilibrium model that allows for noisy information and salience effects. The key feature of their model is that firms' own productivity is salient of aggregate technology to them—a phenomenon which gives rise to a 'false consensus' bias. In line with additional model predictions, firms with a larger 'salience bias' empirically display larger production and forecast-error volatility, as well as lower profits. These systematic differences demonstrate that the measured bias is not the result of random forecast fluctuations. Broer and Kohlhas (2023) put forward a related mechanism. They stress that what they call 'overrevision' of individual forecasts may mask both over- and underreactions to salient public signals, as documented for inflation expectations in the SPF.²²

²¹Similarly, Kučinskas and Peters (2022) document for professional forecasters that their inflation forecasts underreact to aggregate shocks but overreact to idiosyncratic shocks. Using the ifo Survey, Massenot and Pettinicchi (2018) regress, in turn, expectations and forecast errors on past changes of the business situation (rather than on forecast revisions). They find that the regression coefficient is positive and significant, and robustly so, across a number of specifications. They refer to this result as "over-extrapolation".

²²They extend a model of noisy rational expectations by allowing forecasters to be overconfident about the precision of their own information. In this account, absolute overconfidence (perceiving own information

In sum, recent survey evidence shows that firm expectations are responsive to information. Firm-specific information turns out to be more important and impacts expectations more strongly than information about the aggregate economy. This finding emerges from a number of recent contributions and is confirmed once we estimate models (1.3) and (1.4) on our BEP sample. When it comes to the details of the expectation-formation process, the recent literature has put forward a number of promising alternatives to the FIRE benchmark. They go some way to account for the evidence. But further work is required for the profession to be able to settle on a new consensus model.

1.5 Firm expectations and firm decisions

One reason why we care about firm expectations is that they matter for firm decisions—at least according to theory. For the longest time, the link from economic expectations to actions has been taken for granted. At an empirical level, models featuring a key role for expectations that lay the foundation for, e.g., the New Keynesian Phillips curve, have been shown to describe the data reasonably well (e.g., Galí and Gertler 1999). There are also numerous purely empirical studies which suggest that, in general, expectations of economic agents are key for the business cycle (see, for instance, Beaudry and Portier 2006; Born et al. 2019; Enders et al. 2021). These studies, however, do not directly rely on expectations data at the firm level. Only recently has the literature started to explore these data to study the effect of firm decisions on firm actions.

1.5.1 The effect of firm expectations

We revisit some of this work in what follows, with a particular focus on Enders (2020) since their analysis is also based on the BEP. The basic idea of the study is to compare the behavior of firms that report that they expect either an increase or a decrease in production to otherwise very similar firms that expect production to remain unchanged. Because the responses regarding expected production are qualitative, one may think of expectations as a kind of “treatment”: firms may either expect an increase, no change, or a decrease. Of course, expectations are not literally assigned in a random way. By comparing firms that display the same fundamentals but different expectations, however, the assignment can be interpreted as random.

In terms of identification, two features of the ifo Survey are crucial. First, the survey features a fairly large set of control variables, including balance-sheet data and received orders of firms. One may thus approximate the set of fundamentals which matter for firm decisions fairly accurately. Second, the timing of survey responses is key: because the large majority of responses to the survey is filed early in the month, they represent expectations about future periods (namely, for the three months following the current one) at a time when production plans for the current month may be formed but actual demand has not yet been

as more informative than it actually is) makes forecasters overreact to private information while relative overconfidence (perceiving own information as more informative than information of others) makes forecasters underreact to public signals which, in turn, are understood to reflect the response of others to their own forecasts.

observed.²³ Enders (2020) investigate how production expectations impact both production and pricing decisions in the current month. In what follows, we modify the original analysis in three ways. First, for the matching exercise, we use data from 1991–2019, that is, three more years of data. Second, to control for fundamentals we compute the propensity score, that is the likelihood, of a treatment for a given firm-month observation on the basis of model (1.3). In this way, we directly build on the estimates reported in Section 1.4, which allows for macroeconomic control variables, rather than for time-fixed effects as in Enders et al. We use the propensity score to match treated and untreated observations and, eventually, to compute the average treatment effect on the treated (ATT), both for production and pricing decisions. Third, we also report results for various subsets of firms.

Table 1.7 reports the results, separately for firms which report an “increase” and a “decrease” of production expectations. The top row shows the results for the full sample. We observe that expectations of a production increase impact current production and prices positively. Quantitatively our results are very similar to those reported by Enders (2020).²⁴ The effect of an expected production decrease on production and prices is negative and quantitatively comparable to that of an expected production increase. Table 1.7 also reports results for a detailed break-down for different subsets of firms that turn out to be quite similar.

Importantly, expectations may impact current decisions for two reasons. First, expectations may reflect *news* that are not yet incorporated into current fundamentals. According to this interpretation, firm expectations operate as a transmission channel through which future fundamentals impact current decisions. Second, expectations might be fundamentally unwarranted and as such are genuine *noise*. Enders et al. assess the distinct role of news and noise for firm decisions on the basis of forecast errors. Specifically, taking an ex-post perspective, they ask whether firms that expect a change in production behave differently vis-à-vis firms which correctly expect production to remain unchanged, once for firms whose expectations turn out to be correct and once for firms with, in hindsight, incorrect expectations. They find that the treatment effect is present for both correct and incorrect expectations. This finding suggests that expectations impact current firm decisions for both fundamental (news) and non-fundamental reasons (noise).

Other work has also looked into how firm expectations shape firm behavior based on survey evidence. Boneva et al. (2020) study a survey of UK firms and estimate Phillips-curve relationships to capture the effect of firm expectations on firm decisions. Similar to the findings above, they also find an effect on firms’ pricing decisions. Other papers have established a link between firm expectations and firms’ investment decisions. Bachmann and Zorn (2020) do so on the basis of the ifo Investment Survey. Gennaioli et al. (2015), instead, rely on the Duke University Quarterly Survey of Chief Financial Officers. They stress, in particular, that while CFOs’ expectations matter for investment decisions, these expectations cannot be easily accounted for by conventional variables. Ma et al. (2020) establish a relation

²³About 50% of firms answer within the first eight days and another 25% answer in the following week. These figures are calculated for those firms that answer the survey electronically, which is the majority by now.

²⁴This positive effect may reflect a stronger tendency among treated firms to raise production and prices or a reduced tendency to lower production and prices, or both. As they disentangle the two effects, Enders (2020) find that the overall effect is dominated by the increased tendency to raise production and prices.

Table 1.7: Effects of increased and decreased production expectations

Grouped by	Group	Production		Prices	
		increase	decrease	increase	decrease
Full sample		0.152***	-0.193***	0.012***	-0.034***
Number of Employees	Fewer than 50	0.140***	-0.175***	0.025***	-0.040***
	50-199	0.154***	-0.207***	0.003	-0.029***
	200-499	0.183***	-0.149***	0.026**	-0.052***
	500-999	0.186***	-0.245***	-0.009	-0.048*
	More than 1000	0.150***	-0.242***	0.092***	0.006
Employees	First Quartile	0.162***	-0.160***	0.033***	-0.051***
	Second Quartile	0.143***	-0.179***	0.015	-0.017
	Third Quartile	0.140***	-0.229***	-0.005	-0.044***
	Fourth Quartile	0.177***	-0.176***	0.030***	-0.035***
Sales	First Quartile	0.159***	-0.159***	0.038***	-0.028**
	Second Quartile	0.128***	-0.191***	0.005	-0.038***
	Third Quartile	0.139***	-0.163***	-0.002	-0.053***
	Fourth Quartile	0.163***	-0.217***	0.013*	-0.024***
Total Assets	First Quartile	0.153***	-0.151***	0.034***	-0.034***
	Second Quartile	0.132***	-0.225***	0.007	-0.028***
	Third Quartile	0.160***	-0.169***	-0.001	-0.048***
	Fourth Quartile	0.159***	-0.211***	0.016**	-0.029***
Location	Eastern Germany	0.146***	-0.149***	-0.002	-0.025**
	Western Germany	0.144***	-0.183***	0.013**	-0.040***
Sector	Chemical	0.145***	-0.120***	0.011	-0.054***
	Electrical	0.157***	-0.184***	-0.005	-0.052***
	Food	0.154***	-0.224***	-0.002	0.037*
	Furniture	0.114***	-0.160***	-0.018	-0.039**
	Glass	0.122***	-0.192***	0.016	-0.024
	Leather	0.294***	-0.294***	-0.033	0.020
	Machine	0.174***	-0.231***	0.029***	-0.035***
	Metal	0.143***	-0.181***	0.030***	-0.043***
	Oil	0.166*	-0.241*	-0.014	-0.133
	Paper	0.133***	-0.152***	-0.006	-0.065***
	Rubber	0.116***	-0.193***	-0.018	0.015
	Textile	0.247***	-0.223***	0.096***	-0.038
	Vehicle	0.197***	-0.244***	0.004	-0.043*
	Wood	0.147***	-0.197***	0.050*	-0.025

Notes: treatment effect of increased and decreased production expectations. Independent of the sample split, all available observations are used for the matching. The treatment effect is then computed using all observations in a given group. Instead of including time-fixed effects, we use the macro variables introduced in Section 1.4. When grouping by location, we only consider firms that joined the ifo Survey after the German reunification. One, two, and three stars (*) correspond to significance at the 10, 5, and 1 percent significance levels, respectively.

between capital investment and sales forecasts using a business survey of Italian firms run by the Bank of Italy.

1.5.2 Firm-level uncertainty and firm decisions

In theory, not only the first moment of firm expectations matters for firm decisions. The second moment, that is, uncertainty, is important, too. In an influential study, Bloom (2009) emphasized the real option value of delaying an (irreversible) investment decision in the face of increased uncertainty. Whether this matters a lot for aggregate dynamics and the business cycle remains controversial (Bachmann and Bayer 2013, 2014; Bloom et al. 2018). A direct empirical assessment of the effect of uncertainty on firm decisions is thus called for in order to advance our understanding of how firm-level expectations influence firm decisions.

A study by Bachmann et al. (2013) uses the ifo Survey to construct empirical proxies for time-varying business-level uncertainty. They estimate a VAR model to identify uncertainty shocks and find that they induce a temporary contraction of aggregate production in the manufacturing sector as well as of employment and hours—consistent with the notion that uncertainty drives firm decisions. Also, they obtain similar results for the US based on the Business Outlook Survey maintained by the Federal Reserve Bank of Philadelphia. Bachmann et al. (2019), in turn, zoom in on the decisions at the firm level. They find that idiosyncratic firm-level volatility raises the probability of a decision to reset prices (upwards or downwards). This may reflect the fact that firms are exposed to larger shocks as uncertainty (volatility) increases and suggests that the “volatility effect” dominates the “wait-and-see” effect, according to which one would expect a reduced probability to adjust prices. They also establish a fall in the aggregate price level following a shock to average firm-specific volatility.²⁵ Lastly, we note that misperceptions of the extent of uncertainty may also impact firms’ decisions. Ben-David et al. (2013) find for CFOs in the US that more “miscalibrated” (realized returns lie often outside the reported confidence intervals) managers invest more and tolerate higher leverage.

In sum, recent evidence based on survey data suggests that firm expectations matter for firm decisions—as economic theory would suggest. Yet the evidence to date is limited and more research is called for, not least with a view towards assessing the importance of expectations—both its first and its second moment—for firm decisions from a quantitative view. It would be particularly desirable to compare the evidence against predictions from quantitative models which also allow for departures from FIRE in order to account simultaneously for the expectation-formation process (as discussed in Section 1.4 above) and the effect of expectations on firm decisions.

1.6 Conclusion

As more and more survey data on firms’ expectations has become available, the literature has started to explore this data systematically from various angles over the last decade or so. In surveying this work, we have focused on firm expectations about firm-specific

²⁵See also Vavra (2014) for a model-based analysis of how volatility impacts pricing behavior.

developments. We have identified a number of stylized facts and revisited a number of noteworthy insights into the expectation-formation process. Lastly, we have also discussed evidence which illustrates the importance of firm expectations for firm behavior.

More research on firm expectations is called for. The following items feature prominently on our non-exhaustive wish list. First, we need more evidence on firms' forecast errors. While they are not biased unconditionally (Fact 1), they are predictable conditional on some firm-specific variables (Fact 4). Models which account simultaneously for both observations would be important advances. Second, regarding the expectation-formation process of firms, we need to develop a better understanding of how often and how strongly firms update their expectations and what role behavioral features play in this process. Third, we are currently lacking a comprehensive theory which ties together the expectation-formation and decision process of firms. Any advances in these directions are highly welcome. Fourth, while we have made an effort to assemble observations from many countries and surveys, a systematic cross-country comparison of firm-level data on firm expectations is bound to deliver additional valuable insights. While there have been efforts to harmonize firm surveys in the EU, the firm-level data is not available on a common platform. Lastly, we also consider a systematic comparison of qualitative and quantitative survey responses a promising avenue for future research.

1.A Appendices

1.A.1 Expectation errors

Table 1.A.1: Definitions of qualitative expectation errors

Source	Agg. realization $x_{t,h}^i = f(\varsigma_{t,h}^i)$	Expectation error $e_{t,h}^i = f(x_{t,h}^i, x_{t,h t}^i)$	Production		Prices	
			μ	σ	μ	σ
Nerlove (1983)	$\text{sgn}(\varsigma_{t,h}^i)$	$\text{sgn}(x_{t,h}^i - x_{t,h t}^i)$	-0.05	0.65	-0.04	0.65
Bachmann et al. (2013)	$\varsigma_{t,h}^i$	0 if $\text{sgn}(x_{t,h}^i) = \text{sgn}(x_{t,h t}^i)$ $\frac{1}{h}(x_{t,h}^i - x_{t,h t}^i)$ else	-0.03	0.35	-0.02	0.24
Massenot and Pettinicchi (2018)	$\frac{1}{h}\varsigma_{t,h}^i$	$x_{t,h}^i - x_{t,h t}^i$	-0.04	0.53	-0.09	0.41

Notes: schemes for the computation of expectation errors from qualitative surveys like the BEP. Realizations for one month are denoted by $x_{t,1}^i \in \{-1, 0, +1\}$, expectations for h months ahead are denoted by $x_{t,h|t}^i \in \{-1, 0, +1\}$. To account for the difference in reference periods and the qualitative nature, schemes first aggregate monthly realizations over h months and then compare aggregate realizations to expectations. Aggregate realizations $x_{t,h}^i$ are based on the sum of monthly changes over h months $\varsigma_{t,h}^i = \sum_{j=1}^h x_{t+j,1}^i$. Nerlove (1983) and Kawasaki and Zimmermann (1986) set $x_{t,h}^i$ to missing when there are opposite signs in the sum. sgn denotes the sign function and returns 1, 0, or -1 . The last four columns report the mean (μ) and standard deviation (σ) for expectation errors in the BEP.

Table 1.A.1 summarizes the main approaches of earlier work using the ifo Survey. The survey asks for the expected change of a variable (production, prices, business situation, etc.) in the next h months, compared to now. We therefore define as $x_{t,h|t}^i$ the expectation of firm i in month t regarding the change of the firm-specific variable x^i from month t to the period from month $t + 1$ until $t + h$. It can take the values -1 (expected decrease), 0 (no expected change), or 1 (expected increase). The realized change—as reported by the firm—of variable x^i from month $t - 1$ to month t is denoted by $x_{t,1}^i$. Aggregating changes over the h months in question yields $\varsigma_{t,h}^i = \sum_{j=1}^h x_{t+j,1}^i$. Different studies have used different ways how to define a forecast error $e_{t,h}^i$ based on transformations $x_{t,h}^i = f(\varsigma_{t,h}^i)$ of $\varsigma_{t,h}^i$, where $x_{t,h}^i$ is the respective definition of the aggregate realization over the h months. Nerlove (1983) and Kawasaki and Zimmermann (1986) compare the sign of $\varsigma_{t,h}^i$ with that of the expectation $x_{t,h|t}^i$. In their definition, the firm has made no expectation error if the two signs align. Otherwise, there is a forecast error that can be positive or negative (-1 or 1). Bachmann et al. (2013) proceed in a slightly different way. They too assign no expectation error if the sign of the aggregate realization $\varsigma_{t,h}^i$ equals that of the expectation $x_{t,h|t}^i$. In case signs differ, however, they quantify the expectation error by assigning the monthly average of the difference between the aggregate realization $\varsigma_{t,h}^i$ and the expectation $x_{t,h|t}^i$. It can therefore take values between $\pm(h + 1)/h$. Massenot and Pettinicchi (2018) define the expectation error as the difference between the monthly average of the aggregate realization $\varsigma_{t,h}^i/h$ and the expectation $x_{t,h|t}^i$, such that the error may take values between -2 and 2 . Note that with this definition, the error is zero only if the realization of the change takes the expected value in each of the h months.

Yet, the mean and the standard deviation of the expectation errors for production and prices, based on the BEP, are fairly comparable across definition, see the right panels of Table 1.A.1. Moreover, the empirical correlations between the values of the aggregate realization are equal to or above 0.98, while the correlations between expectation errors are at least 0.84. The means of the expectation errors for production and prices, independent of the definition, are close to zero.

1.A.2 Additional figures and tables

Table 1.A.2: Relevant questions from the ifo Survey

Label	Name	Question	Possible answers
Q1	Realized Production	Tendencies in the previous month: Our domestic production activities with respect to product XY have	increased [1] not changed [0] decreased [-1]
Q2	Expected Production	Expectations for the next 3 months: Our domestic production activity regarding good XY will probably	increase [1] not change [0] decrease [-1]
Q3	Realized Prices	Tendencies in the previous month: Taking changes of terms and conditions into account, our domestic sales prices (net) for product XY have been	increased [1] not changed [0] decreased [-1]
Q4	Expected Prices	Expectations for the next 3 months: Taking changes of conditions into account our domestic sales prices (net) for XY will probably be	rising [1] not changing [0] falling [-1]

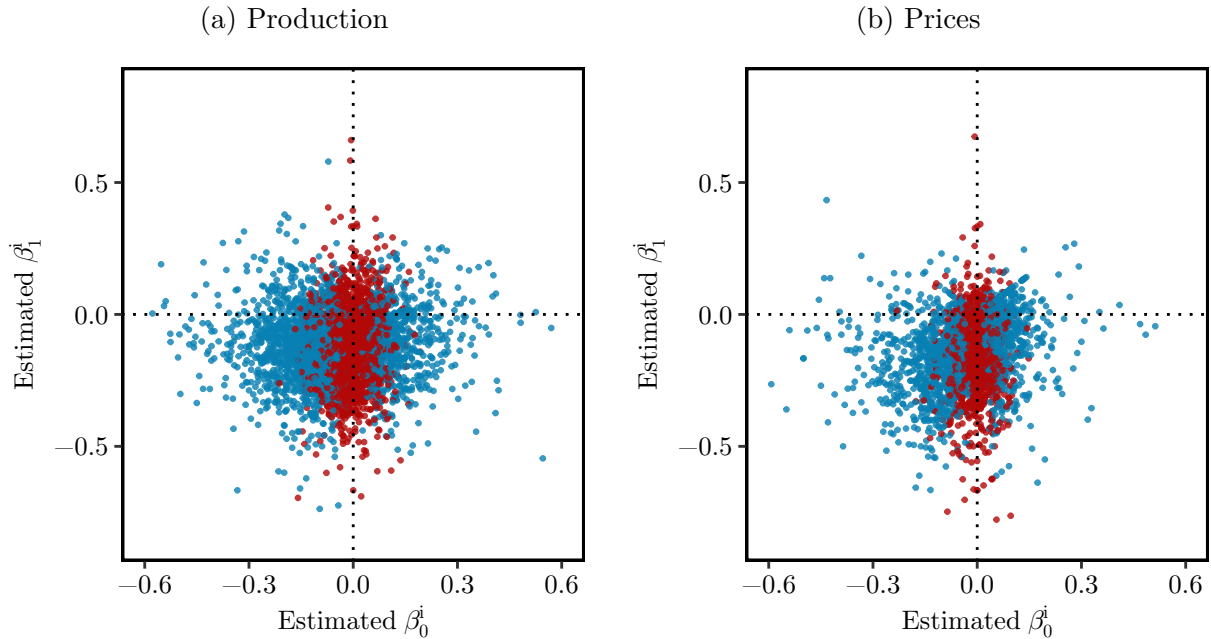
Notes: most recent formulation of the survey questions taken from the EBDC Questionnaire manual.

Table 1.A.3: Summary statistics on firm-level average forecast errors

Grouped by	Group	Production			Prices		
		N	Median	% insig.	N	Median	% insig.
Sector	Chemical	226	-0.0087	83.19	226	-0.0048	78.32
	Electrical	599	-0.0194	78.80	600	-0.0101	82.00
	Food	277	-0.0198	80.51	278	-0.0092	81.29
	Furniture	242	-0.0187	74.79	237	-0.0084	83.97
	Glass	288	-0.0201	76.04	294	-0.0102	79.25
	Leather	63	-0.0111	73.02	62	0.0064	77.42
	Machine	772	-0.0155	80.83	766	-0.0032	84.20
	Metal	612	-0.0129	78.43	583	-0.0104	79.59
	Oil	14	-0.0275	92.86	13	-0.0000	92.31
	Paper	710	-0.0248	75.49	700	-0.0269	72.86
	Rubber	333	-0.0171	76.58	328	-0.0146	79.57
	Textile	315	-0.0261	73.33	329	-0.0108	82.37
	Vehicle	130	0.0031	74.62	128	-0.0021	82.03
	Wood	209	-0.0333	76.56	207	-0.0210	69.08

Notes: estimation of firm-level average forecast errors, entries above provide summary statistics for the estimates for different subgroups of firms. N denotes the number of firms in each group. Sectors are from Bachmann et al. (2019).

Figure 1.A.1: Point estimates for constant and slope



Notes: estimation of equation (1.4) on firm-level observations. Horizontal axis: estimates of β_0^i ; vertical axis: estimates of slope coefficient β_1^i . Colors indicate if the constant is significantly different from 0 (blue) or not (red) at the 5% level. Plot shows values within the 0.025 and 99.75 percent quantiles

Table 1.A.4: Summary statistics on firm-level average squared forecast errors

Grouped by	Group	Production			Prices		
		N	Mean	Median	N	Mean	Median
Sector	Chemical	226	0.1279	0.1152	226	0.0783	0.0498
	Electrical	599	0.1195	0.1083	600	0.0474	0.0332
	Food	277	0.1314	0.1262	278	0.0588	0.0329
	Furniture	242	0.1353	0.1281	237	0.0406	0.0292
	Glass	288	0.1209	0.1073	294	0.0586	0.0349
	Leather	63	0.1127	0.1052	62	0.0490	0.0357
	Machine	772	0.1209	0.1058	766	0.0477	0.0319
	Metal	612	0.1301	0.1156	583	0.0625	0.0354
	Oil	14	0.1054	0.0788	13	0.1557	0.1086
	Paper	710	0.1321	0.1243	700	0.0692	0.0521
	Rubber	333	0.1369	0.1289	328	0.0695	0.0473
	Textile	315	0.1203	0.1118	329	0.0618	0.0330
	Vehicle	130	0.1185	0.1065	128	0.0422	0.0283
	Wood	209	0.1400	0.1299	207	0.0739	0.0496

Notes: Notes: estimation of firm-level average squared forecast errors, entries above provide summary statistics for the estimates for different subgroups of firms. N denotes the number of firms in each group. Sectors are from Bachmann et al. (2019).

Table 1.A.5: Definition of variable blocks

Block	Variable	Description	Frequency	Periods
Survey	Business Situation		monthly	t to t-3
	Realized Production		monthly	t to t-2
	Expected Production		monthly	t-1 to t-3
	Realized Prices		monthly	t to t-2
	Orders		monthly	t to t-3
	Foreign Orders		monthly	t to t-3
	Demand		monthly	t to t-2
	Capacity		monthly	t-1 to t-3
	Expected Prices		monthly	t-1 to t-3
	Employees		annual	
	Avg. Business Situation	two-digit sector level	monthly	t
Sectoral Fixed Effects				
Fundamentals	Financing Coefficient	$\frac{\text{Liabilities} - \text{Provisions}}{\text{Equity} + \text{Provisions}}$	annual	
	Debt Share	$\frac{\text{Total debt}}{\text{Assets}}$	annual	
	Total Assets		annual	
Macro	PPI Growth	versus previous month	monthly	t-2
	CPI Growth	versus previous month	monthly	t-2
	Unemployment		monthly	t-1
	IP Growth	versus previous month	monthly	t-2

Notes: components of the three variable blocks considered as explanatory variables in the ordered probit. The survey and fundamental blocks are taken from Enders (2020).

Table 1.A.6: Overreaction to firm-specific news

Grouped by	Group	Production			Prices		
		N	Mean	Median	N	Mean	Median
Overall		4851	-0.1121	-0.1089	4851	-0.1070	-0.0820
Number of Employees	Fewer than 50	236	-0.1050	-0.1041	236	-0.1029	-0.0777
	50-199	156	-0.0844	-0.0660	156	-0.1108	-0.0827
	200-499	78	-0.0918	-0.0825	78	-0.1059	-0.0739
	500-999	22	-0.1586	-0.1721	22	-0.0826	-0.0693
	More than 1000	5	-0.1433	-0.1833	5	-0.0751	-0.0736
Employees (Quartile)	First Quartile	124	-0.0964	-0.0971	124	-0.1047	-0.0878
	Second Quartile	124	-0.1158	-0.1160	124	-0.1025	-0.0647
	Third Quartile	124	-0.0816	-0.0555	124	-0.1139	-0.0907
	Fourth Quartile	125	-0.1029	-0.1042	125	-0.0978	-0.0667
Sales (Quartile)	First Quartile	107	-0.0989	-0.1029	107	-0.1234	-0.0912
	Second Quartile	112	-0.1016	-0.0846	112	-0.0983	-0.0642
	Third Quartile	109	-0.0999	-0.0903	109	-0.1080	-0.0940
	Fourth Quartile	110	-0.1060	-0.1047	110	-0.1087	-0.0659
Total Assets (Quartile)	First Quartile	130	-0.0962	-0.0955	130	-0.1107	-0.0840
	Second Quartile	131	-0.0979	-0.0987	131	-0.1131	-0.0829
	Third Quartile	130	-0.0954	-0.0870	130	-0.0932	-0.0675
	Fourth Quartile	131	-0.1146	-0.1071	131	-0.1129	-0.0730
Location	Eastern Germany	2203	-0.1121	-0.1099	2203	-0.1060	-0.0806
	Western Germany	1198	-0.1055	-0.1025	1198	-0.1081	-0.0824
Sector	Chemical	271	-0.1113	-0.1105	271	-0.1025	-0.0718
	Electrical	515	-0.1147	-0.1131	515	-0.1078	-0.0876
	Food	358	-0.1092	-0.1108	358	-0.1043	-0.0786
	Furniture	238	-0.1082	-0.1018	238	-0.1117	-0.0817
	Glass	262	-0.1090	-0.0980	262	-0.1170	-0.0931
	Leather	86	-0.1309	-0.1266	86	-0.0880	-0.0523
	Machine	646	-0.1185	-0.1111	646	-0.1088	-0.0813
	Metal	719	-0.1073	-0.1105	719	-0.1052	-0.0773
	Oil	11	-0.0541	-0.0443	11	-0.1178	-0.0508
	Paper	574	-0.1111	-0.1060	574	-0.1102	-0.0892
	Rubber	343	-0.1167	-0.1097	343	-0.1125	-0.0885
	Textile	265	-0.1042	-0.0924	265	-0.1064	-0.0825
	Vehicle	144	-0.1113	-0.1172	144	-0.1042	-0.0755
	Wood	248	-0.1263	-0.1272	248	-0.1042	-0.0862

Notes: estimation of equation (1.4) on firm-level observations. Entries provide summary statistics for the slope estimates based for different subgroups of firms. N denotes the number of firms in each group. When grouping by location, we only consider firms that joined the ifo Survey after the German reunification.

Table 1.A.7: Summary statistics firm-level constant estimates

Grouped by	Group	Production			Prices		
		N	Mean	Median	N	Mean	Median
Overall		4851	-0.0317	-0.0263	4851	-0.0093	0.0056
Number of Employees	Fewer than 50	236	-0.0236	-0.0155	236	0.0005	0.0062
	50-199	156	-0.0237	-0.0274	156	0.0048	0.0133
	200-499	78	0.0068	-0.0023	78	0.0065	0.0094
	500-999	22	-0.0200	-0.0299	22	-0.0004	-0.0071
	More than 1000	5	-0.0148	-0.0344	5	-0.0132	-0.0090
Employees (Quartile)	First Quartile	124	-0.0232	-0.0103	124	-0.0057	0.0053
	Second Quartile	124	-0.0240	-0.0188	124	0.0092	0.0123
	Third Quartile	124	-0.0242	-0.0264	124	0.0011	0.0115
	Fourth Quartile	125	-0.0032	-0.0232	125	0.0058	0.0099
Sales (Quartile)	First Quartile	107	-0.0192	0.0000	107	0.0002	0.0052
	Second Quartile	112	-0.0258	-0.0111	112	0.0032	0.0079
	Third Quartile	109	-0.0127	-0.0150	109	0.0024	0.0103
	Fourth Quartile	110	-0.0311	-0.0325	110	-0.0095	0.0065
Total Assets (Quartile)	First Quartile	130	-0.0196	0.0002	130	-0.0058	0.0051
	Second Quartile	131	-0.0270	-0.0220	131	0.0018	0.0105
	Third Quartile	130	-0.0104	-0.0153	130	0.0107	0.0134
	Fourth Quartile	131	-0.0311	-0.0265	131	-0.0092	0.0093
Location	Eastern Germany	2203	-0.0256	-0.0208	2203	-0.0060	0.0070
	Western Germany	1198	-0.0373	-0.0303	1198	-0.0126	0.0038
Sector	Chemical	271	-0.0446	-0.0291	271	-0.0162	0.0060
	Electrical	515	-0.0435	-0.0315	515	-0.0113	0.0051
	Food	358	-0.0225	-0.0230	358	-0.0070	0.0046
	Furniture	238	-0.0269	-0.0231	238	-0.0113	0.0076
	Glass	262	-0.0343	-0.0128	262	-0.0097	0.0056
	Leather	86	-0.0395	-0.0264	86	-0.0120	0.0113
	Machine	646	-0.0239	-0.0230	646	-0.0052	0.0046
	Metal	719	-0.0303	-0.0231	719	-0.0104	0.0057
	Oil	11	0.0085	-0.0230	11	0.0035	0.0185
	Paper	574	-0.0322	-0.0304	574	-0.0119	0.0032
	Rubber	343	-0.0318	-0.0265	343	-0.0123	-0.0000
	Textile	265	-0.0370	-0.0276	265	-0.0067	0.0065
	Vehicle	144	-0.0360	-0.0332	144	-0.0096	0.0021
	Wood	248	-0.0317	-0.0233	248	-0.0063	0.0100

Notes: summary statistics for the estimates of the constant from the forecaster-by-forecaster regressions in equation (1.4) for different groups of firms. When grouping by location we only consider firms that joined the ifo Survey after the German reunification.

Table 1.A.8: Predictability of expectation errors

Variable	Timing	Production			Prices		
		estimate	t-value	p-value	estimate	t-value	p-value
Constant		0.022	1.22	0.22	0.037***	3.14	0.00
IP growth	real-time	0.424*	1.93	0.05	0.165	1.58	0.11
Unemployment rate	t-1	0.002	1.16	0.24	-0.001	-0.86	0.39
PPI growth	t-2	0.005	0.23	0.82	0.036***	3.61	0.00
CPI growth	t-2	-0.016	-1.07	0.29	-0.007	-1.00	0.32
Expectation about own prices	t	0.012***	3.97	0.00	-0.258***	-81.95	0.00
	t-1	-0.001	-0.39	0.70	0.055***	21.63	0.00
	t-2	-0.010***	-3.87	0.00	0.010***	4.23	0.00
	t-3	-0.010***	-3.30	0.00	0.001	0.26	0.79
Expectation about own production	t	-0.301***	-94.38	0.00	0.002	1.22	0.22
	t-1	0.041***	15.97	0.00	-0.001	-0.78	0.43
	t-2	0.007**	2.52	0.01	-0.001	-0.89	0.37
	t-3	-0.004	-1.26	0.21	0.000	-0.19	0.85
Reported business situation	t	0.007**	2.48	0.01	0.004**	2.21	0.03
	t-1	-0.004*	-1.93	0.05	0.000	-0.17	0.87
	t-2	0.004*	1.77	0.08	0.002	1.21	0.23
	t-3	0.019***	5.95	0.00	0.001	0.74	0.46
Reported backlog of orders	t	-0.020***	-7.01	0.00	-0.011***	-6.23	0.00
	t-1	0.001	0.49	0.63	0.000	0.25	0.80
	t-2	0.004**	1.97	0.05	0.000	0.30	0.76
	t-3	-0.007***	-2.85	0.00	-0.001	-0.73	0.47
most recent reported change in production	t	0.038***	12.22	0.00	0.003	1.58	0.12
	t-1	0.025***	8.98	0.00	0.003*	1.75	0.08
	t-2	0.021***	7.83	0.00	0.004**	2.43	0.02
	t-3	0.023***	7.63	0.00	0.001	0.52	0.60
most recent reported change in prices	t	-0.003	-1.11	0.27	0.060***	16.43	0.00
	t-1	-0.003	-1.35	0.18	0.038***	13.73	0.00
	t-2	0.000	-0.08	0.93	0.033***	12.65	0.00
	t-3	-0.003	-0.91	0.36	0.041***	12.20	0.00
Reported change in demand	t	0.048***	16.18	0.00	0.010***	5.45	0.00
	t-1	0.023***	9.39	0.00	0.004**	2.45	0.01
	t-2	0.014***	6.09	0.00	0.002	1.24	0.22
	t-3	0.006**	2.14	0.03	0.001	0.89	0.38
R^2		0.172			0.170		

Notes: predictive regressions for forecast errors for prices and production. For IP growth we use real-time data for the seasonally and calendar adjusted industrial production and compute monthly growth rates that are also reported in the press releases of DESTATIS. We assume that firms update their information set on the day after the release. Since 2005 firms may complete the survey online. Only for these firms the day of completion is known, which is the sample used for this exercise. One, two, and three stars (*) correspond to significance on the 10, 5, and 1 percent significance levels.

Firm Expectations and News: Micro v Macro

Joint with Benjamin Born, Zeno Enders, Gernot J. Müller, and Manuel Menkhoff

2.1 Introduction

How do firms adjust their expectations to news? Addressing this question yields important insight into their expectation-formation process. *Rational expectations* provide a natural benchmark. In this case, forecast errors are possible but not predictable based on information that is available to the forecaster in real-time—since expectations adjust correctly and instantaneously to news. If, instead, news predicts *positive* forecast errors, expectations adjust too little: they underreact relative to the rational-expectations benchmark. If news predicts *negative* forecast errors, expectations overreact to news. Recent work studies systematically and at different levels of aggregation how news impacts forecast errors, mostly relying on surveys of professional forecasters (Coibion and Gorodnichenko 2015; Bordalo et al. 2020; Broer and Kohlhas 2023).

Against this background, our study offers a new perspective because it relies on a large panel of firm expectations. As a result, we are able to account for heterogeneity in the expectation-formation process along two dimensions. First, we study news of different types. While professional forecasters are surveyed about aggregate indicators, firms in our sample report expectations about firm-specific developments. In this context, we can classify news as either micro or macro, with micro news being information about firm-specific developments and macro news being information about the aggregate economy that, in turn, matters for (expectations about) firm-specific developments, too. Second, by focusing on firm expectations instead of professional forecasters’ expectations, we can exploit a much larger and richer data set and probe into the role of (firm) heterogeneity in the expectation-formation process. Specifically, we rely on the ifo survey of German firms, which features responses from some 1,500 firms each month and covers 15 years of data. In addition, we verify that our main results also hold for the Banca d’Italia’s “Survey on Inflation and Growth Expectations” (SIGE) of Italian firms.

We find that the distinction between micro and macro news is essential: firm expectations overreact to micro news, but simultaneously underreact to macro news. This pattern emerges robustly across a variety of specifications and for all firm types that we consider (e.g., small and large, young and old). It also holds for different measures of expectations and different outcome variables. The variation of overreaction across firms is also systematically related to measures of firm performance. To rationalize these results, we put forward a stylized general-equilibrium model. It builds on the dispersed information model of Lorenzoni (2009), but assumes, in addition, that firms suffer from ‘island illusion’: They systematically underestimate the importance of aggregate developments for their own performance. This

departure from rational expectations allows the model to predict simultaneous over- and underreaction to micro and macro news.

More in detail, the first part of the paper presents new evidence on how firms' expectations change in response to news. This evidence is based on data from the ifo survey of German firms, which is a well-known and widely used survey that has been conducted since 1949 and whose design has since then been adopted by surveys around the world (Becker and Wohlrabe 2008; Born et al. 2022). Our data covers the period from April 2004 to December 2019. We first focus on firms' expectations about their production over the next three months, which are reported in a qualitative manner. This raises some challenges in defining forecast errors, which we address in Section 2.2 below. However, our results are robust once we consider quantitative measures of expectations based on both, the ifo survey and SIGE.

To study how firm expectations respond to news, we adopt the framework of Coibion and Gorodnichenko (2015), which is by now widely used in the literature. The idea is straightforward: we regress firms' forecast errors about the change of production over the next three months on news that is available in the current month. We approximate what is news to firms by their forecast revision, that is, the change in what they report as production expectations. Importantly, these revisions may reflect firm-specific news (micro news) or news about the aggregate economy (macro news). We isolate the effect of the micro component as we purge a firm's forecast revision of the firm-specific impact of a set of macroeconomic indicators that are available in real-time and by controlling for macro news.

