Università degli Studi di Palermo

AREA RICERCA E INNOVAZIONE
SETTORE DOTTORATI E CONTRATTI PER LA RICERCA
U. O. DOTTORATI DI RICERCA

Scienze Economiche e Statistiche – D069
Dipartimento di Scienze Economiche e Statistiche – DSEAS
ECON-01/A

Uncertainty: Natural Language Processing and Dynamic Factor Models to understand and forecast Uncertainty dynamics

IL DOTTORE
Antonio Pietro Maria Morreale

IL COORDINATORE
Prof. Vito Michele Rosario Muggeo

IL TUTOR
Prof. Vito Michele Rosario Muggeo

IL CO-TUTOR
Prof. Davide Furceri

CICLO XXXVII
ANNO CONSEGUIMENTO TITOLO 2026

Alla mia amata ed adorata madre Iole,
alla mia amata ed adorata zia Matilde,
alle mie bellissime e stupende Mamme,
le cui presenze continuano a vivere
in me ogni giorno, in ogni traguardo,
in ogni scelta, ed in tutto ciò che sono.

A voi, che avreste gioito di questo
traguardo più di chiunque altro.

Acknowledgments

I would like to express my sincere gratitude to my supervisors for their guidance, support, and valuable comments throughout my doctoral studies. Their feedback helped me improve the quality, clarity, and structure of this thesis.

I am especially grateful to Professor Davide Furceri for his guidance, insightful suggestions, and continuous support during the development of this thesis. His comments have been particularly important in shaping the research questions and improving the empirical analysis.

I also wish to thank Professor Vito Michele Rosario Muggeo for his important institutional support and for taking on a crucial role at a delicate stage of my doctoral path.

I am deeply grateful to Professor Stefano Soccorsi for his valuable advice, careful reading, and constructive comments. His suggestions were particularly helpful in refining the empirical strategy and improving the overall presentation of the research.

I gratefully acknowledge the Department of Economics at Lancaster University Management School for hosting me as a visiting PhD student from January to May 2025. I am thankful for the stimulating research environment and for the meaningful friendships and human connections I formed in Lancaster, both within and beyond the academic community.

I would like to thank the faculty members and administrative staff of the PhD program for their support during these years. I am also grateful to my fellow PhD colleagues for the discussions, encouragement, and shared moments that made this journey intellectually stimulating and personally meaningful. Their friendship, support, and presence during both joyful and difficult moments have meant a great deal to me.

My heartfelt thanks also go to my closest friends, whose support, patience, and willingness to listen have accompanied me throughout these years. Their presence, encouragement, and affection helped me face many difficult moments and made this journey lighter, more human, and less lonely.

My deepest gratitude goes to my father, my mother and my aunt Matilde, whose love, patience, and support have accompanied me throughout this path. In the final and most difficult part of this journey, after the passing of my mother, their presence proved essential in allowing me to continue and in helping me preserve my emotional well-being. I am profoundly grateful for their strength, care, and unconditional support.

A special thought goes to my mother. Her absence has made the final part of this journey profoundly difficult, but her memory, her love, and her example have remained a source of strength. This thesis is also for her.

Preface

This doctoral thesis studies the measurement, propagation, and macro-financial effects of economic uncertainty through three interconnected applications that combine natural language processing, high-dimensional econometrics, and dynamic factor modelling. Although the three chapters address distinct empirical questions, they are unified by a common methodological concern: how to extract economically meaningful signals from large, noisy, and heterogeneous datasets.

The broader motivation for the thesis lies in the growing transformation of empirical macroeconomics and finance into increasingly data-rich disciplines. Economic information is now dispersed across traditional macroeconomic indicators, financial market variables, textual communication, survey expectations, and high-frequency data. Yet many economically relevant objects — uncertainty, expectations, financial stress, or communication effects — are not directly observable. They must instead be inferred from imperfect proxies and latent structures embedded in large information sets. This creates a central empirical challenge: how to recover informative low-dimensional signals from complex and high-dimensional environments without losing economic interpretability.

Natural language processing and dynamic factor models provide two complementary approaches to this problem. Text-based methods allow researchers to quantify the informational content of language, transforming speeches, news articles, policy statements, and other forms of communication into measurable economic variables. Dynamic factor models, by contrast, provide a statistical framework for summarising co-movements across large panels of variables through a small number of latent common components. While these approaches have often developed separately in the literature, they share a common underlying logic: both seek to recover latent informational structures from high-dimensional data.

This thesis argues that the interaction between these methodologies is particularly useful for the study of uncertainty and financial dynamics. Textual indicators provide rich forward-looking information but often suffer from noise, instability, and context dependence. Factor-based approaches help organise that information within a coherent econometric structure capable of separating common movements from idiosyncratic fluctuations and of studying how informational shocks propagate through the macroeconomy and financial markets.

Main Contributions of the Thesis

The thesis makes three main contributions.

First, it provides an integrated discussion of the literature on uncertainty measurement, natural language processing, and dynamic factor models. The survey chapter organises these strands of research around a common informational perspective, emphasising how different approaches — text-based indicators, econometric uncertainty measures, market-based proxies, and factor-model decompositions — address related empirical problems despite relying on distinct methodologies. Particular attention is devoted to the growing interaction between textual analysis and high-dimensional econometric modelling.

Second, the thesis studies the role of dynamic factor structures in financial volatility forecasting and connectedness analysis. Using the Generalised Dynamic Factor Model within the Euro STOXX 50 universe, the empirical analysis examines whether large cross-sectional information sets improve volatility prediction relative to standard benchmark models. The chapter also investigates the network structure of factor-adjusted idiosyncratic volatility to distinguish common systemic dynamics from asset-specific spillover mechanisms.

Third, the thesis analyses the macroeconomic and international effects of Federal Reserve communication. Sentiment extracted from Federal Reserve speeches using lexicon-based and transformer-based NLP methods is embedded in a structural dynamic factor framework to identify communication-related macro-financial shocks and study their transmission to expectations, uncertainty, real activity, and international spillovers. Rather than treating central bank communication as a purely qualitative object, the chapter approaches it as a measurable informational channel operating within a high-dimensional macro-financial system.

The first chapter, therefore, provides the conceptual and methodological foundation of the thesis. It reviews the main approaches used to measure economic uncertainty, surveys modern NLP methods in economics and finance, and discusses the econometric logic of dynamic factor models and structural identification in large information sets. The chapter places particular emphasis on the links between textual indicators and factor-based frameworks, highlighting both their complementarities and their unresolved methodological challenges.

The second chapter moves from measurement to modelling. It applies the Generalised Dynamic Factor Model to realised volatility measures for Euro STOXX 50 firms, comparing its forecasting performance with standard alternatives and studying the evolution of factor-adjusted network connectedness. The analysis is motivated by the idea that financial volatility is partly driven by pervasive common forces that cannot be adequately captured by purely univariate approaches.

The third chapter studies Federal Reserve communication and its macroeconomic consequences. The analysis combines textual sentiment extraction with structural dynamic factor methods and panel local projections to investigate how communication-related shocks propagate through domestic and international macro-financial variables. The chapter also examines differences between advanced and emerging economies, emphasising the role of global financial linkages in the international transmission of uncertainty and policy communication.

Taken together, the three chapters approach uncertainty, volatility, and communication as informational phenomena whose empirical analysis requires methods capable of handling dimensionality, heterogeneity, and latent dependence structures. The objective of the thesis is not to claim that a single framework can fully resolve the problems of measurement or identification that arise in this literature. Rather, the thesis aims to show that combining textual analysis with high-dimensional econometric methods can provide a richer and more disciplined understanding of how information is generated, transmitted, and absorbed in modern economic and financial systems.

Contents

List of Acronyms

AAII American Association of Individual Investors

DFM Dynamic Factor Model

EPU Economic Policy Uncertainty

FAVAR Factor-Augmented Vector Autoregression

FFR Federal Funds Rate

GARCH Generalized Autoregressive Conditional Heteroskedasticity

GDFM Generalized Dynamic Factor Model

GPR Geo-Political Risk

HAR Heterogeneous Autoregressive Model

IRF Impulse Response Function

LM Loughran–McDonald

NLP Natural Language Processing

PC Principal Component

PCA Principal Component Analysis

SDFM Structural Dynamic Factor Model

SPF Survey of Professional Forecasters

VAR Vector Autoregression

VIX CBOE Volatility Index

WUI World Uncertainty Index

ZLB Zero Lower Bound

1. Uncertainty, Natural Language Processing and Dynamic Factor Models: A Comprehensive Survey

The measurement of economic uncertainty increasingly relies on extracting signals from large-scale textual and quantitative datasets. This chapter surveys the main empirical toolkits used to study uncertainty, with particular emphasis on natural language processing (NLP) methods and dynamic factor models (DFMs). Measuring uncertainty is an economic problem; NLP and DFMs are statistical approaches for extracting uncertainty-relevant information from high-dimensional data. In the literature, uncertainty has been measured using NLP-based indices derived from news articles, such as the Economic Policy Uncertainty Index (EPU), or from professional country reports, such as the World Uncertainty Index (WUI), as well as statistical approaches such as Jurado et al. (2015). By contrast, relatively few contributions combine textual indicators with factor-model decompositions. This chapter reviews foundational contributions, recent methodological advances, and key empirical findings, and discusses how NLP-derived measures and factor-model approaches can be used jointly in macroeconomic and financial applications. The chapter highlights conceptual complementarities between the two approaches and outlines where the literature has begun to bridge them.

1.1 Introduction

Modern economic analysis increasingly confronts a common challenge: extracting meaningful, low-dimensional signals from high-dimensional data. Whether the data consist of hundreds of macroeconomic time series, thousands of financial asset returns, or millions of words from central bank communications, researchers need principled methods to distil actionable information from complex, noisy sources. In this chapter, the central economic problem is the measurement and analysis of uncertainty, while natural language processing (NLP) and dynamic factor models (DFMs) are two methodological toolkits that have been used to address it.

The motivation for bridging these approaches is substantive, not merely taxonomic. Consider a concrete example: the Economic Policy Uncertainty (EPU) index (Baker et al., 2016), one of the most widely used indicators in applied macroeconomics, is constructed by counting keyword occurrences in newspaper articles – a text-as-data approach. The World Uncertainty Index (WUI) (Ahir et al., 2022) provides a complementary measure based on textual information from country reports, offering a broader cross-country perspective. In parallel, econometric approaches such as Jurado et al. (2015) define uncertainty as the common component in forecast error volatility using large information sets. These examples show that uncertainty can be measured using both textual and econometric methods, but the two approaches have typically developed in parallel, with limited integration.

At the same time, incorporating text-derived indicators into macroeconomic models – whether for forecasting, structural identification, or policy evaluation – requires econometric frameworks capable of handling high-dimensional and mixed-frequency data. Dynamic factor models provide precisely such a framework, with applications ranging from business cycle tracking and nowcasting to the identification of structural shocks in large panels (Forni et al., 2009; Stock and Watson, 2016).

These connections run deeper than a simple input-output chain. The intellectual structure of factor analysis – recovering latent common components from observed cross-sectional variation – is shared by both the econometric tradition of DFMs and the NLP tradition of topic modelling, where latent topics are inferred from word co-occurrence patterns across documents (Blei et al., 2003). Recent work has begun to exploit this parallel explicitly, developing *textual factors* that apply factor-analytic techniques directly to text embeddings (Cong et al., 2025). Meanwhile, advances in transformer-based language models such as finBERT (Araci, 2019; Huang et al., 2023) have substantially improved the extraction of sentiment and stance from financial and central bank texts, generating richer inputs for downstream econometric analysis.

Despite these natural complementarities, the literature that jointly integrates NLP and DFMs remains relatively scarce. Most existing work either focuses on constructing textual indicators or on modelling large datasets using factor models, without fully combining the two approaches.

No existing survey covers uncertainty measurement, NLP, and DFMs in an integrated fashion. The comprehensive uncertainty review by Cascaldi-Garcia et al. (2023) catalogues news-based, survey-based, econometric, and market-based uncertainty measures but does not discuss in detail the NLP methods underlying the text-based indices. The surveys of text-as-data methods in economics by Gentzkow et al. (2019), Ash and Hansen (2023), and Lucchetti and Cajueiro (2026) cover the methodological landscape of NLP but do not engage with the factor model frameworks into which text-derived variables are ultimately embedded. Conversely, the DFM literature focuses on estimation, identification, and forecasting without addressing the growing role of textual data as model inputs (Barhoumi et al., 2013; Stock and Watson, 2016; Poncela et al., 2021). This chapter contributes to this discussion by bringing these strands of the literature into a single conceptual review.

The purpose of this survey is to provide the conceptual background for the empirical chapters of the thesis. Chapter 2 applies the Generalised Dynamic Factor Model (GDFM) to daily volatility series for Euro STOXX 50 constituents, using the factor decomposition both to improve volatility forecasting and to separate common from idiosyncratic propagation through factor-adjusted financial networks. Chapter 3 combines NLP-based sentiment measures with a Structural Dynamic Factor Model (SDFM) to study communication-related shocks extracted from Federal Reserve speeches and their domestic and international effects. Within this framework, NLP methods provide textual measures of sentiment, uncertainty, and communication, while dynamic factor models provide the econometric structure used to extract, identify, and interpret common dynamics in high-dimensional data.

We organise the review around three main sections, each self-contained but connected by a unifying thread. Section 1.2 reviews the measurement and macroeconomic effects of uncertainty, covering the four families of uncertainty measures identified by Cascaldi-Garcia et al. (2023) — including text-based indices such as the EPU (Baker et al., 2016) and the World Uncertainty Index (Ahir et al., 2022), econometric approaches based on forecast error volatility (Jurado et al., 2015), and market-based indicators — as well as the growing evidence on how uncertainty shocks affect investment, employment, financial markets, and international spillovers. Section 1.3 surveys NLP methods for economics and finance, from lexicon-based sentiment analysis (Loughran and McDonald, 2011; Shapiro et al., 2022) through machine learning classifiers and topic models (Bybee et al., 2020) to pre-trained transformer models (Araci, 2019; Huang et al., 2023). Section 1.4 reviews the literature on central bank communication and discusses how NLP-derived textual indicators can be combined with dynamic factor models to extract information relevant for macroeconomic analysis and monetary policy transmission (Hansen and McMahon, 2016; Hansen et al., 2019; Bordo and Istrefi, 2023; Istrefi and Piloiu, 2020; Istrefi et al., 2023; Fadda et al., 2025). Section 1.5 reviews dynamic factor models, covering the approximate DFM framework (Stock and Watson, 2002a,b, 2016), the Generalized Dynamic Factor Model of Forni et al. (2000, 2005), structural identification in factor models (Forni et al., 2009; Forni and Gambetti, 2014; Bai and Wang, 2015), and applications to volatility forecasting and financial network analysis (Barigozzi and Hallin, 2016, 2017). Section 1.6 concludes.

1.2 Uncertainty — Concepts, Measurement, and Macroeconomic Effects

1.2.1 Conceptual Foundations

Uncertainty — the inability of economic agents to form precise expectations about the future path of relevant variables — plays a central role in modern macroeconomics. The distinction between risk, understood as quantifiable probability distributions, and genuine uncertainty, understood as unquantifiable ambiguity, dates back to Knight (1921). Empirical macroeconomics, however, has operationalised this distinction through several imperfect proxies rather than through a single observable object. Text-based indices capture perceived uncertainty in public narratives, survey-based measures capture disagreement across forecasters, market-based indicators capture option-implied volatility and risk premia, and econometric measures such as Jurado et al. (2015) capture the volatility of unpredictable components in large information sets. These objects are related, but they are not interchangeable. A central aim of this chapter is therefore to clarify which dimension of uncertainty each measurement approach captures.

The theoretical channels through which uncertainty affects the real economy are well established. The real options channel implies that when uncertainty is high, firms delay irreversible investment and hiring decisions because the option value of waiting increases (Ramey, 2016). The precautionary savings channel predicts that households increase savings and reduce consumption. Financial frictions amplify these effects: heightened uncertainty tightens credit spreads and reduces the supply of external finance, as lenders face greater difficulty in evaluating borrower quality. These channels are not mutually exclusive and tend to reinforce each other during episodes of elevated uncertainty.