To construct macro news, we rely on the ifo *business climate index*, which is an aggregate indicator of the German business cycle compiled on the basis of the ifo survey. This index is widely watched and Bloomberg samples a consensus forecast prior to its release. The difference between the current release of the index and the consensus forecast, both available in real-time, provides us with a natural measure of macro news. Two aspects are important to note. First, the ifo index is constructed by aggregating expectations across firms in the survey such that micro and macro news are directly comparable but differ in the level of aggregation. Second, regarding the timing, we note that macro news is released at the end of the previous month and is thus available as firms report their forecast in the current month—just like micro news. For these reasons, both micro and macro news should not predict the forecast error under rational expectations. And yet, our first key result, based on firm-level and pooled panel regressions, is that they do so robustly.

Our second result is that they do so in systematically different ways. Macro news, or information about the overall economy, tends to lead to positive forecast errors, meaning that actual production ends up exceeding expectations. More concretely, if the current ifo index surprises positively, a firm's production is likely to exceed its expectation over the course of the next three months. 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, that is, an upward revision of production expectations tends to be followed by a worse-than-expected output performance. Firm expectations respond too strongly to micro news: they overreact.

We find that these patterns are a robust feature of our data set. They emerge for alternative definitions of news and forecast errors and also once we consider firms' business expectations which are reported on a quantitative scale and pertain to a 6-month horizon. We also determine whether our findings generalize beyond the ifo survey, which we use as

our main data source. To do this, we turn to the SIGE. This survey provides us with a measure of firms’ quantitative price expectations over a 12-month horizon, and we can use it to measure micro and macro news as we do in the ifo survey. And just like for the ifo survey, we find that firm expectations overreact to micro news but underreact to macro news.

In addition to analyzing the overall response to news using a panel of pooled observations, we also examine how individual firms respond to news by taking advantage of the large number of consecutive observations available for most firms in the ifo survey. We find that overreaction to micro news is a pervasive feature across firms. Firm-level estimates are consistently negative and tightly distributed in a narrow range. There is no economically significant difference in estimates across firm characteristics, such as firm size or firm age. The response to macro news is somewhat more dispersed across firms. Although there is underreaction for most firms, firms differ in how strongly they underreact to macro news. 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.

The estimated response coefficients also vary over time, although they do not change their signs. The underreaction to macro news is strongest during the Great Recession, reflecting a more substantial impact of the macroeconomy in turbulent times. We also find that underreaction and overreaction are persistent over time—forecast errors respond not only to current but also to past news. This finding suggests that our results are not caused by measurement error. Lastly, we establish that the variation in the reaction to news across firms correlates with firm-level outcomes in a systematic way. 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 is associated with higher firm-level production and forecast-error volatility. These findings are consistent with earlier work which shows that firm expectations matter for firm outcomes (Bachmann et al. 2013; Enders et al. 2022).

In the last part of the paper, we put forward a general equilibrium model in order to rationalize our findings. The model builds on Lorenzoni (2009), which in turn is based on Lucas (1972), but can be solved in closed form. In addition to the noisy-information structure of the original model, we assume that firms are prone to ‘island illusion,’ meaning that they tend to underestimate the influence of overall economic conditions on their own performance. We think of island illusion as an instance of salience, which Taylor and Thompson (1982) define as “the phenomenon that when one’s attention is differentially directed to one portion on the environment rather than to others, the information contained in that portion will receive disproportionate weighing in subsequent judgments” (see also Bordalo et al. 2013). Island illusion is hence consistent with the notion that firm-specific developments are salient stimuli to firms because they attract firms’ attention “bottom-up, automatically and involuntarily” (Bordalo et al. 2022). As such, they feature disproportionately in firms’ expectation-formation process—while other sources of information have to be gathered and processed actively.¹

Our model setup differs from earlier work by Bordalo et al. (2020) and Broer and Kohlhas (2023) as we model the response of expectations about firm-level outcomes in a fully specified general-equilibrium setting. This is essential in the context of our analysis because it allows

¹Bianchi et al. (2022) use a machine-learning algorithm to estimate time-varying systematic expectational errors and find that—consistent with our notion of island illusion—survey respondents place too much weight on the private or judgmental component of their forecasts and too little weight on publicly available economic information.

us to account for the cross-equation restrictions which govern the impact of micro and macro news on firm expectations. In the model, information is dispersed across firms. Firms observe their own developments plus a public signal and use this information to forecast sales. Prices are set before actual demand is observed. Firms are then assumed to adjust production in order to meet demand given posted prices. Consequently, the aggregate state of the economy is important for firms when it comes to forecasting their own production. The model is sufficiently stylized so that we can derive our main result in closed form: We show that island illusion causes firm expectations to overreact to micro news and underreact to macro news. It also accounts for how differences in the response to news across firms correlate with firm outcomes, such as profits and forecast-error volatility.

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 2.2 provides details about our data set. In Section 2.3, we introduce our empirical framework and present the results. We develop and solve a general equilibrium model with dispersed information and island illusion in Section 2.4. The final section offers some conclusions.

Related Literature. Our paper builds on three strands of the literature. First, at an empirical level, our work relates to the literature which is concerned with macroeconomic expectations of firms, see, for instance, Andrade et al. (2022), Coibion et al. (2018, 2020), and Savignac et al. (2021), as well as the recent survey by Candia et al. (2022). In contrast, our focus is on firm expectations about firms’ own performance. Here, only a limited number of studies have analyzed firm expectations about firm outcomes (see Born et al. 2022). Massenet and Pettinicchi (2018), in particular, use ifo data as well, regressing expectations and forecast errors on past changes of the business situation (rather than on forecast revisions). They find the regression coefficient is positive and significant, and refer to this result as “over-extrapolation”. Enders et al. (2019), in turn, take a macro perspective and document that the response of firm expectations to monetary policy shocks is non-linear in the size of the shock. Neither of these studies distinguishes between the response to micro and macro news.

Second, our empirical setup builds on a framework that has been popularized by Coibion and Gorodnichenko (2015), see Born et al. (2024) for a survey. Importantly, as in Bordalo et al. (2020), we estimate our model at the level of individual forecasters.² Predictable forecast errors at this level allow us to reject rational expectations. But this does not imply a rejection of 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 et al. 2008; Farmer et al. 2023; Kohlhas and Roberston 2022; Bachmann et al. 2023).³

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*

²See also Angeletos et al. (2021), Broer and Kohlhas (2023), and Kučinskis and Peters (2022) for further 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 (see, again Coibion and Gorodnichenko 2015; Bordalo et al. 2020). This includes models of rational inattention (e.g., Maćkowiak and Wiederholt 2009).

⁴Under certain conditions, behavioral models and incomplete information models give rise to equivalent

can rationalize why there is underreaction to current news (e.g., Farhi and Werning 2019; García-Schmidt and Woodford 2019; Gabaix 2020; Bouchaud et al. 2019; Carroll et al. 2020), while *constrained memory* may account for overreaction (Azeredo da Silveira and Woodford 2019). Ba et al. (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 et al. 2021; Andre et al. 2022). Our model of island illusion is conceptually closely related to *diagnostic expectations* and *overconfidence* (Bordalo et al. 2019, 2020; Broer and Kohlhas 2023). It differs from these approaches in simultaneously accounting for under- and overreactions in a general-equilibrium setting. Such a setting is key because it allows us to model expectations about firm outcomes based on micro and macro news consistently.

2.2 Measuring forecast errors and news

In this section, we first introduce the data set for our empirical analysis. It is centered around the ifo survey of German firms. We also provide details on the construction and descriptive statistics of firms’ forecast errors and the news measures.

2.2.1 The ifo survey

The ifo survey is a mostly qualitative, monthly survey among German firms and representative of the German economy (Hiersemenzel et al. 2022).⁵ While the ifo survey was launched in 1949—and some aggregate statistics based on it were first used by Theil (1955)—the underlying micro data is available for research since 1980. Participation is voluntary and firms only receive non-monetary compensation in the form of sectoral and aggregate results of the survey. 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). Response rates for the ifo survey are generally high: Out of all firms initially contacted in mid-2021, around two-thirds returned at least two surveys. For the comparable Survey of Business Uncertainty in the United States, the response rate is around one-third only (Altig et al. 2022). Response rates remain high also after initial contact, with an average monthly response rate of 82 percent; the sample attrition is moderate (Enders et al. 2022).

Our analysis below relies on measures of firms’ forecast errors and news and builds on three main components: (i) the ifo Business Climate *Survey* in the manufacturing sector (IBS-IND 2020, from now on “ifo survey”), (ii) the ifo Business Climate *Index* (ifo index), and (iii) the Bloomberg consensus forecasts for the ifo index. Our sample is restricted by limited data availability of the Bloomberg forecasts and runs from April 2004 to December 2019.

To measure firm expectations and forecast errors, we rely on the ifo survey. It features a core set of questions, including questions about expected and realized production, prices, and business situation, where firms can report either an increase, no change, or a decrease. While

equilibrium effects (Angeletos and Huo 2021).

⁵Quantitative questions were added in 2005, distributional questions in 2013, see Bachmann et al. (2020, 2021) for details. While the survey is technically at the product level, we follow the literature and treat each respondent as a separate firm (e.g., Bachmann et al. 2013; Born et al. 2022; Enders et al. 2022).

this makes quantitative statements challenging, the qualitative nature arguably reduces the room for measurement error. In our empirical analysis, we rely on time-series data at the level of individual firms. Therefore, we restrict our sample to those firms which are in the survey for at least 30 months and which exhibit some time-series variation in their expectations and expectation errors. In any given month, this leaves us with more than 1,000 responses and often more than 1,500. Panel (a) of Figure 2.1 plots the distribution of firms sorted according to the number of months a firm is in the sample. The median firm is in the survey for around 90 months and 25 percent of firms are in the survey for more than 130 months. We exploit the fact that we have fairly long time series available for individual firms in our analysis in Section 2.3. In particular, it allows us to characterize the heterogeneity of the expectation-formation process systematically.

2.2.2 Forecast errors

To construct firms' forecast errors, we follow the approach of Bachmann et al. (2013) and focus on expected and realized production as reported in the ifo survey. Here, firm j reports for its own production the realized change over the previous month $x_{t,t-1}^j \in \{-1, 0, 1\}$ and the expected change over the following three months $F_t^j(x_{t+3,t}^j) \in \{-1, 0, 1\}$, see Appendix-Table 2.A.1 for the exact wording of the survey questions. To harmonize the time horizons, we aggregate the realized changes over the following three months: $x_{t+3,t}^j = \sum_{s=0}^2 x_{t+s+1,t+s}^j$. Given the aggregated realized and expected changes, we define the forecast error as:

$$x_{t+3,t}^j - F_t^j(x_{t+3,t}^j) = \begin{cases} 0 & \text{if } \text{sign}\{x_{t+3,t}^j\} = \text{sign}\{F_t^j(x_{t+3,t}^j)\}, \\ \frac{1}{3}[x_{t+3,t}^j - F_t^j(x_{t+3,t}^j)] & \text{else.} \end{cases} \quad (2.1)$$

When the signs of the aggregated realized change and the expected change coincide, no error is assigned. In all other cases, the forecast error is equal to the difference between the realized and the expected change, standardized by the forecasting horizon of three months.

Generally, we find forecast errors to be well-behaved. Panel (b) of Figure 2.1 shows the distribution of forecast errors: More than 75 percent of firm-level average forecast errors are not significantly different from zero. And while these forecast errors are based on qualitative rather than quantitative data, Born et al. (2022) show that key facts which characterize firms' forecast errors emerge robustly from qualitative and quantitative data and across countries.

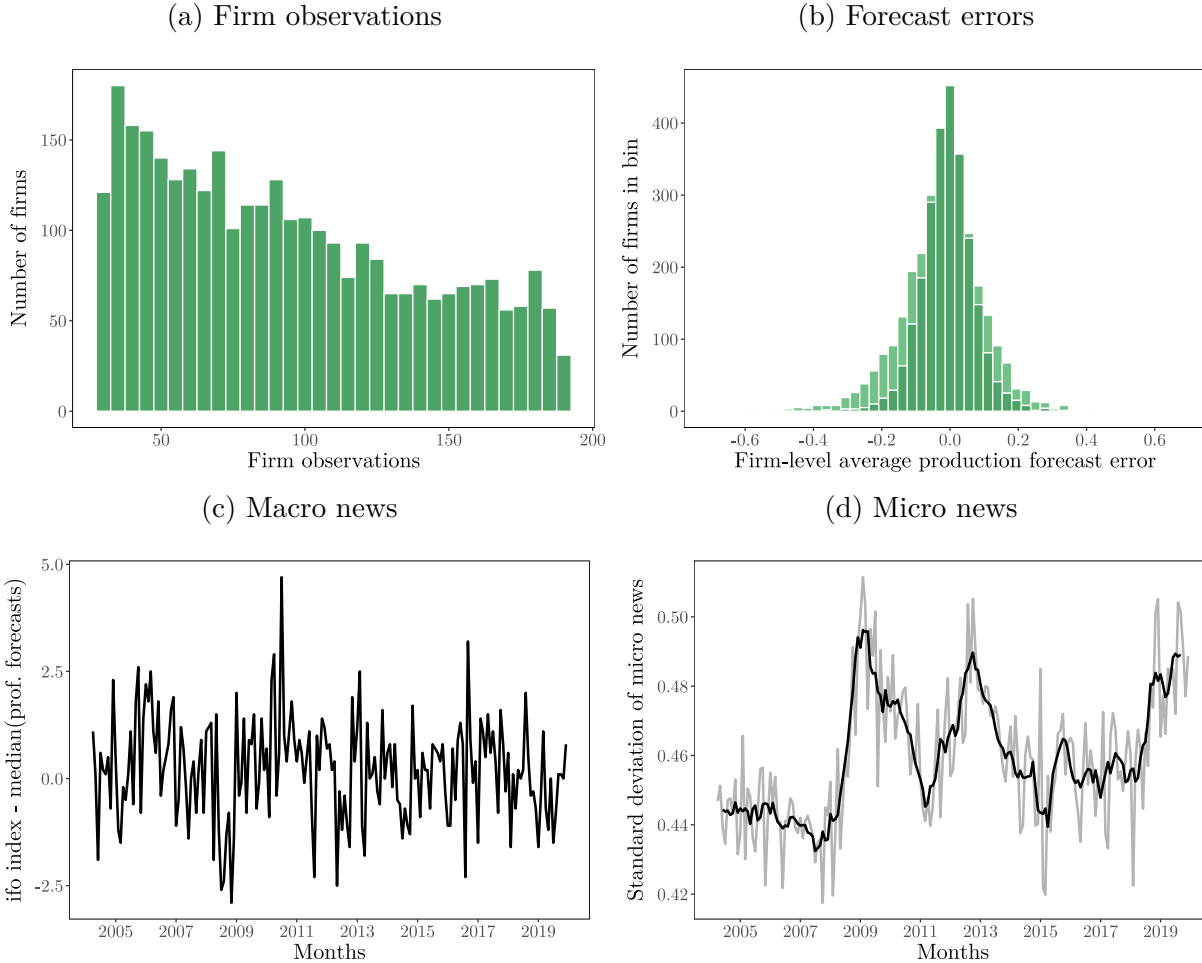
2.2.3 Macro news

To measure macro news, we compute the *surprise component* of the ifo *index*. The ifo index is compiled on the basis of the ifo survey by the ifo Institute and is a widely watched indicator of the German business cycle (Carstensen et al. 2020; Lehmann 2023). The index is based on firms' responses about their current business situation and their business expectations over the next 6 months, see again Appendix-Table 2.A.1 for the wording of the survey question.⁶ The index is compiled as follows:

$$\text{business climate}_t = \sqrt{(\text{business situation}_t + 200)(\text{business expectation}_t + 200)} - 200 ,$$

⁶Since April 2018, the index also includes responses from service-sector firms (Sauer and Wohlrabe 2018).

Figure 2.1: The ifo survey, forecast errors, and news



Notes: Panel (a): distribution of monthly firm observations, i.e., the number of firms for which a firm-specific time series of a certain length is available. Panel (b): histogram of firm-level average forecast errors for production. The color indicates if estimates are significantly different from zero at the five percent level (light green) or not (dark green). Panel (c): macro news over time, defined as the surprise in the ifo index compared to median professional forecasts, see Equation (2.2). Panel (d): cross-sectional standard deviation of micro news over time, defined as the residuals of a regression of forecast revisions on real-time economic indicators, see Equation (2.4). The grey line depicts the standard deviation of micro news at a monthly level and the black line depicts the six-month rolling average.

where $\text{business situation}_t$ and $\text{business expectation}_t$ are balances, that is, the share of positive answers (“increase”) minus the share of negative answers (“decrease”) across firms in month t . For publication, the ifo institute reports the business climate as an index relative to a base year (Sauer and Wohlrabe 2018).

We measure the surprise component in the ifo index based on professional forecasts for the ifo index, available from the Bloomberg consensus survey. In this survey, professional forecasters can submit and update their forecasts of macroeconomic indicators, for example, GDP, employment, and confidence indexes, up until they are released. In the literature, these forecasts have been used to assess the impact of news on long-term treasury bonds (Altavilla

et al. 2017) and stock prices (Elenev et al. 2022; Born et al. 2023; Gilbert et al. 2017; Kurov et al. 2019); see also the construction of uncertainty indexes by Scotti (2016) and the nowcast errors by Enders et al. (2021). For the German ifo index and starting in April 2004, the Bloomberg survey features some 40 professional forecasters.

We construct macro news as the difference between the published ifo index and the median professional forecast of the ifo index from Bloomberg. The timing is key: In the first three weeks of month $t - 1$, firms respond to the survey. Until the last week of month $t - 1$, professional forecasters submit their forecasts for the ifo index in $t - 1$ to Bloomberg. In the last week of month $t - 1$, the ifo institute then publishes the value of the ifo index. In the first three weeks of month t and after observing the macro news, firms again fill out the ifo survey. Formally, we define macro news, as observable at the beginning of month t as follows:

$$\text{macro news}_t = \text{ifo index}_{t-1} - \text{median}(\text{professional forecasts for ifo index}_{t-1}) . \quad (2.2)$$

We display the resulting time series of macro news in Panel (c) of Figure 2.1.

We can be confident that macro news is part of the information set of firms when forecasting their production in t . First, media attention to the index as well as its professional forecasts is high due to its predictive power for the German business cycle. The ifo index is ranked among Bloomberg’s “12 Global Economic Indicators to Watch” and news outlets report on both the realized value and, importantly, the professional forecasts.⁷ Second, information about the aggregate index (as well as the sectoral results) is given to firms as compensation for their participation in the survey by the ifo institute at the end of month $t - 1$.

2.2.4 Micro news

Our measure of micro news is based on forecast revisions. Formally, we define the forecast revision of firm j in month t , FR_t^j , as the first difference of production expectations:

$$FR_t^j = \text{sign}\{F_t^j(x_{t+3,t}^j) - F_{t-1}^j(x_{t+2,t-1}^j)\} , \quad (2.3)$$

which 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 the forecast horizon is fixed at 3 months, the overlap in the monthly forecast revisions is two months. In what follows, we thus assume that forecast revisions reflect mostly news (rather than changes in the forecast horizon).⁸ To assess the informativeness of the forecast revisions, we relate the average forecast revisions over time to German manufacturing production growth, see Figure 2.A.1 in the appendix. It turns out that the average forecast revision is a leading indicator for changes in manufacturing production. This is especially

⁷Examples include leading weekly newspapers *Der Spiegel* and *Die Zeit*. *Der Spiegel* (“Unternehmen sind wegen vierter Coronawelle äußerst besorgt”, 24 November 2021) discusses the November 2021 index value of 96.5 as well as the professional forecast of 96.6. *Die Zeit* (“Geschäftsklimaindex überraschend gestiegen”, 25 January 2022) reports that, contrary to professional forecasts, the January 2022 index value increased by 0.9 points compared to the previous month.

⁸In Section 3, we demonstrate that our findings also hold for alternative specifications where the overlap is more substantial.

visible during the Great Recession and in 2018/2019 when the manufacturing sector cooled down considerably.

Importantly, firms are likely to revise expectations about their own production either because their expectations about the macroeconomy change or because they expect changes in their business conditions due to idiosyncratic developments. Hence, in our analysis below, we control for macro news in order to isolate the effect of micro news which is reflected in the forecast revision. This yields our baseline specification.

In addition, to ensure that forecast revisions are not driven by a macro component, we consider an alternative measure of micro news, which we obtain by purging firms’ forecast revisions of the potential impact of macroeconomic indicators that are observable at the beginning of month t . In this specification, we obtain micro news as the residual of the following regression:

$$FR_t^j = \gamma_j \Gamma_t + \text{micro news}_t^j . \quad (2.4)$$

The vector of macroeconomic indicators Γ_t includes the 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. There are two attractive features of this set-up: i) We only add *observed* changes of the state of the macroeconomy—after correcting for seasonality—in the regression and ii) we run the regressions separately for each firm to allow for firm-specific macro exposure and reactions to the respective changes of macroeconomic states. Panel (d) of Figure 2.1 shows how the cross-sectional dispersion of micro news fluctuates over time. It is largest during the Great Recession, the European debt crisis, and towards the end of our sample period.

Before turning to our main analysis, we verify that macro news impacts forecast revisions significantly. We present results in Table 2.1 for a range of specifications that interact macro news with a number of indicators. Across specifications, we find a significant and positive impact on forecast revisions. The positive sign shows that after receiving positive macro news in the form of a better-than-expected ifo index, firms revise expectations about their own production and business situation upwards as well. This holds across the size distribution of firms, for old and young firms, for firms that have entered the survey more recently and earlier, and for firms where self-reported business-cycle exposure is high and low (see the definition in the table notes). Positive and negative macro news trigger largely symmetric forecast revisions and, last, we find the impact of macro news somewhat stronger during the Great Recession. Generally, however, the economic impact of macro news on forecast revisions is limited. This is in line with our theoretical explanation of a subdued reaction of firms to macro news (see Section 2.4).

2.3 How firm expectations respond to news

In this section, we first introduce our empirical framework, which builds on Coibion and Gorodnichenko (2015). We then report estimates for the average effect of micro and macro news across firms as well as results that account for firm heterogeneity. In addition, we show how the reaction to news is related to real activity. Finally, in Section 2.3.6, we corroborate the results for the ifo survey in the Banca d’Italia’s SIGE.

Table 2.1: 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 < 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		
× Lower macro importance	0.007	0.001
× High macro importance	0.006	0.003
Macro News		
× Positive sign of news	0.012	0.002
× Negative sign of news	0.005	0.001
Macro News		
× outside Great Recession	0.007	0.001
× during Great Recession	0.012	0.002

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 the Great Recession, we rely on a dummy equal to 1 during the years 2007 to 2008 and 0 else. 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. Standard errors are clustered at the firm level.

2.3.1 Empirical framework

Under rational expectations, forecast errors should not be predictable based on information that is available to the forecaster in real time. If one assumes full information in addition to rational expectations, the average forecast error across forecasters should also not be predictable based on average news—a point which Coibion and Gorodnichenko (2015) develop. They test the full-information rational expectations (FIRE) hypothesis based on the following specification:

$$x_{t+h,t} - F_t(x_{t+h,t}) = \beta_0 + \beta_1 \cdot \text{news}_t + \varepsilon_t . \quad (2.5)$$

Here, $x_{t+h,t} - F_t(x_{t+h,t})$ is the average forecast error and news_t is some surprise, typically proxied by the average forecast revisions across forecasters. Under FIRE, we have $\beta_1 = 0$. However, Specification (2.5) is not just simply a test of FIRE. It also points towards specific

alternative models of expectation formation. When positive news tends to be followed by positive forecast errors ($\beta_1 > 0$), the forecast revision turns out to be too weak from an ex-post point of view. Hence, there is an underreaction to news. Conversely, when positive news is on average followed by negative forecast errors ($\beta_1 < 0$), the forecast revision is too strong from an ex-post point of view: There is an overreaction to news.

Earlier work estimates versions of Specification (2.5) using expectations that pertain to macroeconomic outcomes. Coibion and Gorodnichenko (2015), in particular, obtain positive regression coefficients based on the median (consensus) professional forecast for inflation. This result is still consistent with rational expectations: It may simply reflect a failure of the full-information assumption. Yet, and this point is stressed by Coibion and Gorodnichenko (2015), once Specification (2.5) is estimated at the level of individual forecasters, rational expectations imply $\beta_1 = 0$, independently of whether there is full information or not. The key point is that $news_t$ is observed by forecasters in real time. Bordalo et al. (2020) estimate a version of Specification (2.5) based on individual forecasts and find a negative coefficient, that is, they find overreaction to news, rejecting rational expectations, see also Broer and Kohlhas (2023). In sum, once we estimate Specification (2.5) at the level of individual forecasters it provides us with a more stringent test: A test of rational expectations instead of a test of FIRE.

We make three innovations relative to earlier work by estimating a variant of Specification (2.5) on data for individual forecasters. First, we consider firms instead of professional forecasters or households. Second, we focus on firm-level variables, notably production (and prices), rather than macro-level variables (such as aggregate inflation). Last but not least, we distinguish between micro news and macro news regarding firm performance. This distinction takes center stage in our analysis which is based on the following regression equation:

$$x_{t+h,t}^j - F_t^j(x_{t+h,t}^j) = \beta_0 + \beta_1 \cdot \text{micro news}_t^j + \beta_2 \cdot \text{macro news}_t + v_t^j. \quad (2.6)$$

Here, $x_{t+h,t}^j - F_t^j(x_{t+h,t}^j)$ is a firm’s forecast error for its own production defined in Equation (2.1) above. 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 because micro and macro news are part of a firm’s information set, as explained in the previous section. As our baseline, we measure micro news with the forecast revision, defined in Expression (2.3) above, while controlling for macro news, given by the surprise component in the ifo index of the previous month, as in Equation (2.2). Section 2.4 below provides a microfoundation for this specification based on a fully specified structural model. In principle, measurement error may induce a negative correlation between forecast errors and the forecast revisions, a possibility which we consider in Section 2.3.3 below.

2.3.2 Results

To establish our main result, we pool observations across time and firms and estimate Equation (2.6) while allowing for firm-fixed effects. The top panel of Table 2.1 displays the results based on firms’ production expectations. The bottom panel shows results for firm expectations about their business situation which are measured on a quantitative scale. Consider the top panel first. Column (1) on the left reports estimates for a specification

Table 2.1: Over- and underreaction to news

(a) Firms' forecast errors about their production

	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.191*** (0.001)			
Forecast Revision for x_{t+3} net of $\gamma_j\Gamma_t$		-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

(b) Firms' forecast errors about their business situation

	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+6}	-0.441*** (0.004)			
Forecast Revision for x_{t+6} net of $\gamma_j\Gamma_t$		-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 (2.6); observations are pooled across firms, specification includes firm-fixed effects. Panel (a) shows results for the production expectations (3-month horizon, qualitative data), and Panel (b) for the expected business situation (6-month horizon, quantitative data). 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_j (see Section 2.2.4 for more details). All specifications include firm-fixed effects and standard errors clustered at the firm level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

that features forecast revisions and macro news simultaneously. As a result, the forecast revision provides a direct measure of micro news. We find that both types of news induce predictable, statistically significant forecast errors. Hence, we reject rational expectations for firms, consistent with the result of Bordalo et al. (2020) for professional forecasts. In addition, 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 positive forecast errors. This implies, as explained above, that firms overreact to micro news but underreact to macro news. In Section 2.4 below, we offer a

theoretical perspective based on a general-equilibrium model where firms suffer from island illusion.

The remaining columns in the top panel of the table confirm the results reported in Column (1): the micro coefficient remains negative and highly significant when we purge the forecast revision of the impact of real-time macro indicators (second column). The estimate also hardly differs 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 remains virtually unchanged: Column (3). This is to be expected because forecast revisions are purged of the impact of macroeconomic indicators. The macro coefficient remains positive and significant when including only macro news in the regression (fourth column).

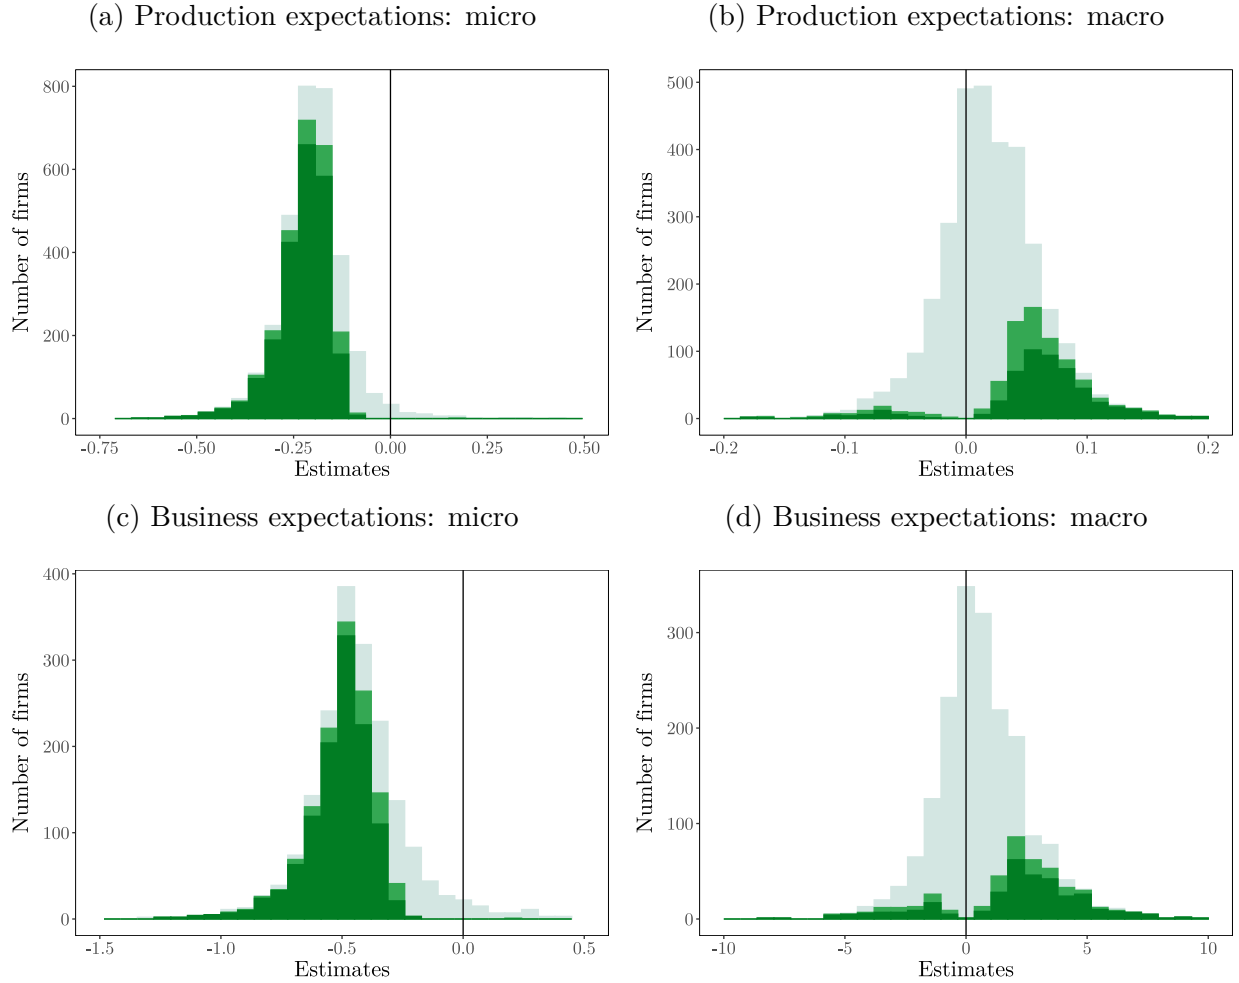
Note that the magnitude of the coefficients in the top panel of Table 2.1 is quantitatively meaningful. In general, the economic importance of the news coefficients is not straightforward to assess due to the qualitative nature of the forecast revisions. However, we may interpret their (relative) importance. Take the specification in Column (2). The average absolute size of micro news is 0.271 and leads to an increase in the absolute value of the forecast error by 0.052 (that is, 0.14 standard deviations of the forecast error). The average absolute size of macro news is 0.971 and leads to an increase in the absolute value of the forecast error by 0.02 (0.05 standard deviations of the forecast error). Hence, the effects on forecast errors are not negligible, and the micro-news effect is about 2-3 times stronger than that of macro news.

The results in Table 2.1 are based on estimates for which we pool observations across firms. But we may exploit the fact that there is a sufficient number of time-series observations for each firm in order to estimate the reaction to news at the level of individual firms. To this end, we re-estimate Specification (2.6) for each of the 3,000 firms in our sample. Throughout, we rely on the forecast revisions purged of macro factors as a measure of micro news and report results in Figure 2.1.⁹ The top panels show the distribution of estimates for β_1 and β_2 based on production expectations. These coefficients capture the response to micro and macro news, respectively. There is a clear pattern: the mass of the estimates for β_1 is concentrated to the left of zero. In fact, as Panel (a) shows, most estimates are significantly smaller than zero (dark green bars). Specifically, for the subset of significant estimates, the micro coefficient is negative for all firms. The estimates for β_2 instead are centered to the right of zero. 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. Overall, our results for the regression which pools observations also hold up when we consider firm-level estimates: the micro coefficient is generally negative while the macro coefficient tends to be positive. In Section 2.3.4 below, we zoom in on how the reactions depend on specific firm characteristics.

A distinct feature of the estimates considered so far is that they are based on qualitative responses of firms: they report whether they expect production to increase, stay the same, or decline. We now turn to a quantitative measure of firm expectations which is also elicited by the ifo survey. It pertains to firms' expected business situation over the next six months and answers are provided in a range from 0 (rather less favorable) to 100 (rather favorable).

⁹As discussed in Section 2.2, our sample includes only firms with at least 30 monthly observations and some variation in their production expectations and forecast errors.

Figure 2.1: Distribution of firm-level responses to news



Notes: Top panels show results for production expectations (3-month horizon, qualitative data), bottom panels 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.

Correspondingly, the survey also asks about the current business situation, with possible answers ranging from 0 (bad) to 100 (good).

We may thus compile forecast errors for the expected business situation over a six-month period, analogously to forecast errors for production expectations.¹⁰ Micro and macro news are measured in the exact same way as above, except that micro news is measured in terms of revisions in business expectations instead of production expectations.

We report results based on firms' business expectations in Panel (b) of Table 2.1 above. As for firm expectations about production reported in Panel (a), we find that firm expectations overreact to micro news but underreact to macro news. Moreover, this holds also across the

¹⁰Link (2020) argues that answers pertain to the level of the expected business situation rather than the change. We report the results of the level interpretation but verify that our results are robust when we consider the alternative interpretation.

alternative specifications in Columns (1) to (4) of the table. This is notable since not only does the nature of responses (qualitative v quantitative) vary across the panels, but also the time horizon (three v six months) and economic concept (production v business situation). With regard to the latter, we note that production expectations are more precisely defined.¹¹ Yet, we also report firm-level estimates based on the business situation in the bottom panels of Figure 2.1 and detect a very similar pattern as in the top panels: when it comes to business expectations, overreaction to micro news is pervasive at the firm level, while firms tend to underreact to macro news.

2.3.3 Measurement error and robustness

In what follows, we show that our results are not likely driven by measurement error, a concern raised by Juodis and Kucinkas (2023) in a related context. In principle, measurement error may indeed induce a mechanical relationship between the forecast revision of period t and the forecast error in period $t + 1$. 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 ε_t^{rep} denote an error term such that the reported expectations amounts to $F_t^{j,rep}(x_{t+h,t}^j) = F_t^j(x_{t+h,t}^j) + \varepsilon_t^{rep}$. The observed forecast error $x_{t+h,t}^j - F_t^j(x_{t+h,t}^j) - \varepsilon_t^{j,rep}$ is then automatically negatively correlated with the reported forecast revision: $FR_t^{j,rep} = FR_t^j + \varepsilon_t^{j,rep} - \varepsilon_{t-1}^{j,rep}$. Hence, taken at face value, measurement error offers an explanation for our results regarding the response to micro news (but not to macro news).¹²

To tackle the issue, we first relate the forecast error in period $t + 1$ to micro news in periods $t - 1$ instead of news in period t . As the first panel of Table 2.2 shows, there is still overreaction to micro news in this case. Second, we consider a fully dynamic specification and regress the forecast error on lagged news in addition to current news.¹³ Specifically, we estimate a model which features 12 lags of both micro and macro news:

$$x_{t+3,t}^j - F_t^j(x_{t+3,t}^j) = \beta_0 + \sum_{p=0}^{12} (\beta_{1,p} \cdot \text{micro news}_{t-p}^j + \beta_{2,p} \cdot \text{macro news}_{t-p}) + \mu_i + v_t^j. \quad (2.7)$$

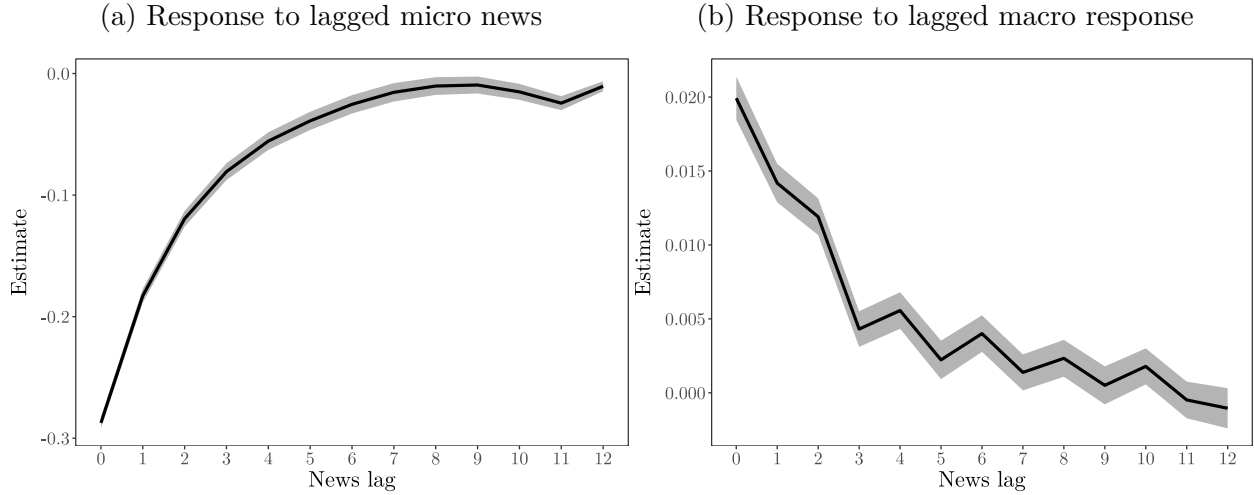
Figure 2.2 displays the results. Note that the overreaction and underreaction is strongest for concurrent news, but it persists over time and declines only gradually. Only after about one year, news ceases to be a cause of forecast errors. This holds both for micro news (left) and macro news (right). This pattern, too, illustrates that our results are not driven by measurement error.

¹¹In addition, the quantitative business situation is only elicited for a subset of firms, starting in September 2005. This accounts for a reduction in the sample size by almost 50 percent.

¹²We note in passing that this kind of measurement error is less of a concern in the case of qualitative data because the answer possibilities of survey participants are limited. Moreover, actual firm decisions are correlated with reported expectations, as we document in Section 2.3.5 below, and the average forecast revision is a leading indicator for changes in manufacturing production, see Section 2.2.4 and Appendix-Figure 2.A.1. This, too, suggests that measurement error is contaminating our data to a negligible extent.

¹³In the context of our analysis, this approach is more suitable than local projections to trace out the effect of news over time because news may be autocorrelated. And indeed, we find that—since micro (macro) news is negatively (positively) autocorrelated—the micro (macro) coefficient on current news is larger (smaller) in this set-up.

Figure 2.2: Response to concurrent and lagged news



Notes: Estimates based on Equation (2.7). Black lines represent point estimates, grey areas correspond to 95% confidence intervals.

That said, the first panel of Table 2.2 provides additional evidence. Turning to the results for the quantitative business situation, we follow Kohlhas and Walther (2021) and exclude outliers of forecast errors and micro news. Again, the estimates show that there is a significant overreaction to micro news, although the estimate is slightly attenuated. We also report estimates that are based on a subsample of observations restricted to firms that revise their qualitative production expectations to zero. In this way, we ensure that the results are not mechanically biased by the qualitative revision scale. The overreaction to micro news is still present. The same holds if we set, in addition, small errors to zero. In the second panel of the table, we report results for a specification in which we again set small forecast errors—potentially driven by measurement error—to zero. We find that results are robust: there is still a significant overreaction to micro news. This also holds when we consider only firms that expect ‘no change’ in production.

In the remainder of the table, we turn to additional robustness tests. So far estimates are based on OLS and the definition of qualitative production forecast errors by Bachmann et al. (2013), see Equation (2.1). The third panel shows that our results also hold when we treat forecast errors qualitatively and use ordered logit rather than OLS for the estimation. Panel 4 reports results for alternative ways to measure macro news. Specifically, we purge firms’ forecast revision by means of time-fixed and time-sector-fixed effects. Again, results are robust to this change.

Lastly, we vary the definition of macro news. We find, in particular, underreaction to 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.

Table 2.2: Alternative specifications

Variation	Details	Micro coeff.	Macro coeff.
1) Micro News (Forecast Revisions)			
Use one month lagged micro news	Table 2.A.2a	-0.021***	0.021***
Business situation (remove outliers)	Table 2.A.2b	-0.387***	0.711***
Use only revisions towards zero	Table 2.A.2c	-0.110***	0.030***
As above and set small errors ($\pm\frac{1}{3}$) to zero	Table 2.A.2d	-0.086***	0.023***
2) Forecast error (Bachmann et al. 2013)			
Set small errors ($\pm\frac{1}{3}$) to zero	Table 2.A.2e	-0.128***	0.018***
Above only for no-change expectations	Table 2.A.2f	-0.192***	0.018***
3) Estimation (OLS)			
Ordered logit	Table 2.A.2g	-1.24***	0.11***
4) Macro component of forecast revision (real-time indicators)			
Fixed effect by time	Table 2.A.2h	-0.194***	0.021***
Fixed effect by time and sector	Table 2.A.2i	-0.196***	0.021***
5) Macro News (surprise component in ifo index)			
Surprise component in manuf. orders	Table 2.A.2j	-0.208***	0.005***
First difference of ifo index	Table 2.A.2k	-0.208***	0.002***
Average forecast revision	Table 2.A.2l	-0.209***	0.345***
Average forecast revision by sector ^a	Table 2.A.2m	-0.211***	0.216***
First difference of stock market index	Table 2.A.2n	-0.208***	0.328***

Notes: Each row corresponds to a variation of the specification for which we report results in Table 2.1, see Appendix 2.A.1 for details. Micro coefficient and Macro coefficient are the estimates on micro and macro news. ^a In this specification, the macro component of forecast revisions is the time and sector average.

*** p<0.01, ** p<0.05, * p<0.1.

2.3.4 Accounting for heterogeneity

Figure 2.1 shows that firms differ in how they react to news. To investigate this more systematically, we zoom in on the determinants of the response to micro and macro news. For this purpose, we re-run the pooled regressions from Table 2.1 while adding interaction terms that capture heterogeneity, both along the cross-sectional and time-series dimensions. We use a Wald test to check if these interaction terms are statistically different from each other. Along the cross-section, we consider the number of employees, firm age, and the duration for which firms participate in the survey. More specifically, for the number of employees, we distinguish between firms in different quartiles; for firm age, we split between firms below 20 years of age and older firms, where a firm's age is measured at the time of the survey based on the year of the reported incorporation; and for the time in the survey, we distinguish between responses submitted during and after the first six months of being in the survey. In addition, we consider heterogeneity regarding the self-reported exposure to the business cycle

Table 2.3: Heterogeneity

Interaction	N	Micro News			Macro News		
		$\hat{\beta}_j$	$SE(\hat{\beta}_j)$	W	$\hat{\beta}_j$	$SE(\hat{\beta}_j)$	W
(1) News Overall (see Table 2.1, (2))	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 \geq 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 \geq 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	
× Medium business-cycle exposure		-0.209***	0.002		0.021***	0.001	
× High business-cycle exposure		-0.208***	0.003		0.022***	0.002	
(6) News	302,737			0.000			0.000
× Positive sign of news		-0.191***	0.002		0.011***	0.001	
× Negative sign of news		-0.232***	0.003		0.035***	0.001	
(7) News	302,737			0.000			0.000
× outside Great Recession		-0.206***	0.001		0.017***	0.001	
× during Great Recession		-0.224***	0.003		0.041***	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, $\hat{\beta}_j$ is the point estimate and $SE(\hat{\beta}_j)$ is 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 from the year of response the year of foundation. For the Great Recession, we rely on a dummy equal to 1 during the years 2007 to 2008 and 0 else. 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]. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

for the firms (see Table 2.A.1 for the wording of the question). Finally, along the time-series dimension, we distinguish between positive and negative news and the period during (outside) the Great Recession.

Table 2.3 displays the results. To facilitate the comparison, we reproduce the results from Table 2.1, Column (2) in the top panel: On average firms overreact to micro news (measured by negative news coefficients) and underreact to macro news (positive news coefficients). We find that this pattern holds across interaction terms. The micro coefficient is robustly negative

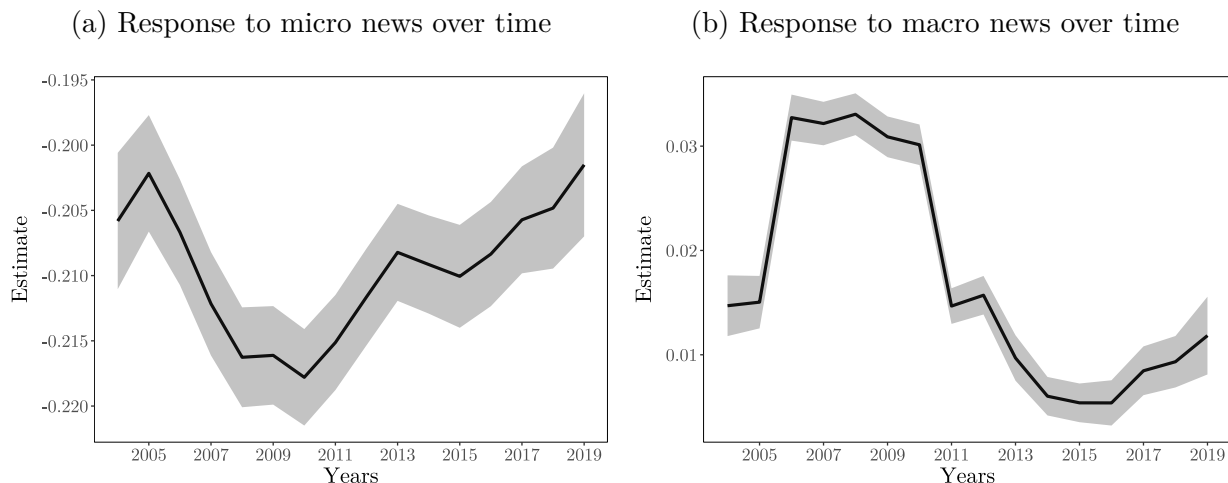
in the cross-section and not significantly different across different levels of firm age, time in the survey, and importance of the business cycle. The overreaction significantly decreases with firm size, but the differences in terms of magnitude are small. This is consistent with the evidence in Panel (a) of Figure 2.1 which shows that the firm-level estimates for β_1 cluster in a fairly tight range. Along the time-series dimension, the micro coefficient is significantly larger for positive news compared to negative news and during the Great Recession compared to other periods.

For the response to macro news, in turn, we find sizeable and significant heterogeneity for firm size, time in the survey, the sign of news, and the Great Recession, again consistent with the more widely distributed estimates of β_2 shown in Panel (b) of Figure 2.1. Looking at firm size (Panel (2) of Table 2.3), the underreaction to macro news is strictly and statistically significantly increasing across employee quartiles. 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 macro economy on the production—and hence the forecast errors—of larger firms. Regarding firm age, reported in Panel (3), there is no statistical difference in the response to macro news between young and old firms. So there is no evidence that firms learn simply by getting older. When comparing the underreaction of firms that recently joined the survey (within six months) to firms with longer tenure, reported in Panel (4), we find evidence for “learning through survey” (Kim and Binder 2023). The underreaction among more tenured firms is about one-third smaller than for firms that recently joined the survey and the difference is statistically significant. This finding is also in line with Massenet and Pettinicchi (2018), who find, for example, that firms’ absolute forecast errors about their own business situation decrease as time since entry in the ifo survey passes. For the exposure to the business cycle, Panel (5), we distinguish between firms that rank the business cycle as very important, important, or less important to them. Here, in line with the heterogeneity by firm size, a high business-cycle exposure is associated with a significantly larger underreaction. Turning to the time-series dimension, we find the underreaction to macro news to be countercyclical. First, the underreaction to negative news is about three times stronger than in the case of positive news, Panel (6), and significantly so. Second, the underreaction is much stronger during the Great Recession, Panel (7), and significantly different from the remaining sample period.

To explore the issue further, we estimate the baseline specification on 5-year rolling windows, following again Coibion and Gorodnichenko (2015). Figure 2.3 shows the results. The left panel shows how the estimated response coefficients for micro news evolve over time, while the right panel does the same for the macro news coefficient. A number of observations are in order. First, firms overreact to micro news and underreact to macro news over the entire sample. Second, the deviations from the rational expectations benchmark are largest during the Great Recession. 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 pattern (in addition to the over- and underreaction to news) conflicts with the notion of rational inattention because one would expect firms to pay more attention to the aggregate economy in times of crisis (see also, Flynn and Sastry 2022). Rather, as argued above, an increased underreaction may simply reflect a stronger impact of macro variables on firm outcomes, without an (sufficiently large) increase in attention.

Finally, we ask what the joint distribution of firm-level response coefficients for micro and

Figure 2.3: Response to news over time



Notes: estimates based on 5-year rolling windows. Black lines represent point estimates, grey areas correspond to 95% confidence intervals.

macro news looks like. To this end, we relate the firm-level estimates of micro and macro news (illustrated in Figure 2.1). Figure 2.A.2 in the appendix displays a binned scatterplot between the micro and macro news coefficients. Indeed, we find a negative relationship that is especially strong if we zoom into the subsample of firms with significant overreaction to micro news and underreaction to macro news ($\rho = -0.35$). Hence, the stronger the underreaction to macro news of a given firm, the stronger is also the overreaction to micro news.

In sum, overreaction to micro news and underreaction to macro news is a robust and pervasive phenomenon—across firms and states of the world.

2.3.5 Reaction to news and firm performance

Expectations matter for firm decisions and firm outcomes, as Enders et al. (2022) establish specifically for the ifo data set. Against this background, we investigate whether over- and underreaction to news is related to measures of firm performance in a systematic way. We will then revisit this evidence in light of our theoretical model below. Specifically, we relate the estimated response coefficients for each firm to their profits, their production volatility, and forecast error volatility. We rely on the firm-level estimates discussed in Section 2.3.2 above and restrict the sample to firms that overreact to micro news and underreact to macro news, in line with the aggregate findings.

Since 2009, the ifo Business Climate Survey includes a quantitative question about the profits in the current year.¹⁴ For each firm, we calculate the average profits and regress them on the micro and macro news coefficients estimated in Section 2.3.2. In addition, we absorb sector- and size-fixed effects. Columns (1) and (2) in Table 2.4 display the results. A stronger

¹⁴Profits are elicited in May and September. We rely on the September wave to capture a larger information set. In addition, we subtract the yearly average profits to ensure that the results are not confounded by heterogeneity over time (in a recession, profits are lower and underreaction stronger, see Section 2.3.4).

Table 2.4: Over- and underreaction to news and real activity

	mean _i (return on sales _{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 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 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-level 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.

overreaction to micro news is associated with a significant decrease in average profits, while a stronger underreaction to macro news is not significantly related to the average profits. In terms of magnitude, a one standard deviation increase in the overreaction to micro news leads to a reduction in profits by on average about 0.14 percentage points.

As a second exercise, we calculate the standard deviation of realized production changes as a proxy for firm-level production volatility. Then, we follow the procedure above and regress it on the estimated response coefficients to micro and macro news, obtained in Section 2.3.2. Columns (3) and (4) in Table 2.4 display the results. The estimates indicate 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 case of macro news. Projecting these cross-sectional estimates on the macro level implies higher micro-level volatility in the presence of over- and underreactions. This is a potential explanation for the high observed idiosyncratic volatility of firm outcome variables (Bachmann et al. 2013; Bloom 2009).

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 2.4 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.

2.3.6 Further evidence for Italian firms

We now turn to an alternative survey of firm expectations in order to assess to what extent our results generalize beyond the ifo survey of German firms. Specifically, we rely on the quarterly “Survey on Inflation and Growth Expectations” (SIGE) operated by the Banca d’Italia, which has also been used by, for example, Coibion et al. (2020). Two features of the SIGE are particularly noteworthy in the context of our analysis. First, it elicits answers in the form of growth rates and, as such, answers are quantitative. Second, it asks firms about their price expectations: not only about their own prices but also about aggregate price developments, that is, inflation.¹⁵

Mimicking our earlier strategy for the ifo survey as closely as possible, we estimate a version of Specification (2.6) on data from the SIGE. Instead of production expectations, we now consider firms’ price expectations: We compute, consistent with the definition of the forecast error in Expression (2.1) above, the one-year-ahead expectation error for firms’ own prices in quarter t by subtracting the expected change reported in quarter t from the actual change, as reported in quarter $t + 4$.

We measure macro news as the surprise component of inflation: we subtract the (average) professional forecast submitted to Consensus Economics up until a month before the publication from the realized inflation rate. To measure micro news, we again rely on forecast revisions, here the first-difference of firms’ expectations about their own prices. As firm expectations are for a twelve-month fixed forecast horizon, the overlap in quarterly forecast revisions is nine months. Since, as above, we include macro news in the regression, the forecast revisions for firms’ own prices allow us to directly estimate the effect of micro news on the forecast error. In an alternative specification, we purge the forecast revision of the change in CPI inflation. Importantly, both news and the change in CPI inflation are in the firm’s information set as the survey question about expected inflation provides firms with the current inflation rate in every quarter.¹⁶

Table 2.5 reports the results. In the first two columns, we proceed in the same fashion as with the data from the ifo survey. Micro news is the forecast revision for a firm’s own prices, both raw and net of aggregate developments. Macro news is the surprise component in the aggregate inflation rate. In line with our findings for the ifo survey, the coefficients for micro news are negative and those for macro news are positive. Both are highly significant.

The third column moves beyond the setup for the ifo survey. Here, we exploit the fact that the SIGE also polls firm expectations about inflation. This allows us to compile firm-specific macro news, namely the forecast revisions of the firm’s aggregate inflation expectations. Also for this specification, coefficients for micro news are negative, those for macro news are positive, and both significantly so.¹⁷ The last three columns then show that these results are

¹⁵For further details on the SIGE, see Appendix 2.A.2 and Grasso and Ropele (2018).

¹⁶See Table 2.A.3 in Section 2.A.2 for the exact wording. For the timing, consider Summer 2022 as an example. On June 13, Consensus Economics polled professional forecasters about their expectations for the inflation rate in the second quarter and published the results on June 16. The Banca d’Italia published the inflation rate on July 8. We use the difference between the realized value and the average professional forecast as a measure of macro news in 2022Q3. Importantly, the SIGE in 2022Q3 ran between August 25 and September 15 and firms are explicitly informed about the current rate of inflation. Macro news is therefore in their information set.

¹⁷In Appendix 2.A.2, Table 2.A.4a shows that this also holds in univariate regressions including either

Table 2.5: Over- and underreaction to news—Italian firms

	Forecast error about firm's own prices					
	(1)	(2)	(3)	(4)	(5)	(6)
Micro News						
Forecast Revision for π_{t+12}^i	-0.478*** (0.022)		-0.457*** (0.020)	-0.405*** (0.016)		-0.376*** (0.013)
FR for π_{t+12}^i net of $\Delta\pi_t$		-0.502*** (0.020)			-0.340*** (0.024)	
Macro News						
Surprise component of π_{t-1}	4.113*** (0.356)	3.758*** (0.470)		2.735*** (0.195)	2.642*** (0.239)	
FR for π_{t+12}			0.242*** (0.058)			0.210*** (0.031)
Drop top and bottom 1%	no	no	no	yes	yes	yes
Observations	21,707	14,030	29,471	21,073	13,610	28,492
R^2	0.103	0.116	0.094	0.074	0.054	0.056
Within R^2	0.127	0.127	0.110	0.097	0.061	0.078

Notes: Regressing firms' forecast errors about their own prices on micro news and macro news. For each type of news, we consider two alternative definitions. For micro news, we consider firms' own forecast revisions for their own prices in their raw form, as well as revisions purged from changes in aggregate inflation for each firm with at least 20 observations. For macro news, we consider the surprise component of inflation in the previous quarter. More specifically, we subtract from the realized value the (mean) professional forecast from Consensus Economics. Alternatively, we consider firms' own forecast revisions about aggregate inflation. Columns (1) to (3) use the full sample, while columns (4) to (6) drop the top and bottom 1% of forecast errors and forecast revisions from the full sample. The sample starts in 2002 (2013 for inflation surprises) and ends in 2022. Firm-fixed effects are always included and standard errors are clustered at the firm level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

robust to dropping the top and bottom 1% of forecast errors and forecast revisions. This again speaks against plain measurement error as a driver of our empirical results. We estimate specification (2.6) at the firm level as well. The resulting coefficients are distributed in a similar way as those for the ifo survey; see Figure 2.A.3 in Section 2.A.2. Looking at the joint distribution of significant micro and macro coefficients, we find that they are negatively and significantly correlated ($\rho = -0.24$), also in line with the results for the ifo survey.

Overall, our results based on the SIGE show that overreaction to micro news and underreaction to macro news is a pertinent feature of firms' expectation formation process. It is not limited to the ifo survey of German firms but also characterizes the expectation formation of Italian firms. This is particularly noteworthy because the SIGE differs from the ifo survey along a number of important dimensions.

micro or macro news. In Table 2.A.4b, we consider as a fourth possible definition for macro news the forecast revision computed by subtracting from the current six-month-ahead inflation expectation the one-year-ahead expectation six months ago, where we also find positive response coefficients for macro news and negative coefficients for micro news, that are both highly significant.

2.4 A model of island illusion

In the following, we develop a stylized model in order to rationalize the evidence established above. Specifically, the model provides a microfoundation for our empirical Specification (2.6) and allows us to establish conditions under which firm expectations overreact to micro news and underreact to macro news. Two aspects set our model apart from related theoretical work, some of which we reference in the introduction above. First, our focus is on expectations about a firm’s own performance and how these, in turn, are shaped by micro and macro news. To represent these news and their interaction in a consistent manner, we need to specify a full-fledged general equilibrium model. Second, the distinct feature of our model is that firms suffer from ‘island illusion’. As a result, firms systematically underestimate the importance of aggregate developments for their own performance. This appears plausible to the extent that for firms firm-specific developments are more salient of economic performance—consistent with findings according to which direct experience impacts (risk) perceptions more strongly than outcomes experienced by others (Smith et al. 2001; Viscusi and Zeckhauser 2015). It is also in line with our results in Section 2.2, which show that firms’ reaction to aggregate news is statistically significant but economically limited.

Our setup relates to Bordalo et al. (2020) where news is overly representative for forecasters and thus triggers an overreaction. Our model, however, accounts for simultaneous over- and underreaction to different types of news at the level of individual forecasters. What sets our model apart from the model of overconfidence put forward by Broer and Kohlhas (2023) is a general-equilibrium perspective that accounts for the cross-equation restrictions regarding the impact of micro and macro news.¹⁸

Formally, we build on the model with dispersed and noisy information put forward by Lorenzoni (2009). We depart from the original model in two ways. First, we assume firms are subject to island illusion. Second, we simplify the original model by assuming predetermined rather than staggered prices in order to solve an approximate model in closed form and to derive analytical results. In what follows, we first describe the structure of the economy, including technology and preferences. Afterward, we specify expectations and policy and present our main result regarding over- and underreaction.

2.4.1 Setup and timing

There is a continuum of islands, indexed by $r \in [0, 1]$, each populated by a representative household and a unit mass of firms, indexed by $j \in [0, 1]$. Each household buys from a subset of all islands, chosen randomly in each period. Specifically, it buys from all firms on n islands included in the set \mathcal{B}_t^r , with $1 < n < \infty$.¹⁹ Households have an infinite planning horizon. Firms manufacture differentiated goods on the basis of island-specific productivity, which is simultaneously driven by a permanent, economy-wide component and a temporary,

¹⁸In related work, 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, however, that asymmetric attention may arise in a fully rational framework.

¹⁹This assumption ensures that households cannot exactly infer aggregate productivity from observed prices. At the same time, individual firms have no impact on the price of households’ consumption baskets.

idiosyncratic component.²⁰ Household-specific demand also features an aggregate and an idiosyncratic stochastic component such that we can write in general terms:

$$\vartheta_t^r = \sqrt{\varpi_\vartheta} \vartheta_t' + \sqrt{1 - \varpi_\vartheta} \bar{\vartheta}_t^{r'} . \quad (2.8)$$

Here ϑ_t^r is either technology a_t^r of a firm on island r or demand q_t^r of the household on the same island, while ϑ_t' and $\bar{\vartheta}_t^{r'}$ are the aggregate and idiosyncratic components, respectively. Both are i.i.d. random variables. The weight ϖ_ϑ determines the importance of aggregate relative to idiosyncratic shocks. Relation (2.8) implies $Var(\vartheta_t^r) = Var(\vartheta_t') = Var(\bar{\vartheta}_t^{r'})$, such that total volatility is divided between the aggregate contribution $\varpi_\vartheta Var(\vartheta_t^r)$ and the idiosyncratic contribution $(1 - \varpi_\vartheta) Var(\vartheta_t^r)$.

The timing of events is as follows: Financial markets are complete such that, assuming identical initial positions, wealth levels of households are equalized at the beginning of each period. Each period consists of three stages. During stage 1 of period t , information about all variables of period $t-1$ is released. Subsequently, nominal wages are determined and the central bank sets the interest rate based on expected inflation.

The aggregate and idiosyncratic components of productivity materialize in the second stage. Concerning technology, firms only observe their own productivity (micro news). Additionally, a noisy public signal about the aggregate demand shock is released to firms and households, based on, say, market research (macro news). Given these information sets, firms set prices.

During the third and final stage, households split up. Workers work for all firms on their island, while consumers allocate their expenditures across differentiated goods based on public information and information reflected in the prices of the goods they purchase. Additionally, individual demand shocks influence their consumption decisions. Because the common productivity component is permanent, demand shocks are purely temporary, and households' wealth and information are equalized in the next period, agents expect the economy to settle on a new steady state from period $t+1$ onward.

2.4.2 Households

A representative household on island r (“household r ”, for short) maximizes lifetime utility

$$U_t^r = E_{t|3}^r \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left(Q_\tau^r \ln C_\tau^r - \frac{(L_\tau^r)^{1+\varphi}}{1+\varphi} \right) \quad \varphi \geq 0, \quad 0 < \beta < 1,$$

where $E_{t|3}^r$ is the expectation operator based on household r 's information set at the time of its consumption decision in stage 3 of period t (see below), while C_t^r denotes the consumption basket of household r . L_t^r is its total labor supply, which aggregates labor the household supplies to individual firms j on island r , $L_t^{j,r}$. As described in Equation (2.8), the demand shock Q_t^r consists of an aggregate and an island-specific component. In linearized form with lower-case letters denoting percentage deviations from steady state, this implies

$$q_t^r = \sqrt{\varpi_q} q_t' + \sqrt{1 - \varpi_q} \bar{q}_t^{r'} \equiv q_t + \bar{q}_t^r,$$

²⁰As argued by Lorenzoni (2009), this setup can account for the empirical observations that the firm-level volatility of productivity is large relative to aggregate volatility and that individual expectations are dispersed.

with $q_t = \sqrt{\varpi_q} q'_t$ and $\bar{q}_t^r = \sqrt{1 - \varpi_q} \bar{q}'_t{}^r$, where q'_t and $\bar{q}'_t{}^r$ are i.i.d. shocks with mean zero and variance $Var(q'_t) = Var(\bar{q}'_t{}^r) = Var(q_t)$. While actual demand, including the shocks, realizes only in stage 3 of the period, a public signal about the (weighted) aggregate component is released to firms and households in the second stage, representing macro news:

$$s_t = q_t + e_t,$$

where e_t is an i.i.d. noise shock with variance σ_e^2 and mean zero. The ratio between the volatility of idiosyncratic demand $Var(q'_t)$ and the volatility $Var(s_t)$ of the signal, which are both observable, is defined as $\bar{v} \equiv Var(q'_t)/Var(s_t)$.

The flow budget constraint of the household is given by

$$E_{t|1} \varrho_{t,t+1}^r \Theta_t^r + B_t^r + \sum_{m \in \mathcal{B}_t^r} \int_0^1 P_t^{j,m,r} C_t^{j,m,r} dj \leq \int_0^1 \Pi_t^{j,r} dj + W_t^r L_t^r + \Theta_{t-1}^r + (1 + r_{t-1}) B_{t-1}^r,$$

where $C_t^{j,m,r}$ denotes the amount bought by household r from firm j on island m and $P_t^{j,m,r}$ is the price for one unit of $C_t^{j,m,r}$. At the beginning of the period, the household receives the payoff Θ_{t-1}^r , given a portfolio of state-contingent securities purchased in the previous period. $\Pi_t^{j,r}$ are the profits of firm j on island r and $\varrho_{t,t+1}^r$ is household r 's stochastic discount factor between t and $t+1$. The period- t portfolio is priced conditional on the (common) information set of stage 1, hence we apply the expectation operator $E_{t|1}$. B_t^r are state non-contingent bonds paying an interest rate of r_t . The complete set of state-contingent securities is traded in the first stage of the period, while state-non-contingent bonds can be traded via the central bank throughout the entire period. The interest rate of the non-contingent bond is set by the central bank. All financial assets are in zero net supply. The bundle C_t^r of goods purchased by household r consists of goods sold in a subset of all islands in the economy²¹

$$C_t^r = \left(\frac{1}{n} \sum_{m \in \mathcal{B}_t^r} \int_0^1 (C_t^{j,m,r})^{\frac{\gamma-1}{\gamma}} dj \right)^{\frac{\gamma}{\gamma-1}} \quad \gamma > 1.$$

While each household purchases a different random set of goods, we assume that all households visit the same number of islands n . The price index of household r is therefore

$$P_t^r = \left(\frac{1}{n} \sum_{m \in \mathcal{B}_t^r} \int_0^1 (P_t^{j,m,r})^{1-\gamma} dj \right)^{\frac{1}{1-\gamma}}.$$

2.4.3 Firms

Firm j on island r produces according to the following production function

$$Y_t^j = A_t^r (L_t^j)^\alpha \quad 0 < \alpha < 1,$$

²¹See, e.g., Enders (2020) for a more detailed treatment of a consumption bundle consisting of a finite number of goods.

featuring labor supplied by the local household as the sole input. $A_t^r = A_t^{j,r}$ denotes the productivity level of firm j , which is the same for all firms on island r .²² During stage 2, the firm sets the optimal price for the current period, conditional on the expectation about the third stage of period t , specified below. Given prices, the level of production is determined by demand during stage 3. Since each island is visited by n consumers, total demand of firm j on island r is given, in linearized form, by

$$q_t^{r,j} = q_t + \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{\bar{q}_t^m}{n}.$$

Log-productivity on each island a_t^r depends on last period's aggregate technology x_{t-1} , an aggregate shock, and an island-specific shock:

$$a_t^r - x_{t-1} = \sqrt{\varpi_a} a_t' + \sqrt{1 - \varpi_a} \bar{a}_t'^r \equiv \varepsilon_t + \eta_t^r,$$

with $\varepsilon_t = \sqrt{\varpi_a} a_t'$ and $\eta_t^r = \sqrt{1 - \varpi_a} \bar{a}_t'^r$, where a_t' and $\bar{a}_t'^r$ are i.i.d. shocks with mean zero and variance $Var(\bar{a}_t'^r) = Var(a_t') = Var(a_t^r - x_{t-1})$. The shock a_t' (and therefore also η_t^r) aggregates to zero across all islands. Idiosyncratic productivity thus contains private information (micro news) about the aggregate level of technology x_t , which follows a random walk

$$\Delta x_t = \sqrt{\varpi_a} a_t' \equiv \varepsilon_t.$$

Firms only observe productivity on their own island a_t^r .

2.4.4 Island illusion

We now turn to the details of the expectation-formation process. To set island illusion apart from rational expectations, we first specify the rational forecasts.

Firms. The rational forecast for Δx_t is given by

$$\bar{E}_{t|2}^j \Delta x_t = \bar{\delta}_x^p (a_t^r - x_{t-1}),$$

where $\bar{E}_{t|2}^j$ is the rational expectation of firm j on island r when setting prices (in stage 2). The coefficient $\bar{\delta}_x^p$ is a function of the structural parameters that capture the informational friction. It is non-negative and smaller than unity:

$$\bar{\delta}_x^p = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\eta^2} = \varpi_a. \quad (2.9)$$

The rational forecast for q_t is given by

$$\bar{E}_{t|2}^j q_t = \bar{\rho}_q^p s_t, \quad \text{with} \quad \bar{\rho}_q^p = \frac{\sigma_q^2}{\sigma_q^2 + \sigma_e^2} = \varpi_q \bar{v}.$$

²²Note that from here on, with a slight abuse of notation, we drop, where the result is unambiguous, the island index r for firm-specific variables in the main text to simplify the expressions: $E_{t|s}^j \equiv E_{t|s}^{j,r}$; $Y_t^j \equiv Y_t^{j,r}$, $L_t^j \equiv L_t^{j,r}$, etc.

Rather than assuming that all expectations are formed in a rational way, however, we suppose that firms are subject to island illusion. Specifically, we assume that firms underestimate the importance of aggregate developments, relative to idiosyncratic developments. Put differently, firms think that their own technology and the demand for their product are driven to a smaller extent by aggregate developments compared to what they would believe under rational expectations. In our setup, island illusion is governed by a single parameter Υ which downweighs the importance of the aggregate component relative to the actual weight:

$$\hat{\varpi}_\vartheta = \Upsilon \varpi_\vartheta.$$

Here $\hat{\varpi}_\vartheta$ is the weight ϖ_ϑ as perceived by firms and Υ measures the degree of the bias. If $\Upsilon = 1$, firms weigh the importance of both components correctly, while $\Upsilon < 1$ reflects island illusion (and $\Upsilon > 1$ the hypothetical case of ‘continent illusion’).²³

Thus, actual firm expectations are formed according to

$$E_{t|2}^j \Delta x_t = \delta_x^p (a_t^r - x_{t-1}) \quad E_{t|2}^j q_t = \rho_q^p s_t,$$

with

$$\begin{aligned} \delta_x^p &= \hat{\varpi}_a = \Upsilon \varpi_a < \varpi_a = \bar{\delta}_x^p \\ \rho_q^p &= \hat{\varpi}_q \bar{v} = \Upsilon \varpi_q \bar{v} < \varpi_q \bar{v} = \bar{\rho}_q^p. \end{aligned}$$

Consumers. Regarding consumers, we assume that they form rational expectations in the following way. While shopping during stage 3, they observe a set of prices. They can hence infer the productivity level of each firm in their sample:

$$E_{t|3}^r \Delta x_t = \delta_x^h \tilde{a}_t^r,$$

where \tilde{a}_t^r is the average over the realizations of $a_t^m - x_{t-1}$ for each island m in household r ’s sample \mathcal{B}_t^r . δ_x^h is equal across households and given in Appendix 2.A.3. Consumers have complete information if $n \rightarrow \infty$. Furthermore, households rationally incorporate the information contained in the public signal concerning the aggregate demand shock into their expectations of the aggregate price level, see Appendix 2.A.3. Note that our results regarding the effects of island illusion on the side of the firms are not affected by a potential bias in the expectation formation process of households, as long as firms have a correct understanding of households’ average reaction to news.

²³The crucial point is that agents misjudge the *relative* contribution of both components to productivity or demand. That is if σ_ε^2 or σ_q^2 is under- or overestimated, agents would still not display a bias if they under- or overestimate σ_η^2 or σ_e^2 by the same degree (i.e., the ‘signal-to-noise ratio’ is correctly assessed, see equations (2.31) and (2.32) in Appendix 2.A.4). Similarly, models of rational inattention assume that agents perceive certain information with noise. Given, however, that they know about this imperfect perception, they have a correct understanding of the signal-to-noise ratio and therefore do not display a bias: Υ would be unity.

2.4.5 Monetary policy and market clearing

The central bank follows an interest-rate feedback rule but sets r_t before observing prices, that is during stage 1 of period t in linearized form:

$$r_t = \psi E_{t|1}^{cb} \pi_t + \nu_t \quad \psi > 1,$$

where π_t is economy-wide net inflation, calculated on the basis of all goods sold in the economy. The expectation operator $E_{t|1}^{cb}$ is conditional on the information set of the central bank. This set consists of information from period $t-1$ only, that is, the central bank enjoys no informational advantage over the private sector.²⁴ ν_t is a monetary policy shock which we include in the model as an example of shock that is observable by firms and households alike.

Goods and labor markets clear in each period:

$$\int_0^1 C_t^{j,m,r} dr = Y_t^{j,m} \quad \forall j, m \quad L_t^r = \int_0^1 L_t^{j,r} dj \quad \forall r,$$

where $C_t^{j,m,r} = 0$ if household r does not visit island m . The asset market clears in accordance with Walras' law.

2.4.6 Accounting for over- and underreaction

In order to account for the evidence presented in Section 2.3 above, we derive a solution of the model based on a linear approximation to the equilibrium conditions around the symmetric steady state; see Appendix 2.A.3 for details. We first define forecast errors and forecast revisions in the model to provide an explicit microfoundation for our empirical specification.