1.2.2 Uncertainty Shocks: The Bloom Framework

The modern empirical literature on uncertainty shocks begins with Bloom (2009), who showed that large temporary uncertainty shocks generate a rapid drop and rebound in employment and output — the so-called *drop and rebound* pattern. Using the VIX as a proxy for uncertainty and a structural VAR framework, Bloom (2009) demonstrated that uncertainty shocks can explain a significant portion of business cycle fluctuations. The analysis was subsequently extended to a DSGE model with heterogeneous firms, confirming that the real options mechanism generates *wait-and-see* dynamics at the firm level — as illustrated in Figure 1.1.

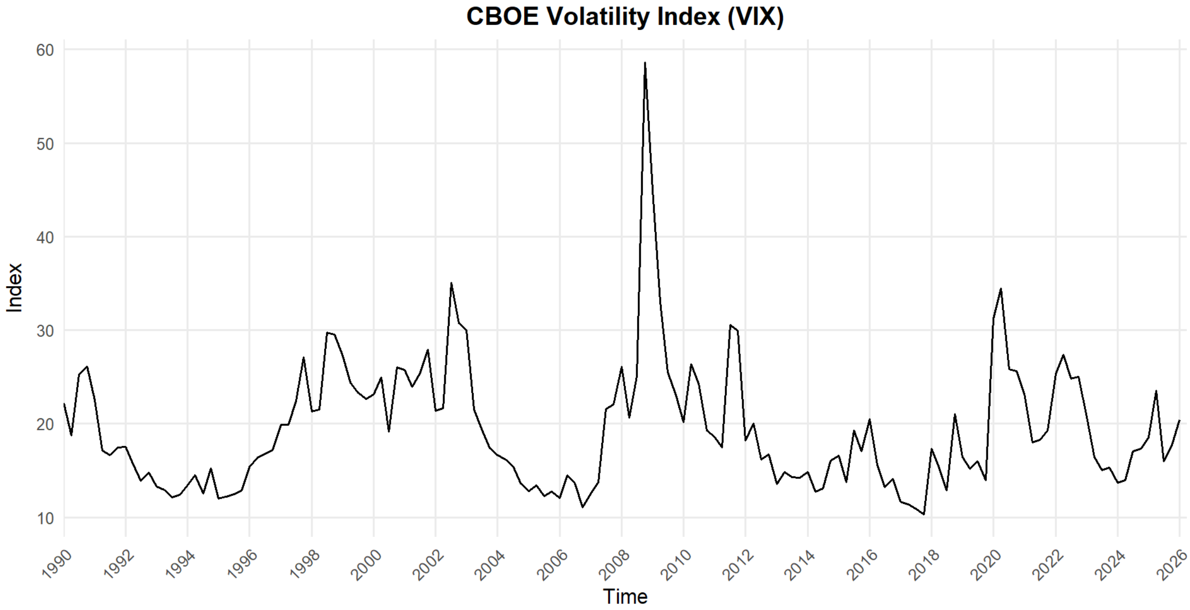

Figure 1.1: CBOE Volatility Index (VIX), 1990Q1–2026Q1.

Notes: The figure reports quarterly averages of daily VIX observations. Shaded areas denote NBER recessions. The VIX measures model-implied volatility extracted from S&P 500 index options. Data are quarterly and span 1990Q1–2026Q1. Source: CBOE via FRED®.

A key question raised by this literature is whether uncertainty shocks cause recessions, or whether recessions endogenously generate uncertainty. Ludvigson et al. (2021) address this identification challenge through a multi-equation framework that distinguishes between financial and macroeconomic uncertainty. Their evidence suggests that financial uncertainty is a more clearly exogenous driver of real downturns, while macroeconomic uncertainty tends to rise as a consequence of deteriorating economic conditions. This asymmetry has important policy implications: if uncertainty is partly endogenous, then policies that stabilise the economy may also reduce uncertainty, creating a virtuous circle.

1.2.3 Text-Based Measures

The most widely used uncertainty indicator in applied economics is the Economic Policy Uncertainty (EPU) index of Baker et al. (2016), which counts the frequency of newspaper articles containing terms related to ”economy”, ”uncertainty”, and specific policy categories. The index is available for 22 countries, and its appeal lies in its simplicity, transparency, and real-time availability. Baker et al. (2016) show that EPU spikes around major political events — elections, government shutdowns, wars — and that elevated EPU is associated with reduced investment and employment growth. The index has been extended to a Global EPU by Davis (2016), constructed as a GDP-weighted average of national EPU indices for 16 countries accounting for two-thirds of global output.

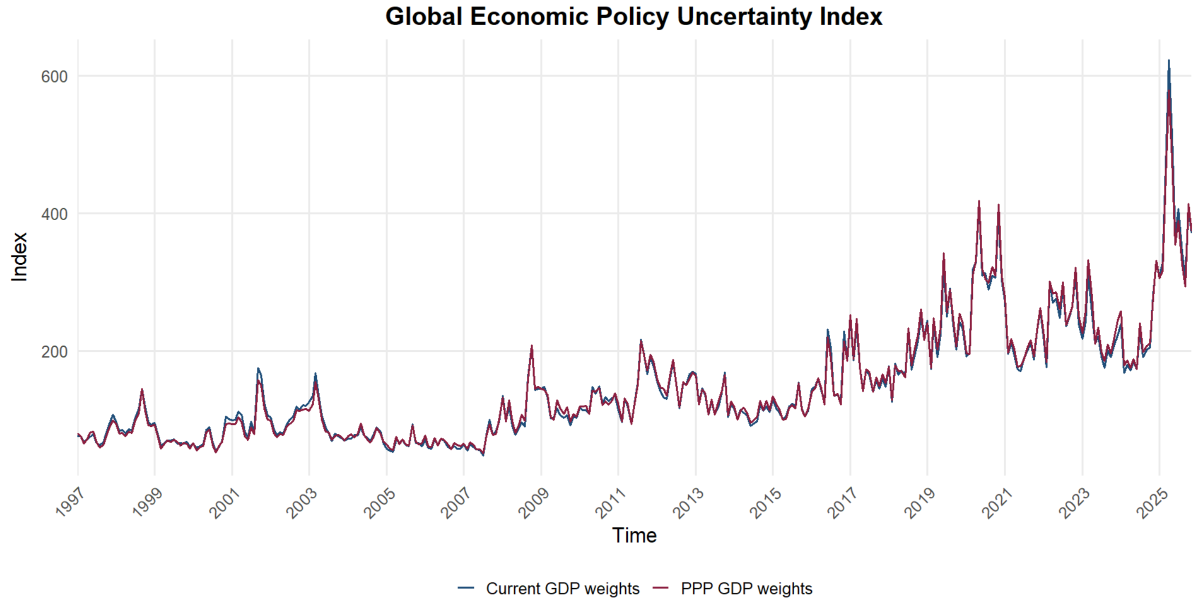

Figure 1.2: Global Economic Policy Uncertainty Index (GEPU), 1997–2026.

Notes: The GEPU index aggregates national EPU indices using GDP weights. Data are monthly. Source: Baker et al. (2016).

A conceptually distinct text-based measure is the World Uncertainty Index (WUI) of Ahir et al. (2022), constructed from the frequency of the word ”uncertainty” in the quarterly Economist Intelligence Unit country reports for 143 economies. Unlike EPU, the WUI uses a single, editorially controlled source, which reduces cross-country variation in media norms. The WUI is particularly relevant for studies of uncertainty spillovers across countries, as it provides the broadest cross-country coverage of any existing uncertainty measure.

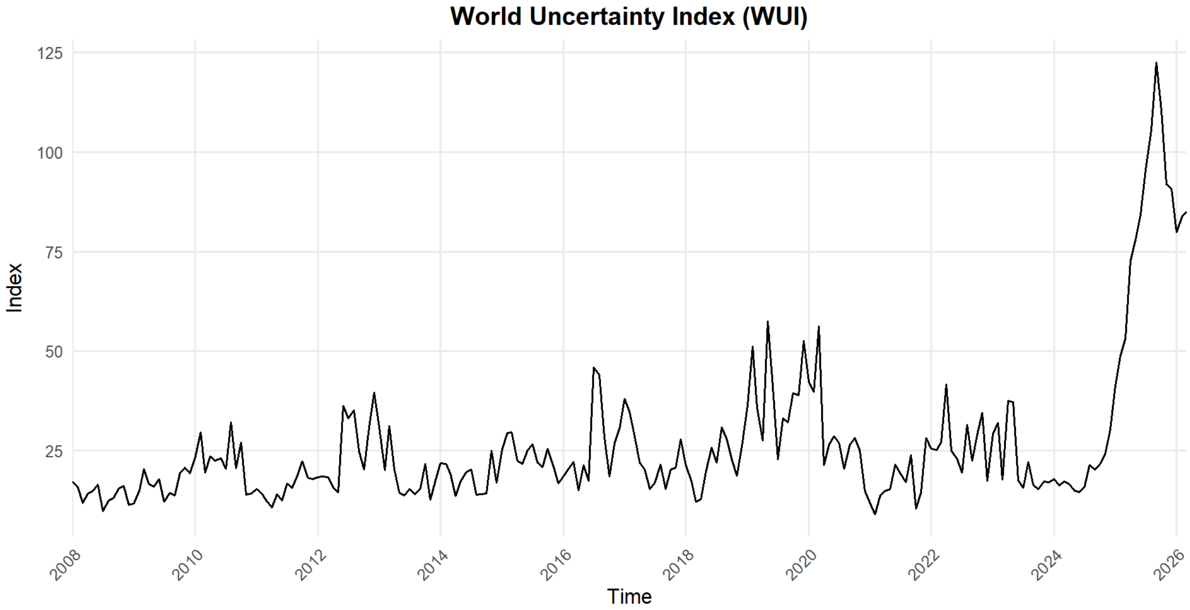

Figure 1.3: World Uncertainty Index (WUI), 2008–2026.

Notes: The WUI measures the frequency of the word *uncertainty* in Economist Intelligence Unit country reports, normalised by total word count. Data are monthly. Source: Ahir et al. (2022).

The Geopolitical Risk (GPR) index of Caldara and Iacoviello (2022) captures a different dimension of uncertainty — one driven by threats and realisations of adverse geopolitical events such as wars, terrorist attacks, and tensions between states. Caldara and Iacoviello (2022) construct the index from automated text searches in major English-language newspapers, distinguishing between ”threats” and ”acts” components. They show that adverse GPR shocks reduce investment and employment and increase the probability of economic disaster. The GPR has been recently extended to country-specific bilateral indices by Alonso Alvarez et al. (2025), who apply a standardised textual analysis methodology to national news sources in 15 languages for 34 countries. The geopolitical fragmentation index of Fernández-Villaverde et al. (2024) offers a complementary approach, measuring the degree to which the global economic system is fragmenting along geopolitical lines.

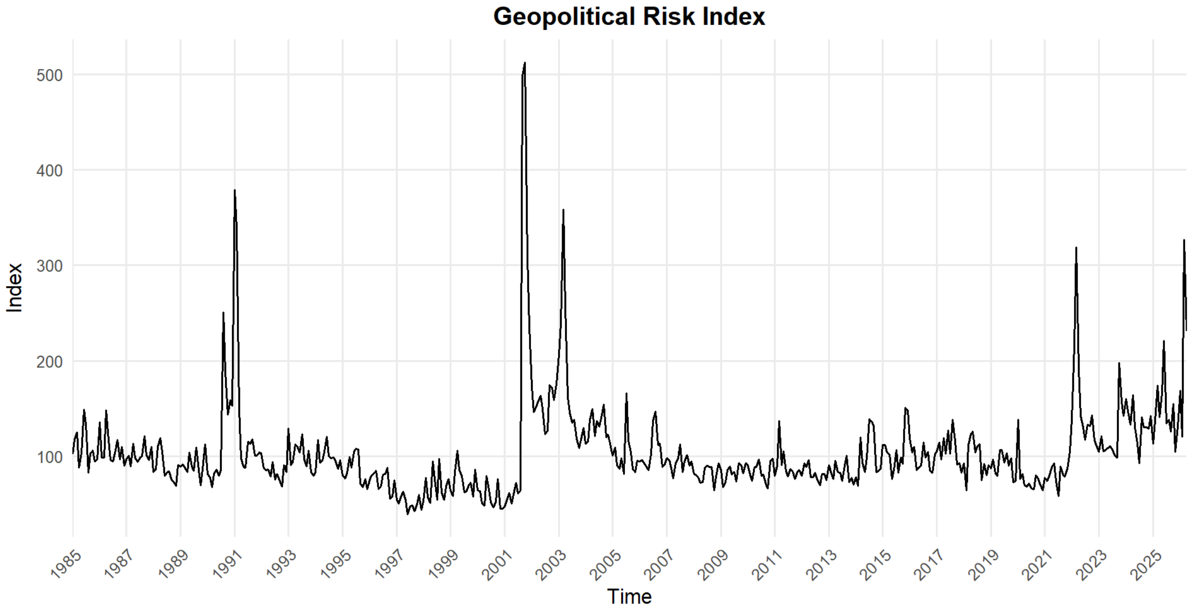

Figure 1.4: Geopolitical Risk Index (GPR), 1990–2026.

Notes: The GPR index captures geopolitical threats and adverse geopolitical events through newspaper-based text analysis. Data are monthly. Source: Caldara and Iacoviello (2022).

A common thread across EPU, WUI, and GPR is their reliance on text-based construction: all three indices ultimately rest on counting or classifying words in journalistic text. This makes them inherently dependent on the quality of the underlying natural language processing — a connection that Section 1.3 of this chapter will explore in detail.

1.2.4 Econometric Measures

An influential alternative to text-based approaches defines uncertainty purely in terms of econometric predictability. Jurado et al. (2015) construct a macroeconomic uncertainty index as the common volatility in the unpredictable component of a large number of economic indicators, estimated through a factor-augmented model — henceforth JLN index. Their approach has two key advantages: it is model-based rather than ad hoc, and it distinguishes between uncertainty (the truly unpredictable) and pure volatility. The resulting estimates reveal that *quantitatively important uncertainty episodes appear far more infrequently than indicated by popular uncertainty proxies* — a finding that challenges the frequent identification of uncertainty shocks based on EPU or VIX spikes. When the JLN index does spike, however, the episodes are larger, more persistent, and more strongly correlated with real activity than those identified by text-based proxies.

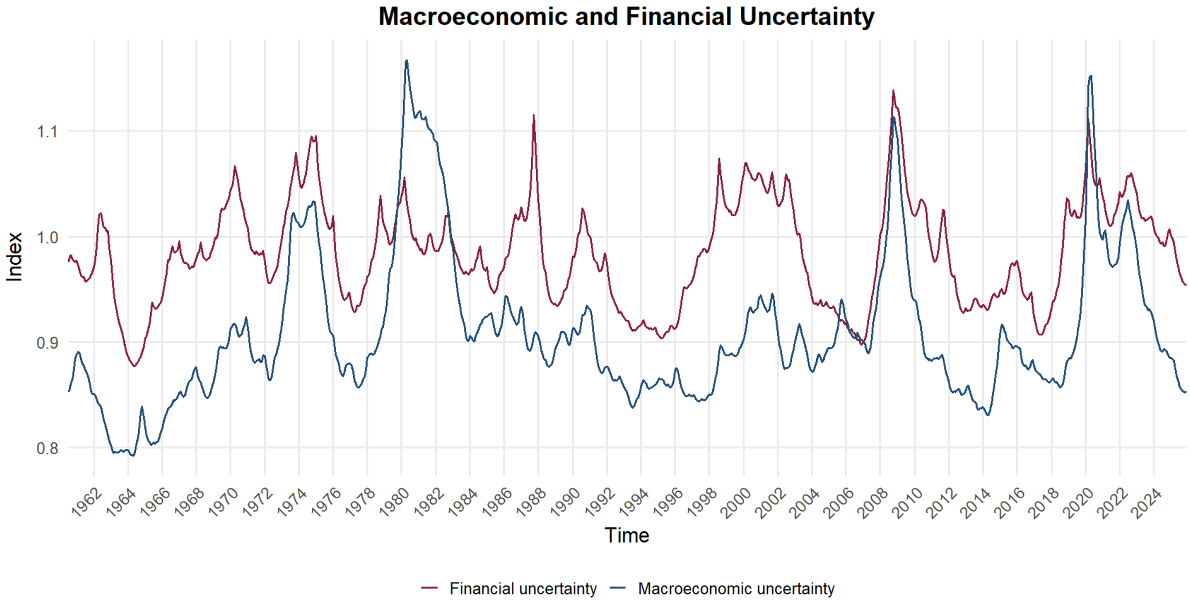

Figure 1.5: Macroeconomic and financial uncertainty indices, 1960–2024.