To map the model to the data, we interpret the intra-period stages of a generic period t as the relevant time units. In what follows we thus drop the time subscript t and index variables only with the stages which define the information flow and the decision-making process within a period t . We can write the forecast error of firm j as follows: $y_3^j - E_2^j(y^j)$, that is, firm j 's actual output in stage 3 relative to its forecast in stage 2. We define the forecast revision accordingly as $FR_2^j = E_2^j(y^j) - E_1^j(y^j)$, that is, the change in the forecast of the same firm between stage 1 and stage 2. This revision reflects the response of firm expectations to the private and the public signal, s , which is common to all firms. Armed with these definitions, we can derive our main result (see Appendix 2.A.4 for the proof):

Proposition 1. *Consider the regression*

$$y_3^j - E_2^j(y^j) = \beta_1 FR_2^j + \beta_2 s_2 + \omega^j, \quad (2.10)$$

where all subscripts refer to different stages of a generic period t . FR^j is the forecast revision of firm j , s is the macro news common to all firms, and ω^j represents a potential error term. In the case of island illusion, that is, for $\Upsilon < 1$, we obtain

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

²⁴Pre-set prices and interest rates allow us to discard the noisy signals about quantities and inflation observed by firms and the central bank in Lorenzoni (2009), simplifying the signal-extraction problem without changing the qualitative predictions of the model. Pre-set wages, on the other hand, guarantee the determinacy of the price level. They do not affect output dynamics after noise and technology shocks, because goods prices may still adjust in the second stage of the period.

where β_1 measures the firm's reaction to micro news and β_2 the reaction to macro news.

Equation (2.10) is the counterpart to our empirical specification (2.6) and thus provides an explicit microfoundation for our empirical analysis. Moreover, Proposition 1 establishes stringent conditions under which our empirical results can be rationalized: In the presence of island illusion, that is, whenever $\Upsilon < 1$, the model predicts simultaneous overreaction to private signals and underreaction to public information by individual firms—based on a single parameter that captures the departure from rational expectations.

Intuitively, in a rational-expectations framework, individual future forecast errors cannot be predicted by current forecast revisions ($\beta_1 = 0$) or public signals ($\beta_2 = 0$), as firms could otherwise easily improve on their forecasts.²⁵ However, given that in our model firms suffer from island illusion and therefore underestimate the importance of aggregate developments, they place too little weight on the private signal ($\delta_x^p < \bar{\delta}_x^p$) when revising their forecast of aggregate technology, relative to the rational-expectations benchmark. Hence, on average, firms attribute too little of a positive surprise in their own technology to a change in aggregate technology. Put differently, after a successful technological innovation at their own firm, managers underestimate the potential of competitors to engineer a similar reduction in prices. Hence they overestimate how much their own production will change, yielding $\beta_1 < 0$.²⁶

Regarding the effect of the public signal on firms' forecast errors, firms also underestimate the role of aggregate developments. That is, they deem aggregate demand disturbances q_t to fluctuate less than they actually do. At the same time, they correctly observe the volatility of the signal, such that they overestimate the contribution of noise to the signal. Consequently, they pay less attention to the signal than under the rational-expectations benchmark ($\rho_x^p < \bar{\rho}_x^p$). Following a positive signal, they hence underestimate the increase in demand for their own and their competitors' products. Hence, firms expect their own demand and the prices of competitors to be lower than what they, on average, turn out to be after a positive signal and, therefore, underestimate their own output, such that $\beta_2 > 0$.

The model allows us to derive a number of additional predictions which conform well with the pattern in the data. We discuss them in turn. As before, proofs are found in Appendix 2.A.4.

Proposition 2. *A higher degree of island illusion (a lower Υ) implies*

- (a) *A stronger overreaction to micro news (a lower β_1) and simultaneously a larger underreaction to the public signal (a larger β_2).*
- (b) *Lower expected profits.*
- (c) *A larger variance of the firm-specific forecast error.*

²⁵To be precise, $\beta_1 = \beta_2 = 0$ as long as agents have a correct estimate of the relative variances of the components of the signals, see the proof of Proposition 1 and Footnote 23.

²⁶In general equilibrium, there are two, partly offsetting effects: On the one hand, firms expect prices of competitors to be on average higher than what they will actually turn out, increasing expected demand for the firms' products. On the other hand, firms expect overall demand to be lower than warranted, reducing expected idiosyncratic demand as well. Overall, the first effect dominates, and firms on average overestimate their future sales after having observed a negative surprise in idiosyncratic technology.

Intuitively, if firms underestimate aggregate developments, they, as explained above, underestimate the information content of the public signal and simultaneously overestimate their technological advantage in case of positive developments in their idiosyncratic technology—which corresponds to the evidence in Figure 2.A.2 in the appendix. Given that the optimal forecast (that achieves an expected forecast error of zero, seen from an econometrician’s view) obtains for $\Upsilon = 1$, any deviations lead to biased forecasts in the profit maximization problem of the firm and hence lower expected profits. Likewise, it raises the forecast-error variance. These predictions are in line with our findings in Section 2.3.5 above.

Proposition 3. *For a given degree of island illusion Υ , a higher business-cycle exposure (a higher ϖ_q) leads to a larger underreaction to macro news (a larger β_2).*

For firms that are more exposed to aggregate demand conditions, island illusion matters more, inducing a stronger underreaction. Intuitively, if demand for a firm’s products is entirely idiosyncratic ($\varpi_q = 0$), island illusion does not play any role as it biases the estimated $\hat{\omega}_q$ towards zero. For those firms, being on an island is no illusion but reality.

2.5 Conclusion


How do expectations adjust to news about the economy? We address this question while zooming in on firms’ expectations about their own performance. This focus sets our study apart from earlier work, as does the distinction between micro and macro news. Analyzing firm surveys from Germany and Italy, we find robustly that firm expectations overreact to micro news and underreact to macro news. We estimate at the level of individual firms and provide detailed evidence which suggests that our results are not driven by measurement error. This allows us to reject rational expectations. But since our estimates show that firm expectations— independent of firm characteristics—respond in a systematically different way to micro and macro news, they directly inform attempts that move beyond rational expectations in modeling the expectation-formation process.

The last part of the paper represents such an attempt. Here we put forward a stylized general equilibrium model which assumes that firms suffer from island illusion: They perceive what’s happening to them as less common than it actually is. We think of island illusion as an instance of salience. In the model, it is governed by a single parameter, representing a disciplined departure from rational expectations. The model provides microfoundation to our empirical specification and shows that island illusion can simultaneously account for overreaction to micro news and underreaction to macro news. Assessing further the validity of island illusion in other contexts of expectation formation seems a promising avenue for future research.

2.A Appendices


2.A.1 Additional figures and tables

Figure 2.A.1: Average forecast revisions and production growth



Notes: The 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 2.A.2: 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 based on the subsample 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 subsample of firms with *significant* negative micro news coefficients and *significant* positive macro news coefficients ($\rho = -0.35$). The firm-level estimates are also displayed in Figure 2.1.

Table 2.A.1: 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.

Table 2.A.2: Alternative specifications

(a) Expectations: use lagged micro news

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.191*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-0.021*** (0.001)	-0.020*** (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	280,583	280,583	302,737
R ²	0.16260	0.09452	0.08988	0.08967
Within R ²	0.08471	0.00580	0.00071	0.00498
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except we use one month lagged micro news. Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

(b) Business Situation: remove outliers (p1, p99)

	Firms' forecast errors about their business situation			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+6}	-0.394*** (0.004)			
Forecast Revision for x_{t+6} net of $\beta_i\Gamma_t$		-0.387*** (0.005)	-0.383*** (0.005)	
Macro News				
Surprise component of the ifo index	0.760*** (0.039)	0.711*** (0.039)		0.615*** (0.040)
Observations	147,226	147,409	147,409	150,166
R ²	0.29231	0.28251	0.27954	0.24130
Within R ²	0.06037	0.04779	0.04384	0.00287
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (b) except that we remove the top and bottom one percent of forecast errors, revisions, and micro news. Firm-fixed effects are always included and standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.2: Alternative specifications, continued.

(c) Expectations: only forecast revisions towards zero

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.091*** (0.003)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except that we only use observations where firms revise their expectations towards zero. Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

(d) Expectations: only forecast revisions towards zero and set small errors to zero

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.072*** (0.002)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except that we only use observations where firms revise their expectations towards zero and set small forecast errors ($\pm\frac{1}{3}$) to zero. Firm-fixed effects are always included and standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.2: Alternative specifications, continued.

(e) Forecast error: set small errors to zero

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.115*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except small forecast errors ($\pm\frac{1}{3}$) are set to zero. Firm-fixed effects are always included and standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

(f) Forecast error: set small errors to zero and no change expected

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.176*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except small forecast errors ($\pm\frac{1}{3}$) are set to zero when expectations are zero. Firm-fixed effects are always included and standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.2: Alternative specifications, continued.

(g) 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.

(h) Micro news: absorb macro comp. of forecast revision with time-fixed effect

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}		-0.191*** (0.001)		
Forecast Revision for x_{t+3} net of $\beta_i \Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except we absorb the macro component from forecast revisions by means of time-fixed effects (see Section 2.2). Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.2: Alternative specifications, continued.

(i) 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 x_{t+3}	-0.191*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except we absorb the macro component from forecast revisions by means of time-sector-fixed effects (see Section 2.2). Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

(j) Macro news: manufacturing orders rather than ifo index

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.190*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except macro news are constructed from the median professional forecast of manufacturing orders. Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.2: Alternative specifications, continued.

(k) 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 x_{t+3}	-0.190*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except macro news is constructed with the first difference of the ifo index. Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

(l) Macro news: average forecast revisions rather than ifo index

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for x_{t+3}	-0.194*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except macro news is constructed with average production forecast revisions. Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.2: Alternative specifications, continued.

(m) 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 x_{t+3}	-0.196*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except macro news is constructed with average production forecast revisions for each sector. Firm-fixed effects are always included and standard errors are clustered at firm level.
 *** p<0.01, ** p<0.05, * p<0.1.

(n) 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 x_{t+3}	-0.190*** (0.001)			
Forecast Revision for x_{t+3} net of $\beta_i\Gamma_t$		-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
Firm FE	✓	✓	✓	✓

Notes: Set-up as in Table 2.1 Panel (a) except macro news is constructed with the first difference of the German stock market index DAX. Firm-fixed effects are always included and standard errors are clustered at firm level.

*** p<0.01, ** p<0.05, * p<0.1.

2.A.2 SIGE Data

The “Survey on Inflation and Growth Expectations” (SIGE) is a quarterly business survey launched in 1999. Until 2011 it features roughly 500 firms per quarter, 1,000 firms between 2011 and 2019, and more than 1,500 since 2021. The median firm responds for 7 quarters and 20 percent of firms respond for more than 23 quarters.²⁷ The questions relevant to our purposes are listed in Table 2.A.3. These questions elicit growth rates in percentage points. The wording of Q3 about expected inflation ensures that firms receive the most recent inflation rates in Italy and the euro area.

Table 2.A.3: Relevant questions from SIGE

Label	Name	Introduced	Wording
Q1	realized change in own prices	2002q4	In the last 12 months, what has been the average change in your firm’s prices?
Q2	expected change in own price	1994q4	For the next 12 months, what do you expect will be the average change in your firm’s prices?
Q3	expected inflation (12 months ahead)	1994q4	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?
Q4	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?

Notes: Wording taken from the September 2022 questionnaire. Starting in 2012q3 alternative wordings for expected inflation (Q3) were used for randomly selected firms. We focus on the traditional wording including information about recent inflation. This wording is shown to 60 percent of the sample.

²⁷For more details on the SIGE, see Grasso and Ropele (2018) and Coibion et al. (2020).

Table 2.A.4: Additional regression results from the SIGE

(a) Univariate regressions

	Forecast error about firm's own prices				
	(1)	(2)	(3)	(4)	(5)
Micro News, firm-level purging	-0.477*** (0.018)				
Micro News, pooled purging		-0.461*** (0.017)			
Micro News, time-fe purging			-0.472*** (0.017)		
Macro News, inflation surprise				3.419*** (0.339)	
Macro News, forecast revision					0.212** (0.083)
Observations	28,561	38,048	38,048	25,420	28,928
R ²	0.0977	0.0904	0.0955	0.0105	0.0025
Within R ²	0.1039	0.1086	0.1128	0.0069	0.0009
Firm FE	✓	✓	✓	✓	✓


Notes: Regressing firms' forecast errors about their own prices on micro news and macro news separately. Micro news is based on firms' forecast revisions (*FR*) net of changes in the aggregate economy. For firm-level purging in column (1), we regress, for each firm separately, *FRs* on the first difference of the inflation rate and use the residuals as micro news. For pooled purging in (2), we run the same regression, but pool observations across firms. For time-fe purging in (3), we regress *FRs* on time-fixed effects and use the residual as micro news. For macro news, we consider two definitions. For inflation surprises in (4), we use the surprise component in the inflation rate of the previous quarter as measured by the difference between the realized rate and the (mean) professional forecast from Consensus Economics. Alternatively, we also consider as macro news, firms' own forecast revisions for 12-month-ahead inflation in columns (5).

(b) Regressions using actual forecast revisions

	Forecast error about firm's own prices		
	(1)	(2)	(3)
Micro News, firm-level purging	-0.495*** (0.023)		
Micro News, pooled purging		-0.454*** (0.027)	
Micro News, time-fe purging			-0.463*** (0.027)
Macro News, forecast revision (6m - L6.12m)	0.431*** (0.142)	0.422*** (0.113)	0.415*** (0.112)
Observations	11,312	14,998	14,998
R ²	0.1024	0.0821	0.0865
Within R ²	0.1189	0.1064	0.1097
Firm FE	✓	✓	✓

Notes: Regressing firms' forecast errors about their own prices on micro news and macro news. Micro news are as defined above. For macro news, we consider the forecast revision computed by subtracting from the current six-months-ahead inflation expectation (Q4 in Table 2.A.3) the twelve-months-ahead inflation expectation six months ago (Q3).

Figure 2.A.3: Firm-level regressions – univariate distribution of news coefficients



Notes: Forecast errors are for firms' own prices. Micro news is a firm's own forecast revision for their own prices purged from changes in inflation. Macro news is the surprise component of inflation in the previous quarter. We run the regression separately for each firm with at least 20 observations. Grey area represents insignificant estimates, light green area represents estimates significant at the 10 % level, dark green area indicates significance at the 5 % level.

2.A.3 Model solution

In Appendix 2.A.4, we provide the proofs for the propositions in Section 2.4. Before that, we outline the model solution and key equilibrium relationships. Throughout, we consider a linear approximation to the equilibrium conditions of the model. Lower-case letters indicate percentage deviations from steady state. We solve the model by backward induction. That is, we start by deriving expectations regarding period $t + 1$. Using the result in the Euler equation of the third stage of period t allows us to determine price-setting decisions during stage 2. Eventually, we obtain the responses of forecasts and realizations to unexpected changes in productivity or the public signal.

Expectations regarding period $t + 1$. Below, E_t^k stands for either $E_{t|2}^{j,r}$, referring to the information set of producer j on island r at the time of her pricing decision, or for $E_{t|3}^r$, referring to the information set of the household on island r at the time of its consumption decision. Variables with only time subscripts refer to economy-wide values. The wage in period $t + 1$ is set according to the expected aggregate labor supply

$$E_t^k \varphi l_{t+1} = E_t^k (w_{t+1} - p_{t+1} - c_{t+1}).$$

This equation is combined with the aggregated production function

$$E_t^k y_{t+1} = E_t^k (x_{t+1} + \alpha l_{t+1}),$$

the expected aggregate labor demand

$$E_t^k (w_{t+1} - p_{t+1}) = E_t^k [x_{t+1} + (\alpha - 1)l_{t+1}],$$

and market clearing $y_{t+1} = c_{t+1}$ to obtain

$$E_t^k x_{t+1} = E_t^k y_{t+1} = E_t^k c_{t+1}. \tag{2.11}$$

Furthermore, the expected Euler equation, together with the Taylor rule, is

$$E_t^k c_{t+1} = E_t^k (c_{t+2} + \pi_{t+2} - \psi \pi_{t+1}).$$

Agents expect the economy to be in a new steady state tomorrow ($E_t^k c_{t+1} = E_t^k c_{t+2}$), given the absence of state variables other than technology, which follows a unit root process, and the demand shock, whose expected value is zero. Ruling out explosive paths yields

$$E_t^k \pi_{t+2} = E_t^k \pi_{t+1} = 0.$$

Stage 3 of period t . After prices are set, each household observes n prices in the economy. Since only productivity is idiosyncratic to firms at the time of setting prices, the productivity level $a_t^{j,r} = a_t^r$ —which is the same for all producers $j \in [0, 1]$ on island r —can be inferred

from each price $p_t^{j,r}$ of the good from producer j on island r . Hence, household r forms its expectations about the change in aggregate productivity according to

$$E_{t|3}^r \Delta x_t = \delta_x^h \hat{a}_t^r,$$

where \hat{a}_t^r is the average over the realizations of $a_{m,t} - x_{t-1}$ for each location m in household r 's sample \mathcal{B}_t^r . The coefficients δ_x^h is equal across households and depend on n, σ_ε^2 , and σ_η^2 in the following way:

$$\delta_x^h = \frac{\sigma_\varepsilon^2}{\underbrace{\sigma_\varepsilon^2 + \sigma_\eta^2/n}_{\rightarrow 1 \text{ if } n \rightarrow \infty}}. \quad (2.12)$$

Furthermore, households rationally incorporate the information contained in the public signal concerning the aggregate demand shock into their expectations of the aggregate price level.

The expectation formation of producers is discussed in the main text. Consumption follows an Euler equation with household-specific inflation, as only a subset of goods is bought. Agents expect no differences between households for $t + 1$, such that expected aggregate productivity and the overall price level impact today's individual consumption. Additionally using $E_{t|3}^r p_{t+1} = E_{t|3}^r p_t$ and $E_{t|3}^r x_{t+1} = E_{t|3}^r x_t$ gives

$$c_t^r = E_{t|3}^r x_t + E_{t|3}^r p_t - p_t^r - r_t + q_t^r. \quad (2.13)$$

Similar to the updating formula for technology estimates, households use all relevant available information to form an estimate about the aggregate price level p_t according to

$$E_{t|3}^r p_t = \delta_p^h \hat{a}_t^r + \kappa_p^h w_t + \tau_p^h x_{t-1} - \eta_p^h r_t + \bar{\rho}_p^h s_t + \bar{\delta}_p^h q_t^r, \quad (2.14)$$

where the undetermined coefficients $\delta_p^h, \kappa_p^h, \tau_p^h, \eta_p^h, \bar{\rho}_p^h$, and $\bar{\delta}_p^h$ represent the impact of the relevant variable on the expected price level. Combining the above gives

$$c_t^r = (1 + \tau_p^h) x_{t-1} + \delta_{xp}^h \hat{a}_t^r + \kappa_p^h w_t - (1 + \eta_p^h) r_t - p_t^r + \bar{\rho}_p^h s_t + (1 + \bar{\delta}_p^h) q_t^r \quad (2.15)$$

where $\delta_{xp}^h = \delta_x^h + \delta_p^h$. We will solve for the coefficients below. Total demand for good j on island r is

$$\begin{aligned} y_t^{j,r} &= -\gamma p_t^{j,r} + \gamma \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{p_t^m}{n} + \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{c_{m,t}}{n} \\ &= -\gamma p_t^{j,r} + \gamma \tilde{p}_t^r + \tilde{y}_t^r, \end{aligned} \quad (2.16)$$

where \tilde{y}_t^r is the average consumption level of customers visiting island r , $1/n$ th of which equals $p_t^{j,r}$. The index \tilde{p}_t^r is the average price index of customers visiting island r . If customers bought on all (that is, infinitely many) islands in the economy, \tilde{p}_t^r would correspond to the overall price level. Given (2.15), we have, with $\kappa^h = (1 + \tau_p^h) x_{t-1} - (1 + \eta_p^h) r_t + \kappa_p^h w_t$,

$$\begin{aligned} \tilde{y}_t^r &= \frac{1}{n} \sum_{\{m|r \in \mathcal{B}_t^m\}} [E_t^m x_t + E_t^m p_t - p_t^m - r_t + q_t^m] \\ &= \kappa^h + \delta_{xp}^h \sum_{m \in \mathcal{B}_t^r} \frac{\hat{a}_t^m}{n} - \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{p_t^m}{n} + (1 + \bar{\delta}_p^h) \left(q_t + \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{\bar{q}_t^m}{n} \right) + \bar{\rho}_p^h s_t. \end{aligned} \quad (2.17)$$

Stage 2 of period t . During the second stage, firms obtain idiosyncratic signals about their productivity. Firms set prices according to

$$\begin{aligned} p_t^{j,r} &= w_t + \frac{1-\alpha}{\alpha} E_{t|2}^{j,r} y_t^{j,r} - \frac{1}{\alpha} a_t^r \\ &\equiv k' + k'_1 E_{t|2}^{j,r} \tilde{p}_t^r + k'_2 E_{t|2}^{j,r} \tilde{y}_t^r - k'_3 a_t^r, \end{aligned}$$

with

$$k' = \frac{\alpha}{\alpha + \gamma(1-\alpha)} w_t \quad k'_1 = \frac{\gamma(1-\alpha)}{\alpha + \gamma(1-\alpha)} \quad k'_2 = \frac{1-\alpha}{\alpha + \gamma(1-\alpha)} \quad k'_3 = \frac{1}{\alpha + \gamma(1-\alpha)}. \quad (2.18)$$

From here onward, expressions that are based on common knowledge only (such as k') are treated like parameters in notation terms, i.e., they lack a time index. This facilitates the important distinction between expressions that are common information and those that are not. Evaluating the expectation of firm j about island-specific demand in period t , using (2.17), results in

$$E_{t|2}^{j,r} \tilde{y}_t^r = \kappa^h + \delta_{xp}^h \left(\frac{1}{n} (a_t^r - x_{t-1}) + \frac{n-1}{n} E_{t|2}^{j,r} \varepsilon_t \right) - \left(\frac{1}{n} p_t^{j,r} + \frac{n-1}{n} E_{t|2}^{j,r} p_t \right) + \left[(1 + \bar{\delta}_p^h) \rho_q^p + \bar{\rho}_p^h \right] s_t. \quad (2.19)$$

where ρ_q^p is the coefficient used by producers to form expectations about the aggregate demand shock based on the signal, and κ^h contains only publicly available information. Furthermore, it is taken into account that the productivity and prices of island r have a non-zero weight in the sample of productivity and price levels observed by consumers visiting island r . Note that producers still take the price index of the consumers as given, since they buy infinitely many goods on the same island.

Inserting the firm expectation (2.19) into the pricing equation (2.18) yields (here, p_t is the average of the prices charged by producers of all other islands, which is the overall price index)

$$p_t^{j,r} \equiv k + k_1 E_{t|2}^{j,r} p_t - k_3 a_t^r + k_4 s_t,$$

with

$$\begin{aligned} \Xi &= 1 - \frac{1}{n} (k'_1 - k'_2) & k &= \frac{1}{\Xi} \left\{ k' + k'_2 \kappa^h + \frac{k'_2 \delta_{xp}^h}{n} [(n-1)(1-\delta_x^p) - 1] x_{t-1} \right\} \\ k_1 &= \frac{n-1}{n\Xi} (k'_1 - k'_2) & k_3 &= \frac{1}{\Xi} \left\{ k'_3 - \frac{k'_2 \delta_{xp}^h}{n} [(n-1)\delta_x^p + 1] \right\} & k_4 &= \frac{k'_2}{\Xi} \left[(1 + \bar{\delta}_p^h) \rho_q^p + \bar{\rho}_p^h \right]. \end{aligned} \quad (2.20)$$

Note that, according to (2.18), $0 < k'_1 - k'_2 < 1$ because $0 < \alpha < 1$ and $\gamma > 1$. Using the definition of k_1 in (2.20), this implies (observe that $n > 1$)

$$0 < k_1 < 1.$$

Aggregating over all producers gives the aggregate price index

$$p_t = k + k_1 \bar{E}_t p_t - k_3 x_t + k_4 s_t,$$

where $\int a_t^r dr = x_t$, and $\bar{E}_t p_t = \iint E_{t|2}^{j,r} p_t dj dr$ is the average expectation of the price level.

The expectation of firm j of this aggregate is therefore

$$\begin{aligned} E_{t|2}^{j,r} p_t &= k - k_3 E_{t|2}^{j,r} x_t + k_1 E_{t|2}^{j,r} \bar{E}_t p_t + k_4 s_t \\ &= k - k_3 \delta_x^p a_t^r - k_3 (1 - \delta_x^p) x_{t-1} + k_1 E_{t|2}^{j,r} \bar{E}_{t|2} p_t + k_4 s_t. \end{aligned} \quad (2.21)$$

Inserting the last equation into (2.20) gives

$$p_t^{j,r} = k + k_1 k - k_1 k_3 (1 - \delta_x^p) x_{t-1} - (k_3 + k_1 k_3 \delta_x^p) a_t^r + k_1^2 E_{t|2}^{j,r} \bar{E}_{t|2} p_t + (k_4 + k_1 k_4) s_t.$$

To find $E_{t|2}^{j,r} \bar{E}_t p_t$, note that firm j 's expectations of the average of (2.21) are

$$E_{t|2}^{j,r} \bar{E}_t p_t = k - k_3 (1 - \delta_x^p) (1 + \delta_x^p) x_{t-1} - k_3 \delta_x^{p2} a_t^r + k_1 E_{t|2}^{j,r} \bar{E}_{t|2}^{(2)} p_t + k_4 s_t.$$

where $\bar{E}^{(2)}$ is the average expectation of the average expectation. The price of firm j is found by plugging the last equation into the second-to-last:

$$\begin{aligned} p_t^{j,r} &= k + k_1 k + k_1^2 k - \left[k_1 k_3 (1 - \delta_x^p) + k_1^2 k_3 (1 - \delta_x^p) (1 + \delta_x^p) \right] x_{t-1} \\ &\quad - \left[k_3 (1 + k_1 \delta_x^p) + k_1^2 k_3 \delta_x^{p2} \right] a_t^r + [k_4 + k_1 k_4 + k_1^2 k_4] s_t + k_1^3 E_{t|2}^{j,r} \bar{E}_{t|2}^{(2)} p_t. \end{aligned}$$

Continuing like this results in some infinite sums

$$\begin{aligned} p_t^{j,r} &= k \left(1 + k_1 + k_1^2 + k_1^3 \dots \right) \\ &\quad - k_1 k_3 (1 - \delta_x^p) \left[1 + k_1 (1 + \delta_x^p) + k_1^2 (1 + \delta_x^p + \delta_x^{p2}) + k_1^3 (1 + \delta_x^p + \delta_x^{p2} + \delta_x^{p3} \dots) \right] x_{t-1} \\ &\quad - k_3 \left(1 + k_1 \delta_x^p + k_1^2 \delta_x^{p2} + k_1^3 \delta_x^{p3} \dots \right) a_t^r + \left[k_4 + k_1 k_4 + k_1^2 k_4 + k_1^3 k_4 + \dots \right] s_t \\ &\quad + k_1^\infty E_{t|2}^{j,r} \bar{E}_{t|2}^{(\infty)} p_t. \end{aligned}$$

This leads to

$$p_t^{j,r} = \frac{k}{1 - k_1} - \frac{k_1 (1 - \delta_x^p)}{1 - k_1} \frac{k_3}{1 - k_1 \delta_x^p} x_{t-1} - \frac{k_3}{1 - k_1 \delta_x^p} a_t^r + \frac{1}{1 - k_1} k_4 s_t + \underbrace{k_1^\infty \bar{E}_{t|2}^{(\infty)}}_{\rightarrow 0} p_t$$

or

$$p_t^{j,r} = \bar{k}_1 + \bar{k}_3 a_t^r + \bar{k}_4 s_t. \quad (2.22)$$

with

$$\bar{k}_1 = \frac{1}{1 - k_1} \left[k - (1 - \delta_x^p) \frac{k_1 k_3}{1 - k_1 \delta_x^p} x_{t-1} \right] \quad \bar{k}_3 = - \frac{k_3}{1 - k_1 \delta_x^p} \quad \bar{k}_4 = \frac{1}{1 - k_1} k_4.$$

The average over all producers yields the aggregate price index as

$$p_t \equiv \bar{k}_1 + \bar{k}_3 x_t + \bar{k}_4 s_t. \quad (2.23)$$

To arrive at qualitative predictions for the impact of the structural shocks ε_t and q_t on output growth and the forecast error, we need to determine the sign and the size of \bar{k}_3 . Note that, according to (2.20),

$$-k_3 = \delta_{xp}^h \frac{k_2' - nk_3'/\delta_{xp}^h + k_2'(n-1)\delta_x^p}{n - (k_1' - k_2')},$$

where the first part of the numerator can be rewritten, by observing (2.18), as

$$k_2' - nk_3'/\delta_{xp}^h = \frac{1 - n/\delta_{xp}^h - \alpha}{\alpha + \gamma(1 - \alpha)}.$$

Using (2.18) and (2.20) thus yields

$$-k_3 = \delta_{xp}^h \frac{(1 - \alpha)[(n - 1)\delta_x^p + 1] - n/\delta_{xp}^h}{(n - 1)[\alpha + \gamma(1 - \alpha)] + 1}.$$

Plugging this into the definition of \bar{k}_3 in (2.23) gives

$$\bar{k}_3 = \delta_{xp}^h \left\{ \frac{(1 - \alpha)[(n - 1)\delta_x^p + 1] - n/\delta_{xp}^h}{(n - 1)[\alpha + \gamma(1 - \alpha)] + 1} \right\} \left\{ 1 - \delta_x^p \frac{(n - 1)(\gamma - 1)(1 - \alpha)}{(n - 1)[\alpha + \gamma(1 - \alpha)] + 1} \right\}^{-1}.$$

To obtain $\delta_{xp}^h = \delta_x^h + \delta_p^h$, we need to find the undetermined coefficients of equation (2.14). Start by comparing this equation with household r 's expectation of equation (2.23):

$$E_{t|3}^r p_t = \underbrace{\bar{k}_1 + \bar{k}_3 x_{t-1}}_{\kappa_p^h w_t + \tau_p^h x_{t-1} - \eta_p^h r_t} + \underbrace{\bar{k}_3 \delta_x^h}_{\delta_p^h} \hat{a}_t^r + \underbrace{\bar{k}_4}_{\bar{\rho}_p^h} s_t, \quad (2.24)$$

with $\bar{\delta}_p^h = 0$, since the household knows that price-setters only have the public signal regarding demand, but not any information about actual demand. Hence, $\delta_{xp}^h = \delta_x^h(1 + \bar{k}_3)$. Inserting this into the above expression for \bar{k}_3 yields

$$\bar{k}_3 \equiv - \frac{n/\Sigma - \delta_x^h \Psi}{\Phi - \delta_x^h \Psi}, \quad (2.25)$$

with

$$\begin{aligned} \Sigma &= (n - 1)[\alpha + \gamma(1 - \alpha)] + 1 > 0 & \Psi &= (1 - \alpha)[(n - 1)\delta_x^p + 1]/\Sigma > 0 \\ \Phi &= 1 - \delta_x^p(n - 1)(\gamma - 1)(1 - \alpha)/\Sigma. \end{aligned}$$

The signs obtain because $n > 1$, $0 < \alpha < 1$, $\delta_x^p > 0$, and $\gamma > 1$. Observe that $\Psi\Sigma < n$ because $\delta_x^p \leq 1$. Hence, $n/\Sigma - \delta_x^h \Psi > 0$ because

$$n - \underbrace{\delta_x^h}_{>0, <1} \underbrace{\Psi\Sigma}_{<n} > 0,$$

implying that the numerator of (2.25) is positive. Turning to the denominator $\Phi - \delta_x^h \Psi$, note that $\Phi - \Psi > 0$. The denominator of (2.25) is therefore positive as well, and we have $\bar{k}_3 < 0$. Next, consider that $n/\Sigma < \Phi$ and we obtain

$$-1 < \bar{k}_3 < 0.$$

This is a key result for the derivation of the propositions in Appendix 2.A.4.

We now turn to \bar{k}_4 . First, observe that

$$\begin{aligned}\Xi &= 1 - \frac{1}{n}(k'_1 - k'_2) \\ &= \frac{[(n-1)\gamma + 1](1-\alpha) + n\alpha}{n[\alpha + \gamma(1-\alpha)]} > 0\end{aligned}$$

and

$$k_1 = \frac{(n-1)\varepsilon(1-\alpha) + (n-1)\alpha + 1 - n}{(n-1)\varepsilon(1-\alpha) + (n-1)\alpha + 1} < 1.$$

Thus,

$$\begin{aligned}\bar{k}_4 &= \frac{1}{1-k_1} \frac{k'_2}{\Xi} [\bar{k}_4 + \rho_q^p] \\ &= \frac{k'_2}{(1-k_1)\Xi - k'_2} \rho_q^p.\end{aligned}$$

Since $k'_2 > 0$, for $\bar{k}_4 > 0$, we need to show that

$$(1-k_1)\Xi > k'_2$$

or

$$n\alpha^2 > -\alpha(1-\alpha)[(n-1)\gamma + 1],$$

which is true, such that $\bar{k}_4 > 0$.

Stage 1 of period t As information sets of agents are perfectly aligned during stage 1, we use the expectation operator $E_{t|1}$ to denote (common) stage-one expectations in what follows. Combining the results regarding expectations about inflation in period $t+1$ with the Euler equation, the Taylor rule, and the random-walk assumption for x_t gives, see equation (2.13),

$$E_{t|1}c_t = E_{t|1}y_t = E_{t|1}x_t + (1-\psi)E_{t|1}\pi_t + E_{t|1}q_t.$$

Remember that the monetary policy shock emerges after wages are set. Its expected value before wage-setting is zero, just like the expected value of the demand shock, as the signal is not yet released. Labor supply is given by

$$\varphi E_{t|1}l_t = E_{t|1}(w_t - p_t - c_t + q_t).$$

This equation can be combined with the aggregated production function

$$E_{t|1}y_t = E_{t|1}(x_t + \alpha l_t),$$

the expected aggregate labor demand

$$E_{t|1}(w_t - p_t) = E_{t|1}[x_t + (\alpha - 1)l_t],$$

and market clearing $y_t = c_t$ to obtain

$$\varphi E_{t|1}l_t = E_{t|1}(x_t + (\alpha - 1)l_t - c_t) + q_t$$

or

$$E_{t|1}y_t = E_{t|1}x_t.$$

Comparing this expression to the Euler equation, we get

$$E_{t|1}\pi_t = 0.$$

Nominal wages are set in line with these expectations. We thus have determinacy of the price level. The central bank then sets its interest rate based on expected inflation.

2.A.4 Proofs

Proof of Proposition 1 Calculating the expectation error of firms for idiosyncratic output, using demand equation (2.16), the island-specific demand (2.17), and the price-level equation (2.23), yields

$$\begin{aligned} FE_{t+1}^{j,r} &= \Delta y_t^{j,r} - E_{t|2}^{j,r} \Delta y_t^{j,r} = \gamma \frac{n-1}{n} (p_t - E_{t|2}^{j,r} p_t) + \tilde{y}_t^r - E_{t|2}^{j,r} \tilde{y}_t^r \\ &= \frac{n-1}{n} [(\gamma - 1)\bar{k}_3 + \delta_x^h(1 + \bar{k}_3)] (\varepsilon_t - E_{t|2}^{j,r} \varepsilon_t) + q_t - E_{t|2}^{j,r} q_t + \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{\bar{q}_t^k}{n} \\ &\equiv \Lambda (\varepsilon_t - E_{t|2}^{j,r} \varepsilon_t) + q_t - E_{t|2}^{j,r} q_t + \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{\bar{q}_t^k}{n}, \end{aligned} \quad (2.26)$$

where the Euler equations (2.15) of customers of island r is used in the second equation. The effect Λ of the expectation error regarding aggregate technology innovations, $\varepsilon_t - E_{t|2}^{j,r} \varepsilon_t$, on the expectation error regarding own output is negative if

$$\gamma - 1 > -\delta_x^h \frac{1 + \bar{k}_3}{\bar{k}_3}. \quad (2.27)$$

Since

$$-\frac{1 + \bar{k}_3}{\bar{k}_3} = \frac{(n-1)(1-\alpha)(\gamma-1)(1-\delta_x^p)}{n - \delta_x^h(1-\alpha)[(n-1)\delta_x^p + 1]},$$

inequality (2.27) is fulfilled if

$$1 > \delta_x^h(1-\alpha),$$

which is correct, such that $\Lambda < 0$. The gap between expected own and aggregate output can be calculated using (2.16), (2.19), (2.22), and (2.23):

$$\begin{aligned} E_{t|2}^{j,r} y_t^{j,r} - E_{t|2}^{j,r} y_t &= -\gamma \frac{n-1}{n} (p_t^{j,r} - E_{t|2}^{j,r} p_t) + E_{t|2}^{j,r} \tilde{y}_t^r - E_{t|2}^{j,r} y_t \\ &= \frac{1}{n} \left[-\gamma(n-1)\bar{k}_3 + \delta_x^h(1 + \bar{k}_3) - \bar{k}_3 \right] E_{t|2}^{j,r} \eta_t^r \equiv K_1 E_{t|2}^{j,r} \eta_t^r. \end{aligned} \quad (2.28)$$

Aggregating individual Euler equations (2.13) over all individuals, using (2.23), and (2.24) gives aggregate output as

$$y_t = E_{t|3}^r x_t + E_{t|3}^r p_t - p_t - r_t + q_t = x_{t-1} + \underbrace{\left[\delta_x^h - \bar{k}_3(1 - \delta_x^h) \right]}_{>0} \varepsilon_t + q_t - \underbrace{\frac{\alpha}{\alpha + \psi(1 - \alpha)}}_{<0} \nu_t.$$

Note that, if households have full information ($n \rightarrow \infty$), we get $\delta_x^h \rightarrow 1$ and $y_t = x_t - \nu_t \alpha / (\alpha + \psi(1 - \alpha))$. The signs indicated above result from $0 < -\bar{k}_3 < 1$ (derived above).

Forecast revisions are then given by the change in expectations between before and after receiving the private and public signals (that is, between stage one and stage two). The last equation implies

$$E_{t|2}^{j,r} y_t - x_{t-1} = \left[\delta_x^h - \bar{k}_3(1 - \delta_x^h) \right] E_{t|2}^{j,r} \varepsilon_t + \rho_q^p s_t - \frac{\alpha}{\alpha + \psi(1 - \alpha)} \nu_t.$$

Using this equation together with equation (2.28) in the forecast revision gives

$$\begin{aligned} FR_t^{j,r} &= E_{t|2}^{j,r} (y_t^{j,r} - y_{t-1}^{j,r}) - E_t (y_t^{j,r} - y_{t-1}^{j,r}) = E_{t|2}^{j,r} y_t^{j,r} - E_{t|2}^{j,r} y_t + E_{t|2}^{j,r} y_t - E_t y_t \\ &= K_1 E_{t|2}^{j,r} \eta_t^r + \left[\delta_x^h - \bar{k}_3(1 - \delta_x^h) \right] E_{t|2}^{j,r} \varepsilon_t + \rho_q^p s_t - \frac{\alpha}{\alpha + \psi(1 - \alpha)} \nu_t. \end{aligned}$$

Since

$$E_{t|2}^{j,r} \varepsilon_t = \delta_x^p (\varepsilon_t + \eta_t^r) \quad E_{t|2}^{j,r} \eta_t^r = (1 - \delta_x^p) (\varepsilon_t + \eta_t^r) \quad (2.29)$$

we can write the above as

$$\begin{aligned} FR_t^{j,r} &= K_1 (1 - \delta_x^p) (\varepsilon_t + \eta_t^r) + \left[\delta_x^h - \bar{k}_3(1 - \delta_x^h) \right] \delta_x^p (\varepsilon_t + \eta_t^r) + \rho_q^p s_t - \frac{\alpha}{\alpha + \psi(1 - \alpha)} \nu_t \\ &\equiv X_1 \varepsilon_t + X_1 \eta_t^r + X_1^q q_t + X_1^e e_t + K_\nu \nu_t. \end{aligned}$$

with

$$X_1 = K_1 (1 - \delta_x^p) + \left[\delta_x^h - \bar{k}_3(1 - \delta_x^h) \right] \delta_x^p \quad X_1^q = \rho_q^p \quad K_\nu = -\frac{\alpha}{\alpha + \psi(1 - \alpha)}.$$

Similarly, making use of (2.29), the forecast error (2.26) can be written as

$$FE_{t+1}^{j,r} = \Lambda \left[(1 - \delta_x^p) \varepsilon_t - \delta_x^p \eta_t^r \right] + (1 - \rho_q^p) q_t - \rho_q^p e_t - \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{\bar{q}_t^k}{n}. \quad (2.30)$$

The sign of β_1 of regression (2.10) can then be determined in two steps. Since both independent variables, forecast revisions and the signal, are correlated, we first regress forecast revisions on the signal, yielding the regression coefficient

$$Coe f_1 = \frac{Cov(FR_t^{j,r}, s_t)}{Var(s_t)} = \frac{X_1^q \sigma_q^2 + X_1^q \sigma_e^2}{\sigma_q^2 + \sigma_e^2} = X_1^q.$$

The residual of this regression can therefore be written as $FR_t^{j,r} - Coe f_1 s_t$. The sign of the coefficient β_1 of regression (2.10) then depends on the sign of

$$\begin{aligned} Cov(FE_{t+1}^{j,r}; FR_t^{j,r} - Coe f_1 s_t) &= Cov(FE_{t+1}^{j,r}; FR_t^{j,r}) - Coe f_1 Cov(FE_{t+1}^{j,r}, s_t) \\ &= \underbrace{(X_1^q - Coe f_1)}_{=0} R_e^q + \underbrace{\Lambda X_1}_{<0} \underbrace{R_\eta}_{>0} < 0, \end{aligned}$$

with

$$R_e^q = (1 - \rho_q^p) \sigma_q^2 - \rho_q^p \sigma_{e,q}^2 \quad R_\eta = (1 - \delta_x^p) \sigma_\varepsilon^2 - \delta_x^p \sigma_\eta^2.$$

The signs obtain from $\Lambda < 0$ and

$$K_1 = \frac{1}{n} \left[-\gamma(n-1) \bar{k}_3 + \delta_x^h (1 + \bar{k}_3) - \bar{k}_3 \right] > 0 \quad X_1 = K_1 (1 - \delta_x^p) + \left[\delta_x^h - \bar{k}_3 (1 - \delta_x^h) \right] \delta_x^p > 0,$$

as well as

$$R_\eta > 0 \quad \text{if} \quad \frac{\hat{\sigma}_\eta^2}{\hat{\sigma}_\varepsilon^2} > \frac{\sigma_\eta^2}{\sigma_\varepsilon^2}, \quad (2.31)$$

that is

$$R_\eta > 0 \quad \text{if} \quad \frac{1 - \Upsilon \varpi_a}{\Upsilon \varpi_a} > \frac{1 - \varpi_a}{\varpi_a},$$

which results from the assumption of island illusion, $\Upsilon < 1$. Hence, $\beta_1 < 0$.

The sign of the coefficient β_2 of regression (2.10) can equivalently be derived by first regressing the forecast revision on the signal, which gives the coefficient

$$Coe f_2 = \frac{Cov(FR_t^{j,r}, s_t)}{Var(FR_t^{j,r})} = \frac{X_1^q \sigma_q^2 + X_1^q \sigma_e^2}{X_1^2 \sigma_\varepsilon^2 + X_1^2 \sigma_\eta^2 + (X_1^q)^2 \sigma_q^2 + (X_1^q)^2 \sigma_e^2 + (K_\nu)^2 \sigma_\nu^2},$$

which is positive since $X_1^q > 0$. The sign of β_2 in regression (2.10) then depends on the sign of

$$\begin{aligned} Cov(FE_{t+1}^{j,r}; s_t - Coe f_2 (FR_t^{j,r})) &= Cov(FE_{t+1}^{j,r}; s_t^q) - Coe f_2 Cov(FE_{t+1}^{j,r}, FR_t^{j,r}) \\ &= \underbrace{(1 - Coe f_2 X_1^q)}_{>0} \underbrace{R_e^q}_{>0} - \underbrace{Coe f_2}_{<0} \underbrace{\Lambda X_1}_{<0} R_\eta. \end{aligned}$$

The signs obtain because

$$1 - Coe f_2 X_1^q = \frac{X_1^2 \sigma_\varepsilon^2 + X_1^2 \sigma_\eta^2 + (K_\nu)^2 \sigma_\nu^2}{X_1^2 \sigma_\varepsilon^2 + X_1^2 \sigma_\eta^2 + (X_1^q)^2 \sigma_q^2 + (X_1^q)^2 \sigma_e^2 + (K_\nu)^2 \sigma_\nu^2},$$

which is positive but smaller than unity, and

$$R_e^q > 0 \quad \text{if} \quad \frac{\hat{\sigma}_e^2}{\hat{\sigma}_q^2} > \frac{\sigma_e^2}{\sigma_q^2}, \quad (2.32)$$

that is

$$R_e^q > 0 \quad \text{if} \quad \frac{1/\bar{v} - \Upsilon\varpi_q}{\Upsilon\varpi_q} > \frac{1/\bar{v} - \varpi_q}{\varpi_q},$$

which results from the assumption of island illusion. Hence, $\beta_2 > 0$. ■

Proof of Proposition 2

A higher degree of island illusion (a lower Υ) implies...

a) A stronger overreaction to micro news (a lower β_1) and simultaneously a larger underreaction to the public signal (a larger β_2).

The coefficient β_1 of regression (2.10) is, where results from the proof of Proposition 1 are inserted in the first line

$$\begin{aligned} \beta_1 &= \frac{\text{Cov}(FE_{t+1}^{j,r}; FR_t^{j,r} - \text{Coe}f_1 s_t)}{\text{Var}(FR_t^{j,r} - \text{Coe}f_1 s_t)} = \frac{\overbrace{((X_1^q - \text{Coe}f_1) R_e^q + \Lambda X_1 R_\eta)}^{=0}}{\text{Var}(X_1 \varepsilon_t + X_1 \eta_t^r + X_1^q q_t + X_1^q e_t + K_\nu \nu_t - X_1^q s_t)} \\ &= \frac{\Lambda[\sigma_\varepsilon^2 - \delta_x^p \sigma_a^2]}{X_1 \sigma_a^2 + (K_\nu)^2 \sigma_\nu^2 / X_1}. \end{aligned}$$

First note that the derivative of X_1 with respect to δ_x^p equals

$$\frac{\partial X_1}{\partial \delta_x^p} = \frac{\partial K_1}{\partial \delta_x^p} (1 - \delta_x^p) - K_1 + \delta_x^h - \bar{k}_3 (1 - \delta_x^h) - (1 - \delta_x^h) \delta_x^p \frac{\partial \bar{k}_3}{\partial \delta_x^p}.$$

Since

$$\frac{\partial K_1}{\partial \delta_x^p} = \frac{1}{n} [-\gamma(n-1) + \delta_x^h - 1] \frac{\partial \bar{k}_3}{\partial \delta_x^p}$$

we have

$$\begin{aligned} \frac{\partial X_1}{\partial \delta_x^p} &= -K_1 + \delta_x^h - \bar{k}_3 (1 - \delta_x^h) + \left\{ \frac{1}{n} [-\gamma(n-1) + \delta_x^h - 1] (1 - \delta_x^p) - (1 - \delta_x^h) \delta_x^p \right\} \frac{\partial \bar{k}_3}{\partial \delta_x^p} \\ &= \bar{k}_3 \left[\frac{1}{n} \gamma(n-1) + \frac{1}{n} - (1 - \delta_x^h) \right] + \delta_x^h \left[1 - \frac{1}{n} (1 + \bar{k}_3) \right] + \\ &\quad \left\{ \frac{1}{n} [-\gamma(n-1)(1 - \delta_x^p) + \delta_x^h - 1] + \delta_x^p \frac{1}{n} - \delta_x^p \left[\frac{1}{n} \delta_x^h + 1 - \delta_x^h \right] \right\} \frac{\partial \bar{k}_3}{\partial \delta_x^p} \\ &= \Lambda + \frac{n-1}{n} [-\gamma(1 - \delta_x^p) + \delta_x^p (\delta_x^h - 1) + (\delta_x^h - 1)/(n-1)] \frac{\partial \bar{k}_3}{\partial \delta_x^p}. \end{aligned}$$

Because

$$\begin{aligned}\frac{\partial \bar{k}_3}{\partial \delta_x^p} &= \frac{\delta_x^h}{\Phi - \delta_x^h \Psi} \frac{\partial \Psi}{\partial \delta_x^p} + \frac{n/\Sigma - \delta_x^h \Psi}{(\Phi - \delta_x^h \Psi)^2} \left(\frac{\partial \Phi}{\partial \delta_x^p} - \delta_x^h \frac{\partial \Psi}{\partial \delta_x^p} \right) = \left[\delta_x^h + \bar{k}_3 \left((\gamma - 1) + \delta_x^h \right) \right] \frac{(n-1)(1-\alpha)}{\Sigma(\Phi - \delta_x^h \Psi)} \\ &= n\Lambda \frac{1-\alpha}{\Sigma(\Phi - \delta_x^h \Psi)}\end{aligned}$$

with

$$\Sigma(\Phi - \delta_x^h \Psi) = (n-1)(1-\alpha) \left[(\gamma-1)(1-\delta_x^p) - \delta_x^p \delta_x^h \right] + n - \delta_x^h(1-\alpha),$$

such that

$$\frac{\partial \bar{k}_3}{\partial \delta_x^p} = \frac{\Lambda}{[-1 + \gamma(1 - \delta_x^p) + (1 - \delta_x^h)\delta_x^p](n-1)/n + 1/(1-\alpha) - \delta_x^h/n} < 0,$$

we can also write

$$\frac{\partial X_1}{\partial \delta_x^p} = \Lambda \frac{n\alpha/(1-\alpha)}{(n-1)[\gamma(1-\delta_x^p) + (1-\delta_x^h)\delta_x^p] + n\alpha/(1-\alpha) + 1 - \delta_x^h} \equiv \Lambda K_4 < 0,$$

with $K_4 > 0$. The derivative of β_1 with respect to δ_x^p is then positive if

$$\begin{aligned}\frac{\partial \Lambda}{\partial \delta_x^p} R_\eta - \Lambda \sigma_a^2 &> \Lambda R_\eta \frac{(\sigma_a^2 - K_\nu^2 \sigma_\nu^2 / X_1^2)}{X_1 \sigma_a^2 + (K_\nu)^2 \sigma_\nu^2 / X_1} \frac{\partial X_1}{\partial \delta_x^p} \\ \frac{X_1 K_5 R_\eta - \sigma_a^2}{K_4 \Lambda R_\eta} &> \frac{\sigma_a^2 - K_\nu^2 \sigma_\nu^2 / X_1^2}{\sigma_a^2 + (K_\nu)^2 \sigma_\nu^2 / X_1} < 1,\end{aligned}$$

with

$$K_5 = \frac{n-1}{n} \frac{\gamma-1 + \delta_x^h}{\Lambda} \frac{\partial \bar{k}_3}{\partial \delta_x^p}.$$

The above is fulfilled if

$$\begin{aligned}-\sigma_a^2 &< \left(\frac{K_4}{X_1} \Lambda - K_5 \right) R_\eta \\ \text{or } -1 &< \left(\frac{K_4}{X_1} \Lambda - K_5 \right) (\varpi_a - \delta_x^p).\end{aligned}\tag{2.33}$$

Since

$$\frac{K_4}{X_1} \Lambda - K_5 = \frac{\frac{\alpha}{1-\alpha} \frac{\Lambda}{X_1} - \frac{n-1}{n} (\gamma-1 + \delta_x^p)}{[-1 + \gamma(1 - \delta_x^p) + (1 - \delta_x^h)\delta_x^p](n-1)/n + 1/(1-\alpha) - \delta_x^h/n}$$

inequality (2.33) can be written as

$$1 - \gamma(1 - \delta_x^p) - (1 - \delta_x^h)\delta_x^p / (n-1) / n - 1/(1-\alpha) + \delta_x^h/n < \left[\frac{\alpha}{1-\alpha} \frac{\Lambda}{X_1} - \frac{n-1}{n} (\gamma-1 + \delta_x^p) \right] (\varpi_a - \delta_x^p)$$

or

$$(\varpi_a - 1)(\gamma - 1) \frac{n-1}{n} + \frac{\delta_x^p}{n} [\varpi_a(n-1) + 1] - 1 < \frac{\alpha}{1-\alpha} \left[(\varpi_a - \delta_x^p) \frac{\Lambda}{X_1} + 1 \right].$$

We start with the left-hand side, which can be expressed as

$$(\varpi_a - 1)(\gamma - 1 + \delta_x^p) \frac{n-1}{n} + \delta_x^p - 1 < 0,$$

where the inequality follows from $\varpi_a, \delta_x^p < 1$. The right-hand side is positive if

$$(\varpi_a - \delta_x^p) \frac{\Lambda}{X_1} + 1 > 0. \quad (2.34)$$

Substituting X_1 and then Λ yields

$$\begin{aligned} \gamma \frac{\bar{k}_3}{\Lambda} &> \frac{1}{n-1} + \varpi_a \\ \gamma &> \frac{n-1}{n} \left[(\gamma - 1) + \delta_x^h \left(1 + \frac{1}{\bar{k}_3} \right) \right] \left(\frac{1}{n-1} + \varpi_a \right) \\ \underbrace{\gamma(1 - \varpi_a)}_{>0} &> \underbrace{\left[\delta_x^h - 1 + \frac{\delta_x^h}{\bar{k}_3} \right]}_{<0} \underbrace{\left(\frac{1}{n-1} + \varpi_a \right)}_{>0}, \end{aligned}$$

such that inequality (2.33) is fulfilled and hence

$$\frac{\partial \beta_1}{\partial \Upsilon} = \underbrace{\frac{\partial \beta_1}{\partial \delta_x^p}}_{>0} \underbrace{\frac{\partial \delta_x^p}{\partial \Upsilon}}_{>0} > 0,$$

demonstrating that a larger degree of ‘island illusion’ (a lower Υ) leads to a stronger overreaction to micro news (a lower β_1).

Concerning the effect of Υ on β_2 ,

$$\beta_2 = \frac{(1 - \text{Coe}f_2 X_1^q) R_e^q - \text{Coe}f_2 \Lambda X_1 R_\eta}{\text{Var}(s_t - \text{Coe}f_2 FR_t^{j,r})}$$

$$\beta_1 = \frac{\Lambda X_1 R_\eta}{\text{Var}(X_1 \varepsilon_t + X_1 \eta_t^r + X_1^q q_t + X_1^q e_t + K_\nu \nu_t - X_1^q s_t)} \equiv \frac{\Lambda X_1 R_\eta}{V_{\beta_1}},$$

such that, also substituting X_1^q ,

$$\beta_2 = \frac{(1 - \text{Coe}f_2 \rho_q^p) R_e^q - \text{Coe}f_2 \beta_1 V_{\beta_1}}{\text{Var}(s_t - \text{Coe}f_2 FR_t^{j,r})}.$$

Since

$$R_e^q = (1 - \rho_q^p) \sigma_q^2 - \rho_q^p \sigma_{e,q}^2 = (1 - \Upsilon \varpi_q \bar{v}) \varpi_q \bar{v} \text{Var}(s_t) - \Upsilon \varpi_q \bar{v} (1 - \varpi_q \bar{v}) \text{Var}(s_t) = (1 - \Upsilon) \varpi_q \bar{v} \text{Var}(s_t).$$

and, see the proof of Proposition 1,

$$Coe f_2 = \frac{Cov(FR_t^{j,r}, s_t)}{Var(FR_t^{j,r})} = \frac{X_1^q \sigma_q^2 + X_1^q \sigma_e^2}{X_1^2 \sigma_\varepsilon^2 + X_1^2 \sigma_\eta^2 + (X_1^q)^2 \sigma_q^2 + (X_1^q)^2 \sigma_e^2 + (K_\nu)^2 \sigma_\nu^2},$$

such that

$$Var(s_t - Coe f_2 FR_t^{j,r}) = (1 - Coe f_2)^2 Var(s_t) + Coe f_2^2 V_{\beta_1} = Var(s_t) \frac{V_{\beta_1}}{Var(FR_t^{j,r})},$$

as well as

$$1 - Coe f_2 \rho_q^p = \frac{X_1^2 \sigma_a^2 + (K_\nu)^2 \sigma_\nu^2}{Var(FR_t^{j,r})} = \frac{V_{\beta_1}}{Var(FR_t^{j,r})}$$

we obtain

$$\begin{aligned} \beta_2 &= \frac{\frac{V_{\beta_1}}{Var(FR_t^{j,r})} (1 - \Upsilon) \varpi_q \bar{v} Var(s_t) - \frac{\rho_q^p Var(s_t)}{Var(FR_t^{j,r})} \beta_1 V_{\beta_1}}{Var(s_t) \frac{V_{\beta_1}}{Var(FR_t^{j,r})}} \\ &= \varpi_q \bar{v} [1 - \Upsilon(1 + \beta_1)]. \end{aligned}$$

The derivative of β_2 w.r.t. Υ is therefore

$$\frac{\partial \beta_2}{\partial \Upsilon} = -\varpi_q \bar{v} \left(1 + \beta_1 + \Upsilon \frac{\partial \beta_1}{\partial \Upsilon} \right),$$

where $\frac{\partial \beta_1}{\partial \Upsilon} > 0$ was derived above. Regarding the size of β_1 , note that

$$\begin{aligned} \beta_1 &= \frac{\Lambda X_1 \sigma_a^2 \varpi_a (1 - \Upsilon)}{X_1^2 \sigma_a^2 + (K_\nu)^2 \sigma_\nu^2} > -1 \\ X_1 \sigma_a^2 [X_1 + \Lambda \varpi_a (1 - \Upsilon)] &> -(K_\nu)^2 \sigma_\nu^2. \end{aligned}$$

Since we have shown that inequality (2.34) holds, we also know that $X_1 + \Lambda \varpi_a (1 - \Upsilon) > 0$, such that $\beta_1 > -1$ and

$$\frac{\partial \beta_2}{\partial \Upsilon} < 0.$$

Hence, a higher degree of island illusion (a lower Υ) leads to a larger underreaction to macro news (a higher β_2). ■

(b) Lower expected profits

As usual, the firm's maximization problem states that profits are maximized if the price is a fixed markup over marginal costs. In linearized form

$$p_t^{j,r} = mc_t^{j,r},$$

where $mc_t^{j,r}$ are marginal costs, given by

$$\begin{aligned} mc_t^{j,r} &= w_t - a_t^r + \frac{1-\alpha}{\alpha}(y_t^{j,r} - a_t^r) \\ &= w_t + \frac{1-\alpha}{\alpha}y_t^{j,r} - \frac{1}{\alpha}a_t^r. \end{aligned}$$

Since the wage w_t and technology a_t^r are known at the time when prices are set (and independent of Υ), we have

$$mc_t^{j,r} - E_{t|2}^{j,r} mc_t^{j,r} = \frac{1-\alpha}{\alpha}(y_t^{j,r} - E_{t|2}^{j,r} y_t^{j,r}) = \frac{1-\alpha}{\alpha} FE_{t+1}^{j,r}.$$

The forecast error $FE_{t+1}^{j,r}$ is given by equation (2.30). Its expected value is zero and its variance is minimal at $\Upsilon = 1$, see below. Hence, expected profits are also at their maximum at $\Upsilon = 1$. Furthermore, given that the profit function (at the point of approximation) is concave in $P_t^{j,r}$, the larger the distance to the optimal price, the lower realized profits. ■

(c) *A larger variance of the firm-specific forecast error*

The forecast error $FE_{t+1}^{j,r}$ is given by equation (2.30). Its variance results as

$$\begin{aligned} \text{Var}(FE_{t+1}^{j,r}) &= \\ &\Lambda^2 \sigma_a^2 \left[(1 - \delta_x^p)^2 \varpi_a + (\delta_x^p)^2 (1 - \varpi_a) \right] + \text{Var}(s_t) \left[(1 - \rho_q^p)^2 \varpi_q \bar{v} + (\rho_q^p)^2 (1 - \varpi_q \bar{v}) \right] + \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{\bar{q}_t^k}{n} \\ &= \Lambda^2 \sigma_a^2 \varpi_a \left[(1 - \Upsilon)^2 \varpi_a + 1 - \varpi_a \right] + \text{Var}(s_t) \varpi_q \bar{v} \left[(1 - \Upsilon)^2 \varpi_q \bar{v} + 1 - \varpi_q \bar{v} \right] + \sum_{\{m|r \in \mathcal{B}_t^m\}} \frac{\bar{q}_t^k}{n}, \end{aligned}$$

such that

$$\frac{\partial \text{Var}(FE_{t+1}^{j,r})}{\partial \Upsilon} = -2(1 - \Upsilon) \left[\Lambda^2 \sigma_a^2 \varpi_a^2 + \text{Var}(s_t) (\varpi_q \bar{v})^2 \right].$$

Hence, $\text{Var}(FE_{t+1}^{j,r})$ is minimal at $\Upsilon = 1$ and rises as $|1 - \Upsilon|$ increases. ■

Proof of Proposition 3

As shown in the proof of Proposition 2 a), β_2 can be written as

$$\beta_2 = \varpi_q \bar{v} [1 - \Upsilon(1 + \beta_1)],$$

such that

$$\frac{\partial \beta_2}{\partial \varpi} = \bar{v} [1 - \Upsilon(1 + \beta_1)] > 0,$$

where we have used the result $\beta_1 > -1$ from the same proof. That is, a higher attachment to the business cycle (a higher ϖ) leads to a larger underreaction to macro news (a larger β_2). ■

Political Distance and International Trade

3.1 Introduction

Geopolitical tensions are rising worldwide, and some foresee a New Cold War.¹ As examples, consider the Russo-Ukrainian war starting in 2014, which culminated in the Russian invasion of Ukraine in February 2022, the tensions surrounding the legal status of Taiwan relative to China, and the trade war between China and the United States (US). These events prompted considerations among policymakers to decouple international trade, that is, to reorganize trade towards politically close countries and away from politically distant ones. For the US, Yellen (2022) established “Friend-Shoring” as the ideal for the future path of US trade policy. The EU introduced “Open Strategic Autonomy” (Le Maire 2020; Ioannou et al. 2023) and Truss (2021) considers a “Network of Liberty” for the UK. In this paper, I ask: what was the role of political distance for international trade in the past, and how much would we expect international trade to change in case of such a decoupling scenario?

I answer this question in three steps. First, I introduce the gravity model as a well-established framework for trade flows and trade costs. For trade flows, I entertain two complementary panels that each feature intra-national and international trade in goods. One is the well-established TradeProd panel of CEPII. It covers up to 167 countries from 1966 until 2018 and features nine aggregate sectors. While its coverage in terms of countries and periods is substantial, the sectoral breakdown is relatively coarse and difficult to align with other classification schemes. The other panel is novel, and I construct it to have a detailed sectoral breakdown (33 sectors) that can easily be aligned with other classification schemes at the expense of coverage across time (2012 to 2020) and countries (85). For trade costs, I employ traditional measures, that is, indicators for economic integration agreements and tariffs, and complement them with a novel measure: political distances derived from countries’ voting behavior at the United Nations General Assembly (UN GA). These distances have not yet been used to study international trade flows.

As a second step, I incorporate these trade costs in a standard empirical gravity model and find that, within a country pair and on average, an increase in political distance predicts a significant decrease in bilateral trade. More specifically, an average-sized increase in bilateral political distance² predicts a decrease in bilateral trade in the same year by 4 percent. This finding is robust to various alternative specifications, but there is heterogeneity across countries, sectors, and time. The predicted decrease in trade is more than twice as large for

¹The Economist discussed “A new kind of cold war” (16 May 2019). Der Spiegel, a leading German newspaper, featured the headline “Superpower Posturing – Fears Grow of New Cold War Between U.S. and China” (9 March 2023) and Le Monde diplomatique discusses “A cold war by any other name” (June 2023).

²For political distance, I think of an average-sized increase as the median pair-wise standard deviation.

trade flows involving the US, the EU, or the UK and for trade flows in strategic sectors. In the 1970s, it was 22 percent; between 2000 and 2009, it was insignificant, and in 2018 it was 4 percent and statistically significant again.

As a final step, I use these estimates to quantify the degree of trade diversion in a counterfactual decoupling scenario that mimics a New Cold War in 2018. Specifically, I use political distances observed during the Cold War to compute counterfactual trade flows for 2018. The median value of a counterfactual trade flow is not much different from the observed one in 2018, as countries trade in the counterfactual more with politically near countries and less with distant ones. This hides, however, substantial reshuffling of trade: the median absolute change of a trade flow is 56 percent of its actual value in 2018.

More in detail, the first part of the paper introduces the gravity model as the lens of the analysis and the data used for its estimation. I use two panels on intra-national and international trade flows that also include zeros for no trade as required for theory-consistent estimation of the gravity equation. The first panel is the well-established TradeProd panel of the CEPII. It is also used by, for example, Baier et al. (2019), Bergstrand et al. (2015), and Larch et al. (2019). It covers 167 countries from 1966 until 2018 at six industrial sectors that are difficult to align with other classification schemes. I address this issue by constructing a novel panel on sectoral trade. It is computed from data on industrial production at the sectoral level collected by the United Nations, and international trade flows at the goods level from the CEPII's BACI panel. The sectoral panel covers only 85 countries from 2012 to 2020, but its sectoral resolution is higher, as it features 33 sectors. In this paper, the sectoral panel is particularly useful for zooming into the heterogeneity across strategic and non-strategic sectors, as well as trade in global value chains (GVC).

For trade costs, I employ traditional measures, that is, tariffs from the TradeProd panel and data on economic integration agreements from Baier and Bergstrand (2021). I complement these trade costs with a novel one: political distance. To measure it, I consider countries' voting behavior at the UN GA, where countries vote on 80 resolutions per year on average.³ I compute the political distance between two countries based on the squared distance between their votes. This measure is based on Cohen (1960) and uses the voting data of Voeten (2013). The resulting political distances vary substantially over time and in the expected way. During the Cold War, Russia and China were at record distances to the US, while France and Canada were as close to the US as they ever were. With the end of the Cold War, this fragmentation decreased. This also holds when looking beyond the US and considering all political distances. Since 2018, however, only the fragmentation of political distances to the US has increased, but not the fragmentation of all political distances. Overall fragmentation may, however, increase if political tensions stay high or continue growing.

I combine political distances, tariffs, data on economic integration agreements,⁴ and trade flows from the TradeProd panel. The combined panel features 155 countries as exporters and importers and covers 53 years from 1966 to 2018. This panel is the basis of most of

³Häge (2011) and Bailey et al. (2017) discuss details on the construction of these measures specifically for UN GA voting data.

⁴I also estimate traditional specifications that feature log distance, indicators for contiguity, common languages, colonial links, and economic integration agreements, as well as tariffs. The estimates for these specifications are in Table 3.1 and are well in line with those in the meta-analysis of Head and Mayer (2014) and Yotov et al. (2016).

the empirical analysis. I use it to estimate a gravity model that features importer-year- and exporter-year-fixed effects and pair-fixed effects, in line with the recommendations in the literature (Baldwin and Taglioni 2006; Head and Mayer 2014; Yotov et al. 2016). The key variable that is novel to gravity models in international trade is political distance.⁵

The result is clear cut: increased political distance predicts significantly less bilateral trade. To quantify the change, consider the median of pair-wise standard deviations of political distance as an average-sized change in political distance. The baseline specification suggests that an average-sized increase in political distance predicts a significant decrease in trade of 4 percent on average when holding economic integration agreements and tariffs constant.

The result also arises from several alternative specifications. It holds when computing political distance based on resolutions on human rights only and when computing political distance based on Scott (1955), but not when using the ideal point distance of Bailey et al. (2017). It holds when using alternative data on economic integration agreements and when including a full set of indicators for different levels of economic integration. It also arises from using only every third, fourth, and fifth year in the panel as suggested by Yotov et al. (2016) and from using only positive trade flows. Lastly, it is also robust to accounting for globalization by including indicators for international trade flows each year as proposed by Baier et al. (2019).

Next, I investigate the heterogeneity across countries, sectors, and time and find significant differences across all dimensions. Concerning countries, I distinguish three groups. The first group comprises the US, the EU, and the UK. Their political leaders have announced or implemented plans to reorganize their trade relations, in particular with respect to global value chains (GVC). The second group consists of all other advanced economies as classified by the IMF, and the third group comprises all other economies and includes China. For imports into the US, the EU, or the UK, an increase in political distance by the mean pair-specific standard deviation predicts a decrease in trade of 12 percent. For other advanced economies, such an increase in political distance predicts an increase in imports by 2 percent, and for all other countries, it predicts a decrease of 2.5 percent. For exports, the picture is quite similar, except for other countries, where political distance is not a significant predictor.

To analyze the heterogeneity across sectors, I use the sectoral panel that features 33 sectors at the two-digit level of the Central Product Classification (CPC, United Nations 2015).⁶ Here, I distinguish sectors that belong to GVCs, as well as sectors that produce final goods and other goods. For trade in final goods and other goods, political distance is not a significant predictor. For trade in GVC goods, however, it is a significant predictor. Here, an increase in the political distance by the mean pair-specific standard deviation predicts a decrease in trade by 9 percent. Lastly, I distinguish strategic and non-strategic sectors. Based on the list of Tran (2022), I define sectors, such as communications equipment and medical appliances, as strategic and all others as non-strategic. An average-sized increase in political distance within a country pair predicts no significant change in trade flows in non-strategic sectors, but for strategic sectors, it predicts a decrease in trade by 9 percent

⁵The IMF (2023) uses a gravity model and political distance to study foreign direct investment.

⁶This classification is unique in that it can be matched to the Harmonized System (HS, World Customs Organization 2022), the Classification by Broad Economic Categories (BEC, United Nations 2016), and the sectoral classification of the International Standard Industrial Classification of All Economic Activities (ISIC Rev. 4, United Nations 2008).

Also in the time-series dimension, I find substantial variation. An increase in the political distance by the mean pair-specific standard deviation during the Cold War predicted a decrease in trade by 22 percent, while it predicted no significant change between 2000 and 2009, consistent with the “distance puzzle” in Borchert and Yotov (2017) and Yotov (2012). In 2018 the predicted decrease in trade was 4 percent and statistically significant again.

To address endogeneity concerns in empirical gravity models, Baier and Bergstrand (2007) and Baier et al. (2019) propose a simple test for strict exogeneity of economic integration agreements: under strict exogeneity, future economic integration agreements should have no significant impact on trade today (see also Wooldridge 2008). Applying this test to political distance, I find that, in contrast to existing results for economic integration agreements, future values of political distance do predict current trade. Hence, the estimates should not be interpreted causally. Instead, I interpret them as coefficients that approximate the conditional expected value of bilateral trade.

This interpretation motivates the final exercise, where I consider a counterfactual decoupling scenario similar to a New Cold War in 2018. In this counterfactual, I replace political distances in 2018 with those observed during the Cold War, more precisely in 1963, at the time of the Cuban Missile Crises. Afterwards, I compare the trade flows in the counterfactual to the ones observed in 2018 to assess how much trade is reshuffled in the decoupling scenario. To that end, I would ideally observe every country pair which is in the panel in 2018, also during the Cold War. Based on the difference between their political distances and using the predictive coefficients from the baseline specification, I could easily compute counterfactual trade and compare it to observed trade in 2018. However, not all countries in the sample in 2018 were UN members during the Cold War, and some did not even exist at that time. For 20 percent of country pairs in 2018, neither country existed or was a UN member. They account for less than 7 percent of all trade in 2018 and are dropped from the exercise. When only one country was a UN member, I replace the non-member with his political predecessor or the average of the ten countries the non-member was closest to in 2018 (“10-nearest-neighbors-approach”). When both countries in a pair were UN members during the Cold War, I simply compute their political distance at that time. This procedure yields 126 countries or 16,000 country pairs covering 93 percent of all trade in 2018 for the counterfactual analysis.

Using the political distances of the Cold War, I compute counterfactual bilateral trade flows. I compare these counterfactual values to the ones observed in 2018. The median relative difference is quite small, as countries now trade more with politically close countries and less with distant ones. In fact, the median relative difference to the observed values is +5 percent, so trade increases a little, as political distances are slightly lower on average in the counterfactual than in actuality. However, this modest increase hides substantial reshuffling of trade: the median absolute change of a trade flow is 56 percent of its actual value in 2018. This holds across different country groups, but it is largest for other economies, that include, for example, China. For these countries, the median absolute change of a trade flow is 68 percent of its actual value in 2018.

The rest of the paper is organized as follows. In the remainder of the introduction, I place the paper’s contribution in the context of the literature. Section 3.2 introduces the gravity model as the lens I use for the analysis and the data used for the estimation. In Section 3.3, I show the results from the baseline and alternative specifications. Afterwards, I investigate

the heterogeneity across countries, sectors, and time and discuss the interpretation of the estimates. Section 3.4 uses the baseline estimates and quantifies the degree of trade diversion in a counterfactual New Cold War. The final section concludes.

Related Literature. This paper relates to the literature on empirical gravity models and the literature studying geopolitical fragmentation and its impact on economic outcomes.

The literature on empirical gravity models deals with the theory-consistent estimation of gravity models and the causal effects of trade costs, particularly free trade agreements (FTAs). The recent literature is built on three main ideas. First is the contribution of Baier and Bergstrand (2007), who identify the causal effect of FTAs using pair-fixed effects. These pair-fixed effects, along with importer-time- and exporter-time-fixed effects introduced by Fally (2015), have now become a standard component of empirical gravity equations. Baier and Bergstrand (2007) also introduce the test for strict exogeneity of Wooldridge (2008) into the gravity framework and find that in the presence of pair-fixed effects, FTAs are indeed exogenous. For the estimation, however, they rely on a log-linearized gravity equation. The second idea is about the estimator used for gravity equations. Silva and Tenreyro (2006) show that using Ordinary Least Squares on log-linearized gravity equations yields biased estimates when the residual is heteroscedastic. As a remedy, they propose using the Poisson pseudo-maximum likelihood (PPML) estimator on the multiplicative version of the gravity equation. This setup can also incorporate trade flows with a value of zero. The third idea is about the inclusion of intra-national (domestic) trade flows. Dai et al. (2014) and Anderson and Yotov (2016) argue that in the gravity model, consumers choose from domestically and internationally produced goods and that this should be reflected in the data as well.

This trinity has since become standard in modern gravity econometrics. Bergstrand et al. (2015) use it to re-estimate the causal effects of economic integration agreements (EIAs) while also accounting for the effects of globalization. Larch et al. (2019) use it to study the effects of GATT membership, and Baier et al. (2019) study the heterogeneous but mostly positive causal effects of FTAs. I also adopt the trinity but move beyond the existing literature by introducing a new trade cost, political distance, that is significant even when controlling for EIAs and bilateral tariffs already.

The second strand of literature I contribute to studies the (potential) geopolitical fragmentation and its implications for economic outcomes, for example, international trade, foreign direct investment, and welfare. In this spirit, the IMF (2023) is closest to this paper. They study foreign direct investment (FDI) and also use a gravity model with political distance. They find that an increase in political distance predicts a significant decrease in FDI. This aligns with the findings in this paper covering international trade.