Notes: The indices measure the common volatility of the unpredictable component in large macroeconomic and financial datasets. Data are monthly. Source: Jurado et al. (2015).

The connection between uncertainty measurement and factor models is tight and deserves emphasis in this survey. The Jurado et al. (2015) procedure is itself a dynamic factor model: it extracts common factors from a large panel, uses them to construct forecasts of each series, and then treats the common volatility of the resulting forecast errors as the uncertainty measure. Carriero et al. (2018) pursue a related approach using a large Bayesian VAR with stochastic volatility, finding substantial commonality in uncertainty across many variables — again pointing to the factor structure of uncertainty itself. These methods illustrate a broader point: uncertainty measurement at the frontier increasingly requires the high-dimensional econometric tools reviewed in Section 1.5.

1.2.5 Market-Based Measures

Financial market prices provide high-frequency, forward-looking uncertainty proxies. The Chicago Board Options Exchange Volatility Index (VIX), which measures the implied volatility of S&P 500 index options, is the most prominent example. Its advantage is timeliness and market-based pricing; its disadvantage is that it conflates uncertainty with risk premia and investor sentiment. Bauer et al. (2021) develop a market-based monetary policy uncertainty (MPU) index derived from interest rate derivatives, which captures uncertainty specifically about the future path of policy rates.

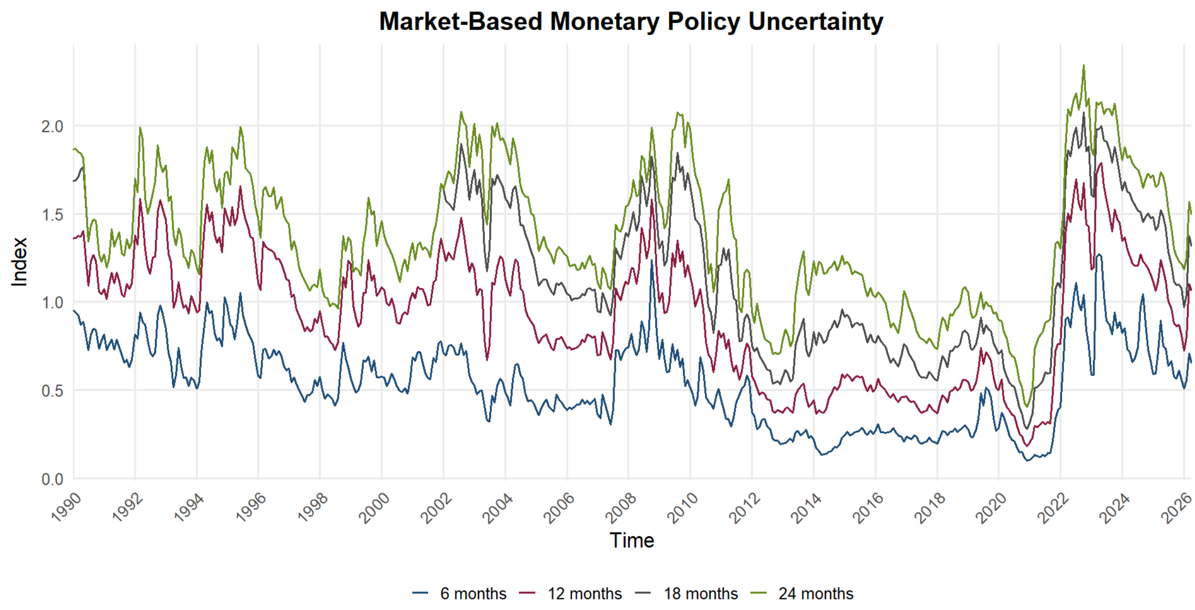

Figure 1.6: Market-based monetary policy uncertainty across alternative horizons, 1990–2026.

Notes: The figure reports the monetary policy uncertainty (MPU) index at horizons of 6, 12, 18, and 24 months. Series are monthly and reported in levels. Source: Bauer et al. (2021).

They show that MPU is distinct from broader financial uncertainty and that it spikes around FOMC announcements. Lakdawala et al. (2021) demonstrate that US monetary policy uncertainty has significant international spillover effects, reducing industrial production and increasing financial stress abroad — a finding that underscores how uncertainty generated by a single central bank's decisions can propagate globally through financial channels.

1.2.6 A Taxonomy and Open Questions

The comprehensive survey by Cascaldi-Garcia et al. (2023) organises the uncertainty landscape into four families: news-based, survey-based, econometric, and market-based measures. Each captures a different facet of the phenomenon, and the correlations across families are often surprisingly low — suggesting that "uncertainty" is not a single latent variable but a multidimensional concept. Survey-based measures, such as the dispersion of forecasts in the Survey of Professional Forecasters (Sill, 2012, 2014; Legerstee and Franses, 2015), capture disagreement among experts; this is related to, but distinct from, the unpredictability concept of Jurado et al. (2015). Gambetti et al. (2023) formalise this distinction between *agreed* and *disagreed* uncertainty, showing that they have different macroeconomic effects.

The macroeconomic effects of uncertainty have been documented along several dimensions. On investment and employment, the evidence consistently shows that uncertainty shocks depress both, though the magnitude and persistence depend on the measure used and the identification strategy employed (Ramey, 2016). On asset prices, Birru and Young (2022) examine the interaction between sentiment and uncertainty, finding that the two concepts are related but distinct in their pricing implications. Bianchi et al. (2022) show that belief distortions — systematic deviations of subjective beliefs from rational expectations — interact with uncertainty to generate macroeconomic fluctuations that standard models miss.

On international spillovers, US uncertainty propagates globally through trade and financial channels. Miranda-Agrippino and Rey (2020) document a *global financial cycle* driven largely by US monetary policy and risk appetite, of which uncertainty is a key component. Dedola et al. (2017) show that when the Fed "sneezes", other economies catch a cold, with the transmission depending on the degree of financial integration. Understanding these spillover mechanisms is essential for evaluating how textual signals from central bank communication — the subject of Section 1.3 — may propagate beyond national borders.

Two open questions stand out. First, the ***identification problem***: most uncertainty measures are endogenous to the business cycle, making it difficult to establish whether uncertainty causes recessions or the reverse (Ludvigson et al., 2021). Structural factor models, reviewed in Section 1.5, offer a promising avenue for addressing this challenge by exploiting the information content of large panels to achieve identification. Second, the ***measurement problem***: the fact that different uncertainty indices often behave differently across time and episodes raises the question of which measure is appropriate for which purpose. Both problems point toward the need for combining text-based inputs with structural econometric models — the central argument of this chapter.

1.3 Natural Language Processing

1.3.1 Text as Data: The Paradigm

Economic agents produce vast quantities of text: central bank statements, corporate filings, news articles, legislative records, and social media posts. This text contains forward-looking information, subjective assessments, and nuanced signals that are absent from standard quantitative data. The *text as data* paradigm recognises that these documents, once appropriately processed, can serve as inputs to economic models on the same footing as traditional numerical indicators.

The foundational methodological references are Gentzkow et al. (2019), who provide a formal framework for econometric analysis of text, and Ash and Hansen (2023), who survey the algorithmic toolkit available to economists, organised around four core empirical tasks: measurement, prediction, causal inference, and description. Lucchetti and Cajueiro (2026) offers the most recent comprehensive survey, covering the full pipeline from raw text to economic application with a focus on finance. On the applied side, Hassan et al. (2025) illustrate how textual analysis of earnings conference call transcripts can provide insights into firm-level responses to economic shocks that elude traditional non-text data sources.

A useful distinction separates three generations of methods: (i) dictionary-based or lexicon approaches, which count pre-defined words; (ii) ML classifiers and topic models, which learn representations from data; and (iii) pre-trained transformer models, which encode deep contextual understanding of language. Each generation subsumes its predecessor in flexibility, but also introduces new challenges of interpretability, computational cost, and validation. The following subsections trace this progression.

1.3.2 Lexicon-Based Approaches

The earliest and still widely used approach to extracting quantitative information from text relies on dictionaries: curated word lists in which each term is assigned a sentiment label (positive, negative, uncertain) or a thematic category. A document's sentiment score is then computed as the frequency-weighted sum of dictionary terms it contains.

The seminal contribution in finance is Loughran and McDonald (2011), who demonstrated that general-purpose negative word lists (such as the Harvard General Inquirer) dramatically misclassify common financial terms. Words like "liability", "tax", and "cost" are negative in everyday English but neutral or even positive in accounting contexts. Loughran and McDonald developed a finance-specific dictionary that has since become the standard in the literature, and they provided an influential survey of the field's progress and pitfalls (Loughran and McDonald, 2016).

Shapiro et al. (2022) developed a daily news sentiment index by applying a combination of lexicons to economic and financial newspaper articles from 1980 to 2015. They show that combining existing lexicons and accounting for negation substantially improves classification accuracy relative to off-the-shelf approaches, and that daily news sentiment has predictive power for macroeconomic outcomes such as industrial production and the unemployment rate. Their work illustrates both the promise and the fragility of lexicon methods: results are sensitive to the choice of dictionary and to how context modifiers (negation, intensifiers) are handled.

The EPU index of Baker et al. (2016), the WUI of Ahir et al. (2022), and the GPR index of Caldara and Iacoviello (2022), discussed in Section 1.2, all rely on variants of the lexicon approach. This is their strength (simplicity, transparency, reproducibility) and their limitation (context-blindness, inability to distinguish sarcasm, irony, or conditional statements). The question of whether more sophisticated NLP methods can improve text-based uncertainty measurement remains surprisingly underexplored.

1.3.3 Machine Learning and Topic Models

Supervised ML classifiers — support vector machines (SVMs), random forests, naive Bayes — represent the next step in sophistication. Given a labelled training set, these models learn to map text features (typically bag-of-words representations) to categories such as positive/negative sentiment or policy-relevant/irrelevant. Their advantage over dictionaries is that they can capture complex, non-linear relationships between words and outcomes; their disadvantage is that they require labelled training data, which is expensive to produce in specialised domains.

Unsupervised methods, particularly topic models, have proven especially influential in economics. Latent Dirichlet Allocation (Blei et al., 2003) treats each document as a mixture of latent topics, where each topic is a probability distribution over words. Applied to economic text, LDA and its variants can discover interpretable thematic structures without requiring prior categorisation.

Bybee et al. (2020) provide a striking application: from the full text of $800,000$ Wall Street Journal articles spanning 1984-2017, they estimate a topic model that summarises business news as interpretable topical themes and quantifies the proportion of news attention allocated to each theme over time. They show that these text-based inputs accurately track a wide range of economic activity measures and have incremental forecasting power for macroeconomic outcomes above and beyond standard numerical predictors. Their approach reveals the *text-based narratives* that underlie shocks identified in numerical economic data — a finding that connects directly to the uncertainty literature's reliance on text-based indicators.

A recent and conceptually important development is the work of Cong et al. (2025), who introduce *textual factors* — a framework that applies factor-analytic techniques directly to text embeddings. By representing documents as vectors using neural word embeddings, clustering them, and then extracting spanning factors through topic modelling, they create a factor structure of unstructured data suitable for downstream econometric regressions. This approach makes explicit the deep structural parallel between factor models for quantitative data — Section 1.5 — and latent topic structures for text: in both cases, low-dimensional latent variables explain the co-movement of high-dimensional observables.

1.3.4 Pre-trained Transformer Models

The introduction of BERT (Devlin et al., 2019) — Bidirectional Encoder Representations from Transformers — marked a paradigm shift in NLP. Unlike previous approaches that processed words independently or sequentially, BERT uses self-attention mechanisms to encode the full bidirectional context of each word in a sentence. Pre-trained on massive general-language corpora, BERT and its variants can be fine-tuned on small domain-specific datasets to achieve state-of-the-art performance on classification, sentiment analysis, and question answering tasks.

The financial domain has seen several adaptations. Araci (2019) introduces finBERT, a BERT model further pre-trained on financial text corpora. They show that even with a smaller training set and fine-tuning only part of the model, finBERT outperforms traditional ML methods in financial sentiment classification. Huang et al. (2023) develop a separate finBERT variant, pre-trained specifically on corporate filings and analyst reports, and demonstrate that it substantially outperforms both the Loughran-McDonald dictionary and other Machine Learning algorithms. Their key finding is that finBERT excels precisely where other methods fail: in identifying the positive or negative sentiment of sentences that simpler algorithms mislabel as neutral, likely because it uses contextual information in financial text. Yang et al. (2020) present a third variant, pre-trained on financial communications, confirming the robustness of the domain-adaptation approach.

The rapid improvement from dictionaries to transformers raises an important question for the uncertainty literature reviewed in Section 1.2: would uncertainty indices constructed with finBERT or similar models differ materially from those based on keyword counting? To date, surprisingly little work has addressed this question systematically.

An increasingly important strand of the literature combines NLP-derived indicators with large-dimensional macroeconomic models. In this framework, textual information extracted from central bank speeches, policy statements, or news articles is transformed into quantitative time series — such as sentiment, uncertainty, or topic-based indicators — and then incorporated into dynamic factor models or factor-augmented VARs. The advantage of this approach is that it allows qualitative information contained in unstructured text to be linked to the common macroeconomic dynamics extracted from large panels of economic and financial variables. In this sense, NLP methods provide information extraction from text, while DFMs provide a statistical representation of the underlying macro-financial structure generating the observed co-movements.

GPT-family models More recently, the rapid diffusion of generative large language models (LLMs), including GPT-family models, has further expanded the scope of text-as-data methods in economics. Unlike discriminative transformer models such as BERT, which are typically fine-tuned for classification tasks, generative LLMs can be used for summarisation, information extraction, topic labelling, reasoning over textual narratives, and constructing qualitative indicators from large corpora. These tools are increasingly being explored in macroeconomic and central bank applications, including the analysis of policy communication, narratives, and expectations. For instance, Silva et al. (2025) use a fine-tuned large language model to classify central bank communication across topic, stance, sentiment, and audience. Relatedly, Gambacorta et al. (2024) develop domain-specific language models for central banking, showing how central-bank-specific models can improve the analysis of monetary policy communication.

The present thesis focuses on discriminative models, and in particular on finBERT, for two main reasons. First, the empirical objective is not to generate text or produce open-ended classifications, but to construct a reproducible and interpretable sentiment indicator at the level of individual speeches. Second, encoder-based models such as BERT and finBERT are well-suited to supervised sentiment classification and allow the textual signal to be consistently aggregated over time and subsequently embedded in a structural dynamic factor model. Generative LLMs, therefore, represent an important extension of the literature, but the use of finBERT is more closely aligned with the measurement and identification objectives of this thesis.

1.4 Central Bank Communication and Textual Information

Central bank communication has become a primary channel of monetary policy, particularly since the zero lower bound constrained conventional interest rate tools (Filardo and Hofmann, 2014; McKay et al., 2016; Blinder et al., 2024). The volume and variety of central bank text — policy statements, minutes, press conferences, speeches, research papers, and increasingly social media (Gorodnichenko et al., 2024) — make it a natural laboratory for NLP methods.

The pioneering application is Hansen and McMahon (2016), who combine computational linguistics with a factor-augmented VAR (FAVAR) to study the macroeconomic effects of FOMC communication. By decomposing FOMC statements into *current economic conditions* and *forward guidance* components, they identify communication shocks and trace their impact on the macroeconomy. Their central finding is that shocks to forward guidance are more important than the communication of current economic conditions in explaining macroeconomic outcomes — a result with direct implications for the conduct of monetary policy. Hansen et al. (2019) extend this analysis to examine the long-run information effects of central bank communication, showing that FOMC language contains persistent signals about the committee's policy intentions.

Subsequent work has expanded the scope in several directions. Ahrens et al. (2024) study market responses to central bank speeches using high-frequency data, documenting significant intraday financial market reactions to the textual content of speeches. Gorodnichenko et al. (2023) analyze the *voice* of monetary policy, showing that the vocal tone of Fed Chair press conferences contains information beyond the text itself. Cieslak et al. (2023) study policymakers' own uncertainty, extracting measures of how uncertain FOMC members are from the language of meeting transcripts.