Góes and Bekkers (2022) also consider countries' voting behavior at the UN GA and construct two blocks, one led by the US and the other by China. Using a multi-sector multi-region general equilibrium model, they find that welfare losses in a decoupling scenario, that is, an increase in trade costs between the two country blocks, can be drastic, as large as 12 percent in some regions. Eppinger et al. (2021) consider a similar notion of decoupling, where trade costs for international trade in inputs increase. They find that the welfare losses from a decoupling far exceed the benefits of reduced exposure to foreign shocks.

3.2 Gravity and political distance

In this section, I introduce the gravity model and the data used for its estimation. More specifically, I present a simple gravity model to lay the groundwork in Section 3.2.1. Estimation of this model requires a panel of intra-national and international trade flows and trade costs. For such trade flows, I use a workhorse panel as well as a novel sectoral panel that I introduce in Section 3.2.2. In Section 3.2.3, I discuss trade costs. I combine different established panels, for example, on tariffs and economic integration. I complement these panels with a novel type of trade costs: political distance measured based on countries' voting behavior at the United Nations General Assembly. I conclude by providing summary statistics of the panels that combine trade flows and trade costs in Section 3.2.4.

3.2.1 A simple gravity model

Gravity equations arise from various models, for example, Anderson (1979), Arkolakis et al. (2012), and Eaton and Kortum (2002). I present a simple demand-side derivation of the gravity equation to introduce relevant concepts and the data required for the estimation. It is based on Yotov et al. (2016) and does not provide any contribution beyond the original paper.

The starting point is a world of N countries, each producing a good differentiated by place of origin. The supply of each good is fixed at Q_i , and its factory price is p_i . Hence, the value of domestic production of country i is $Y_i = p_i Q_i$, which is also equal to nominal income. Its aggregate expenditures are $E_i = \phi_i Y_i$, where ϕ_i is exogenous and positive. In each country j , there are representative consumers with preferences over their consumption of variety i , denoted c_{ij} :

$$U_j = \left(\sum_i \alpha_i^{\frac{1-\sigma}{\sigma}} c_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}.$$

Here, α_i is an exogenous preference shifter, and $\sigma > 1$ is the constant elasticity of substitution among varieties. Consumers maximize their utility by choosing a consumption over goods produced domestically and internationally subject to their budget constraint $E_j = \sum_j p_{ij} c_{ij}$. The delivered price of a variety i in country j is $p_{ij} = p_i t_{ij}$, where $t_{ij} \geq 1$ are bilateral trade costs defined as iceberg costs. Hence, to deliver one unit of its variety to country j , country i must ship t_{ij} units. Optimality then implies that the total expenditures of country j on the variety of country i are

$$X_{ij} = c_{ij} p_{ij} = \left(\frac{\alpha_i p_i t_{ij}}{P_j} \right)^{1-\sigma} E_j,$$

where $P_j = (\sum_i (\alpha_i p_i)^{1-\sigma})^{\frac{1}{1-\sigma}}$ is the CES consumer price index of country j . Imposing market clearing ($Y_i = \sum_j X_{ij}$) lastly yields the structural gravity system:

$$X_{ij} = \frac{Y_i E_j}{Y} \left(\frac{t_{ij}}{\prod_i P_j} \right)^{1-\sigma} \quad (3.1)$$

where $\Pi_i^{1-\sigma} = \sum_j \left(\frac{t_{ij}}{P_j}\right)^{1-\sigma} \frac{E_j}{Y}$ and $P_j^{1-\sigma} = \sum_i \frac{Y_i}{Y} \left(\frac{t_{ij}}{\Pi_i}\right)^{1-\sigma}$ are structural terms that Anderson and Van Wincoop (2003) refer to as outward and inward multilateral resistances. Based on the gravity equation in Equation 3.1, two terms drive bilateral trade. The first term, $\frac{Y_i E_j}{Y}$, captures size effects and incorporates the idea that all else equal, larger economies import and export more. The size term also equals bilateral trade in a world without trade costs with $t_{ij} = 1 \forall i, j$. The second term, $\left(\frac{t_{ij}}{\Pi_i P_j}\right)^{1-\sigma}$, captures trade costs and explains deviations from the hypothetical, frictionless benchmark by bilateral trade costs t_{ij} and multilateral resistance terms P_j and Π_i , that capture, for example, the ease of market access and openness to trade. Assuming that the gravity equation holds across time t , modern textbook estimation, as in Yotov et al. (2016), uses panel data on intra-national and international trade flows, as well as measures of trade costs, to estimate the trade cost elasticities. I discuss the data I use for trade flows and trade costs in the following two sections.

3.2.2 Trade flows

Modern textbook estimation of gravity models requires panel data on intra-national and international trade flows. I use two such panels that complement each other. One is the well-established CEPII Trade and Production Database (TradeProd) panel, and the other is the panel on sectoral trade that I compile.⁷ For each trade flow, TradeProd consolidates the values reported by the importer and the exporter. Consolidated values are generally based on the records of the importing country. When such records from the importing country are not available, the records of the exporting country are used instead. The consolidated values in TradeProd exclude transportation costs. Mayer et al. (2023) report that TradeProd traces over 90 percent of world manufacturing production from 2010 to 2018 and provide further details on the panel. TradeProd is also used by, for example, Bergstrand et al. (2015), Baier et al. (2019), and Larch et al. (2019).

The sectoral panel is based on production data from UNIDO and international trade data in the CEPII's BACI panel, both of which are publicly available. For more details on the construction of the sectoral panel, see Section 3.A.2.

Initially, both panels are at the sectoral level, which I only use for a later heterogeneity analysis. Nonetheless, it is helpful to compare the dimensions of both panels. I provide summary statistics on both panels' dimensions in Table 3.1a.

The TradeProd panel features only nine large sectors. In exchange, it is extensive in its coverage of countries and years. It features 167 countries as importers and exporters and covers the period from 1966 until 2018. The total number of observations is just below ten million observations at the sectoral level. The sectoral panel is more detailed in its sectoral resolution as it features 33 sectors.⁸ In exchange for this increased resolution, it features

⁷The TradeProd panel is available at http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=5

⁸These sectors are CPC Version 2.1 two-digit codes (United Nations 2015). As such, they can be conveniently mapped to other classifications, such as the Harmonized System (HS) or the International Standard Industrial Classification of All Economic Activities (ISIC). The sectoral panel is based on industrial production taken from the United Nations Industrial Statistics Database (INDSTAT 4) at the level of three- and four-digit codes of ISIC Revision 4 and international trade flows at the goods level from the CEPII BACI panel.

Table 3.1: Two panels of trade flows

(a) Panel dimensions at the sectoral level

Summary statistic	TradeProd	Sectoral panel
Number of sectors	9	33
Number of importing countries	167	88
Number of exporting countries	167	88
Start of sample	1966	2012
End of sample	2018	2020
Number of observations	9,968,175	2,299,968

(b) Distribution of aggregate trade flows

Trade flow	Summary statistic	TradeProd	Sectoral panel
International trade flows			
	Mean	1,277	1,641
	10 % Quantile	0	0
	50 % Quantile	20	36
	90 % Quantile	1,691	2,330
	Interquartile range	260	376
Intra-national trade flows			
	Mean	341,512	337,284
	10 % Quantile	963	0
	50 % Quantile	27,021	19,692
	90 % Quantile	575,953	587,644
	Interquartile range	111,671	118,027

Notes: The upper panel shows panel dimensions of the TradeProd and the sectoral panel. The lower panel shows summary statistics for the value of trade flows that are featured both in the TradeProd and the sectoral panel. All values are in millions of current USD. A list of all countries in the TradeProd panel is in Table 3.A.2, and countries in the sectoral panel are in Table 3.A.1.

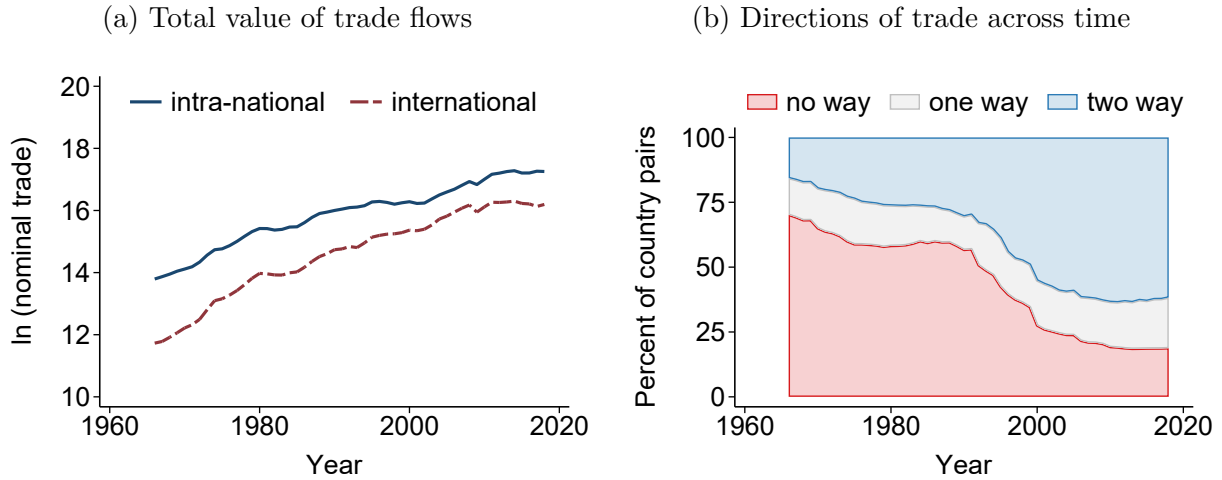
fewer countries than the TradeProd panel (88) and a shorter sample period from 2012 to 2020. The high sectoral resolution also dramatically increases the number of observations. The sectoral panel features just over two million observations at the sectoral level.

Except for a later heterogeneity analysis based on the sectoral panel, I sum up across sectors and focus on aggregate trade. Table 3.1b shows that the distribution of aggregate trade flows is quite similar for trade flows featured in both panels.

International trade flows are heavily skewed, as the mean is far larger than the median. In TradeProd, the median value of trade flows is 20 million US-Dollars (USD), while the mean is almost 1,300 million USD. The sectoral panel's median value is 36 million USD, and the mean is 1,641 million USD. In both panels, the values of intranational trade are outsized by intra-national (domestic) trade. The median value of an intra-national trade flow is almost three orders of magnitude larger than that of a median international trade flow, both in the TradeProd and sectoral panels.

For most of the following analysis, I use the TradeProd panel due to its extensive coverage of countries and years, and use the sectoral panel to study the sectoral heterogeneity.

Figure 3.1: Fifty years of trade in TradeProd



Notes: The left panel shows (log) total intra-national and international trade. The right panel shows how many percent of country pairs do not trade at all (no way, bottom), only trade in one direction (one way, middle), or trade in both directions (two way, top).

To gauge the relative sizes of intra-national and international trade, consider Figure 3.1a. It plots the total value of all intra-national and international trade flows over the sample period. Two observations stand out: (i) intra-national trade outsizes international trade and (ii) until the Great Financial Crisis, international trade grew faster than intra-national trade. Since then, however, growth slowed down, and intra-national trade grew only moderately while international trade stagnated.⁹ This evidence is in line what *The Economist* (2019) dubbed “Slowbalisation”.

Figure 3.1b examines international trade more closely and considers the directions in which country pairs trade with each other. Following Helpman et al. (2008), it plots the share of country pairs across time that fall into one of three categories: pairs not trading with each other at all (no-way), pairs trading in only one direction (one-way) or pairs trading in both directions (two-way). Over time, the share of one-way trade is constant at roughly 15%. However, there have been substantial changes in no-way and two-way over time. In the 1970s, the most common outcome was no-way trade, accounting for just below 75% of all country pairs. Since then, however, this share steadily decreased while one-way or two-way trade increased. This trend continued through the mid-1990s when two-way trade became the most common outcome until 2010. Since then, the shares of no-way, one-way, and two-way trade remain unchanged, with two-way trade between approximately 60% of all country pairs and the remaining pairs evenly split between no-way and one-way trade.

⁹More specifically, between 1966 and 2008, international trade grew by 10.6% per year on average, while intra-national trade grew at 7.4%. Between 2009 and 2018, intra-national trade grew by 3.2% per year on average, while international trade stagnated with an average growth rate below 0.2%.

3.2.3 Trade costs

For traditional measures of trade costs, I combine two datasets. The first is the CEPII Gravity panel. For any pair of countries from 1948 to 2020, it provides, for example, measures of geographical distance, contiguity, common languages, colonial links, and information on regional trade agreements. For more details on the construction of the Gravity dataset, see Conte et al. (2022). For regional trade agreements, I also use the NSF-Kellogg Institute Database on Economic Integration Agreements of Baier and Bergstrand (2021). It classifies bilateral levels of economic integration on a scale from 0 (no agreement) to 6 (Economic Union) for each country pair and year between 1950 and 2018.

Traditional gravity models do not feature political distance explicitly. Instead, political distance enters implicitly, for example, via economic integration agreements. However, since the current discussion emphasizes political distance as key in reorganizing trade connections, it is essential to consider political distance explicitly. To measure political distance, I consider countries' voting behavior on resolutions at the United Nations General Assembly (UN GA) as covered in Voeten et al. (2009).

Figure 3.2a shows the total number of resolutions each year. On average, countries vote on roughly 80 resolutions per year. To measure political distance, I generally consider votes on resolutions of all topics. Later on, I also use votes on resolutions related to human rights and, more generally, other measures in robustness checks.

For every resolution, countries can vote yes (coded as 1), abstain or be absent (2), or vote no (3). Figure 3.2b shows the average number of ballots per year for a given vote (yes, no, or absent). From the beginning of the sample period, the number of UN members steadily increased from around 120 in 1966 to 165 in 1991. In 1992, after the dissolution of the Soviet Union, 15 post-Soviet states joined the UN. At the end of the sample period, 193 countries are members of the UN. The distribution of votes over time is relatively constant. On average, 67 % of countries vote yes, 8% abstain or are absent at the time of voting, and the remaining 25 % vote against a resolution.

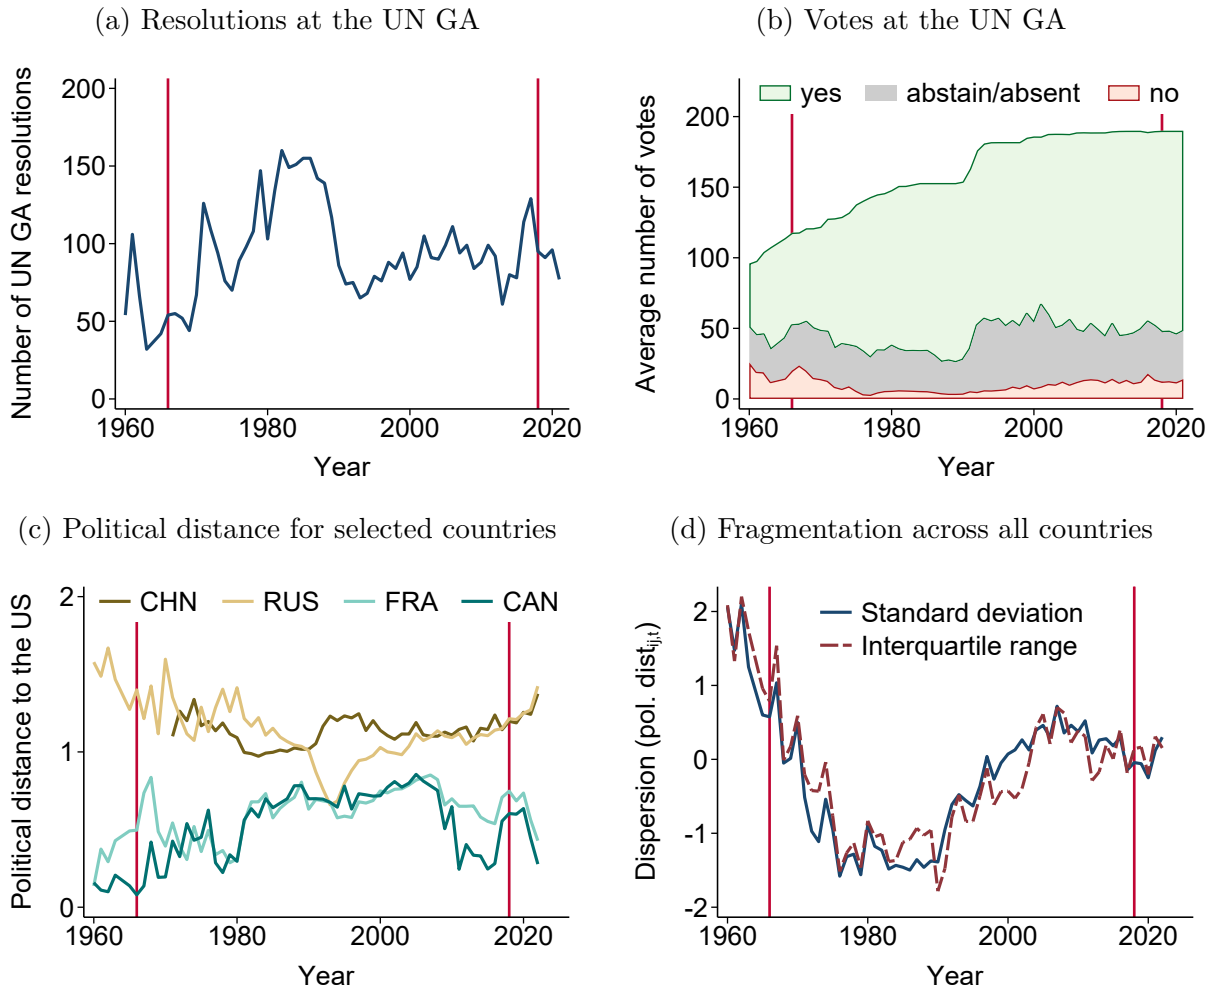
To quantify political distance based on countries' votes at the UN GA, I build upon existing measures of political similarity. I adapt them to reflect political distance, parallel to geographical distance, which already features prominently in gravity models. More specifically, I start with κ as proposed by Cohen (1960).¹⁰ It measures political similarity and has an upper bound of 1. To obtain a measure of political distance with 0 as a lower bound, I take κ , subtract one and then flip the sign of the result. Hence, I calculate the political distance between countries i and j in year t as

$$\text{political distance}_{ij,t} = -(\kappa_{ij,t} - 1) = \frac{\sum_k (v_{i,k} - v_{j,k})^2}{\sum_k (v_{i,k} - \bar{v}_i)^2 + \sum_k (v_{j,k} - \bar{v}_j)^2 + \sum (\bar{v}_i - \bar{v}_j)^2}.$$

Here $v_{i,k}$ and $v_{j,k}$ are the votes (1 for yes, 2 abstain/absent, 3 no) of countries i and j on resolution k in year t , and \bar{v}_i and \bar{v}_j are their averages. The first equality follows from my definition of political distance. The second equality simply replaces κ by its definition as in Häge (2011). As examples of time series of political distance, consider Figure 3.2c.

¹⁰In later robustness checks, I also consider the absolute ideal point distance provided by Bailey et al. (2017) and π proposed by Scott (1955).

Figure 3.2: Resolutions at the UN GA, votes, and political distance



Notes: The panel in the top left shows the total number of resolutions voted on in the UN GA each year. The panel in the top right shows the average number of ballots for a given vote. I sum up the number of ballots of each type of vote across resolutions each year and divide by the number of resolutions. The panel in the bottom left shows the political distance of China, Russia, France, and Canada from the US. The panel in the bottom right shows the standard deviation and the interquartile range of political distance between all country pairs countries each year. Both time series are standardized. In the bottom two panels, Russia refers to the Soviet Union until 1991. China refers to the People's Republic of China, which was admitted to the UN in 1971, replacing Taiwan (the Republic of China). The vertical lines indicate the start of the sample period (1966) and its end (2018).

Table 3.2: Descriptive statistics for the combined TradeProd panel

(a) Panel dimensions – number of unique values

Variable	Overall	1966 to 1980	1981 to 2000	2000 to 2018
Exporting countries	155	118	153	154
Importing countries	155	118	153	154
Year	53	15	20	18
Number of observations	948,843	179,620	347,037	422,186

(b) Main variables

Variable	Source	Mean	p10	p50	p90	IQR	SD
Trade flows (Million USD)	TradeProd	862.32	0.00	0.10	109.11	7.39	47687.35
Any EIA	Baier and Bergstrand (2021)	0.22	0.00	0.00	1.00	0.00	0.41
Ad-valorem tariff	TradeProd	0.07	0.00	0.02	0.20	0.12	0.11
Political distance	Voeten et al. (2009)	0.80	0.36	0.88	1.11	0.40	0.30

Notes: The upper panel shows the number of unique values of the variables that identify observations in the panel, both overall and broken down by periods. The lower panel shows descriptive statistics for the main variables. p10, p50, and p90 denote the 10, 50, and 90 percent quantiles, IQR is the interquartile range, and SD is the standard deviation. Any EIA is an indicator equal to 1 for any bilateral economic integration agreement.

It shows the political distance of China, Russia, France, and Canada towards the US from 1966 until 2018. France and Canada are politically closer to the US than China and Russia, particularly during the Cold War. Towards the end of the Cold War, fragmentation, which I define as the dispersion of political distance, decreased. Political distance towards the US decreased for Russia, while it increased for France and Canada. This also holds when considering the political distance across all country pairs, as in Figure 3.2d.¹¹ Since 2018, however, only the fragmentation of political distances to the US has increased, but not the fragmentation of all political distances. Overall fragmentation may, however, increase if political tensions stay high or continue growing.

3.2.4 Descriptive statistics

In this section, I briefly discuss the panel that combines trade flows from TradeProd and trade costs from previously discussed sources. I provide descriptive statistics for the dimensions of the panel in Table 3.2a and the main variables it covers in Table 3.2b.

Overall, the panel features 155 countries as importers or exporters, where coverage increases over time. The panel's time dimension is also substantial and covers 53 years, from 1966 until 2018. Trade flows are from TradeProd and are measured in millions of USD. They include intra-national and international trade and values of zero in case of no reported trade. Data on economic integration agreements are from Baier and Bergstrand (2021). Following the literature, I consider an indicator equal to one for any economic integration agreement

¹¹Figure 3.A.1 shows that this finding is robust to using alternative measures of political distance.

between a country pair.¹² More than 20 percent of observations in the panel are for country pairs with an economic integration agreement in place at the time. Ad-valorem tariffs are from TradeProd as well. Starting at the sectoral level of TradeProd, I set them equal to zero whenever no tariffs are reported. Then, I compute the average tariff across all sectors. The mean ad-valorem tariff across time and country pairs is 7 percent, but half of all ad-valorem tariffs in the panel are less than 0.2 percent, and 90 percent of all ad-valorem tariffs are less than 20 percent. Besides these traditional trade costs, the panel features political distance. I compute it based on countries' voting behavior at the UN GA. It ranges from 0 (within countries, for example) to 2, but 90 percent of political distances are less than 1.11. Their mean is 0.80, and the median is 0.88.

3.3 Bilateral trade and political distance

This section first introduces the estimation framework and presents results for the baseline specification in Section 3.3.2 and alternative specifications in Section 3.3.3. Afterwards, I consider the heterogeneity across sectors, countries, and time in Section 3.3.4 and discuss identification in Section 3.3.5.

3.3.1 Estimation framework

To estimate gravity equations, such as Equation 3.1, several studies, including Baldwin and Taglioni (2006), Yotov et al. (2016), and Head and Mayer (2014), established a set of best practices. Adopting these best practices and adding political distance as a novel trade cost yields the empirical gravity equation:

$$X_{ij,t} = \exp\left(\pi_{i,t} + \pi_{j,t} + \pi_{ij} + \beta_1 \text{poldist}_{ij,t} + \beta_2 \text{EIA}_{ij,t} + \beta_3 \log(1 + \text{tariff}_{ij,t})\right) \times \varepsilon_{ij,t}. \quad (3.2)$$

Bilateral trade flows from country i to country j in year t , $X_{ij,t}$, consist of intra-national and international trade flows and include trade flows with zero value when no trade is reported. The right-hand side of the empirical gravity equation features fixed effects and time-varying bilateral trade costs. Time-fixed effects for each importer and exporter, $\pi_{j,t}$ and $\pi_{i,t}$, correspond to multilateral resistance terms as shown by Fally (2015). At an empirical level, these fixed effects control for any factors that equally affect a country's imports or exports to all destinations. Pair-fixed effects, π_{ij} , absorb all time-constant bilateral trade costs, such as geographical distance, contiguity, common languages, and colonial links. Hence estimation focuses on time-varying, bilateral trade costs. Here, political distance is the key component and is not featured in standard specifications. Furthermore, I also control for the level of economic integration using an indicator variable for the presence of any bilateral economic integration agreements and the mean ad-valorem tariff. Lastly, $\varepsilon_{ij,t}$ is the residual.

For the estimation of Equation 3.2, I employ the standard Poisson pseudo-maximum likelihood (PPML) estimator (see also Silva and Tenreyro 2006; Yotov 2022; Baier et al. 2019).¹³ At a conceptual level, the setup in Equation 3.2 is quite similar to the IMF (2023),

¹²As a later robustness test, I also consider indicator variables for each type of economic integration agreement. This leaves the results unchanged, however.

¹³I use the Stata package `ppmlhdfe` by Correia et al. (2020).

Table 3.1: Political distance and bilateral trade

	(1)	(2)	(3)	(4)
Political distance				
based on UN GA resolutions			-0.220*** (0.055)	-0.276*** (0.057)
Traditional trade costs				
Log distance	-0.751*** (0.027)			
Contiguity	0.357*** (0.054)			
Common language	0.203*** (0.061)			
Colonial links	-0.067 (0.103)			
Economic integration agreements	0.180*** (0.060)	0.148*** (0.050)	0.205*** (0.044)	
Log (1+tariff)	-2.926*** (0.385)	-4.161*** (0.315)	-4.057*** (0.329)	
Observations	1,004,353	986,892	859,192	859,192
Pair-fixed effects	No	Yes	Yes	Yes
Intra-national-fixed effects	Yes	No	No	No

Notes: Results for aggregate trade in goods based on Tradeprod. The dependent variable is the trade flow from country i to country j in year t . Trade flows are intra-national and international and have a value of 0 in the case of no reported trade. Political distance is based on countries' votes on all UN GA resolutions. Log distance, dummies for contiguity, common spoken languages, and colonial links are from the Gravity database. The indicator for economic integration agreements is based on Baier and Bergstrand (2021). Mean bilateral tariffs are from Tradeprod. All specifications include importer-year- and exporter-year-fixed effects. Standard errors are clustered at importer-exporter level, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

who also estimate a gravity model that features political distance. There are two key differences, however. The IMF (2023) studies foreign direct investment from 2003 to 2021. In contrast, I study international trade from 1966 until 2018, that is, over the last 50 years. Nonetheless, the IMF (2023) provides a valuable comparison I draw later.

3.3.2 Results

I first show a specification without political distance to establish a baseline and compare the results to the literature. Afterwards, I add political distance. The results are in Table 3.1, where all specifications feature importer-time-fixed and exporter-time-fixed effects.

The first two columns are standard specifications, and the estimates align with the literature. The last two columns augment the standard specifications by political distance. Column (1) reports estimates for a specification that includes traditional trade costs only: distance, indicators for contiguity, common language, colonial links, and regional trade

agreement (RTA), as well as mean bilateral tariffs, as well as country-specific fixed-effects for domestic trade. The significance and sign of the estimated parameters are very much in line with the literature (see, for example, Yotov et al. (2016) and Head and Mayer (2014)). Column (2) adds, following best practices, pair-fixed effects, which absorb time-constant bilateral trade costs (distance, contiguity, common language, and colonial links), and, more generally, any unobserved bilateral and time-constant characteristics. Hence, estimation focuses on bilateral and time-varying trade costs, that is, economic integration agreements and tariffs, whose predictive coefficients increase moderately in absolute value. Column (3) additionally features the key variable, political distance. It turns out that it is a highly significant predictor of trade.

To illustrate its magnitude, consider an average change in the political distance within a given country pair. Computing the time-series standard deviation of political distance for each country pair and taking its mean yields a value of 0.17. This is quite close to, for example, the increase in political distance between the US and Russia from 2021 to 2022 of 0.15. Given the point estimate of -0.23 , such a change in the political distance would predict a decrease in trade by 3.8 percent, while holding economic integration agreements and tariffs constant and any characteristics that affect all trade partners of the US and Russia equally.¹⁴

Column (4) then drops economic integration agreements and tariffs and focuses on political distance only. The resulting predictive semi-elasticity increases, now predicting a drop in trade by 25 percent when political distance increases by one unit. In a comparable empirical setup, but studying foreign direct investment, the IMF (2023) also finds that an increase in political distance predicts significantly less trade.

3.3.3 Robustness and alternative specifications

Next, I show that the baseline results also arise from a number of alternative specifications. To that end, I run a battery of alternative specifications that vary specific aspects of the baseline specification and collect the main results in Table 3.2.¹⁵

The first row replicates the baseline results from Column (3) of Table 3.1: An increase in political distance by one unit predicts a significant drop of bilateral trade by 20 percent, even in the presence of pair-fixed effects and when controlling for economic integration agreements and tariffs. This specification measures political distance based on Cohen (1960, κ) and uses countries' votes in the UN GA on all resolutions. Furthermore, it uses data from Baier and Bergstrand (2021) for economic integration agreements. For the estimation, I use annual data and include trade flows with positive values and values of 0 in case of no bilateral trade. I vary all of these aspects in the remainder of this section.

The first set of alternative specifications is about the measure of political distance and the data used to compute it. Within this set, I use the voting data of Voeten (2013) for the first four alternatives. I consider, alternatively, votes on resolutions on human rights issues only, and also consider Scott (1955, π) as an alternative measure. The fourth and final alternative specification features the ideal point distance proposed by Bailey et al. (2017) as political distance. This measure is also used by the IMF (2023). For all measures, except for ideal

¹⁴The exact change is $\exp(\widehat{\beta}_1 \cdot \sigma_{\text{pol. dist}}) - 1 = \exp(-0.23 \cdot 0.17) - 1 = -0.038$.

¹⁵For the full results, see Table 3.A.3.

Table 3.2: Alternative specifications

Aspect (Baseline)	Alternative specification	Details	$\hat{\beta}_1$	SE($\hat{\beta}_1$)
Reminder: Baseline specification				
	See Column (3) in Table 3.1		-0.220***	0.055
1) Political distance (based on all UN GA resolutions and κ)		Table 3.A.3a		
	Human rights resolutions, κ		-0.218***	0.037
	All resolutions, π		-0.141***	0.045
	Human rights resolutions, π		-0.195***	0.033
	Ideal point distance, Bailey et al. (2017)		-0.018	0.025
2) Economic integration (indicator based on Baier and Bergstrand (2021))		Table 3.A.3b		
	Indicator based on WTO data		-0.187***	0.054
	Interaction term for levels of EIAs		-0.078*	0.045
3) Data (use all available years)		Table 3.A.3c		
	Three year intervals		-0.174***	0.059
	Four year intervals		-0.183**	0.072
	Five year intervals		-0.300***	0.066
4) Trade flows (zero and positive)		Table 3.A.3d		
	Only positive trade flows		-0.233***	0.056
5) Explicitly account for globalisation (no)		Table 3.A.3e		
	Include international flows by year indicators		-0.148***	0.042

Notes: Alternative specifications relative to the baseline in Column (3) of Table 3.1. Each column varies the baseline specification by one aspect. Full results are in the tables listed under Details. $\hat{\beta}_1$ is the coefficient for political distance and SE($\hat{\beta}_1$) is the standard error. UN GA resolutions related to Human Rights as classified by Voeten et al. (2009). κ indicates that political distance is based on Cohen (1960), and π indicates that political distance is based on Scott (1955). Bailey et al. (2017) use all UN votes to measure political distance. Besides importer-year-fixed, exporter-year-fixed, and pair-fixed effects, all specifications also include measures of bilateral economic integration and the mean bilateral tariff. Full results are in Table 3.A.3. Standard errors are clustered at importer-exporter level. *** p<0.01, ** p<0.05, * p<0.1.

point distance, the coefficients for political distance are negative and highly significant. One potential explanation for the lack of significance for ideal point distance is that, unlike other measures, it does not detect an increase in fragmentation since 2010, as visible in Figure 3.2d. Figure 3.A.1 shows the standard deviation of political distance to the US (fragmentation) each year for different measures of political distance. Here, ideal point distance stands out since it is the only measure that suggests a steady decrease in fragmentation. At the same time, all other measures show increased fragmentation since 2010, in line with increased political tensions since then, particularly in the recent past.

The second set of alternative specifications varies the data used for economic integration and how they enter the regression equation. In the baseline, I follow the literature (Yotov et al. 2016, for example) and use a simple indicator for economic integration agreements of any kind based on Baier and Bergstrand (2021). Here, the first alternative uses data on regional trade agreements from the WTO, as reported in the Gravity dataset. The second alternative uses interaction terms for different levels of economic integration rather than a simple indicator. The interaction terms are again based on Baier and Bergstrand (2021).

In both alternatives, the coefficient for political distance is negative and highly significant. Using interaction terms, however, the size of the coefficient is roughly cut in half.

For the third set of alternative specifications, I experiment with the frequency of the data. In the baseline, I use every available year in the sample, as Egger et al. (2021) advocated. Yotov et al. (2016) and Olivero and Yotov (2012) suggest experimenting with interval data rather than the annual panel itself to allow for the adjustment of trade flows. I follow this suggestion and report results for specifications that use every third, fourth, or fifth year of data in the sample. The coefficient for political distance remains unchanged in these alternative specifications.

In the fourth alternative specification, I address concerns about including trade flows with a value of zero in the estimation. These concerns are not warranted, as using positive trade flows only has no impact on the sign and significance of the estimates and has a negligible effect on the numeric value of the estimates. This is well in line with existing theoretical and empirical results of, for example, Silva and Tenreyro (2006).

The last alternative specification explicitly controls for the effects of globalization. Baier et al. (2019) propose to include indicators equal to one for international trade for each year. They argue that the resulting coefficients capture globalization, as over time, all countries trade more with each other and less within their own domestic markets. Including these globalization terms reduces the estimated semi-elasticity of political distance by seven percentage points but leaves its significance unchanged.

3.3.4 Heterogeneity

In the previous section, I estimated a single parameter for political distance. In doing so, I averaged across countries, sectors, and time. In this section, I investigate the heterogeneity along these dimensions. To that end, I expand the baseline specification and interact political distance with variables that capture heterogeneity and modify the sample where necessary. I use the TradeProd panel introduced in Section 3.2 for most of this. TradeProd covers a long time series and many countries at the cost of a relatively coarse sectoral breakdown. For the heterogeneity across sectors, I use the sectoral panel summarized in Section 3.2 and Table 3.1a. It covers a far shorter period (2012 to 2020) and countries (21) but, in exchange, features 33 CPC two-digit sectors. I present the results for heterogeneity across countries and sectors in Table 3.3.

First, I consider heterogeneity across destination or importing countries. The US, the United Kingdom, and the European Union form the first group. Their political leaders have already announced plans to reorganize the trade relations towards politically close countries.¹⁶ The second group consists of all remaining advanced economies as defined by the IMF, and the third group consists of all other economies. The US, the EU, and the UK drive the negative and significant overall estimate from the baseline specification. Their semi-elasticity is -0.78 , highly significant, and almost ten times larger than that of other advanced economies (-0.07), where it is significant at the 10 percent level only. For other importing countries, political distance has no significant predictive power.

¹⁶For the US, see, for example, Yellen (2022) and Biden (2021). For the United Kingdom, see Truss (2021) and for European Union, see Störmer et al. (2021) and Le Maire (2020).

Table 3.3: Elasticities differ meaningfully across countries and sectors

Interaction	Details	N	$\widehat{\beta}_{1,j}$	$SE(\widehat{\beta}_{1,j})$	Wald
Reminder: baseline result		859,192			
Overall			-0.220***	0.055	
(1) Political Distance	Table 3.A.4a	859,192			0.000
× Destination is USA, EU, or UK			-0.758***	0.086	
× Destination other advanced economies			0.127**	0.062	
× Destination other economies			-0.143**	0.065	
(2) Political Distance	Table 3.A.4b	859,192			0.000
× Origin is USA, EU, or UK			-0.749***	0.078	
× Origin other advanced economies			0.110*	0.065	
× Origin other economies			-0.104	0.068	
(3) Political Distance	Table 3.A.4c	147,147			0.000
× Sectors intensive in final goods			0.016	0.054	
× Sectors intensive in GVC goods			-0.539***	0.064	
× Sectors intensive in other goods			0.016	0.053	
(4) Political Distance	Table 3.A.4d	98,098			0.000
× Non-strategic sectors			0.080	0.051	
× Strategic sectors			-0.526***	0.081	

Notes: I estimate the baseline specification while adding interaction terms and adjusting the sample accordingly. All specifications feature political distance based on Cohen (1960) measured by countries' votes on all resolutions at the UN GA. Furthermore, all specifications include a dummy for economic integration agreements based on Baier and Bergstrand (2021) and mean bilateral tariffs from TradeProd, as well as importer-year-fixed, exporter-year-fixed, and pair-fixed effects. N is the number of observations, $\widehat{\beta}_j$ is the point estimate for an interaction term and $SE(\widehat{\beta}_j)$ is its standard error. W is the p-value for a Wald test where H_0 is that all interaction terms are identical. For destinations, I distinguish three groups. The first group consists of the United States, members of the European Union, and the United Kingdom. Political leaders in these countries have announced plans to reorganize their trade connections. Other advanced economies are as defined by the IMF. All remaining countries are collected in other economies. At the sectoral level, I consider two dimensions of heterogeneity. For different types of sectors or goods (GVC, final, and others, based on the Classification by Broad Economic Categories (United Nations 2016)), I consider the most common type of good for each sector. GVC goods are processed and specific intermediate goods. Final goods are goods for final consumption or capital formation. The "sectors to watch" of Tran (2022) form the starting point for strategic sectors. Since direct correspondence between these sectors to CPC sectors is not available, I classify CPC sectors 43 to 48 as strategic. These sectors cover, for example, machinery, communications equipment, and medical appliances. All other sectors are considered non-strategic. For these two dimensions, I rely on the sectoral panel (see Table 3.1a and Section 3.A.2) rather than TradeProd. For the estimation, I aggregate the sectoral data across all sectors of a given characteristic, for example, strategic sectors, and then estimate the baseline specification. Standard errors are clustered at the importer-exporter level.