Within this literature, Fadda et al. (2025) study the communication of uncertainty by central banks; Istrefi et al. (2023) analyze Fed speeches to recover signals about financial stability concerns and their link to monetary policy decisions; Istrefi and Piloiu (2020) examine how public opinion interacts with central banks when economic policy uncertainty is high; and Bordo and Istrefi (2023) study how perceptions of FOMC members map into the categories of hawks, doves, and swingers. Taken together, these papers suggest that central bank communication should not be treated as a single undifferentiated signal, but rather as a multidimensional information set whose textual content affects beliefs, uncertainty, and the perceived policy stance in distinct ways.

On the sentiment side, Ehrmann and Talmi (2020) introduce a measure of semantic similarity in central bank communication, showing that departures from established language patterns increase financial market volatility. Petropoulos and Siakoulis (2021) use an adaptive NLP sentiment index derived from central bank speeches to predict financial market turbulence. Yu et al. (2023) construct a central bankers' sentiment index and link it to the global financial cycle. Beyond the Fed, Tillmann and Walter (2019) study the effect of diverging communication between the ECB and the Bundesbank, and Campiglio et al. (2025) examine central bank communication on climate change.

The identification of monetary policy shocks from text represents a particularly active frontier. Aruoba and Drechsel (2024) propose a method based on applying NLP to the documents that Federal Reserve staff prepare for FOMC meetings, capturing the committee's information set at the time of policy decisions and using ML to predict rate changes conditional on that information set. The residual constitutes a monetary policy shock. Swanson and Jayawickrema (2023) argue that speeches by the Fed Chair are more important than FOMC announcements in moving financial markets, and develop an improved high-frequency measure of monetary policy shocks that incorporates speech content.

Chapter 3 of this thesis contributes to this line of research by combining sentiment indicators extracted from Federal Reserve speeches with a structural dynamic factor model estimated on a large panel of macroeconomic and financial variables.

1.4.1 Bridges to Factor Models and Open Questions

The NLP literature reviewed in this section produces high-dimensional textual indicators: sentiment scores, topic proportions, embedding vectors, and uncertainty classifications. Converting these into useful inputs for macroeconomic analysis requires methods capable of handling dimensionality, co-movement, and dynamic structure. Dynamic factor models, reviewed in Section 1.5, provide precisely such a framework.

The connections are substantive, not merely methodological. Hansen and McMahon (2016) embed their text-derived variables into a FAVAR — a factor-augmented VAR that is a special case of the DFM frameworks discussed in the next section. Cong et al. (2025) explicitly constructs *textual factors* using the same dimensionality reduction logic as econometric factor models. Bybee et al. (2020) show that the topic proportions extracted from news articles behave like latent factors that drive economic fluctuations. These examples illustrate a broader pattern: NLP produces the high-dimensional inputs, and factor models provide the econometric structure to use them.

Several open questions merit attention. First, the *validation problem*: how should text-derived measures be validated against economic outcomes, and what is the appropriate benchmark? Second, the *domain adaptation problem*: pre-trained models like finBERT perform well in the domains on which they are fine-tuned, but their performance on out-of-domain text (e.g., central bank speeches from emerging markets) is less studied. Third, the *temporal stability problem*: the meaning of economic language evolves, and models trained on one period may perform poorly on another. Finally, and most relevant to this survey, the *integration problem*: how should text-derived indicators be formally incorporated into structural econometric frameworks, and what identification assumptions are required? The dynamic factor models reviewed in the next section provide a natural toolkit for addressing this question.

1.4.2 Limitations of NLP-Based Uncertainty Measures

NLP-based measures have expanded the empirical toolkit available to economists, but they should be interpreted as noisy proxies for latent informational content rather than as direct observations of uncertainty or communication shocks. Dictionary-based indicators depend on the choice of word lists, the treatment of negation and context, and the composition of the underlying textual corpus. Transformer-based models improve contextual classification but introduce additional concerns related to domain adaptation, training data, temporal stability, and reproducibility.

These limitations are particularly relevant for central bank communication. A speech, statement, or press conference may simultaneously convey policy intentions, assessments of current macroeconomic conditions, concerns about financial stability, and information about future risks.

Textual sentiment extracted from such documents can therefore overlap with information shocks, news shocks, or changes in broader macro-financial expectations. For this reason, text-derived indicators should be embedded in econometric frameworks that make their identifying assumptions explicit.

This point motivates the role of dynamic factor models in the empirical chapters of the thesis. Factor models do not remove the identification problem, but they provide a disciplined structure for organising high-dimensional textual and numerical information, separating common from idiosyncratic dynamics, and studying how latent informational components propagate across variables.

1.5 Dynamic Factor Models — Theory, Estimation, and Structural Identification

1.5.1 Motivation: The Curse and the Blessing of Dimensionality

Macroeconomists and financial analysts routinely face datasets in which the number of observed time series n is large relative to — or even exceeds — the number of available time periods T. National statistical offices publish hundreds of indicators; financial markets generate thousands of asset prices at intraday frequency; and, as Section 1.3 has documented, text-based methods now produce additional high-dimensional inputs (topic proportions, sentiment scores, embedding dimensions). Traditional multivariate methods such as unrestricted VARs become infeasible in this setting: the number of parameters grows quadratically in n, leading to severe overfitting and unreliable inference.

Dynamic factor models (DFMs) resolve this dimensionality problem by postulating that the co-movement of a large panel of time series is driven by a small number of latent common factors. If $\mathbf{x}_t$ is an n-dimensional vector of observables, the basic DFM assumes

$$x_{it} = \boldsymbol{\chi}_{it} + \boldsymbol{\xi}_{it}, \qquad \chi_{it} = \boldsymbol{\lambda}_i' \mathbf{f}_t, \tag{1.1}$$

where $\mathbf{f}_t$ is an $r \times 1$ vector of static factors and $\boldsymbol{\lambda}_i$ is the corresponding vector of factor loadings, and $\boldsymbol{\xi}_t$ is the idiosyncratic component.

Stacking across variables, the model can be written in matrix form as

$$\mathbf{x}_t = \boldsymbol{\Lambda} \mathbf{f}_t + \boldsymbol{\xi}_t, \tag{1.2}$$

where $\mathbf{x}_t$ is an $n \times 1$ vector, $\boldsymbol{\Lambda}$ is an $n \times r$ loading matrix, and $\boldsymbol{\xi}_t$ collects the idiosyncratic components.

The factors themselves follow a VAR process:

$$\mathbf{f}_t = \mathbf{A}(L)\mathbf{f}_{t-1} + \mathbf{R}\mathbf{u}_t, \tag{1.3}$$

where $\mathbf{A}(L)$ is a lag polynomial, $\mathbf{u}_t$ is a q-dimensional vector of reduced-form innovations with covariance matrix $\mathbf{\Sigma}_u$, and $\mathbf{R}$ is an $r \times q$ matrix mapping the dynamic shocks into the space of static factors. The innovations $\mathbf{u}_t$ represent the fundamental common shocks driving the co-movement in the large-dimensional panel of macroeconomic and financial variables.

This framework is simultaneously a dimension-reduction device (summarising n series in q factors), a forecasting tool (using factors as predictors), and a structural model (identifying shocks through the factor dynamics). The three roles correspond to three generations of the DFM literature.

1.5.2 Static Factor Models and Principal Components

The first generation of large-dimensional factor analysis in macroeconomics is associated with Stock and Watson (2002a) and Stock and Watson (2002b), who showed that principal components (PC) extracted from large panels provide powerful forecasts of key macroeconomic variables. Their approach is nonparametric in spirit: factors are estimated as the eigenvectors of the sample covariance matrix of the standardized data, with no distributional or dynamic assumptions.

Formally, the PC estimator solves

$$\hat{\mathbf{F}} = \arg\min_{\mathbf{F},\mathbf{\Lambda}} \frac{1}{nT} \sum_{i=1}^{n} \sum_{t=1}^{T} \left(x_{it} - \boldsymbol{\lambda}_i' \mathbf{f}_t\right)^2 , \tag{1.4}$$

subject to the normalization $\mathbf{F}'\mathbf{F}/T = \mathbf{I}_q$. The solution is given by the q eigenvectors associated with the q largest eigenvalues of the $T \times T$ matrix $\mathbf{X}\mathbf{X}'/n$.

Bai (2003) provide the asymptotic theory for these estimators under the "large n, large T" framework, establishing that the estimated factors converge to the true factors (up to rotation) at rate $\min(\sqrt{n}, \sqrt{T})$. The question of how many factors to retain is addressed by the information criteria of Bai and Ng (2002), which penalise the residual variance by a function of n and T:

$$IC(q) = \ln V(q) + q \cdot g(n, T), \tag{1.5}$$

where $V(q)$ is the average residual variance with q factors and $g(n, T)$ is a penalty function — e.g., $g(n, T) = (n + T)/(nT) \cdot \ln(nT/(n + T))$.

The PC-based approach has two key strengths: computational simplicity and robustness to mild misspecification of the factor dynamics. Its main limitation is that it treats the factor model as static — the loadings do not distinguish between contemporaneous and lagged effects of the factors, and the resulting factors need not correspond to economically interpretable latent variables.

1.5.3 Dynamic Factor Models and the GDFM

The Generalized Dynamic Factor Model (GDFM), originally developed by (Forni et al., 2000, 2005), provides a representation of high-dimensional time series in which each variable is decomposed into a common component driven by a small number of dynamic shocks and an idiosyncratic component. Differently from static factor models, the GDFM is best understood as a representation result under minimal assumptions on the data-generating process, rather than as a model imposing specific parametric restrictions. In particular, under suitable conditions, a large class of stationary processes admits a dynamic factor representation in which the common component captures the pervasive co-movements across variables (Hallin and Lippi, 2013).

Formally, let x_{it} denote the i-th variable in a large panel. The GDFM representation writes

$$x_{it} = \chi_{it} + \xi_{it}, \qquad \chi_{it} = \sum_{j=1}^{q} c_{ij}(L) u_{jt}, \tag{1.6}$$

where u_{jt} are q orthogonal dynamic shocks and $c_{ij}(L)$ are lag polynomials. The common component χ_{it} is thus generated by a small number q of orthogonal dynamic shocks u_t through a possibly infinite distributed lag structure. ξ_{it} is the idiosyncratic component.

This highlights the key difference with the static factor model in Section 1.5.2: while the latter imposes a finite-dimensional representation in terms of contemporaneous factors, the GDFM allows for a richer dynamic structure and can accommodate situations in which no finite-dimensional static factor representation exists.

Number of Dynamic Factors

A key issue in the empirical implementation of the GDFM is the determination of the number of dynamic factors q. In parallel with the static case discussed in Section 1.5.2, several criteria have been proposed in the literature.

Among the most widely used is the Hallin–Liska criterion (Hallin and Liška, 2007), which is based on the eigenvalues of the spectral density matrix and provides a consistent estimator of the number of dynamic factors under general conditions.

It is important to emphasise that, similarly to the information criteria used for static factor models, the Hallin–Liska criterion represents one of several available approaches. Alternative methods and tests have been proposed, and the choice among them may depend on the specific features of the data and the empirical application.

Finally, Forni et al. (2017) propose an estimator that does not require the existence of a finite-dimensional static factor representation.

GDFM for Volatility and Networks

An important extension of the GDFM framework addresses not the levels but the *volatilities* of large panels. Barigozzi and Hallin (2016) apply the GDFM to realised volatility measures, extracting common volatility factors that capture the systematic component of risk across a large cross-section of assets. Their model can be written as

$$v_{it} = \chi_{it}^{(v)} + \eta_{it}, \qquad \text{with} \qquad \chi_{it}^{(v)} = \sum_{j=1}^{m} c_{ij}(L) g_{jt}, \tag{1.7}$$

where v_{it} denotes the (log-)realized volatility of asset i, $\chi_{it}^{(v)}$ is its common volatility component, g_{jt} are the common volatility factors, and η_{it} is the idiosyncratic volatility component. This decomposition is relevant for financial risk management because it separates the portion of an asset's volatility driven by market-wide forces from the asset-specific component.

A related strand of the literature studies financial connectedness. Following the seminal contribution of Diebold and Yılmaz (2014), several approaches have been proposed to measure the transmission of shocks across financial variables. Within the GDFM framework, Barigozzi and Hallin (2017) show how dynamic factor models can be used to recover network structures in large panels. More recently, Barigozzi et al. (2021) develop time-varying GDFMs to measure financial connectedness and its evolution over time, while the FNETS methodology of Barigozzi et al. (2024) combines factor-adjusted estimation with network topology inference.

It is important to distinguish between connectedness measures and network-based approaches. Connectedness measures typically capture the transmission of systematic risk through common risk factor components and forecast error variance decompositions, whereas network approaches such as FNETS focus on the propagation of idiosyncratic shocks and firm-specific interactions after removing the pervasive common factor structure.

1.5.4 Parametric Estimation: Kalman Filter and Maximum Likelihood

The third approach to DFM estimation casts the model in state-space form and estimates it by maximum likelihood via the Kalman filter and smoother. Doz et al. (2011) propose a computationally efficient two-step procedure: first, PC are used to initialise the factors; then, quasi-maximum likelihood (QML) estimation refines the factor estimates by using a Gaussian likelihood as a quasi-likelihood criterion, without requiring the full distribution of the data to be correctly specified. Doz et al. (2012) establish the asymptotic properties of the QML estimator, showing that it is consistent and asymptotically normal as both n and T grow.

Poncela et al. (2021) provide a comprehensive survey of factor extraction via Kalman filter and smoothing, comparing the performance of PC and state-space methods across Monte Carlo designs and empirical applications. Their main conclusion is that *Kalman filter and smoother methods tend to perform better than PC when the cross-sectional dimension is moderate relative to T and when the idiosyncratic components are serially correlated*, but that the differences diminish as n grows large.

Bayesian methods offer an alternative framework that naturally accommodates uncertainty about the number of factors, the lag structure, and time variation in parameters. Bai and Wang (2015) develop identification and estimation methods for structural DFMs in a Bayesian setting, addressing the well-known rotational indeterminacy problem that plagues factor models.

1.5.5 Structural Dynamic Factor Models and Identification

For policy analysis and causal inference, extracting factors is not enough — one needs to identify *structural shocks* within the factor dynamics. This is the domain of Structural Dynamic Factor Models (SDFMs). The important point is that identification in DFMs, as in VARs, still requires assumptions; the advantage of DFMs is that they allow the researcher to condition on a much larger information set.

The identification problem in factor models has two layers. The first is the classical rotation indeterminacy: if $\mathbf{f}_t$ and $\mathbf{\Lambda}$ solve the factor model, so do $\mathbf{H}\mathbf{f}_t$ and $\mathbf{\Lambda}\mathbf{H}^{-1}$ for any invertible matrix $\mathbf{H}$. The second is the structural identification problem familiar from SVAR analysis: recovering economically meaningful shocks from the reduced-form factor innovations $\mathbf{u}_t$. Structural shocks $\boldsymbol{\varepsilon}_t$ are related to reduced-form innovations via

$$\mathbf{u}_t = \mathbf{B}\boldsymbol{\varepsilon}_t, \qquad \mathbb{E}[\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t'] = \mathbf{I}_q, \tag{1.8}$$

where $\mathbf{B}$ is the $q \times q$ impact matrix to be identified.

A key conceptual point is that non-fundamentalness is not a property of the model in the abstract, but of the structural shocks relative to a given information set. With a sufficiently rich information set, structural shocks can be fundamental. In this sense, DFMs help because they enlarge the information set available to the econometrician, thereby reducing the scope for omitted-information problems. This is the sense in which the factor-model framework can address non-fundamentalness: not by eliminating identification assumptions, but by making the conditioning set much richer.

Identification via the Cross-Sectional Dimension Forni et al. (2009) address both layers simultaneously. Their key insight is that in a large panel, the response of each of the n observables to each structural shock provides over-identifying restrictions that are simply not available in small-scale SVARs. Formally, define the impulse response function of observable i to structural shock j as

$$\frac{\partial x_{i,t+h}}{\partial \varepsilon_{jt}} = \boldsymbol{\lambda}_i'\mathbf{A}^h\mathbf{b}_j, \tag{1.9}$$

where $\boldsymbol{\lambda}_i'$ is the i-th row of $\mathbf{\Lambda}$, $\mathbf{A}$ is the VAR companion matrix, and $\mathbf{b}_j$ is the j-th column of $\mathbf{B}$.

As $n \to \infty$, the cross-sectional variation in these responses provides information that can be used to pin down $\mathbf{B}$ under suitable restrictions, but not without restrictions. The point is that the large dimensionality of the panel supplies many more moments and cross-equation restrictions than a standard small VAR.

Frequency-Domain Identification: Forni, Gambetti, Granese, Sala and Soccorsi (2025) A recent contribution by Forni et al. (2025) extends the frequency-domain SDFM methodology and applies it to the identification of supply and demand shocks in US macroeconomic data. The key intuition is that identification can be achieved by imposing economically meaningful restrictions on the contribution of structural shocks over specific frequency bands of the spectral density of the observed data. In this framework, long-run restrictions are interpreted as restrictions over low-frequency components rather than as point restrictions at frequency zero. Closely related identification ideas are also employed in Chapter 3.

The Factor-Augmented VAR (FAVAR) The FAVAR of Bernanke et al. (2005) offers a complementary identification strategy. Rather than identifying shocks within the factor structure itself, the FAVAR treats the estimated factors as observable proxies for the state of the economy and augments a standard VAR with these factors:

$$\begin{pmatrix} \widehat{\mathbf{f}}_t \\ y_t \end{pmatrix} = \mathbf{\Phi}(L) \begin{pmatrix} \widehat{\mathbf{f}}_{t-1} \\ y_{t-1} \end{pmatrix} + \mathbf{e}_t, \tag{1.10}$$

where y_t is an observed policy variable (e.g., the federal funds rate) and $\widehat{\mathbf{f}}_t$ are the estimated factors. Identification then proceeds using standard SVAR methods — Cholesky orderings, sign restrictions, or external instruments — applied to the augmented system. The FAVAR's appeal lies in its simplicity and flexibility: it can accommodate any identification scheme that works in a VAR, while exploiting the information content of a large panel through the estimated factors.

The FAVAR has become the workhorse model for evaluating the effects of monetary policy in a data-rich environment. It also provides the natural econometric framework for incorporating text-derived communication variables into structural analysis: Hansen and McMahon (2016) embed their NLP-derived FOMC communication measures directly into a FAVAR to identify communication shocks and trace their macroeconomic transmission.

1.5.6 Other Applications in Macroeconomics and Finance

Dynamic factor models have found application across a wide range of domains, and their usefulness is clearest when the researcher faces a noisy, high-dimensional, and time-sensitive data environment. Their central advantage is that they convert a large and heterogeneous information set into a small number of latent components that can be updated in real time and interpreted economically. This makes them particularly well-suited to tasks in which the goal is not only forecasting, but also tracking the current state of the economy, distinguishing common from idiosyncratic movements, and linking reduced-form variation to structural interpretations.

Nowcasting and Business Cycle Monitoring. The ability of DFMs to handle mixed-frequency data and missing observations makes them the backbone of real-time macroeconomic monitoring. In practice, this is the setting in which the appeal of factor methods is most immediate: official macroeconomic indicators are published with delays, at different frequencies, and often with substantial revision, while policy institutions require timely assessments of current conditions. The nowcasting literature shows how a large information set can be used to reconstruct the present more effectively than a small set of aggregate indicators. Giannone et al. (2008) show that a large panel of macroeconomic series can be used to nowcast GDP in real time, even when the data arrive with ragged edges and heterogeneous publication lags. In a related application, Altissimo et al. (2010) develop the New Eurocoin indicator, which extracts the medium-term component of euro-area GDP growth in real time using a GDFM. Aruoba et al. (2009) provide a comparable real-time business-cycle indicator for the US economy. Together, these contributions show that the practical value of DFMs lies not only in asymptotic theory but in their ability to summarise a noisy data flow into a timely measure of current economic conditions.

Macroeconomic Forecasting and Monetary Policy Analysis. A second major application concerns forecasting and the transmission of monetary policy. Stock and Watson (2002a) establish the benchmark result that factor-augmented forecasts outperform autoregressive and small-VAR benchmarks for a wide range of US macroeconomic variables. Subsequent work has explored whether gains persist out of sample, whether they vary across forecast horizons, and whether nonlinear factor models can improve further. Forni and Gambetti (2010) apply a structural factor model to study the dynamic effects of monetary policy, while Forni et al. (2018) show that infinite-dimensional factor spaces can improve forecasting performance. These contributions are especially relevant for the present thesis because they show how factor methods can move from measurement to structural interpretation. The FAVAR framework of Bernanke et al. (2005) has been widely adopted for studying the transmission of monetary policy. By extracting factors from large panels, the FAVAR avoids the "price puzzle" and other anomalies that afflict small-scale VARs.

Financial Volatility and Risk. The GDFM extensions of Barigozzi and Hallin (2016, 2017) demonstrate that factor models can also be used when the object of interest is not the level of activity, but its volatility or risk component. In this setting, the common component is interpreted as the systematic part of volatility shared across assets, while the idiosyncratic component captures asset-specific movements. This is valuable for financial risk management because it separates market-wide volatility from firm-level or asset-level noise. The broader implication is that DFMs are not confined to macroeconomic aggregates: they can also be used to study the propagation of financial stress, the time variation of systemic risk, and the co-movement patterns that underpin financial connectedness. In that sense, the GDFM extends the same logic of dimensionality reduction to a different economic object.

What these applications have in common is not the particular target variable, but the econometric structure that makes them possible. In each case, the relevant economic object is latent, dynamic, and embedded in a high-dimensional environment. DFMs provide a way to recover that object from the data without forcing the researcher to choose a narrow information set a priori. This is why the same framework can be used for nowcasting, monetary policy analysis, uncertainty measurement, and financial volatility. The applications are different, but the underlying logic is the same: extract common variation, separate it from idiosyncratic noise, and use the resulting representation for forecasting, monitoring, or structural interpretation.

1.6 Conclusion

This chapter has reviewed the literature on economic uncertainty, natural language processing, and dynamic factor models to clarify how different empirical approaches recover uncertainty-related information from high-dimensional data. The main point is not that these literatures have already converged into a single unified framework, but rather that they address related measurement and identification problems using complementary tools.

Three lessons emerge from the review. First, uncertainty is a multidimensional empirical object. Text-based indicators capture uncertainty as reflected in public narratives and institutional reports; survey-based measures capture disagreement across forecasters; market-based indicators capture option-implied volatility and risk premia; and econometric measures such as Jurado et al. (2015) capture the volatility of unpredictable components in large information sets. These measures are related, but they should not be treated as interchangeable.

Second, NLP methods have made textual information increasingly usable in empirical macroeconomics and finance. They allow researchers to quantify sentiment, uncertainty, policy tone, and communication in documents that would otherwise remain qualitative. At the same time, text-derived measures remain sensitive to source selection, dictionary choices, model training, domain adaptation, and temporal instability. Their economic interpretation, therefore, requires caution and validation against relevant outcomes.

Third, dynamic factor models provide a natural econometric architecture for organising high-dimensional numerical and textual information. Factor models reduce dimensionality, extract latent common components, and provide a framework for forecasting and structural analysis in large panels. This is why they play a central role in the empirical chapters of the thesis: Chapter 2 uses the GDFM to study common and idiosyncratic volatility dynamics in European equity markets, while Chapter 3 combines NLP-derived communication measures with an SDFM to analyse Federal Reserve communication and uncertainty dynamics.

Taken together, the literature reviewed in this chapter suggests that uncertainty should not be interpreted as a single measurable object, but as a family of related phenomena that can be proxied through textual information, forecast disagreement, financial volatility, and latent econometric components. The empirical chapters build on this insight by combining NLP techniques with dynamic factor models to study how uncertainty-related information is measured, transmitted, and propagated across macroeconomic and financial systems.

2. EuroSTOXX 50: Does the GDFM help predict periods of high volatility?

This chapter studies whether dynamic factor information improves short-horizon volatility forecasting for EuroSTOXX 50 constituents. Using daily Parkinson volatility for 46 firms over 2002–2022, I decompose volatility into common and idiosyncratic components through the Generalized Dynamic Factor Model (GDFM) and compare forecast accuracy against HAR and GARCH(1,1) benchmarks. The GDFM delivers positive portfolio-level MSE gains, with stronger improvements for sectors and firms whose volatility is more exposed to common variation. I then use factor-adjusted idiosyncratic networks to examine how residual volatility linkages evolve during periods of market stress. Network modularity declines sharply during the Global Financial Crisis and the Covid-19 episode, suggesting that idiosyncratic propagation becomes less sectorally segmented when systemic volatility dominates. The results indicate that factor-based volatility forecasts are most informative when common volatility components are empirically important. At the same time, network diagnostics provide complementary evidence on the structure of residual risk transmission.

2.1 Introduction

The introduction of the euro in 2002 created, for the first time, a structurally integrated European equity market in which exchange-rate risk among member states was eliminated and cross-border investment flows deepened. Since then, a large body of research has documented substantial co-movement in European equity volatility, driven by macroeconomic shocks, monetary policy, and global financial conditions. However, it is not obvious that stronger cross-sectional dependence during turbulent periods automatically translates into superior out-of-sample forecasting performance for multivariate models relative to univariate alternatives. Forecasting gains depend on the persistence and predictability of the common component, the signal-to-noise ratio, and the horizon of interest. For this reason, the empirical relevance of cross-sectional information remains an open question.

A natural framework for addressing this question is the Generalized Dynamic Factor Model (GDFM). In textbook finance, asset returns and volatilities can be decomposed into a market (systematic) component and an idiosyncratic component, where only the former represents non-diversifiable risk. The GDFM formalises this decomposition in a high-dimensional time-series setting by extracting a small number of latent factors that capture the common dynamics across assets, while the idiosyncratic components represent firm-specific variation. This structure reduces dimensionality while preserving the cross-sectional information that univariate models discard. The GDFM has been successfully applied to macroeconomic and financial data (Forni et al., 2000, 2005), and recent work shows that it can recover market volatility shocks and provide consistent prediction intervals (Barigozzi and Hallin, 2016, 2020). Yet, despite its theoretical appeal, the GDFM has not been systematically evaluated as a practical forecasting tool for realised volatility at the individual-stock level in European equity markets.

Beyond forecasting, understanding how firm-specific volatility propagates across the market requires a complementary perspective. Even though idiosyncratic risk is diversifiable in principle, shocks to large or highly interconnected firms can have aggregate consequences through balance-sheet linkages, liquidity spirals, or correlated trading behaviour. This insight motivates the use of network-based methods to study financial spillovers. A growing literature interprets forecast error variance decompositions from VAR models as weighted directed networks (Diebold and Yılmaz, 2014), and several studies have applied this approach to European markets (Mensi et al., 2018; Aslam et al., 2021; Huynh et al., 2021). However, these analyses typically apply network methods directly to observed returns or volatilities, without first removing the common factor structure. This omission can confound market-wide shocks with firm-specific propagation channels, as emphasised by Giglio and Xiu (2021) and Bekaert et al. (2025), who show that failing to model common components inflates estimated correlations and distorts risk assessments.

This chapter contributes to the literature by combining GDFM-based forecasting with a factor-adjusted network analysis of idiosyncratic volatility for the EuroSTOXX 50 universe. The contribution is threefold. First, we provide a comprehensive out-of-sample evaluation of the GDFM for forecasting realised volatility of European blue-chip equities, comparing its performance with standard univariate benchmarks such as HAR (Corsi, 2009) and GARCH(1,1) (Engle, 1982; Bollerslev, 1986). Second, we construct idiosyncratic Granger-causality networks using the FNETS methodology of Barigozzi et al. (2024), applied to the idiosyncratic components extracted from the GDFM. This yields a sparse directed network that isolates firm-specific propagation channels, free from the confounding influence of market-wide shocks. Third, we relate the network structure to the behaviour of the common component, showing how the two perspectives jointly characterise periods of market stress and stability.

Relative to the existing literature, this chapter differs in two important respects. First, while several studies analyse volatility spillovers in Europe (Mensi et al., 2018; Aslam et al., 2021; Huynh et al., 2021), none combines a dynamic factor decomposition with a factor-adjusted network analysis. Second, while Barigozzi and Hallin (2017) apply a GDFM-based approach to US markets, this chapter focuses on European equities and jointly examines forecast performance and the topology of the idiosyncratic volatility network. This dual perspective provides new insights into how common and idiosyncratic components shape volatility dynamics in an integrated financial market.

The link between forecast performance and idiosyncratic network topology is not mechanical, but follows from the same common–idiosyncratic decomposition. If a large share of volatility is driven by pervasive common shocks, factor-based forecasts are expected to perform relatively better than purely univariate models, since they exploit information from the cross-section. At the same time, once common volatility is removed, the residual idiosyncratic network can be used to assess whether firm-specific volatility propagation remains organised along sectoral communities or becomes more diffuse during stress episodes. The chapter, therefore, uses forecasting results and network diagnostics as complementary evidence: the former evaluates the predictive content of common volatility components, while the latter describes the residual propagation structure after those common components have been filtered out.

The remainder of the chapter is organised as follows. Section 2.2 reviews the related literature. Section 2.3 describes the dataset and volatility proxies. Section 2.4 presents the GDFM forecasting procedure and the factor-adjusted network estimation. Section 2.5 reports the empirical findings. Section 2.6 concludes.

2.2 Related Literature

This chapter connects three strands of the literature: (i) volatility forecasting with factor models, (ii) network-based connectedness analysis in financial markets, and (iii) empirical studies of European equity volatility and contagion. This section reviews each strand and highlights how the present work contributes to them.

2.2.1 Volatility Forecasting and Factor Models

The econometrics of volatility forecasting has been dominated by two classes of models. Univariate conditional volatility models — notably the GARCH family initiated by Engle (1982) and Bollerslev (1986) — remain the standard tool for modelling time-varying variance at the individual asset level. The Heterogeneous Autoregressive (HAR) model of Corsi (2009), which captures the long-memory properties of realised volatility through a parsimonious combination of daily, weekly, and monthly components, has become a leading alternative, particularly when realised or range-based volatility measures are available.

Both approaches, however, treat each asset in isolation. When applied to large portfolios, they ignore the cross-sectional dependence among assets — dependence that intensifies during crises and is central to risk management. Multivariate GARCH specifications address this limitation in principle, but the number of parameters grows quadratically with the panel dimension, making estimation infeasible for panels of 40 or more series (Lütkepohl, 2005). Symitsi et al. (2018) compare popular covariance forecasting models on a portfolio of major European equity indices and find that a parsimonious VHAR model achieves the best performance, but even this approach operates at the index level rather than at the individual stock level.

Generalized Dynamic Factor Models (GDFMs) provide a scalable alternative. The foundational theory was developed by Forni et al. (2000) and Forni and Lippi (2001), who showed that a large panel of time series can be decomposed into a common component, driven by a small number of latent dynamic factors, and an idiosyncratic component specific to each series. Barigozzi and Hallin (2016) apply this framework to volatilities, demonstrating that the GDFM can recover unobserved market volatility shocks in a fully nonparametric and model-free manner. Barigozzi and Hallin (2020) extend the analysis by deriving consistency results and constructing one-step-ahead conditional prediction intervals, showing that the GDFM outperforms componentwise GARCH(1,1) on the S&P 100 panel. Barigozzi et al. (2019) further generalise the framework to identify both global and local shocks in international financial markets, including European ones, though their focus is on shock identification rather than systematic forecasting. Alessi et al. (2009) combine the GDFM with GARCH innovations for multivariate volatility forecasting on large panels, confirming the superiority of the factor-based approach.

More recently, Giglio and Xiu (2021) have shown that omitting latent factors from asset pricing models biases the estimation of risk premia, because the unmodelled common variation leaks into the idiosyncratic component. Bekaert et al. (2025) provide complementary evidence that correctly estimating the idiosyncratic variance component is essential for accurate volatility forecasts, especially during periods of market stress. These findings motivate the use of factor-based decompositions not merely as a dimensionality-reduction technique but as a substantive requirement for unbiased risk measurement.

A further strand of the volatility forecasting literature has explored the role of cross-market information. Bunčić and Gisler (2016) show that US equity volatility information substantially improves out-of-sample forecasts in 17 international markets, and Tang et al. (2022) confirm this finding specifically for major European indices (CAC, DAX, FTSE). These results underscore the importance of cross-sectional information for European volatility forecasting, but the information is incorporated through US volatility as a predictor variable rather than through a structural factor decomposition of the European panel itself.

Despite this body of work, the GDFM has not been systematically tested as a volatility forecasting tool for European equities at the individual stock level. The present chapter fills this gap by conducting the first rolling-window, out-of-sample comparison of GDFM-FNETS against HAR and GARCH(1,1) on a panel of 46 EuroSTOXX 50 constituents, using data from 2002-2022 and focusing on the GFC and *Covid-19* crisis episodes.