*** p<0.01, ** p<0.05, * p<0.1.

Second, I consider heterogeneity across sectors. Here, I use the sectoral panel introduced in Section 3.2 and described in more detail in Section 3.A.2. This panel is at the level of 33 two-digit codes of the Central Production Classification (CPC, United Nations 2015). For heterogeneity across the type of goods produced across sectors, I use the Classification by Broad Economic Categories (BEC, United Nations 2016). Following the BEC, I distinguish processed and specific intermediate goods (GVC goods), goods for final consumption and capital formation (final goods), and other goods.¹⁷ After aggregating trade flows from the sectoral level to the level of GVC, final, and other goods, I estimate the baseline specification interacting political distance with the goods type. For sectors that are intensive in final goods and other goods, the semi-elasticity for political distance is positive and, for final goods only, significant at the ten percent level. In contrast, for sectors that are intensive in GVC goods, the elasticity is negative, highly significant, and largest in absolute value among all interaction terms.


Lastly, I turn toward strategic and non-strategic sectors. I use the list of "sectors to watch" of Tran (2022) also used by the IMF (2023). Tran (2022) identifies five strategic sectors: semiconductors, telecommunications and 5G infrastructure, equipment needed for the green energy transition, active pharmaceutical ingredients, and strategic and critical minerals. These sectors roughly correspond to general-purpose machinery (CPC code 43), special-purpose machinery (44), office, accounting and computing machinery (45), electrical machinery and apparatus (46), radio, television and communication equipment and apparatus (47), and medical appliances, precision and optical instruments, watches and clocks (48). These sectors then constitute the group of strategic sectors. All other sectors are considered non-strategic. For the estimation, I aggregate values across all strategic and all non-strategic sectors. The interaction term for political distance is statistically significant for both types of sectors. For non-strategic goods, it is positive, while for strategic sectors, it is negative. Put differently: political distance predicts significantly more trade for non-strategic sectors while it predicts significantly less trade for strategic sectors. This finding is very much in line with the IMF (2023), who find that political distance is a stronger predictor for FDI in strategic sectors than for non-strategic sectors.

Now, I turn to heterogeneity across time and ask how the elasticity of political distance has changed over time. The TradeProd panel shines here since it starts in 1966, allowing insights from the Cold War up until 2018. I estimate the elasticity of political distance for every year by interacting political distance and time. Figure 3.2 shows the resulting time-series. The semi-elasticity for political distance changed dramatically between 1966 and 2018. From 1966 until 1980, it was close to unity and highly significant. Between 1980 and 2000, it shrunk towards zero and turned insignificant in 2000. Until 2009 it was borderline significant, and since then, it has again become highly significant and negative. This finding is consistent with what the IMF (2023) report for FDI using a comparable specification.

It is also consistent with the "distance puzzle" discussed by Yotov (2012) and Borchert and Yotov (2017). They study the elasticity of trade with respect to *geographical* distance over time and estimate a gravity equation that features intra-national and international trade

¹⁷BEC is defined at the level of six-digit codes of the Harmonized System (HS, World Customs Organization 2022). Using the correspondence table of the United Nations, I assign to each CPC sector the most common type of good in the BEC.

Figure 3.2: Elasticity of political distance across time



Notes: The figure shows the prediction coefficients for political distance interacted with years and otherwise using the baseline specification. Point estimates are red lines, and 95% confidence intervals are grey areas.

without pair-fixed effects. Their main result is that the elasticity of trade with respect to geographical distance decreases over time. Bergstrand et al. (2015) and Baier et al. (2019) use pair-fixed effects and arrive at a similar conclusion. They rationalize this finding as evidence of globalization, as all countries trade more internationally and less intra-nationally.¹⁸ At least for political distance, and since 2009, this trend seems to reverse, and international trade costs, in this case, political distance, are increasing again.

3.3.5 Identification

Can the estimates discussed in the previous section be interpreted causally? I rely on a simple test proposed by Wooldridge (2008) to answer this question empirically. He argues that, under strict exogeneity conditional on fixed-effects, future values of independent variables have no partial effect on the dependent variable. Baier and Bergstrand (2007) adapt this idea for gravity models. For FTAs, they include lags to allow for their phasing-in and, more importantly, include leads to test for strict exogeneity. I follow their approach and test for exogeneity by adding leads and lags of political distance to the baseline specification. The results are in Table 3.4.

Column (1) is the baseline specification. Columns (2) to (4) add the first, second, and third lead and lag of political distance separately. Looking at lags of political distance first, I find that the coefficients are negative and significant. This is very much in line with “phasing-in” effects for economic integration agreements, as in, for example, Baier and Bergstrand (2007).

The key to the question of causality, however, are the coefficients for leads of political distance. Under strict exogeneity conditional on fixed-effects, coefficients for leads of political distance should not be significant. Columns (2) to (4) show this is not the case, as the

¹⁸Specification (5) in Table 3.2 shows that political distance remains a significant predictor of trade even when accounting for the effects of globalization by means of indicator variables for international trade flows and each year.

Table 3.4: Causality

	(1)	(2)	(3)	(4)
Political distance				
in year $t + 3$				-0.14*** (0.04)
in year $t + 2$			-0.10*** (0.03)	
in year $t + 1$		-0.14*** (0.03)		
in year t	-0.22*** (0.06)	-0.12*** (0.03)	-0.14*** (0.03)	-0.13*** (0.03)
in year $t - 1$		-0.05 (0.04)		
in year $t - 2$			-0.06 (0.04)	
in year $t - 3$				-0.11** (0.05)
Traditional trade costs				
Economic integration agreements	0.20*** (0.04)	0.20*** (0.04)	0.19*** (0.04)	0.19*** (0.05)
Log(1+tariff)	-4.06*** (0.33)	-4.02*** (0.33)	-4.01*** (0.33)	-4.02*** (0.33)
Observations	859,192	852,101	836,501	827,416

Notes: Results for aggregate trade in goods based on Tradeprod. The dependent variable is the trade flow from country i to country j in year t . Trade flows feature intra-national and international flows and flows with a value of 0 in the case of no reported trade. Political distance is based on countries' votes for all UN GA resolutions. The dummy for any economic integration is based on the EIA database. Mean bilateral tariffs are from Tradeprod. All specifications include importer-year- and exporter-year-fixed effects. Standard errors are clustered at importer-exporter level, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

coefficients for leads of political distance are highly significant and similar in size to the coefficient of current political distance. Hence, I refrain from adding a causal interpretation of these results. Instead, I think of the coefficients as parameters describing the conditional expected value of bilateral trade flows. This idea motivates the counterfactual analysis in the next section.

3.4 Counterfactuals

The previous section established that political distance is a significant predictor of trade. I use this insight to study how much international trade would change in case of a decoupling scenario. More specifically, I consider a counterfactual decoupling of international trade in

2018 by setting all political distances to the values observed during 1962, at the time of the Cuban Missile Crisis. After computing (or approximating) these political distances, I compute counterfactual trade flows for 2018, compare them to trade flows observed in 2018, and measure the degree of trade reshuffling.

However, simply computing political distances for every country pair in the data in 2018 is not feasible as not all countries I observe in 2018 were UN members during the Cold War, let alone existed in the first place. How I proceed depends on whether no, one, or both countries in a pair were UN members. In 20 percent of all country pairs I observe in 2018, neither country was a UN member during the Cold War. In 2018, these country pairs amounted to less than 7 percent of all trade and are dropped for this exercise. For country pairs, where only one country was a UN member during the Cold War, I first check if the missing country was a former member of the Soviet Union¹⁹. If it was, I replace the missing country with the Soviet Union and use the resulting political distance. Otherwise, I adopt a “10-nearest-neighbors-approach” and compute, for the top 10 closest countries to the missing country, the distance to the UN member during the Cold War and use the resulting political distance. This leaves 126 countries or some 16,000 country pairs that cover 93 percent of all trade in 2018 for the counterfactual analysis.

Figure 3.A.2 in the Appendix shows that the differences between actual and counterfactual distances are generally centered around 0 and less than 1 in absolute value for 98 percent of all bilateral pairs. For example, the distance between the US and Russia in the counterfactual scenario increases by 0.47 or 40 percent of its actual value in 2018.

For this exercise, I use the predictive coefficient estimated over the entire sample period.²⁰ Based on the counterfactual political distances, I compute counterfactual trade in 2018 as follows:

$$X_{ij,2018}^{CF} = \hat{\mathbb{E}} \left[X_{ij,2018}^{CF} \mid \text{political distance}_{ij,2018} = \text{political distance}_{ij,1962}^{CF} \right] \quad (3.3)$$

Here, $X_{ij,2018}^{CF}$ is the counterfactual trade flow. After replacing the actual political distance with the counterfactual one, I compute $X_{ij,2018}^{CF}$ as the conditional expected value of bilateral trade. To measure the change in a given trade flow, I compute the relative difference between counterfactual and actual value as

$$rd_{ij} = X_{ij,2018}^{CF} / X_{ij,2018} - 1. \quad (3.4)$$

The median of rd_{ij} across all bilateral trade flows describes whether trade flows increase or decrease on average. Taking the absolute value of rd_{ij} and then taking the median summarizes the absolute change. Here, positive and negative differences do not cancel out, so this median describes how much trade flows change in absolute value. It measures how much trade is reshuffled to a different trade partner in the counterfactual. Table 3.1 shows these values.

The median change in bilateral trade flows is small but positive at 5 percent. This aligns with the moderate decrease in political distances illustrated in Figure 3.A.2. Put differently: in the counterfactual scenario, countries are closer on average and therefore trade more with

¹⁹I also replace China with the Soviet Union in this procedure. This seems reasonable, as Figure 3.2c shows that, upon joining the UN, China was politically as distant from the US as Russia.

²⁰Figure 3.2 shows that this average value ($\beta_1 = -0.22$) is quite close to its counterpart estimated specifically for 2018.

Table 3.1: Counterfactual v. actual exports

Countries	Country group	Median(rd_{ij})	Median($ rd_{ij} $)
(1) Overall		0.05	0.56
(2) By destination			
	US, EU, UK	0.03	0.49
	Other advanced economies	-0.04	0.46
	Other economies	0.08	0.60
(3) By origin			
	US, EU, UK	-0.01	0.35
	Other advanced economies	0.15	0.38
	Other economies	0.07	0.68

Notes: Table shows compares actual and counterfactual trade. Bilateral counterfactual trade is computed as in Equation 3.3, replacing political distance in 2018 with its value during the Cold War. I compute the relative difference between counterfactual and actual exports as in Equation 3.4. The table reports, for the country groups in Columns (1) and (2), the median relative difference (Column (3)) and the median absolute difference (Column (4)). Other advanced economies and other economies as defined by the IMF.

each other. This also holds when considering trade flows from certain country groups of destinations or origins. Looking at how much of these trade flows is reshuffled to a different trade partner, consider the median absolute relative change, which is substantial. The median trade flow changes by an absolute value of 56 percent of its value in 2018. There is only some heterogeneity across country groups, such as importers and exporters, but reshuffling is largest when it involves other economies, for example, China, as destination or origin.

3.5 Conclusion

What was the role of political distance for international trade in the past, and how much would we expect international trade to change in case of a decoupling scenario?

I answered this question through the lens of a standard gravity model that features indicators for economic integration agreements and tariffs. Moving beyond the existing literature, however, I additionally considered political distance a novel trade cost. I computed political distance from countries' voting behavior at the UN GA and characterized the distribution of political distances across years. Incorporating it into a standard gravity model, I found that an increase in political distance predicts a significant decrease in bilateral trade. This finding also arose from numerous alternative specifications, but I also documented the heterogeneity across countries, sectors, and time. Using standard tests in the literature, I found that political distance is not exogenous conditional on fixed effects. Hence I thought of the estimates as parameters that approximate the conditional expected value of bilateral trade. I applied this idea in a counterfactual, where I considered a New Cold War at the political level. I set political distances in 2018 to those observed during the Cold War and compared the counterfactual values to those observed in 2018. While the median change in trade flows was moderate, it hid substantial trade reshuffling, as the median absolute

difference was 56 percent of its value in 2018.

Some notes and qualifications to these results are in place that provide avenues for future research. While I did use a gravity model as a lens of analysis, it was intentionally stylized. Here, future research could push further by considering welfare to make qualified statements about the potential winners and losers of the decoupling scenario. This way, one could derive clear policy implications. Currently, the model does not provide these additional insights. Building on the existing literature, however, the model and the counterfactual analysis can easily be modified to allow these insights. Another avenue for future research is the identification of causal effects. While establishing causality in gravity models is notoriously difficult, one could look at other measures or consider shocks to political distances. To that end, data beyond the UN GA voting data might be helpful. The sectoral panel I used for the heterogeneity analysis might also help identify the heterogeneity of political distance's (causal) effects. The sectoral panel may also be helpful for a more detailed counterfactual analysis that considers each country's and sector's specific exposure to political distance. Alternatively, I focused on international trade's past, that is, until 2018, and its potential future. Future research could use more recent trade data and alternative identification schemes to directly study the impact of recent political shocks.

Lastly, gravity models extend beyond trade flows and have also been used to study, for example, migration flows (see Beine et al. (2016) for an overview). Hence, future research could study the role of political distance in migration flows. To do this, intuitively speaking, one might use data on migration flows instead of trade flows and estimate a model that is comparable to the one presented in this paper.

3.A Appendices

3.A.1 Additional figures and tables

Table 3.A.1: List of countries in the sectoral panel

ISO	N	I	X	ISO	N	I	X	ISO	N	I	X
AFG	765	0.22	0.16	DEU	765	0.00	0.00	PAN	765	0.15	0.11
ALB	765	0.14	0.10	GRC	765	0.02	0.00	PER	765	0.02	0.03
AGO	765	0.09	0.25	HUN	765	0.01	0.00	PHL	765	0.09	0.06
AZE	765	0.05	0.09	ISL	765	0.05	0.05	POL	765	0.00	0.00
AUS	765	0.00	0.00	IND	765	0.03	0.03	PRT	765	0.01	0.00
AUT	765	0.00	0.00	IDN	765	0.00	0.00	QAT	765	0.06	0.08
BGD	765	0.13	0.04	IRN	765	0.11	0.06	ROU	765	0.01	0.00
ARM	765	0.06	0.12	IRQ	765	0.16	0.24	RUS	765	0.01	0.00
BEL	765	0.01	0.01	ISR	765	0.07	0.05	RWA	765	0.13	0.28
BIH	765	0.01	0.02	ITA	765	0.00	0.00	SAU	765	0.02	0.02
BWA	765	0.18	0.28	JPN	765	0.00	0.00	SEN	765	0.12	0.12
BRA	765	0.00	0.00	KAZ	765	0.03	0.06	SGP	765	0.00	0.00
BGR	765	0.01	0.00	JOR	765	0.04	0.00	SVK	765	0.01	0.01
BLR	765	0.05	0.01	KEN	765	0.07	0.07	SVN	765	0.02	0.01
CAN	765	0.00	0.00	KOR	765	0.00	0.00	ZWE	765	0.15	0.17
LKA	765	0.06	0.03	KGZ	765	0.16	0.20	ESP	765	0.00	0.00
CHN	765	0.00	0.00	LVA	765	0.06	0.01	SWE	765	0.00	0.00
COL	765	0.01	0.02	LTU	765	0.04	0.01	CHE	765	0.00	0.00
CRI	765	0.03	0.04	LUX	765	0.02	0.02	ARE	765	0.02	0.02
HRV	765	0.04	0.01	MYS	765	0.03	0.03	TUR	765	0.00	0.00
CYP	765	0.05	0.01	MLT	765	0.05	0.02	UKR	765	0.05	0.04
CZE	765	0.00	0.00	MUS	765	0.03	0.04	MKD	765	0.04	0.07
DNK	765	0.00	0.00	MEX	765	0.02	0.02	GBR	765	0.00	0.00
ECU	765	0.07	0.06	MNG	765	0.11	0.20	TZA	765	0.06	0.07
EST	765	0.04	0.02	MDA	765	0.13	0.10	USA	765	0.00	0.00
FJI	765	0.16	0.23	OMN	765	0.08	0.04	URY	765	0.09	0.03
FIN	765	0.00	0.00	NLD	765	0.01	0.01	UZB	765	0.19	0.18
FRA	765	0.00	0.00	NZL	765	0.00	0.00				
GEO	765	0.06	0.02	NOR	765	0.00	0.00				

Notes: N: number of observations. I and X: share of non-zero imports and exports. All countries are in the panel from 2012 to 2020.

Table 3.A.2: List of countries in TradeProd

ISO	Min(t)	N	I	X	ISO	Min(t)	N	I	X	ISO	Min(t)	N	I	X
AFG	1966	7,021	0.56	0.51	GAB	1966	7,021	0.57	0.50	NAM	2000	2,909	0.26	0.21
AGO	1976	5,963	0.42	0.59	GBR	1966	7,021	0.01	0.09	NER	1966	6,998	0.44	0.57
ALB	1966	7,021	0.65	0.51	GEO	1992	4,058	0.40	0.32	NGA	1966	7,021	0.23	0.38
ARE	1971	6,512	0.33	0.33	GHA	1966	7,021	0.30	0.36	NIC	1966	7,021	0.40	0.47
ARG	1966	7,021	0.25	0.15	GMB	1966	7,014	0.51	0.64	NLD	1966	7,021	0.01	0.09
ARM	1992	4,061	0.40	0.44	GRC	1966	7,021	0.09	0.12	NOR	1966	7,021	0.09	0.10
AUS	1966	7,021	0.09	0.09	GTM	1966	7,021	0.38	0.37	NPL	1966	7,021	0.55	0.50
AUT	1966	7,021	0.04	0.09	HND	1966	7,021	0.37	0.41	NZL	1966	7,021	0.16	0.15
AZE	1992	4,061	0.39	0.41	HRV	1992	4,061	0.10	0.15	OMN	1971	6,507	0.38	0.40
BDI	1966	7,011	0.56	0.67	HTI	1966	7,021	0.80	0.56	PAK	1966	7,021	0.15	0.12
BEL	1966	7,021	0.02	0.09	HUN	1966	7,021	0.27	0.17	PER	1966	7,021	0.26	0.26
BEN	1966	7,021	0.42	0.63	IDN	1966	7,021	0.19	0.17	PHL	1966	7,021	0.25	0.17
BFA	1966	7,021	0.45	0.64	IND	1966	7,021	0.14	0.09	PNG	1975	6,077	0.62	0.60
BGD	1974	6,187	0.37	0.19	IRL	1966	7,021	0.06	0.10	POL	1966	7,021	0.38	0.16
BGR	1966	7,021	0.54	0.23	IRN	1966	7,021	0.36	0.28	PRT	1966	7,021	0.12	0.11
BHR	1971	6,507	0.28	0.39	IRQ	1966	6,972	0.52	0.58	PRY	1966	7,021	0.46	0.45
BHS	1973	6,296	0.45	0.42	ISL	1966	7,021	0.31	0.30	QAT	1971	6,512	0.35	0.43
BIH	1992	4,049	0.45	0.36	ISR	1966	7,021	0.26	0.21	ROU	1966	7,021	0.46	0.18
BLR	1992	4,061	0.36	0.28	ITA	1966	7,021	0.00	0.09	RUS	1992	4,061	0.18	0.11
BLZ	1981	5,381	0.46	0.51	JAM	1966	7,021	0.37	0.34	RWA	1966	7,020	0.64	0.65
BOL	1966	7,016	0.34	0.49	JOR	1966	7,021	0.27	0.35	SAU	1966	7,016	0.18	0.27
BRA	1966	7,021	0.19	0.11	JPN	1966	7,021	0.02	0.09	SDN	2011	1,232	0.12	0.25
BRB	1966	7,021	0.34	0.48	KAZ	1992	4,061	0.27	0.35	SEN	1966	7,020	0.30	0.41
BRN	1984	5,024	0.43	0.53	KEN	1966	7,021	0.37	0.26	SGP	1966	7,021	0.15	0.14
BWA	2000	2,909	0.35	0.38	KGZ	1992	4,043	0.44	0.51	SLV	1966	7,021	0.40	0.44
CAF	1966	6,980	0.55	0.62	KHM	1966	7,011	0.58	0.51	SOM	1966	6,969	0.88	0.68
CAN	1966	7,021	0.06	0.10	KOR	1991	4,184	0.03	0.07	SUN	1966	2,588	0.39	0.31
CHE	1997	3,305	0.01	0.05	KWT	1966	7,021	0.29	0.35	SUR	1975	6,075	0.51	0.53
CHL	1966	7,021	0.25	0.24	LAO	1966	7,021	0.67	0.61	SVK	1993	3,922	0.09	0.11
CHN	1971	6,512	0.31	0.13	LBN	1966	7,021	0.26	0.31	SVN	1992	4,061	0.07	0.14
CIV	1966	7,021	0.30	0.34	LBR	1966	6,973	0.84	0.56	SWE	1966	7,021	0.03	0.09
CMR	1966	7,021	0.32	0.41	LBY	1966	7,021	0.42	0.56	SWZ	2000	2,909	0.38	0.27
COG	1966	7,020	0.41	0.50	LCA	1979	5,616	0.47	0.69	SYR	1966	7,021	0.54	0.39
COL	1966	7,021	0.22	0.23	LKA	1966	7,021	0.31	0.19	THA	1966	7,019	0.14	0.13
CPV	1975	6,077	0.61	0.70	LSO	2000	2,909	0.59	0.52	TJK	1992	4,052	0.66	0.56
CRI	1966	7,021	0.31	0.33	LTU	1992	4,061	0.26	0.21	TON	1999	3,047	0.71	0.74
CUB	1966	7,021	0.71	0.38	LUX	1999	3,056	0.11	0.07	TTO	1966	7,016	0.38	0.35
CYP	1966	7,021	0.27	0.25	LVA	1992	4,061	0.32	0.24	TUN	1966	7,021	0.23	0.28
CZE	1993	3,922	0.03	0.07	MAR	1966	7,021	0.20	0.21	TUR	1966	7,021	0.17	0.18
DEU	1991	4,184	0.00	0.06	MDA	1992	4,061	0.33	0.40	TZA	1966	7,021	0.40	0.39
DNK	1966	7,021	0.04	0.09	MDG	1966	7,021	0.38	0.43	UGA	1966	7,021	0.47	0.47
DOM	1966	7,013	0.39	0.43	MDV	1966	7,017	0.74	0.71	UKR	1992	4,061	0.21	0.16
DZA	1966	7,021	0.25	0.39	MEX	1966	7,021	0.15	0.15	URY	1966	7,021	0.39	0.30
ECU	1966	7,021	0.33	0.34	MKD	1993	3,922	0.25	0.36	USA	1966	7,021	0.02	0.09
EGY	1966	7,021	0.18	0.22	MLT	1966	7,021	0.29	0.25	UZB	1992	4,046	0.95	0.50
ERI	1993	3,922	0.92	0.64	MMR	1966	7,021	0.55	0.47	VEN	1966	7,021	0.37	0.32
ESP	1966	7,021	0.03	0.09	MNE	2006	1,997	0.16	0.50	VNM	1977	5,848	0.51	0.30
EST	1992	4,061	0.23	0.23	MNG	1966	7,021	0.72	0.65	YEM	1991	4,184	0.34	0.44
ETH	1993	3,922	0.19	0.29	MOZ	1975	6,077	0.54	0.49	ZAF	2000	2,909	0.00	0.05
FIN	1966	7,021	0.09	0.09	MUS	1968	6,819	0.34	0.39	ZMB	1966	7,021	0.40	0.49
FJI	1970	6,615	0.46	0.60	MWI	1966	7,020	0.45	0.52	ZWE	1980	5,499	0.37	0.36
FRA	1966	7,021	0.00	0.09	MYS	1966	7,021	0.16	0.12					

Notes: Min(t): first year of observation. N: number of observations. I and X: share of non-zero imports and exports. The Soviet Union (SUN) is in the panel until 1991, all other countries until 2018.

Table 3.A.3: Details on alternative specifications

(a) Alternative measures of political distance

	(1)	(2)	(3)	(4)	(5)
Political distance					
based on UN GA resolutions	-0.220*** (0.055)				
based on human rights resolutions, κ		-0.218*** (0.037)			
based on all resolutions, π			-0.141*** (0.045)		
based on human rights resolutions, π				-0.195*** (0.033)	
ideal point distance (Bailey et al. 2017)					-0.018 (0.025)
Other trade costs					
Economic integration agreements	0.205*** (0.044)	0.187*** (0.043)	0.207*** (0.044)	0.187*** (0.043)	0.211*** (0.043)
Log (1+tariff)	-4.057*** (0.329)	-3.766*** (0.353)	-4.048*** (0.332)	-3.785*** (0.353)	-4.080*** (0.345)
Observations	859,192	849,412	859,192	849,412	846,208

Notes: Results for aggregate trade in goods based on TradeProd. The dependent variable are trade flows from country i to country j in year t . Trade flows feature intra-national and international flows and flows with a value of 0 in the case of no reported trade. Column (1) is the baseline specification and measures political distance based on κ (Cohen 1960) using UN GA resolutions related to human rights as classified by Voeten et al. (2009). Column (2) uses all UN GA resolutions to measure political distance. Columns (3) and (4) measure political distance based on π (Scott 1955). Column (3) uses only human rights resolutions, and Column (4) uses all resolutions. Column (5) uses the ideal point distance of Bailey et al. (2017) for political distance. The measure is based on all UN GA resolutions. All columns additionally include a dummy for economic integration agreements based on the EIA database, and mean bilateral tariffs are taken from TradeProd. Furthermore, all specifications include importer-year-, exporter-year-, and pair-fixed effects. Standard errors are clustered at importer-exporter level, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.A.3: Details on alternative specifications, continued

(b) Alternative measures of economic integration

	(1)	(2)	(3)
Political distance			
based on UN GA resolutions	-0.220*** (0.055)	-0.187*** (0.054)	-0.078* (0.045)
Other trade costs			
Economic integration agreements	0.205*** (0.044)		
Economic integration agreements (WTO data)		0.218*** (0.050)	
Log (1+tariff)	-4.057*** (0.329)	-3.865*** (0.326)	-3.705*** (0.306)
No Agreement			0.000 (.)
Non-Reciprocal PTA			0.045 (0.044)
Preferential Trade Agreement			0.268*** (0.084)
Free Trade Agreement			0.216*** (0.051)
Customs Union			0.579*** (0.061)
Common Market			0.770*** (0.056)
Economic Union			0.945*** (0.066)
Observations	859,192	859,192	859,192

Notes: Results for aggregate trade in goods based on TradeProd. The dependent variable are trade flows from country i to country j in year t . Trade flows feature intra-national and international flows and flows with a value of 0 in the case of no reported trade. Column (1) is the baseline specification and measures economic integration using a dummy for any economic integration agreement based on data from Baier and Bergstrand (2021). Column (2) uses a dummy for economic integration agreements based on WTO data as reported in the Gravity dataset. Column (3) uses fixed-effects for different levels of economic integration agreements. In all columns, political distance is based on UN GA votes on all UN GA resolutions and based on κ (Cohen 1960), and mean bilateral tariffs are taken from TradeProd. Furthermore, all specifications include importer-year-, exporter-year-, and pair-fixed effects. Standard errors are clustered at importer-exporter level, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.A.3: Details on alternative specifications, continued

(c) Interval data rather than annual data

	(1)	(2)	(3)	(4)
Political distance				
based on UN GA resolutions	-0.220*** (0.055)	-0.174*** (0.059)	-0.183** (0.072)	-0.300*** (0.066)
Other trade costs				
Economic integration agreements	0.205*** (0.044)	0.208*** (0.046)	0.208*** (0.047)	0.236*** (0.045)
Log (1+tariff)	-4.057*** (0.329)	-4.027*** (0.332)	-3.707*** (0.319)	-4.431*** (0.382)
Observations	859,192	283,072	219,020	169,099

Notes: Results for aggregate trade in goods based on TradeProd. The dependent variable are trade flows from country i to country j in year t . Trade flows feature intra-national and international flows and flows with a value of 0 in the case of no reported trade. Column (1) is the baseline specification and uses data for all years featured in TradeProd. Column (2) uses data for every third year starting with 1966. Column (3) uses data for every fourth year starting with 1966. Column (4) uses data for every fifth year starting with 1966. In all columns, political distance is based on UN GA votes on all UN GA resolutions and based on κ (Cohen 1960). Columns additionally include a dummy for economic integration agreements based on Baier and Bergstrand (2021) and mean bilateral tariffs are from TradeProd. Furthermore, all specifications include importer-year-, exporter-year-, and pair-fixed effects. Standard errors are clustered at importer-exporter level, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

(d) Positive trade flows only

	(1)	(2)
Political distance		
based on UN GA resolutions	-0.220*** (0.055)	-0.233*** (0.056)
Other trade costs		
Economic integration agreements	0.205*** (0.044)	0.199*** (0.044)
Log (1+tariff)	-4.057*** (0.329)	-4.001*** (0.329)
Observations	859,192	630,946

Notes: Results for aggregate trade in goods based on TradeProd. Dependent variable are intra-national and international flows from country i to country j in year t . Column (1) is the baseline and features observations with a value of 0 when no trade is reported. Column (2) uses only observations where some (positive) trade value is reported. In all columns, political distance is based on UN GA votes on all UN GA resolutions and based on κ (Cohen 1960). Columns additionally include a dummy for economic integration agreements based on Baier and Bergstrand (2021) and mean bilateral tariffs are from TradeProd. Furthermore, all specifications include importer-year-, exporter-year-, and pair-fixed effects. Standard errors are clustered at importer-exporter level, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.A.3: Details on alternative specifications, continued

(e) Accounting for globalisation

	(1)	(2)
Political distance		
based on UN GA resolutions	-0.220*** (0.055)	-0.148*** (0.042)
Other trade costs		
Economic integration agreements	0.205*** (0.044)	0.096*** (0.034)
Log (1+tariff)	-4.057*** (0.329)	-1.371*** (0.289)
Observations	859,192	630,946

Notes: Results for aggregate trade in goods based on TradeProd. The dependent variable are trade flows from country i to country j in year t . Trade flows feature intra-national and international flows and flows with a value of 0 in the case of no reported trade. Column (1) is the baseline specification and uses data for all years featured in TradeProd. Column (2) additionally features, for each year, an indicator variable for international trade to capture the effects of globalisation (Baier et al. 2019). In all columns, political distance is based on UN GA votes on all UN GA resolutions and based on κ (Cohen 1960). Columns additionally include a dummy for economic integration agreements based on Baier and Bergstrand (2021) and mean bilateral tariffs are from TradeProd. Furthermore, all specifications include importer-year-, exporter-year-, and pair-fixed effects. Standard errors are clustered at importer-exporter level, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.A.4: Full results for heterogeneity

(a) Heterogeneity across destinations

	(1)	(2)
Political distance		
Overall	-0.220*** (0.055)	
× Destination is US, EU, or UK		-0.749*** (0.078)
× Destination are other advanced economies		0.110* (0.065)
× Destination are other economies		-0.104 (0.068)
Other trade costs		
Economic integration agreements	0.205*** (0.044)	0.190*** (0.042)
Log (1+tariff)	-4.057*** (0.329)	-4.076*** (0.325)
Observations	859,192	859,192
Wald statistic for identical interaction terms		81.22
p-value		0.00

Notes: Column (1) is the baseline specification. Column (2) adds interaction terms for country groups with advanced economies as classified by the IMF. Standard errors are clustered at the importer-exporter level.

*** p<0.01, ** p<0.05, * p<0.1.

Table 3.A.4: Full results for heterogeneity, continued

(b) Heterogeneity across origins

	(1)	(2)
Political distance		
Overall	-0.220*** (0.055)	
× Origin is US, EU, or UK		-0.758*** (0.086)
× Origin are other advanced economies		0.127** (0.062)
× Origin are other economies		-0.143** (0.065)
Other trade costs		
Economic integration agreements	0.205*** (0.044)	0.195*** (0.042)
Log (1+tariff)	-4.057*** (0.329)	-4.044*** (0.328)
Observations	859,192	859,192
Wald statistic for identical interaction terms		78.58
p-value		0.00

Notes: Column (1) is the baseline specification. Column (2) adds interaction terms for country groups with advanced economies as classified by the IMF. Standard errors are clustered at the importer-exporter level.

*** p<0.01, ** p<0.05, * p<0.1.

Table 3.A.4: Full results for heterogeneity, continued

(c) Heterogeneity across types of goods

	(1)	(2)
Political distance		
Overall	-0.220*** (0.055)	
× Final goods		0.016 (0.054)
× GVC goods		-0.539*** (0.064)
× Other goods		0.016 (0.053)
Other trade costs		
Economic integration agreements	0.205*** (0.044)	0.043* (0.026)
Log (1+tariff)	-4.057*** (0.329)	-0.964 (0.738)
Observations	859,192	147,147
Wald statistic for identical interaction terms		70.41
p-value		0.00

Notes: Column (1) is the baseline specification. Column (2) adds interaction terms for types of goods associated most frequently with each sector according to United Nations (2016, BEC,). Standard errors are clustered at the importer-exporter level. *** p<0.01, ** p<0.05, * p<0.1.

Table 3.A.4: Full results for heterogeneity, continued
(d) Heterogeneity across strategic and non-strategic sectors


	(1)	(2)
Political distance		
Overall	-0.220*** (0.055)	
× Non-strategic sectors		0.080 (0.051)
× Strategic sectors		-0.526*** (0.081)
Other trade costs		
Economic integration agreements	0.205*** (0.044)	0.045 (0.030)
Log (1+tariff)	-4.057*** (0.329)	-0.728 (0.763)
Observations	859,192	98,098
Wald statistic for identical interaction terms		36.47
p-value		0.00

Notes: Column (1) is the baseline specification. Column (2) adds interaction terms for strategic or non-strategic sectors based on Tran (2022). Column (2) uses the panel on sectoral trade discussed in Section 3.A.2. Standard errors are clustered at the importer-exporter level. *** p<0.01, ** p<0.05, * p<0.1.

Figure 3.A.1: Standardized dispersion of political distance to the US

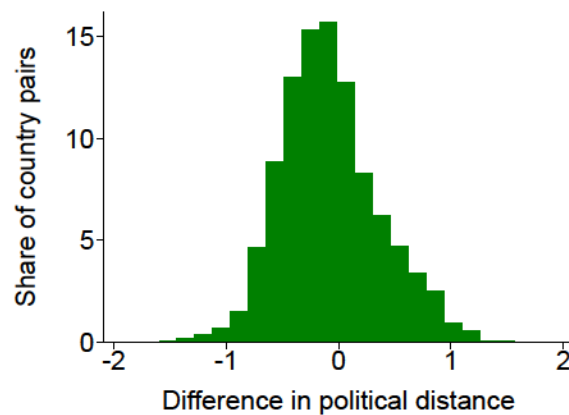
(a) based on Scott (1955) and Cohen (1960)

(b) based on Bailey et al. (2017)



Notes: Both panels show the standard deviation of political distance to the US each year using different measures of political distance. All time series are standardized. Solid lines are for measures based on countries' votes on resolutions related to human rights. Dashed lines are for measures based on all ballots. The left panels show measures based on κ (Cohen 1960) using blue lines and measures based on π (Scott 1955) using orange lines. The right panel shows political distance by the ideal point distance of Bailey et al. (2017).

Figure 3.A.2: Distribution of political distances: counterfactual vs. 2018



Notes: Histogram of the difference in political distances computed as the counterfactual distance (1972) minus the observed political distance in 2018. Positive values correspond to more distance in the counterfactual scenario than in actuality.

3.A.2 Constructing a sectoral trade panel

The TradeProd panel I use for most of the empirical analysis maximizes its coverage in terms of years and countries at the expense of a coarse sectoral resolution. In this appendix, I discuss the construction of a complementary panel on sectoral trade that maximizes sectoral coverage at the expense of coverage across countries and time. For this panel, I build upon Heid et al. (2021). They construct a panel of aggregated intra-national and international trade by combining international trade flows from the United Nations Comtrade database with domestic production from the United Nations Industrial Development Organizations Industrial Statistics Database (INDSTAT). In contrast to Heid et al. (2021), however, I focus on trade flows at the sectoral level.

To that end, I combine international trade flows at the goods level from the CEPII BACI database with data on domestic production at the sectoral level from the INDSTAT database.²¹ Both BACI and INDSTAT feature multiple versions that differ in their sample periods and the definition of goods or sectors. Out of all possible combinations of BACI and INDSTAT datasets, only one combination can be expressed using a single sectoral classification. This combination overlaps from 2012 until 2020, which is the sample period for the panel on sectoral trade. The panel on sectoral trade then features intra-national and international trade between these countries at the level of 33 sectors.