2.2.2 Network Connectedness in Financial Markets

Diebold and Yılmaz (2014) show that the generalised forecast error variance decomposition (GFEVD) of a VAR model can be interpreted as a weighted directed network, where each edge represents the share of forecast error variance that one variable transmits to another. This insight has given rise to a large literature on financial connectedness, with applications to systemic risk monitoring, identification of systemically important institutions, and measurement of cross-market spillovers.

Subsequent methodological advances have extended the framework in several directions: Chen et al. (2025) develop methods for estimating time-varying networks in high-dimensional settings; Martin et al. (2024) propose network-informed restricted VAR models; Haeran Cho and Fearnhead (2024) address high-dimensional time series segmentation via factor-adjusted VAR modelling; Beyhum and Striaukas (2024) develop tests for sparse idiosyncratic components in factor-augmented regressions; and Katsouris (2024) study robust estimation in network VAR models with nonstationary regressors.

A key innovation in the network estimation literature is the Factor-Adjusted Network Estimation for Time Series (FNETS) procedure of Barigozzi et al. (2024), implemented in the `R` package described in Owens et al. (2023). FNETS first removes the common factor component from the observed panel via the GDFM, then estimates a sparse VAR on the idiosyncratic residuals using LASSO-type regularisation. This procedure yields a Granger causality network that reflects only the firm-specific interactions among series, net of the pervasive market-wide shocks. Priola et al. (2021) propose weighting such Granger causality networks by the GFEVD via element-wise multiplication, so that only statistically significant causal linkages receive spillover weights. Greenwood-Nimmo et al. (2021) develop a related connectedness measure based on the total spillover index.

The topology of financial networks is commonly characterised by summary measures from graph theory. Barabási and Pósfai (2016) provide standard definitions of network diameter, average path length, and modularity. Modularity, in particular, measures how strongly a network is divided into distinct communities. The dissolution of community structure during crises is theoretically predicted by the liquidity spiral mechanism of Brunnermeier and Pedersen (2009) and the fire-sale contagion channel of Cont and Schaanning (2017) and Shleifer and Vishny (2011): when margin calls and forced liquidations propagate across sectors, the boundaries between communities erode.

This chapter contributes to the network connectedness literature by applying the FNETS-based Granger causality network, weighted by GFEVD, to the idiosyncratic component of European blue-chip volatilities — a combination that, to the best of my knowledge, has not been attempted before.

2.2.3 European Equity Volatility and Contagion

The empirical literature on financial connectedness in Europe is richer than is sometimes acknowledged, but it has developed along two separate tracks that leave a gap precisely where the present chapter contributes.

The first track focuses on *financial institutions and sovereign risk*. Grilli et al. (2015) study market connectivity and financial contagion using simulation-based models. Paltalidis et al. (2015) analyse the transmission channels of systemic risk in the European financial network, with a focus on banking institutions. Magkonis and Tsopanakis (2020) examine the financial connectedness between eurozone core and periphery countries at a disaggregated level. Deev and Lyócsa (2020) construct cross-quantilogram networks for 205 European financial institutions, finding strong interconnectedness across all quantiles. Foglia and Angelini (2020) map tail-risk spillovers between banks, insurance companies, and shadow banking in the Eurozone using the TENET model.

The second track applies the Diebold–Yılmaz connectedness framework *directly to equity market volatilities*. Mensi et al. (2018) investigate volatility spillovers among GIPSI and global stock markets, finding that crises intensify spillovers. Aslam et al. (2021) estimate intraday volatility spillovers across twelve European stock markets during *Covid-19*, reporting that approximately 78% of forecast error variance comes from cross-market spillovers. Most closely related to the present chapter, Huynh et al. (2021) analyse tail-risk networks for the 46 largest Eurozone firms over 2006-2020 — a panel nearly identical to mine — but use the TENET model rather than a factor-based decomposition.

What is missing from both tracks is a *structural decomposition* of equity volatility before network construction. All the studies cited above apply connectedness measures to observed returns or volatilities, without first separating the common factor component from the idiosyncratic one. As a consequence, the resulting networks conflate market-wide shocks with firm-specific propagation channels. The GDFM-FNETS framework used in this chapter addresses this limitation by first extracting the common component and then constructing the idiosyncratic Granger causality network from the idiosyncratic residuals. In this way, the resulting network captures spillovers that are not driven by market-wide shocks, offering a complementary perspective on firm-level propagation mechanisms.

2.3 Data

The dataset comprises 46 of the 50 constituents of the EuroSTOXX 50 index. Four stocks were excluded because their price history on Yahoo Finance does not cover the full sample period, either due to later listing dates or to corporate events (mergers, delistings) that interrupt the series. The remaining 46 stocks represent the largest and most liquid equities in the euro area, spanning all major sectors; Table A.2 reports the full list with ticker, company name, sector, and country.

Daily data — consisting of open, high, low, and close prices — were retrieved from Yahoo Finance and cover every trading day from 2 January 2002 to 30 December 2022, yielding approximately $5,400$ observations per stock. The sample begins with the introduction of euro banknotes and coins, a deliberate choice: the euro's physical launch created a structurally integrated European equity market for the first time, eliminating exchange-rate risk among member states and fostering deeper cross-border investment flows. Starting the sample at this date ensures that the estimated common factor structure reflects genuine economic integration rather than currency-driven co-movement.

2.3.1 Volatility Proxy

Following Diebold and Yılmaz (2014), I measure daily volatility using the Parkinson range-based estimator (Parkinson, 1980):

$$\hat{\sigma}_{it}^2 = \frac{1}{4\ln 2}\left(\ln P_{it}^{\text{high}} - \ln P_{it}^{\text{low}}\right)^2, \tag{2.1}$$

where P_{it}^{high} and P_{it}^{low} are the highest and lowest transaction prices of stock i on day t. The scaling factor $1/(4\ln 2)$ corrects for the fact that, under a geometric Brownian motion, the expected squared log-range exceeds the true variance by a factor of $4\ln 2$ (Parkinson, 1980).

As shown by Alizadeh et al. (2002), range-based estimators are substantially more efficient than squared-return proxies, approximately Gaussian, and robust to microstructure noise — properties that are particularly relevant for a panel of 46 stocks over two decades, where high-frequency tick data may not be uniformly available.

2.3.2 Descriptive Statistics

Figure 2.1 displays the daily Parkinson volatility series for all 46 stocks over the sample period. The strong co-movement during the GFC (2008–2009) and the *Covid-19* outbreak (March 2020) is clearly visible, motivating the use of a factor model to capture this common variation. Table 2.3.2 reports descriptive statistics (mean, standard deviation, skewness, kurtosis, minimum, and maximum) for the volatility series, both at the individual stock level and aggregated by sector.

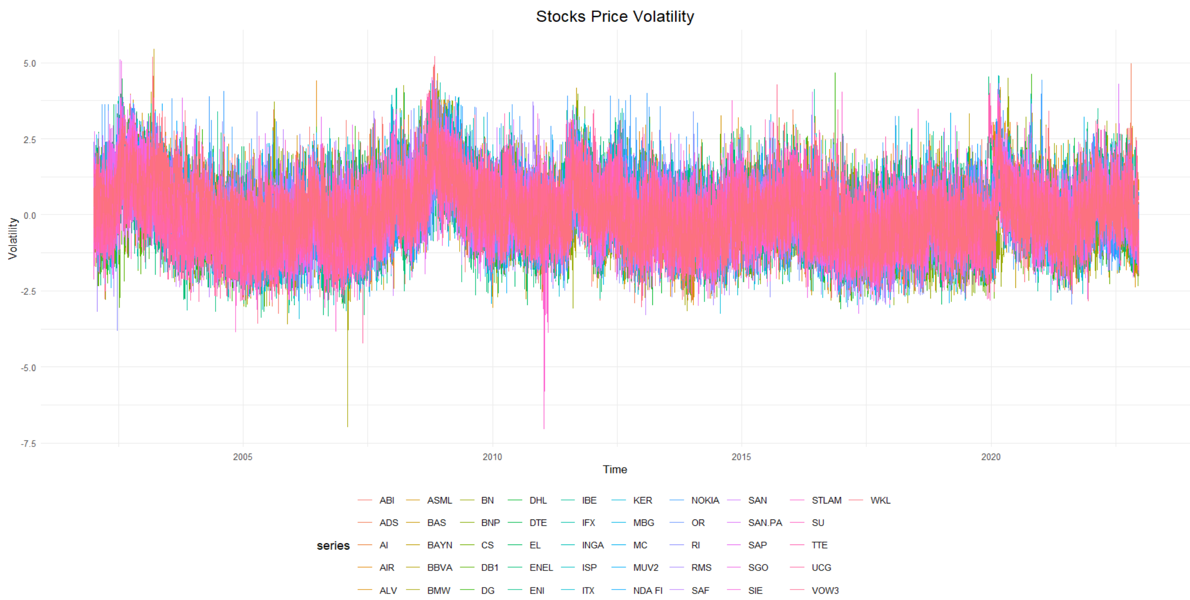

Figure 2.1: Daily Parkinson volatility for the 46 EuroSTOXX 50 constituents, 2002–2022.

Notes: Volatility is computed using the Parkinson (1980) estimator in Eq 2.1. The figure reports daily volatility for all 46 EuroSTOXX 50 constituents over 2002–2022. Each coloured line represents the Parkinson volatility series of one stock. The y-axis is in log-volatility units. Source: Yahoo Finance and author's calculations.

2.4 Methodology

The empirical analysis proceeds in two steps. Section 2.4.1 describes the forecasting exercise, presenting each competing model and the forecast evaluation procedure. Section 2.4.2 describes the network analysis conducted on the idiosyncratic component of the GDFM-FNETS. All rolling-window estimations use a window of $W = 256$ trading days (approximately one calendar year). This choice balances two competing requirements: the window must be large enough to ensure stable estimation of the factor structure and the sparse VAR coefficients, yet short enough to track time-varying changes in network topology during crisis episodes. Feng and Zhang (2025) show that, for pure volatility forecasting, the loss function is U-shaped in window size with a minimum between 1,000 and 2,000 observations; however, their analysis does not account for the need to capture structural shifts in network connectivity, which requires shorter windows.

2.4.1 Forecasting Exercise

I compare the out-of-sample forecasting performance of three models: the GDFM-FNETS, the HAR, and the GARCH(1,1). For each forecast origin t, the models are estimated using the rolling window $[t - W + 1, t]$ and are used to generate multi-step-ahead forecasts for horizons $h = 1, \ldots, H$, where $H = 10$. Forecast accuracy is then evaluated by comparing the horizon-specific predictions with the corresponding realised Parkinson volatility at $t + h$.

GDFM

The Generalized Dynamic Factor Model decomposes the p-dimensional observed volatility vector $\mathbf{x}_t = (x_{1t}, \ldots, x_{pt})'$ into a common component $\boldsymbol{\chi}_t$ and an idiosyncratic component $\boldsymbol{\xi}_t$ (Forni et al., 2000; Forni and Lippi, 2001):

$$\mathbf{x}_t = \boldsymbol{\chi}_t + \boldsymbol{\xi}_t, \tag{2.2}$$

where $\boldsymbol{\chi}_t = \mathbf{B}(L)\,\mathbf{u}_t$ is driven by a q-dimensional vector of common shocks $\mathbf{u}_t$ through the dynamic loading matrix $\mathbf{B}(L)$, and $\boldsymbol{\xi}_t$ captures the stock-specific variation not explained by the common factors.

Estimation of the GDFM. Estimation proceeds in the following steps. The frequency-domain representation of the GDFM follows the seminal contribution of Forni et al. (2000) and Forni and Lippi (2001). In this framework, the observed panel is decomposed into a pervasive common component and an idiosyncratic component. For forecasting purposes, however, it is important to distinguish between the two-sided reconstruction of the common component, which exploits information from both past and future observations, and a genuine out-of-sample forecasting exercise, which must rely only on the information available at the forecast origin.

Using the frequency-domain factor extraction of the observable spectral density, it is possible to:

1. estimate the spectral density matrix of $\mathbf{x}_t$ on a grid of Fourier frequencies using a suitable smoothing window;

2. at each frequency, extract the leading r eigenvectors of the spectral density to obtain frequency-specific factor loadings and factor scores;

3. recover the covariance structure of the common and idiosyncratic components and reconstruct the time-domain common component $\hat{\boldsymbol{\chi}}_t$ and the idiosyncratic residuals $\hat{\boldsymbol{\xi}}_t = \mathbf{x}_t - \hat{\boldsymbol{\chi}}_t$.

The forecasting exercise in this chapter is implemented through the GDFM–FNETS forecasting framework of Barigozzi et al. (2024). Hence, the forecast is not presented as a direct implementation of the constrained-projection predictor of Forni et al. (2005). Rather, the GDFM decomposition provides the factor-adjusted representation on which the FNETS forecasting procedure is built. At each rolling-window origin, the model is re-estimated using only the information available up to time t, so that the resulting forecasts do not rely on future observations.

It is useful to clarify how the one-sided forecasting approach of Forni et al. (2005) differs from the forecasting implementation used below. In the Forni et al. (2005) framework, the forecast of the common component is not obtained by fitting a VAR to the estimated factors. Rather, the predictor is constructed through a constrained projection that exploits the estimated covariance structure of the common and idiosyncratic components.

Following the notation in Giovannelli et al. (2021), let $\hat{\Gamma}_0^{\chi}$ and $\hat{\Gamma}_0^{\xi}$ denote the estimated contemporaneous covariance matrices of the common and idiosyncratic components, respectively. The generalised principal components are obtained by solving the generalised eigenvalue problem:

$$\hat{\mathbf{z}}_j^g \hat{\Gamma}_0^{\chi} = \hat{\nu}_j^g \hat{\mathbf{z}}_j^g \hat{\Gamma}_0^{\xi}, \qquad j = 1, \ldots, r, \tag{2.3}$$

$$\text{subject to} \quad \begin{cases} \hat{\mathbf{z}}_j^g \hat{\Gamma}_0^{\xi} \hat{\mathbf{z}}_j^{g\prime} = 1, \\ \hat{\mathbf{z}}_i^g \hat{\Gamma}_0^{\xi} \hat{\mathbf{z}}_j^{g\prime} = 0, \quad i \neq j. \end{cases} \tag{2.4}$$

Let $\hat{\mathbf{Z}}_h^g = (\hat{\mathbf{z}}_1^g, \ldots, \hat{\mathbf{z}}_{r_h}^g)'$ denote the matrix collecting the r_h generalized eigenvectors. In contrast to Giovannelli et al. (2021), who focus on forecasting a single target variable, the present analysis constructs forecasts for the entire cross-section of log-volatilities. Accordingly, $\hat{\mathbf{Z}}_h^g$ is used in the forecasting step for all 46 variables in the system, rather than only for the first component.

$$\hat{f}_t = \hat{\mathbf{Z}}^g \mathbf{x}_t. \tag{2.5}$$

In the Forni et al. (2005) constrained-projection predictor, the h-step-ahead forecast of the common component can be written as $\hat{\Delta}_h \hat{f}_t$ where $\hat{\Gamma}_h^{\chi} \hat{\mathbf{Z}}^{g\prime} \left(\hat{\mathbf{Z}}^g \hat{\Gamma}_0 \hat{\mathbf{Z}}^{g\prime} \right)^{-1}$.

Here, $\hat{\Gamma}_h^{\chi}$ denotes the estimated lag-h covariance matrix of the common component, while $\hat{\Gamma}_0$ denotes the estimated contemporaneous covariance matrix of $\mathbf{x}_t$. In applications focusing on a single target variable, only the corresponding row of $\hat{\Delta}_h$ is used. In the present chapter, instead, the object of interest is the full volatility vector of the p firms in the panel, so the common-component forecast is naturally a p-dimensional vector.

This discussion is included to clarify the role of Forni et al. (2005) and to distinguish the Forni et al. (2005) constrained-projection predictor from the forecasting implementation used in this chapter. The empirical forecasts below are not obtained by directly implementing Eq. (2.3). Instead, they follow the GDFM–FNETS forecasting framework of Barigozzi et al. (2024), which builds on the factor-adjusted decomposition of the panel into common and idiosyncratic components. Therefore, the GDFM provides the common–idiosyncratic representation, while the operational forecasting procedure is the one implemented in GDFM–FNETS.

Constructing the h-step ahead forecast of the common component. Following the GDFM–FNETS forecasting procedure, after obtaining the estimated factor time series $\{\hat{f}_t\}_{t=1}^T$ within each rolling window, the forecast of the common component is constructed as follows:

1. fit a low-dimensional VAR(p) to the estimated factors:

$$\hat{f}_t = \sum_{\ell=1}^{p} \hat{A}_\ell \hat{f}_{t-\ell} + \hat{\varepsilon}_t^{(f)}, \tag{2.6}$$

   where p is selected by information criteria on the factor VAR, and $\hat{\varepsilon}_t^{(f)}$ denotes the estimated VAR residual;

2. obtain the h-step ahead forecast of the factors by iterating the estimated VAR:

$$\hat{f}_{t+h|t} = \Phi_h(\hat{A}_1, \ldots, \hat{A}_p)\, \hat{f}_t,$$

   where $\Phi_h(\cdot)$ denotes the companion-matrix iteration;

3. reconstruct the h-step ahead forecast of the common component using the estimated dynamic loadings, or their finite-lag approximation:

$$\hat{\boldsymbol{\chi}}_{t+h|t} = \hat{\mathbf{B}}(L)\, \hat{f}_{t+h|t}. \tag{2.7}$$

In practice, $\hat{\mathbf{B}}(L)$ is implemented using the finite-lag representation implied by the spectral estimator. The VAR on the estimated factors captures the dynamic persistence of the common component and is the key step that produces multi-horizon forecasts for $\hat{\boldsymbol{\chi}}_{t+h|t}$. This step should therefore be interpreted as the GDFM–FNETS forecast of the common component, rather than as the constrained-projection formula of Forni et al. (2005). Importantly, the forecasting exercise is implemented recursively: at each rolling-window origin, the factors, their VAR dynamics, and the common-component forecast are obtained using only the information available at that point in time.

Forecasting the idiosyncratic component. Following the GDFM–FNETS framework of Barigozzi et al. (2024), the forecasting procedure is based on a factor-adjusted VAR representation. After removing the pervasive common component from the volatility panel, the remaining idiosyncratic component is modelled as a sparse vector autoregressive process. Starting from the decomposition in Eq. (2.2), the idiosyncratic dynamics are approximated by:

$$\boldsymbol{\xi}_t = \sum_{\ell=1}^{p} A_\ell \boldsymbol{\xi}_{t-\ell} + \mathbf{e}_t, \tag{2.8}$$

where the coefficient matrices A_ℓ are sparse. This specification allows the model to capture residual lead–lag dependence across firms after common market-wide volatility shocks have been filtered out. The h-step-ahead forecast is then obtained by combining the forecast of the common component with the forecast of the idiosyncratic VAR component:

$$\hat{\mathbf{x}}_{t+h|t} = \hat{\boldsymbol{\chi}}_{t+h|t} + \hat{\boldsymbol{\xi}}_{t+h|t}. \tag{2.9}$$

The model is re-estimated at each rolling-window origin.

Importantly, this approach does not assume that the idiosyncratic volatility components are serially uncorrelated. On the contrary, the sparse VAR component is included precisely to account for residual serial and cross-sectional dependence that remains after factor adjustment.

HAR

The Heterogeneous Autoregressive model of Corsi (2009) captures long-memory dynamics in volatility through a parsimonious regression on daily, weekly, and monthly components. For each stock i, the one-step-ahead forecast is:

$$\hat{\sigma}^2_{i,t+1|t} = \hat{\beta}_{i,0} + \hat{\beta}_{i,D}\,\hat{\sigma}^2_{i,t} + \hat{\beta}_{i,W}\,\overline{\sigma}^2_{i,t}(5) + \hat{\beta}_{i,M}\,\overline{\sigma}^2_{i,t}(22), \tag{2.10}$$

where $\overline{\sigma}^2_{i,t}(h) = h^{-1}\sum_{j=0}^{h-1}\hat{\sigma}^2_{i,t-j}$ denotes the average volatility over the past h days. The daily ($h = 1$), weekly ($h = 5$), and monthly ($h = 22$) components are designed to capture heterogeneous trading horizons. The coefficients $\hat{\beta}_{i,0}, \hat{\beta}_{i,D}, \hat{\beta}_{i,W}, \hat{\beta}_{i,M}$ are estimated by OLS within each rolling window, separately for each stock.

For horizons $h = 2, \ldots, H$, forecasts are generated recursively, replacing lagged realised volatilities with their forecasted counterparts.

GARCH(1,1)

The GARCH(1,1) model of Bollerslev (1986) specifies the conditional variance of returns r_{it} as:

$$h_{i,t+1} = \hat{\omega}_i + \hat{\alpha}_i\, r^2_{it} + \hat{\beta}_i\, h_{it}, \tag{2.11}$$

where $h_{it} = \mathrm{Var}(r_{it} \mid \mathcal{F}_{t-1})$ is the conditional variance, $\hat{\omega}_i > 0$, $\hat{\alpha}_i \geq 0$, $\hat{\beta}_i \geq 0$, and $\hat{\alpha}_i + \hat{\beta}_i < 1$. The parameters are estimated by quasi-maximum likelihood within each rolling window, separately for each stock. The conditional variance forecast serves as the volatility prediction to be compared against the Parkinson realised volatility.

For horizons $h = 2, \ldots, H$, the conditional variance forecast is obtained recursively.

Forecast Evaluation

Forecasts from the GDFM are compared to standard univariate benchmarks (HAR and GARCH(1,1)) as described in Section 2.4.1. Statistical comparisons use the conditional predictive ability framework of Giacomini and White (2006) to test whether the GDFM forecasts outperform the benchmarks at various horizons. Standard errors for average loss differentials are computed with Newey–West HAC estimators to account for serial dependence and heteroskedasticity (Newey and West, 1986; Andrews, 1991; Zeileis, 2004). For each model $m \in \{\text{HAR, GARCH}\}$ and each forecast origin t, I compute the loss function at each horizon $h = 1, \ldots, H$ as the squared forecasting error. At the individual stock level, the horizon-specific loss for stock i is:

$$L^{(m)}_{i,t}(h) = \left(\hat{\sigma}^2_{i,t+h|t} - \sigma^2_{i,t+h}\right)^2. \tag{2.12}$$

The multi-horizon loss for stock i is then defined as the average across horizons:

$$\bar{L}_{i,t}^{(m)} = \frac{1}{H} \sum_{h=1}^{H} L_{i,t}^{(m)}(h). \tag{2.13}$$

At the sector level, $\bar{L}_{s,t}^{(m)}$ is the average of $\bar{L}_{i,t}^{(m)}$ over all stocks i belonging to sector s. At the portfolio level, $\bar{L}_{\text{port},t}^{(m)}$ is the average across all 46 stocks.

The relative MSE gain of GDFM over the benchmark m is:

$$\Delta_{m,\ell,t}^{rel} = \frac{\left(\bar{L}_{\ell,t}^{(m)} - \bar{L}_{\ell,t}^{\text{GDFM}}\right)}{\bar{L}_{\ell,t}^{(m)}} = 1 - \frac{\bar{L}_{\ell,t}^{\text{GDFM}}}{\bar{L}_{\ell,t}^{(m)}}, \tag{2.14}$$

where $\ell \in \{i, s, \text{port}\}$ indexes the evaluation level. Thus, $\Delta_{m,\ell,t}^{rel} > 0$ indicates superior performance of the GDFM.

Multiple testing and the portfolio-level hypothesis. Inference is conducted at three aggregation levels with distinct inferential goals:

- **Ticker level:** we test 46 individual nulls (one per stock). Because this is a multiple-comparison setting, p-values are adjusted within the ticker family using the Benjamini–Hochberg false discovery rate (FDR) procedure (Benjamini and Hochberg, 1995).
- **Sector level:** we test one null per sector (averaging across tickers in the sector) and apply the Benjamini–Hochberg correction within the family of sector tests.
- **Portfolio level:** the objective is different: here we assess a single, aggregate hypothesis that the GDFM improves forecast accuracy for the whole portfolio. For this purpose we report a single joint statistic (e.g., the mean relative MSE reduction across all stocks) and compute its standard error using HAC robust methods. When a formal joint test of the vector of stock-level differences is required, we implement a block bootstrap (or a multivariate bootstrap) to obtain a joint p-value for the portfolio null. The distinction is intentional: multiple testing controls FDR when screening many individual series, while the portfolio test evaluates an overall economic effect with a single hypothesis.

Statistical significance is assessed using the Giacomini-White test (Giacomini and White, 2006), which evaluates whether the mean loss difference is significantly different from zero. Under the null hypothesis

$$H_0\colon \mathbb{E}\left[\bar{L}_{\ell,t}^{(m)} - \bar{L}_{\ell,t}^{\text{GDFM}}\right] = 0,$$

the models have equal predictive accuracy. The standard error of the average loss differential is computed using a Newey-West HAC estimator (Newey and West, 1986; Andrews, 1991) to account for potential serial dependence and heteroskedasticity.

Since the test is applied simultaneously across 46 individual stocks and multiple sectors, p-values are adjusted for multiple comparisons using the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). Specifically, the family of null hypotheses tested is:

$$\text{Ticker level:} \qquad H_{0,i} \colon \mathbb{E}\left[\bar{\Delta}^{rel}_{m,i,t}\right] = 0, \quad i = 1, \ldots, 46, \tag{2.15}$$

$$\text{Sector level:} \qquad H_{0,s} \colon \mathbb{E}\left[\bar{\Delta}^{rel}_{m,s,t}\right] = 0, \quad s = 1, \ldots, S, \tag{2.16}$$

$$\text{Portfolio level:} \qquad H_{0,\text{port}} \colon \mathbb{E}\left[\bar{\Delta}^{rel}_{m,\text{port},t}\right] = 0, \tag{2.17}$$

where S is the number of sectors in the panel. The Benjamini-Hochberg procedure orders the p-values within each family as $p_{(1)} \leq p_{(2)} \leq \cdots \leq p_{(K)}$ and rejects $H_{0,(j)}$ for all $j \leq j^*$, where

$$j^* = \max\left\{j \colon p_{(j)} \leq \frac{j}{K}\,\alpha\right\}, \tag{2.18}$$

and $\alpha = 0.05$ is the target false discovery rate.

2.4.2 Network Analysis

The network analysis follows a strict and explicit sequence that must be respected in empirical implementation:

1. **Estimate the GDFM and recover idiosyncratic components.** Using the procedure in Section 2.4.1, obtain $\hat{\boldsymbol{\chi}}_t$ and the idiosyncratic residuals $\hat{\boldsymbol{\xi}}_t = \mathbf{x}_t - \hat{\boldsymbol{\chi}}_t$. *This step is performed before any network estimation.* FNETS does not extract factors; it operates on the residuals produced by the factor model.

2. **Fit a sparse VAR to the idiosyncratic residuals.** Apply the FNETS factor-adjusted VAR estimation to $\hat{\boldsymbol{\xi}}_t$ to estimate a high-dimensional sparse VAR:

$$\hat{\boldsymbol{\xi}}_t = \sum_{k=1}^{\ell} \mathbf{A}_k \, \hat{\boldsymbol{\xi}}_{t-k} + \boldsymbol{\eta}_t, \tag{2.19}$$

   where sparsity is induced via LASSO-type regularisation and tuning is selected by cross-validation or stability selection (Barigozzi and Brownlees, 2019; Barigozzi et al., 2024).

3. **Construct the directed, weighted idiosyncratic network.** The sparsity pattern of the estimated transition matrices $\{\hat{\mathbf{A}}_k\}$ defines the Granger-causal adjacency structure; economic weights are assigned using GFEVD computed from the VAR moving-average representation and combined with the binary adjacency via element-wise multiplication (Diebold and Yılmaz, 2014; Priola et al., 2021).

Idiosyncratic Granger Causality and Weighted Network

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote the directed, weighted graph representing idiosyncratic spillovers, where $\mathcal{V} = \{1, \dots, p\}$ is the set of nodes (stocks) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of directed edges. The network is constructed from the sparse VAR fitted to the idiosyncratic residuals $\hat{\boldsymbol{\xi}}_t$ (see Section 2.4.2). The construction proceeds in three steps: (i) binary adjacency from LASSO-estimated VAR coefficients; (ii) economic weighting via Generalized Forecast Error Variance Decompositions (GFEVD); (iii) combination of significance and economic weight into a final weighted adjacency matrix.

Binary adjacency from sparse VAR. From the estimated sparse VAR in Equation (2.19) we obtain estimated transition matrices $\{\hat{\mathbf{A}}_k\}_{k=1}^{\ell}$. A directed binary edge from node j to node i is present if at least one lag coefficient from j to i is nonzero:

$$G_{ij} = \begin{cases} 1 & \text{if } [\hat{\mathbf{A}}_k]_{ij} \neq 0 \text{ for some } k \in \{1, \dots, \ell\}, \\ 0 & \text{otherwise,} \end{cases} \tag{2.20}$$

The LASSO (or other sparsity) estimator and the chosen tuning procedure ensure that retained links correspond to stable, predictive relationships; additional stability selection or bootstrap filtering can be applied to reduce false positives (Barigozzi and Brownlees, 2019; Barigozzi et al., 2024).

Moving-average representation and GFEVD. Assuming stationarity of the fitted idiosyncratic VAR, the process admits an infinite MA representation:

$$\hat{\boldsymbol{\xi}}_t = \sum_{h=0}^{\infty} \boldsymbol{\Psi}_h \, \boldsymbol{\eta}_{t-h}, \tag{2.21}$$

with $\boldsymbol{\Psi}_0 = \mathbf{I}_p$ and $\{\boldsymbol{\Psi}_h\}$ obtained recursively from the estimated $\{\hat{\mathbf{A}}_k\}$.

Let $\boldsymbol{\Sigma}_\eta = \mathrm{Var}(\boldsymbol{\eta}_t)$ and denote by $\sigma_{jj} = [\boldsymbol{\Sigma}_\eta]_{jj}$ the innovation variance of series j. The H-step Generalized Forecast Error Variance Decomposition entry measuring the contribution of shocks in j to the forecast error variance of i is:

$$\theta_{ij}(H) \;=\; \frac{\sigma_{jj}^{-1} \sum_{h=0}^{H-1} \left(\mathbf{e}_i' \boldsymbol{\Psi}_h \boldsymbol{\Sigma}_\eta \mathbf{e}_j\right)^2}{\sum_{h=0}^{H-1} \mathbf{e}_i' \boldsymbol{\Psi}_h \boldsymbol{\Sigma}_\eta \boldsymbol{\Psi}_h' \mathbf{e}_i}, \tag{2.22}$$

where $\mathbf{e}_i$ is the ith canonical vector. The matrix $\boldsymbol{\Theta}(H) = [\theta_{ij}(H)]_{i,j}$ collects these shares for all pairs (i, j) (Diebold and Yılmaz, 2014). In practice we set H to the same horizon used in the forecasting exercise (e.g., $H = 10$ trading days) to ensure comparability.

Weighted adjacency matrix. To combine statistical significance (Granger causality) with economic magnitude (GFEVD), we form the Hadamard product of the GFEVD matrix and the binary adjacency:

$$\mathbf{W} \;=\; \boldsymbol{\Theta}(H) \odot \mathbf{G}, \tag{2.23}$$

so that $W_{ij} = \theta_{ij}(H)$ if a Granger-causal link from j to i is retained, and $W_{ij} = 0$ otherwise (Priola et al., 2021). The matrix $\mathbf{W}$ is the final weighted, directed adjacency matrix used for all subsequent network analyses.

Normalization and row/column sums. For some summary measures we normalise rows of $\mathbf{W}$ so that outgoing spillovers sum to one:

$$\tilde{W}_{ij} = \frac{W_{ij}}{\sum_j W_{ij}} \quad \text{if } \sum_j W_{ij} > 0, \quad \tilde{W}_{ij} = 0 \text{ otherwise.}$$

Row sums and column sums of $\mathbf{W}$ (or $\tilde{\mathbf{W}}$) provide node-level measures of out- and in-spillovers respectively.

Network topology measures. We summarise the topology of the directed, weighted network with standard measures.