In what follows, I first introduce the individual datasets. Then I discuss how to express both using the Central Product Classification (CPC). Lastly, I provide summary statistics for the coverage of the panel on sectoral trade in terms of sectors, countries, and years, and show that, at an aggregate level, the results are in line with those obtained from TradeProd panel I use for most of the empirical analysis.

Domestic production I use data from the UN Industrial Development Organization (UNIDO) for domestic production. Their INDSTAT4 database reports, at the level of four-digit ISIC Rev. 4 codes, domestic production for up to 107 countries from 1991 to 2020. Countries report their production at different levels of aggregation (two-, three- or four-digit codes). I use data reported at the two-digit level whenever it is available. In all other cases, I compute production at the two-digit level by summing up production across all corresponding three- or four- digit codes to increase coverage. There are cases, where production is reported only for a combination of ISIC codes. In these cases, I distribute the total value of that record evenly across all ISIC codes it refers to, before aggregating. This procedure results in a country-level panel of domestic production at the level of 24 ISIC Rev. 4 two-digit codes.

International trade flows For international trade flows, I use the CEPII BACI database. It contains bilateral trade flows at the goods level (six-digit codes of the Harmonized System (HS, World Customs Organization 2022)). The CEPII offers versions based on different iterations of the HS, which is updated roughly every five years. While older HS versions allow for a longer sample period, it is difficult to accurately translate the HS codes to their more recent counterparts, since some HS codes are dropped between iterations. I use the HS

²¹The BACI database builds on UN Comtrade data as well but standardizes the definition of goods that changes every five years in UN Comtrade.

2012 version of the BACI, which covers international trade flows at the level of some 5,100 harmonized HS 2012 six-digit codes.

A common classification for domestic production and international trade flows

At this stage, domestic production is at the level of ISIC two-digit codes, while international trade flows are at the level of HS six-digit codes. I use the Central Product Classification (CPC), Revision 2.1 to bring both together. It is the only classification scheme that features correspondence tables to ISIC Rev. 4 codes (domestic production) and HS 2012 codes (international trade). In both cases, however, the correspondence is not always unique (1:1, 1:m, or m:1), and frequently more than one ISIC or HS code is assigned to multiple codes in CPC. Table 3.A.5 illustrates how I proceed, when m different ISIC codes correspond to n different CPC codes. For HS codes, I use the same procedure. Using the correspondence key, I first count, for each ISIC code, the number of assigned CPC codes (Table 3.A.5a) and divide the value recorded for the ISIC code by the number of assigned CPC codes to obtain scaled values (Table 3.A.5b). Then, I merge the scaled values to the correspondence key (Table 3.A.5c). Lastly, I sum up the scaled values for each CPC code to obtain the corresponding values in the CPC classification (Table 3.A.5d). For each sector in each country and each year, this procedure ensures that the total value recorded across all m ISIC codes is equal to the total value across all n CPC codes. The example emphasizes this: The total value using ISIC-codes is $(20 + 10 =) 30$ (Table 3.A.5b) and exactly equal to the total value using CPC-codes $(15 + 15 =) 30$.

Computing intra-national trade With international trade flows and domestic production, both by CPC two-digit codes, at hand, I follow Yotov et al. (2016) and Heid et al. (2021) and compute intra-national trade by subtracting from domestic production the sum of exports for a given year, country and sector. For international trade, I use the values from the BACI database.

Table 3.A.5: An illustration of an $m : n$ conversion

(a) ISIC-CPC correspondence			(b) Record using ISIC codes			
ISIC	CPC	n_{CPC}	ISIC	value	n_{CPC}	scaled value
1	a	2	1	20	2	10
1	b	2	2	10	2	5
2	a	2				
2	b	2				

(c) Merging (a) and (b)			(d) Sum over CPC codes	
ISIC	CPC	scaled value	CPC	sum(scaled value)
1	a	10	a	15
1	b	10	b	15
2	a	5		
2	b	5		

Notes: Table 3.A.5 illustrates the steps I follow for a $m : n$ merge, using a correspondence of two ISIC codes to two CPC codes as an example. Table 3.A.5a is the correspondence provided by the UN Statistics Division at <https://unstats.un.org/unsd/classifications/Econ/CPC.cshtml>. n_{CPC} denotes, for a given ISIC code, the number of corresponding CPC codes. Table 3.A.5b shows a generic trade record based on ISIC codes after merging n_{CPC} from Table 3.A.5a and computing the scaled value by dividing the value by n_{CPC} . Table 3.A.5c then shows the result of merging the correspondence (Table 3.A.5a) and the ISIC coded record (Table 3.A.5b). As a final step, I sum the scaled values for a given CPC code as shown in Table 3.A.5d. This procedure ensures that the total value across ISIC codes (20+10=30) is equal to the total value across CPC codes (15+15=30).

Conclusion

This dissertation analyzed two distinct fields of research in empirical macroeconomics. The first two chapters focused on firm expectations about their own variables with a particular focus on their expectation formation. In these chapters, we consolidated what we know about firm expectations about their own variables, as well as their determinants and effects. We showed that firms react differently to micro and macro news and explained this finding in a general-equilibrium model. The third chapter focused on international trade and informed the discussion about decoupling.

More specifically, in Chapter 1, we synthesized the empirical evidence on firm expectations about their own variables. To illustrate our results, we used the German ifo survey. First, we illustrated our results using the German ifo survey and established six stylized facts about firm expectations about their own variables. Next, we considered the determinants of firm expectations and found that firm-specific variables are particularly important for firms' expectation formation. Lastly, we summarized the causal effects of firms' expectations on their decisions.

Chapter 2 zoomed in on the expectation formation of firms about their own prices and production and how they react to different types of news. We distinguished two types of news. Micro news was about firms' own developments, and macro news was about aggregate developments. We showed that firms overreact to micro news and underreact to macro news. This finding was robust across numerous alternative specifications, and it arose from the German ifo survey and the "Survey on Inflation and Growth Expectations" of the Banca d'Italia. We proposed a general-equilibrium model where firms suffered from "island illusion" to rationalize these findings.

Chapter 3 picked up the ongoing discussion about rising geopolitical tensions and a potential decoupling of international trade towards politically close countries and away from politically distant ones. To study the role of political distance in international trade, I introduced it as a novel trade cost to the gravity model. I computed political distance based on the squared difference between countries' votes at the United Nations General Assembly. I found that an increase in political distance by its mean pair-specific standard deviation predicts a significant decrease in aggregate bilateral trade by 4 percent on average. To investigate the heterogeneity across sectors, I constructed a new sectoral panel of intra-national and international trade. I found the predicted decrease by an increase in political distance to be 8 percent—twice as large as the average predicted decrease—for trade flows involving the US, the EU, or the UK, and for trade flows in strategic sectors. Lastly, I used these novel insights to study a counterfactual decoupling scenario comparable to a New Cold War in 2018. In this scenario, political distances in 2018 are set to the values during the Cold War. Overall, countries are closer to each other, so the median value of a trade flow increases slightly by 5 percent. This small median change, however, hides substantial reshuffling of trade, as the median absolute change of a trade flow is 56 percent of its actual value in 2018.

In conclusion, this dissertation made two contributions that advance our understanding of macroeconomics from a theoretical and empirical perspective. The first two chapters contributed to efforts to converge to a new paradigm for rational expectations. The third chapter informed the discussion about decoupling international trade by quantifying how much trade reshuffling would occur in such a scenario.

References

- Abberger, Klaus and Klaus Wohlrabe (2006). “Einige Prognoseeigenschaften des ifo Geschäftsklimas – Ein Überblick über die neuere wissenschaftliche Literatur”. *ifo Schnelldienst* 59, 19–26.
- Altavilla, Carlo, Domenico Giannone, and Michele Modugno (2017). “Low frequency effects of macroeconomic news on government bond yields”. *Journal of Monetary Economics* 92, 31–46.
- Altig, Dave, Scott Baker, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Scarlet Chen, Steven J. Davis, Julia Leather, Brent Meyer, Emil Mihaylov, Paul Mizen, Nicholas Parker, Thomas Renault, Pawel Smietanka, and Gregory Thwaites (2020). “Economic uncertainty before and during the COVID-19 pandemic”. *Journal of Public Economics* 191, 104274.
- 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.
- Anderson, James E. (1979). “A Theoretical Foundation for the Gravity Equation”. *American Economic Review* 69 (1), 106–116.
- Anderson, James E. and Eric Van Wincoop (2003). “Gravity with gravitas: A solution to the border puzzle”. *American Economic Review* 93 (1), 170–192.
- Anderson, James E. and Yoto V. Yotov (2016). “Terms of trade and global efficiency effects of free trade agreements, 1990–2002”. *Journal of International Economics* 99, 279–298.
- Anderson, Oskar, Rainald K. Bauer, Hellmuth Führer, and Jens-Peter Petersen (1956a). “Ursachen und Typen kurzfristiger Produktions- und Preisplanrevisionen der Unternehmer”. *ifo-Studien* 2, 1–26.
- Anderson, Oskar, Hildegard Fürst, and Willi Schulte (1956b). “Zur Analyse der unternehmerischen Reaktionsweise”. *ifo-Studien* 2, 129–156.
- Anderson, Oskar and Werner H. Strigel (1960). “Empirische Untersuchungen des Unternehmerverhaltens an Hand von Konjunkturtest-Daten”. *ifo-Studien* 6, 143–156.
- 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, Ingar Haaland, Chris Roth, and Johannes Wohlfart (2022). “Narratives about the Macroeconomy”. Mimeo. Briq Institute.
- 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.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare (2012). “New Trade Models, Same Old Gains?” *American Economic Review* 102 (1), 94–130.

- Azeredo da Silveira, Rava and Michael Woodford (2019). “Noisy Memory and Over-Reaction to News”. *AEA Papers and Proceedings* 109, 557–61.
- Ba, Cuimin, J. Aislinn Bohren, and Alex Imas (2023). “Over- and Underreaction to Information”. Mimeo, University of Chicago.
- 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, Kai Carstensen, Stefan Lautenbacher, and Martin Schneider (2020). “Uncertainty is more than risk - Survey evidence on Knightian and Bayesian firms”. Mimeo. Stanford University.
- (2021). “Uncertainty and Change: Survey Evidence of Firms’ Subjective Beliefs”. NBER Working Paper 29430.
- Bachmann, Rüdiger and Steffen Elstner (2015). “Firm optimism and pessimism”. *European Economic Review* 79, 297–325.
- Bachmann, Rüdiger and Christian Bayer (2013). “‘Wait-and-see’ business cycles?” *Journal of Monetary Economics* 60 (6), 704–719.
- (2014). “Investment Dispersion and the Business Cycle”. *American Economic Review* 104 (4), 1392–1416.
- Bachmann, Rüdiger, Kai Carstensen, Manuel Menkhoff, and Martin Schneider (2023). “Firm expectations and uncertainty in normal times and times of crisis”. Mimeo. Stanford University.
- 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), 217–249.
- Bachmann, Rüdiger and Peter Zorn (2020). “What drives aggregate investment? Evidence from German survey data”. *Journal of Economic Dynamics and Control*, 103873.
- Baier, Scott L. and Jeffrey H. Bergstrand (2007). “Do free trade agreements actually increase members’ international trade?” *Journal of International Economics* 71 (1), 72–95.
- (2021). “NSF-Kellogg Institute Database on Economic Integration Agreements”. Version 2021.
- Baier, Scott L., Yoto V. Yotov, and Thomas Zylkin (2019). “On the widely differing effects of free trade agreements: Lessons from twenty years of trade integration”. *Journal of International Economics* 116, 206–226.
- Bailey, Michael A., Anton Strezhnev, and Erik Voeten (2017). “Estimating Dynamic State Preferences from United Nations Voting Data”. *The Journal of Conflict Resolution* 61 (2), 430–456.
- Balduzzi, Pierluigi, Emanuele Brancati, Marco Brianti, and Fabio Schiantarelli (2020). “The Economic Effects of COVID-19 and Credit Constraints: Evidence from Italian Firms’ Expectations and Plans”. IZA Discussion Paper 13629.
- Baldwin, Richard and Daria Taglioni (2006). “Gravity for Dummies and Dummies for Gravity Equations”. NBER Working Paper 12516.

- Balleer, Almut, Sebastian Link, Manuel Menkhoff, and Peter Zorn (2020). “Demand or Supply? Price Adjustment during the Covid-19 Pandemic”. *Covid Economics: Vetted and Real-Time Papers*, 31, 59–102.
- Banca d’Italia (2019). “Survey on inflation and growth expectations”. *Methods and Sources: Methodological Notes*.
- Bank of Japan, Research and Statistics Department (2020). “Tankan (short-term economic survey of enterprises in Japan) explanation”. Mimeo. Bank of Japan.
- (2022). “Frequently asked questions on Tankan (short-term economic survey of enterprises in Japan)”. www.boj.or.jp/en/statistics/outline/exp/tk/faqtk02.htm/#p0106, accessed on 01/12/22.
- Barrero, Jose Maria (2021). “The micro and macro of managerial beliefs”. *Journal of Financial Economics forthcoming*.
- Baumeister, Christiane (2023). “Measuring market expectations”. *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Elsevier. Chap. 14, 413–441.
- Beaudry, Paul and Franck Portier (2006). “Stock Prices, News, and Economic Fluctuations”. *American Economic Review* 96 (4), 1293–1307.
- 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.
- Beine, Michel, Simone Bertoli, and Jesús Fernández-Huertas Moraga (2016). “A practitioners’ guide to gravity models of international migration”. *The World Economy* 39 (4), 496–512.
- Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey (2013). “Managerial Miscalibration”. *Quarterly Journal of Economics* 128 (4), 1547–1584.
- Bennett, Adam (1984). “Output expectations of manufacturing industry”. *Applied Economics* 16 (6), 869–879.
- Bergstrand, Jeffrey H., Mario Larch, and Yoto V. Yotov (2015). “Economic integration agreements, border effects, and distance elasticities in the gravity equation”. *European Economic Review* 78, 307–327.
- Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma (2022). “Belief distortions and macroeconomic fluctuations”. *American Economic Review* 112 (7), 2269–2315.
- Biden, Joseph (2021). “Executive Order on America’s Supply Chains”.
- Bloom, Nicholas (2009). “The impact of uncertainty shocks”. *Econometrica* 77 (3), 623–685.
- Bloom, Nicholas, Philip Bunn, Scarlet Chen, Paul Mizen, Pawel Smietanka, and Gregory Thwaites (2019). “The impact of Brexit on UK firms”. Staff Working Paper 818. Bank of England.
- Bloom, Nicholas, Steven J. Davis, Lucia Foster, Brian Lucking, Scott Ohlmacher, and Itay Saporta Eksten (2020). “Business-level expectations and uncertainty”. NBER Working Paper 28259.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry (2018). “Really Uncertain Business Cycles”. *Econometrica* 86 (3), 1031–1065.
- Bloom, Nicholas, Takafumi Kawakubo, Charlotte Meng, Paul Mizen, Rebecca Riley, Tatsuro Senga, and John Van Reenen (2021). “Do well managed firms make better forecasts?” NBER Working Paper 29591.

- Bloom, Nicholas and John Van Reenen (2007). “Measuring and explaining management practices across firms and countries”. *Quarterly Journal of Economics* 122, 1351–1408.
- Boneva, Lena, James Cloyne, Martin Weale, and Tomasz Wieladek (2020). “Firms’ Price, Cost and Activity Expectations: Evidence from Micro Data”. *The Economic Journal* 130 (627), 555–586.
- Borchert, Ingo and Yoto V. Yotov (2017). “Distance, globalization, and international trade”. *Economics Letters* 153, 32–38.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020). “Overreaction in macroeconomic expectations”. *American Economic Review* 110 (9), 2748–82.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2019). “Diagnostic Expectations and Stock Returns”. *The Journal of Finance* 74 (6), 2839–2874.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2013). “Salience and Consumer Choice”. *Journal of Political Economy* 121 (5), 803–843.
- (2022). “Salience”. *Annual Review of Economics* 14, 521–544.
- Born, Benjamin, Jonas Dovern, and Zeno Enders (2023). “Expectation dispersion, uncertainty, and the reaction to news”. *European Economic Review* 154, 104440.
- Born, Benjamin, Zeno Enders, and Gernot J. Müller (2024). “On FIRE, news, and expectations”. *Handbook of Economic Expectations in Historical Perspective*. Ed. by Ingo Köhler, Laetitia Lenel, Alexander Nützenadel, and Jochen Streb. Routledge, forthcoming.
- Born, Benjamin, Zeno Enders, Gernot J. Müller, and Knut Niemann (2022). “Firm expectations about production and prices: Facts, determinants, and effects”. *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Elsevier. Chap. 12, 355–383.
- Born, Benjamin, Gernot J Müller, Moritz Schularick, and Petr Sedláček (2019). “The costs of economic nationalism: evidence from the brexit experiment”. *The Economic Journal* 129 (623), 2722–2744.
- Bouchaud, Jean-Philippe, Philipp Krüger, Augustin Landier, and David Thesmar (2019). “Sticky expectations and the profitability anomaly”. *The Journal of Finance* 74 (2), 639–674.
- Broer, Tobias and Alexandre Kohlhas (2023). “Forecaster (Mis-)Behavior”. *Review of Economics and Statistics* (forthcoming).
- Bruine de Bruin, Wändi, Alycia Chin, Jeff Dominitz, and Wilbert van der Klaauw (2023). “Household surveys and probabilistic questions”. *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Academic Press. Chap. 1, 3–31.
- Buchheim, Lukas, Carla Krolage, and Sebastian Link (2022). “Sudden stop: When did firms anticipate the potential consequences of COVID-19?”. *German Economic Review* 23 (1), 79–119.
- Bundesbank, Deutsche (2021). “Assessments and expectations of firms in the pandemic: findings from the Bundesbank Online Panel Firms”. Monthly report, April.
- Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko (2022). “The Macroeconomic Expectations of Firms”. *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Elsevier. Chap. 11, 321–353.

- Carlsson, Mikael and Oskar N. Skans (2012). “Evaluating Microfoundations for Aggregate Price Rigidities: Evidence from Matched Firm-Level Data on Product Prices and Unit Labor Cost”. *American Economic Review* 102 (4), 1571–1595.
- 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 and Rüdiger Bachmann (2023). “Firm surveys”. *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Academic Press. Chap. 2, 33–70.
- 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, Cheng, Takahiro Hattori, and Yulei Luo (2021). “Information rigidity and elastic inattention: evidence from Japan”. CREPE Discussion Paper 96.
- Chen, Cheng, Tatsuro Senga, Chang Sun, and Hongyong Zhang (2020). “Uncertainty, Imperfect Information, and Expectation Formation over the Firm’s Life Cycle”. Mimeo. Queen Mary University.
- Clements, Michael P., Robert W. Rich, and Joseph S. Tracy (2023). “Surveys of professionals”. *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Academic Press. Chap. 3, 71–106.
- Cohen, Jacob (1960). “A coefficient of agreement for nominal scales”. *Educational and psychological measurement* 20 (1), 37–46.
- Coibion, Olivier and Yuriy Gorodnichenko (2015). “Information rigidity and the expectations formation process: A simple framework and new facts”. *American Economic Review* 105 (8), 2644–2678.
- 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.
- Conte, Maddalena, Pierre Cotterlaz, and Thierry Mayer (2022). “The CEPII Gravity Database”. *CEPII Working Paper* (2022-05).
- Correia, Sergio, Paulo Guimarães, and Tom Zylkin (2020). “Fast Poisson estimation with high-dimensional fixed effects”. *The Stata Journal* 20 (1), 95–115.
- Dai, Mian, Yoto V. Yotov, and Thomas Zylkin (2014). “On the trade-diversion effects of free trade agreements”. *Economics Letters* 122 (2), 321–325.
- Dovern, Jonas, Lena Müller, and Klaus Wohlrabe (2020). “How Do Firms Form Expectations of Aggregate Growth? New Evidence from a Large-Scale Business Survey”. CESifo Working Paper 8179.
- Eaton, Jonathan and Samuel Kortum (2002). “Technology, geography, and trade”. *Econometrica* 70 (5), 1741–1779.

- EBDC-BEP (2019). “Business Expectations Panel 01/1980 – 06/2019, LMU-ifo Economics & Business Data Center, Munich”. *doi: 10.7805/ebdc-bep-2019*.
- Egger, Peter H., Mario Larch, and Yoto V. Yotov (2021). “Gravity Estimations with Interval Data: Revisiting the Impact of Free Trade Agreements”. *Economica* 89 (353), 44–61.
- Elenev, Vadim, Tzuo Hann Law, Dongho Song, and Amir Yaron (2022). “Fearing the Fed: How Wall Street reads Main Street”. Mimeo. Wharton School.
- 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.
- Eminidou, Snezana and Marios Zachariadis (2022). “Firms’ expectations and monetary policy shocks in the euro area”. *Journal of International Money and Finance* 122, 102556.
- Enders, Zeno (2020). “Heterogeneous Consumers, Segmented Asset Markets, and the Effects of Monetary Policy”. *Economic Journal* 130, 1031–1056.
- 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.
- Enders, Zeno, Michael Kleemann, and Gernot J. Müller (2021). “Growth Expectations, Undue Optimism, and Short-Run Fluctuations”. *Review of Economics and Statistics* 103 (5), 905–921.
- Eppinger, Peter, Gabriel Felbermayr, Oliver Krebs, and Bohdan Kukharsky (2021). “Decoupling Global Value Chains”. *CEifo Working Paper Series* (9079).
- Fally, Thibault (2015). “Structural gravity and fixed effects”. *Journal of International Economics* 97 (1), 76–85.
- Farhi, Emmanuel and Iván Werning (2019). “Monetary Policy, Bounded Rationality, and Incomplete Markets”. *American Economic Review* (11) (109), 887–392.
- Farmer, Leland, Emi Nakamura, and Jon Steinsson (2023). “Learning about the long run”. *Journal of Political Economy* (forthcoming).
- Flynn, Joel P. and Karthik A. Sastry (2022). “Attention cycles”. Mimeo. Massachusetts Institute of Technology.
- FRB Atlanta (2021). “Survey of Business Uncertainty - Methodology”.
- Freuding, Julia, Raffaella Seitz, and Klaus Wohlrabe (2021). “Was steckt hinter dem ifo Geschäftsklima? Einschätzungen der Unternehmen zu ihrer aktuellen Lage und Erwartungen”. *ifo Schnelldienst* 74 (08), 40–45.
- Gabaix, Xavier (2020). “A Behavioral New Keynesian Model”. *American Economic Review* (8) (110), 2271–2327.
- Galí, Jordi and Mark Gertler (1999). “Inflation dynamics: A structural econometric analysis”. *Journal of Monetary Economics* 44 (2), 195–222.
- García-Schmidt, Mariana and Michael Woodford (2019). “Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis”. *American Economic Review* (1) (109), 86–120.
- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer (2015). “Expectations and Investment”. *NBER Macroeconomics Annual 2015, Volume 30*. NBER Chapters. National Bureau of Economic Research, 379–431.

- Gilbert, Thomas, Chiara Scotti, Georg Strasser, and Clara Vega (2017). “Is the intrinsic value of a macroeconomic news announcement related to its asset price impact?” *Journal of Monetary Economics* 92, 78–95.
- Giustinelli, Pamela (2023). “Expectations in education”. *Handbook of economic expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Elsevier. Chap. 7, 193–224.
- Glas, Alexander and Matthias Hartmann (2021). “Uncertainty measures from partially rounded probabilistic forecast surveys”. Mimeo.
- Glynn, D. R. (1969). “The CBI industrial trends survey”. *Applied Economics* 1 (3), 183–196.
- Grasso, Adriana and Tiziano Ropele (2018). “Firms’ inflation expectations and investment plans”. Banca d’Italia Working Paper 1203.
- Guiso, Luigi and Giuseppe Parigi (1999). “Investment and demand uncertainty”. *Quarterly Journal of Economics* 114 (1), 185–227.
- Góes, Carlos and Eddy Bekkers (2022). “The Impact of Geopolitical Conflicts on Trade, Growth, and Innovation”. *WTO Staff Working Papers*.
- Head, Keith and Thierry Mayer (2014). “Gravity Equations: Workhorse, Toolkit, and Cookbook”. *Handbook of international economics*. Ed. by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff. Vol. 4. Handbook of International Economics. Elsevier. Chap. 3, 131–195.
- Heid, Benedikt, Mario Larch, and Yoto V. Yotov (2021). “Estimating the effects of non-discriminatory trade policies within structural gravity models”. *Canadian Journal of Economics/Revue canadienne d'économique* 54 (1), 376–409.
- Helpman, Elhanan, Marc Melitz, and Yona Rubinstein (2008). “Estimating Trade Flows: Trading Partners and Trading Volumes”. *Quarterly Journal of Economics* 123 (2), 441–487.
- Henzel, Steffen R. and Sebastian Rast (2013). “Prognoseeigenschaften von Indikatoren zur Vorhersage des Bruttoinlandsprodukts in Deutschland”. *ifo Schnelldienst* 66, 39–46.
- Hiersemenzel, Magdolna, Stefan Sauer, and Klaus Wohlrabe (2022). “On the representativeness of the ifo business survey”. CESifo Working Paper 9863.
- Hirshleifer, David, Anieg Low, and Siew Hong Teoh (2012). “Are Overconfident CEOs Better Innovators?” *The Journal of Finance* 67 (4), 1457–1498.
- Häge, Frank M. (2011). “Choice or Circumstance? Adjusting Measures of Foreign Policy Similarity for Chance Agreement”. *Political Analysis* 19 (3), 287–305.
- IBS-IND (2020). *Ifo Business Survey Industry 1/1980 - 06/2020*. LMU-ifo Economics & Business Data Center, Munich, doi: 10.7805/ebdc-ibs-ind-2020b.
- IMF (2023). “Goeconomic Fragmentation and Foreign Direct Investment”. *World Economic Outlook April 2023*. International Monetary Fund.
- INSEE (2007). “The French business survey on the situation and outlook in industry: methodology”. *Insee Méthodes* 117.
- Ioannou, Demosthenes, Javier J Pérez, Hans Geeroms, Isabel Vansteenkiste, Pierre-François Weber, Ana M Almeida, Irina Balteanu, Iván Kataryniuk, Maria Grazia Attinasi, Kristel Buysse, et al. (2023). “The EU’s Open Strategic Autonomy from a Central Banking Perspective. Challenges to the Monetary Policy Landscape from a Changing Geopolitical Environment”. *ECB Occasional papers* (311).

- Juodis, Artūras and Simas Kucinskas (2023). “Quantifying noise in survey expectations”. *Quantitative Economics* 14 (2), 609–650.
- Kaplan, Steven N., Mark M. Klebanov, and Morten Sorensen (2012). “Which CEO characteristics and abilities matter?” *Journal of Finance* 67, 973–1007.
- Kawasaki, Seiichi and Klaus F. Zimmermann (1986). “Testing the rationality of price expectations for manufacturing firms”. *Applied Economics* 18 (12), 1335–1347.
- Kim, Gwangmin and Carola Binder (2023). “Learning-through-Survey in Inflation Expectations”. *American Economic Journal: Macroeconomics* 15 (2), 254–278.
- Kohlhas, Alexandre and Donald Roberston (2022). “Cautions expectations”. Mimeo. Oxford University.
- Kohlhas, Alexandre and Ansgar Walther (2021). “Asymmetric Attention”. *American Economic Review* 111 (9), 2879–2925.
- Krüger, Fabian and Lora Pavlova (2020). “Quantifying subjective uncertainty in survey expectations”. KIT Working Paper in Economics 139.
- Kukuvec, Anja and Harald Oberhofer (2020). “The propagation of business expectations within the european union”. CESifo Working Paper Series 8198. CESifo Group Munich.
- Kurov, Alexander, Alessio Sancetta, Georg Strasser, and Marketa Halova Wolfe (2019). “Price drift before U.S. macroeconomic news: Private information about public announcements?” *Journal of Financial and Quantitative Analysis* 54 (1), 449–479.
- Kučinskas, Simas and Florian S. Peters (2022). “Measuring Under- and Overreaction in Expectation Formation”. *Review of Economics and Statistics*, 1–45.
- König, Heinz, Marc Nerlove, and Gilles Oudiz (1981). “On the formation of price expectations”. *European Economic Review* 16 (1), 103–138.
- Larch, Mario, José-Antonio Monteiro, Roberta Piermartini, and Yoto Yotov (2019). “On the Effects of GATT/WTO Membership on Trade: They are Positive and Large After All”. *SSRN Electronic Journal* (2019-4).
- Le Maire, Bruno (2020). “Strengthening the EU’s resilience and strategic autonomy”. *The European Files*.
- Lehmann, Robert (2023). “The Forecasting Power of the ifo Business Survey”. *Journal of Business Cycle Research* 19 (1), 43–94.
- Link, Sebastian (2020). “Harmonization and interpretation of the ifo business survey’s micro data”. *Journal of Economics and Statistics* 240 (4), 543–555.
- Lorenzoni, Guido (2009). “A Theory of Demand Shocks”. *American Economic Review* 99 (5), 2050–2084.
- Lucas, Robert E. (1972). “Expectations and the Neutrality of Money”. *Journal of Economic Theory* 4 (2), 103–124.
- (1973). “Some International Evidence on Output-Inflation Tradeoffs”. *The American Economic Review* 63 (3), 326–334.
- Lui, Silvia, James Mitchell, and Martin Weale (2010). “Qualitative business surveys: signal or noise?” *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 174 (2), 327–348.
- Ma, Yueran, Tiziano Ropele, David Sraer, and David Thesmar (2020). “A Quantitative Analysis of Distortions in Managerial Forecasts”. NBER Working Paper 26830.
- Maćkowiak, Bartosz and Mirko Wiederholt (2009). “Optimal sticky prices under rational inattention”. *American Economic Review* 99 (3), 769–803.

- Malmendier, Ulrike and Geoffrey Tate (2005a). “CEO overconfidence and corporate investment”. *The Journal of Finance* 60 (6), 2661–2700.
- (2005b). “Does overconfidence affect corporate investment? CEO overconfidence measures revisited”. *European Financial Management* 60 (11), 649–659.
- Mankiw, N. Gregory and Ricardo Reis (2002). “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve”. *Quarterly Journal of Economics* 117 (4), 1295–1328.
- Massenot, Baptiste and Yuri Pettinicchi (2018). “Can firms see into the future? Survey evidence from Germany”. *Journal of Economic Behavior & Organization* 145, 66–79.
- Mayer, Thierry, Gianluca Santoni, and Vicent Vicard (2023). “The CEPII Trade and Production Database”. *CEPII Working Paper* (2023-01).
- McFadden, Daniel (1974). “Conditional logit analysis of qualitative choice behavior”. *Frontiers in Econometrics*.
- McIntosh, James, Fabio Schiantarelli, and William Low (1989). “A qualitative response analysis of UK firms’ employment and output decisions”. *Journal of Applied Econometrics* 4 (3), 251–264.
- Meyer, Brent H., Nicholas B. Parker, and Xuguang S. Sheng (2021a). “Unit Cost Expectations and Uncertainty: Firms’ Perspectives on Inflation”. FRB Atlanta Working Paper 2021-12a.
- Meyer, Brent H., Brian Prescott, and Xuguang Simon Sheng (2021b). “The impact of the COVID-19 pandemic on business expectations”. *International Journal of Forecasting*.
- Morikawa, Masayuki (2016). “Business uncertainty and investment: Evidence from Japanese companies”. *Journal of Macroeconomics* 49, 224–236.
- (2019). “Firms’ Subjective Uncertainty and Forecast Errors”. Discussion papers 19055. Research Institute of Economy, Trade and Industry (RIETI).
- Nerb, Gernot (1987). “Der Konjunkturtest im Lichte neuerer wirtschaftstheoretischer Ansätze”. *ifo Studien* 32, 27–40.
- Nerb, Gernot and Stefan Sauer (2020). “Einführung in die ifo Umfragen”. *Handbuch der umfragebasierten konjunkturforschung*. Ed. by Georg Goldrian. Vol. 88. ifo Beiträge zur Wirtschaftsforschung. ifo Institut für Wirtschaftsforschung. Chap. 200, 1–8.
- Nerlove, Marc (1983). “Expectations, plans, and realizations in theory and practice”. *Econometrica* 51 (5), 1251–1279.
- Olivero, María Pía and Yoto V. Yotov (2012). “Dynamic gravity: endogenous country size and asset accumulation”. *The Canadian Journal of Economics / Revue canadienne d’Economie* 45 (1), 64–92.
- Rosewell, B. C. (1987). “The CBI industrial trends survey and capacity working”. *Working below capacity*. Ed. by Derek Bosworth and David F. Heathfield. London: Palgrave Macmillan UK, 3–22.
- Sauer, Stefan and Klaus Wohlrabe (2018). “The new ifo business climate index for Germany”. *CESifo Forum* 19 (2), 59–64.
- (2019). “CEO or intern – Who actually answers the questionnaires in the ifo business survey?” *CESifo Forum* 20 (2), 29–31.
- Savignac, Frédérique, Erwan Gautier, Yuriy Gorodnichenko, and Olivier Coibion (2021). “Firms’ inflation expectations: New evidence from France”. NBER Working Paper 29376.

- Schwarz, Norbert, Hans-J. Hippler, Brigitte Deutsch, and Fritz Strack (1985). “Effects of category range on reported behavior and comparative judgments”. *The Public Opinion Quarterly* 49 (3), 388–395.
- Scott, William A. (1955). “Reliability of Content Analysis: The Case of Nominal Scale Coding”. *Public Opinion Quarterly* 19 (3), 321.
- Scotti, Chiara (2016). “Surprise and uncertainty indexes: Real-time aggregation of real-activity macro-surprises”. *Journal of Monetary Economics* 82, 1–19.
- Seiler, Christian and Klaus Wohlrabe (2013). “Das ifo Geschäftsklima und die deutsche Konjunktur”. *ifo Schnelldienst* 66 (18), 17–21.
- Shiller, Robert J. (2017). “Narrative Economics”. *American Economic Review* 107 (4), 967–1004.
- Silva, João Santos and Silvana Tenreyro (2006). “The log of gravity”. *The Review of Economics and Statistics* 88 (4), 641–658.
- Sims, Christopher A. (2003). “Implications of rational inattention”. *Journal of Monetary Economics* 50 (3). Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information, 665–690.
- Smith, Jeremy and Michael McAleer (1995). “Alternative procedures for converting qualitative response data to quantitative expectations: an application to Australian manufacturing”. *Journal of Applied Econometrics* 10 (2), 165–185.
- 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.
- Strasser, Georg (2013). “Exchange rate pass-through and credit constraints”. *Journal of Monetary Economics* 60 (1). Carnegie-NYU-Rochester Conference, 25–38.
- Störmer, Eckhard, Stefan Muench, Lucia Vesnic-Alujevic, Fabiana Scapolo, and Christiano Cagnin (2021). *Shaping and securing the EU’s open strategic autonomy by 2040 and beyond*. Publications Office of the European Union.
- Taylor, Shelley E. and Suzanne C. Thompson (1982). “Stalking the elusive “vividness” effect”. *Psychological Review* (89) (2), 155–181.
- The Economist (2019). “The steam has gone out of globalisation”. *The Economist*.
- Theil, Henri (1955). “Recent experiences with the Munich business test: An expository article”. *Econometrica* 23 (2), 184–192.
- Thomas, D. G. (1995). “Output expectations within manufacturing industry”. *Applied Economics* 27 (5), 403–408.
- Tran, Hung (2022). “Our guide to friend-shoring: Sectors to watch”. *Issue brief*. Atlantic Council.
- Trebing, Michael and Caroline B. Fenske (2018). “The “Philly Fed Index” turns 50 with steadfast success”. *Economic Insights* 3 (4).
- Trieb, Thomas P and Justin Tumlinson (2013). “Learning to forecast the hard way—evidence from German reunification”. NBER Working Paper 19209.
- Truss, Elizabeth (2021). “Building the Network of Liberty”. *Speech at Chatham House*.
- United Nations (2008). “International Standard Industrial Classification of All Economic Activities Revision 4”. *Statistical Paper Series M*.
- (2015). “Central Product Classification (CPC), Version 2.1”. *UN Statistical Papers*.
- (2016). “Classification by Broad Economic Categories Rev. 5”. *UN Statistical Papers*.

- 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.
- Voeten, Erik (2013). “Data and Analyses of Voting in the UN General Assembly”. *Routledge handbook of international organization*.
- Voeten, Erik, Anton Strezhnev, and Michael Bailey (2009). “United Nations General Assembly Voting Data”. Version 29.
- Woodford, Michael (2002). “Imperfect Common Knowledge and the Effects of Monetary Policy”. Princeton: Princeton University Press, 25–58.
- Wooldridge, Jeffrey M. (2008). *Econometric Analysis of Cross Section and Panel Data*. The MIT Press, 644.
- World Customs Organization (2022). *Harmonized System Nomenclature 2022*. World Customs Organization.
- Yellen, Janet L. (2022). “Remarks by Secretary of the Treasury Janet L. Yellen on Way Forward for the Global Economy”. *US Department of the Treasury Press Release*.
- Yotov, Yoto V. (2012). “A simple solution to the distance puzzle in international trade”. *Economics Letters* 117 (3), 794–798.
- (2022). “On the role of domestic trade flows for estimating the gravity model of trade”. *Contemporary Economic Policy* 40 (3), 526–540.
- Yotov, Yoto V., Roberta Piermartini, and Mario Larch (2016). *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*. WTO iLibrary.