- **Diameter.** Let $\delta(u, v)$ denote the length of the shortest directed path from node u to node v (set $\delta(u, v) = \infty$ if v is not reachable from u). The diameter of the network is the maximum finite shortest-path distance:

$$\operatorname{diam}(\mathcal{G}) = \max_{\substack{u,v \\ \delta(u,v)<\infty}} \delta(u, v). \tag{2.24}$$

  A smaller diameter indicates that shocks can traverse the network in fewer steps.

- **Average path length.** The average shortest directed path length over all reachable ordered pairs is

$$\bar{\delta} = \frac{1}{|\{(u, v) : u \neq v,\ \delta(u, v) < \infty\}|} \sum_{\substack{u \neq v \\ \delta(u,v)<\infty}} \delta(u, v). \tag{2.25}$$

- **Modularity.** To measure community structure we compute directed modularity adapted to the binary adjacency $\mathbf{G}$ (or to a thresholded version of $\mathbf{W}$). Given a partition $\{c_1, \ldots, c_C\}$ of nodes into C communities, directed modularity is

$$Q = \frac{1}{m} \sum_{i,j} \left[G_{ij} - \frac{k_i^{\text{out}}\, k_j^{\text{in}}}{m} \right] \not\Vdash(c_i = c_j), \tag{2.26}$$

  where $m = \sum_{i,j} G_{ij}$ is the total number of directed edges, $k_i^{\text{out}} = \sum_j G_{ij}$ and $k_j^{\text{in}} = \sum_i G_{ij}$. The partition is chosen to maximise Q (e.g., via Louvain or other community detection algorithms) and $Q \in [-1, 1]$ with larger values indicating stronger community structure (Barabási and Pósfai, 2016).

Interpretation. Because the network is constructed from idiosyncratic residuals obtained after removing the GDFM common component, links in $\mathbf{W}$ represent firm-specific predictive spillovers that are not attributable to market-wide shocks. This separation provides a cleaner measure of idiosyncratic contagion and facilitates comparisons between the behaviour of the common component (forecasting gains) and the topology of firm-level propagation channels.

Spillover Index

To validate the GDFM-FNETS decomposition, I compute a total spillover index from the common component following Greenwood-Nimmo et al. (2021). Both the idiosyncratic network modularity and the common-component spillover index are calculated on daily rolling windows of $W = 256$ observations. If the factor decomposition is internally consistent, periods of high common-component spillover (systemic stress transmitted through market-wide factors) should coincide with low idiosyncratic modularity (dissolution of the community structure in the residual network). I test this prediction by examining the correlation between the two series.

2.5 Empirical Results

This section reports the empirical findings from the two exercises described in Section 2.4: the out-of-sample forecasting comparison (Section 2.5.1) and the factor-adjusted network analysis of the idiosyncratic component (Section 2.5.2). The presentation is strictly limited to the procedures actually implemented: Giacomini–White tests for forecast comparisons with Benjamini–Hochberg correction for multiple testing, and correlation tests between idiosyncratic network modularity and the common-component spillover index.

2.5.1 Forecasting Performance

Forecasts are produced on rolling windows of length $W = 256$ trading days and evaluated at horizons $h = 1, \ldots, H$ with $H = 10$. For each forecast origin and horizon we compute the squared forecasting loss and the horizon-averaged loss measures in Equations (2.12)–(2.13). The primary performance metric is the relative MSE gain defined in Equation (2.14).

Statistical inference for pairwise forecast comparisons uses the Giacomini–White conditional predictive ability test (Giacomini and White, 2006). Because tests are performed across many tickers and across sectors, reported p-values are adjusted within each family of hypotheses using the Benjamini–Hochberg false discovery rate procedure (Benjamini and Hochberg, 1995). At the portfolio level we report an aggregate statistic (the cross-sectional mean relative MSE reduction) together with HAC robust standard errors where appropriate (Newey and West, 1986; Andrews, 1991).

Ticker-level Results

Table 2.1 (main text) reports, for each of the 46 constituents, the horizon-averaged relative MSE gain of the GDFM versus the HAR and GARCH(1,1) benchmarks, the Giacomini–White statistic, and the Benjamini–Hochberg adjusted p-value. The empirical regularity is twofold. First, the GDFM produces positive relative MSE gains for a majority of tickers versus HAR and for an even larger subset versus GARCH(1,1). Second, after FDR correction, a non-negligible fraction of tickers exhibit statistically significant improvements at conventional levels. These findings are consistent with the view that factor-based aggregation can improve volatility signal extraction when common variation is informative for the forecast target (Barigozzi and Hallin, 2016, 2020).

The results show substantial heterogeneity across firms. The GDFM delivers positive relative MSE gains against HAR for slightly more than half of the stocks and performs more consistently against GARCH(1,1). This pattern suggests that cross-sectional information is useful for volatility forecasting, but not uniformly so. Firms whose volatility is more strongly linked to the common market component benefit more from the factor-based specification, whereas firms with dominant idiosyncratic persistence display weaker or even negative gains. This heterogeneity is consistent with the view that factor models improve forecasting performance primarily when omitted common components are empirically relevant.

Sector-level Results

Table 2.2 aggregates ticker losses by sector and reports sector-level relative MSE gains and adjusted p-values. Sectors characterised by stronger common variation (for example, Information Technology and Financials in our sample) show larger average gains, consistent with textbook intuition that factors capture systematic drivers that affect groups of firms simultaneously (Forni et al., 2000, 2005).

The GDFM outperforms the benchmarks in most sectors, although the gains are heterogeneous and not uniformly positive. This pattern is consistent with the common–idiosyncratic decomposition underlying the GDFM-FNETS framework. When volatility is driven to a greater extent by pervasive market-wide forces, factor-based forecasts are expected to perform better because they exploit cross-sectional information that is unavailable to purely univariate benchmarks. By contrast, when residual firm-specific or sector-specific persistence remains important, the relative advantage over HAR is naturally more limited, since HAR is designed to capture own-series volatility persistence. The sectoral evidence therefore supports a conditional interpretation of the forecasting results: the GDFM is most useful when common volatility dynamics are relevant, but its gains are not expected to be uniform across all sectors.

Portfolio-level Results

At the portfolio level, the equally weighted relative MSE gain is positive and statistically significant against both HAR and GARCH(1,1). This result indicates that, although ticker-level gains are heterogeneous, the factor-based specification delivers aggregate forecasting improvements for

Table 2.1: Ticker-level relative MSE gains of the GDFM relative to HAR and GARCH(1,1) benchmarks

	GDFM vs. HAR			GDFM vs. GARCH(1,1)		
Ticker	$\bar{\Delta}$	GW stat	Adj. p	$\bar{\Delta}$	GW stat	Adj. p
ABI	0.47	5.13	0.00	0.07	0.54	0.63
ADS	0.25	4.65	0.00	0.17	2.64	0.02
AI	-1.30	0.95	0.51	0.14	2.15	0.06
AIR	0.06	3.82	0.00	0.11	1.23	0.30
ALV	0.44	2.63	0.02	0.06	0.37	0.75
ASML	0.18	4.00	0.00	0.28	6.46	0.00
BAS	0.23	2.25	0.05	0.17	3.24	0.00
BAYN	-0.05	0.65	0.64	0.13	2.87	0.01
BBVA	-0.50	-0.70	0.63	0.17	2.72	0.02
BMW	0.07	3.40	0.00	0.21	4.05	0.00
BN	-0.40	-0.07	0.97	-0.04	-0.88	0.43
BNP	-0.38	0.23	0.90	0.19	3.36	0.00
CS	-1.27	0.29	0.88	-0.06	-1.32	0.27
DB1	0.08	2.58	0.02	0.19	3.00	0.01
DG	0.04	3.34	0.00	0.12	1.41	0.24
DHL	-0.61	-0.87	0.52	0.17	3.34	0.00
DTE	-0.73	0.68	0.64	0.16	2.81	0.01
EL	0.41	6.50	0.00	0.09	1.03	0.39
ENEL	-0.97	-0.63	0.64	0.03	-0.15	0.88
ENI	0.26	4.65	0.00	0.08	0.98	0.41
IBE	0.28	3.14	0.00	0.15	1.61	0.18
IFX	-1.03	0.16	0.93	0.31	7.72	0.00
INGA	0.27	4.45	0.00	0.24	4.02	0.00
ISP	-0.52	0.26	0.89	0.16	2.99	0.01
ITX	-0.72	-1.04	0.46	0.06	0.75	0.50
KER	-0.29	-0.13	0.94	0.16	1.90	0.10
MBG	0.46	5.18	0.00	0.21	3.65	0.00
MC	-0.74	-0.02	0.98	0.13	2.18	0.06
MUV2	0.44	6.35	0.00	0.19	3.07	0.01
NDA FI	-1.41	-1.56	0.21	0.15	1.54	0.19
NOKIA	-0.70	0.88	0.52	0.12	1.17	0.33
OR	0.04	2.51	0.03	0.10	0.89	0.43
RI	0.30	3.60	0.00	0.16	2.11	0.07
RMS	0.03	3.61	0.00	0.19	4.82	0.00
SAF	0.44	5.57	0.00	0.08	0.81	0.47
SAN.FR	-0.14	1.09	0.45	0.12	1.46	0.22
SAN.SP	0.08	2.33	0.04	0.10	0.92	0.43
SAP	0.02	3.20	0.00	0.26	4.91	0.00
SGO	0.26	4.10	0.00	0.15	2.01	0.08
SIE	0.01	3.44	0.00	0.22	3.78	0.00
STLAM	-0.49	-1.79	0.14	0.13	2.44	0.03
SU	-0.52	-0.50	0.73	0.04	0.25	0.82
TTE	-1.10	-1.12	0.45	-0.01	-1.03	0.39
UCG	-0.54	-1.05	0.46	0.28	4.33	0.00
VOW3	0.35	4.88	0.00	0.13	2.06	0.07
WKL	-0.62	0.92	0.51	0.17	3.31	0.00

Notes: $\bar{\Delta}$ denotes the horizon-averaged relative MSE gain defined in Equation (2.14). Positive values indicate lower MSE for the GDFM relative to the benchmark. Adj. p reports Benjamini–Hochberg adjusted p-values from Giacomini–White tests. Values in red denote significance at the 5% level.

Table 2.2: Sector-level relative MSE gains of GDFM over benchmark models

	GDFM vs. HAR			GDFM vs. GARCH(1,1)		
Sector	$\bar{\Delta}_s$	GW stat	Adj. p	$\bar{\Delta}_s$	GW stat	Adj. p
Consumer Staples	0.10	3.85	0.00	0.07	0.68	0.56
Consumer Discretionary	0.14	3.00	0.01	0.16	3.00	0.01
Materials	-0.77	1.46	0.16	0.15	2.81	0.01
Industrials	0.14	3.37	0.00	0.13	2.06	0.06
Financials	0.38	2.62	0.01	0.15	2.85	0.01
Information Technology	0.39	2.48	0.02	0.24	5.70	0.00
Health Care	0.07	3.67	0.00	0.12	1.82	0.10
Communication	-0.73	0.68	0.50	0.16	2.81	0.01
Utilities	0.35	2.27	0.03	0.09	0.67	0.56
Energy	0.42	3.02	0.01	0.04	-0.36	0.72

Notes: $\bar{\Delta}_s$ is the sector-level mean loss difference, computed as the cross-sectional average of ticker-level losses within each sector. Other notation as in Table 2.1.

Table 2.3: Portfolio-level relative MSE gains of GDFM over benchmark models

	GDFM vs. HAR			GDFM vs. GARCH(1,1)		
EuroSTOXX 50	$\bar{\Delta}_p$	GW stat	Adj. p	$\bar{\Delta}_p$	GW stat	Adj. p
	0.2470	3.0185	0.0025	0.8369	2.4890	0.0128

Notes: $\bar{\Delta}_p$ denotes the portfolio-level relative MSE gain defined in Equation (2.14), computed using equally weighted portfolio losses across the 46 constituents. Positive values indicate lower MSE for the GDFM relative to the benchmark. Adj. p reports Benjamini–Hochberg adjusted p-values from Giacomini–White tests. Values in red denote significance at the 5% level.

a diversified basket of EuroSTOXX 50 constituents. Economically, this suggests that common volatility components contain useful predictive information for portfolio-level risk monitoring, even when their contribution varies across individual firms and sectors.

$$L^{(m)}_{\text{port},t} = \frac{1}{46}\sum_{i=1}^{46} L^{(m)}_{i,t}. \tag{2.27}$$

This corresponds to an equally weighted portfolio of the 46 EuroSTOXX 50 constituents, where each stock receives a weight of 1/46. The choice of equal weights ensures that the aggregate forecasting metric is not dominated by a small number of high-volatility stocks.

The aggregate relative MSE gain of the GDFM is statistically significant at the 5% level against both HAR and GARCH(1,1), indicating that the cross-sectional information captured by the factor decomposition translates into measurable forecasting gains at the portfolio level.

The economic interpretation of this result is that the usefulness of the GDFM does not arise from uniformly improving each volatility forecast. Rather, the model extracts a common volatility signal that is valuable at higher levels of aggregation. This is especially relevant for portfolio risk monitoring, where the objective is not necessarily to forecast every idiosyncratic movement, but to capture systematic changes in the volatility environment affecting many assets simultaneously.

2.5.2 Network Analysis

All network results are constructed from the idiosyncratic residuals $\hat{\boldsymbol{\xi}}_t$ obtained after GDFM extraction and follow the FNETs procedure for estimating a sparse VAR on idiosyncratic terms and constructing a weighted directed network via GFEVD combined with Granger-adjacency (Barigozzi et al., 2024; Diebold and Yılmaz, 2014; Priola et al., 2021). The network measures reported below are computed on one-year (256 trading day) estimation windows unless otherwise stated.

Static Network Analysis: Crisis VS Economically Stable Windows

Estimating idiosyncratic networks on representative pre-crisis and crisis windows around the GFC and the Covid-19 episode reveals a consistent empirical pattern: in stable windows, the idiosyncratic network exhibits high modularity and sectoral clustering; in crisis windows, modularity declines sharply, and the network becomes denser. Figures 2.2 and 2.3 display the estimated networks for selected pre-crisis and crisis windows. In pre-crisis windows, nodes cluster by sector with few cross-sector edges; during crisis windows, the networks become markedly denser, and the sector-coded communities lose their spatial coherence.

Table 2.4 summarises diameter, average path length and modularity for selected windows. This empirical pattern aligns with prior empirical work documenting increased connectedness during stress episodes in European markets (Mensi et al., 2018; Aslam et al., 2021; Huynh et al., 2021) and is consistent with theoretical mechanisms of liquidity spirals and fire-sale contagion (Brunnermeier and Pedersen, 2009; Cont and Schaanning, 2017; Shleifer and Vishny, 2011).

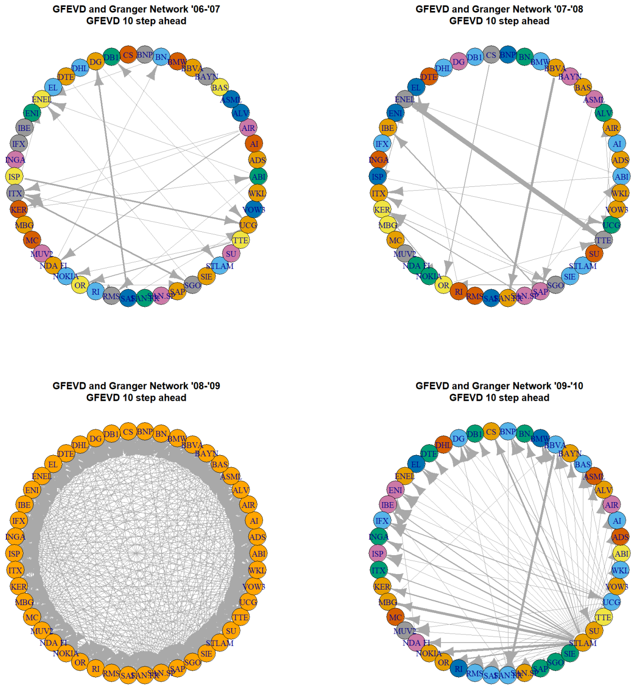

Figure 2.2: Factor-adjusted idiosyncratic Granger Causality volatility networks during the Global Financial Crisis, 2006–2010.

Notes: Networks are estimated from the idiosyncratic residuals obtained after GDFM extraction. Directed edges are identified via sparse VAR Granger causality and weighted using 10-step-ahead GFEVD contributions. Node colours indicate communities detected using the Walktrap algorithm. Each panel corresponds to a rolling one-year estimation window.

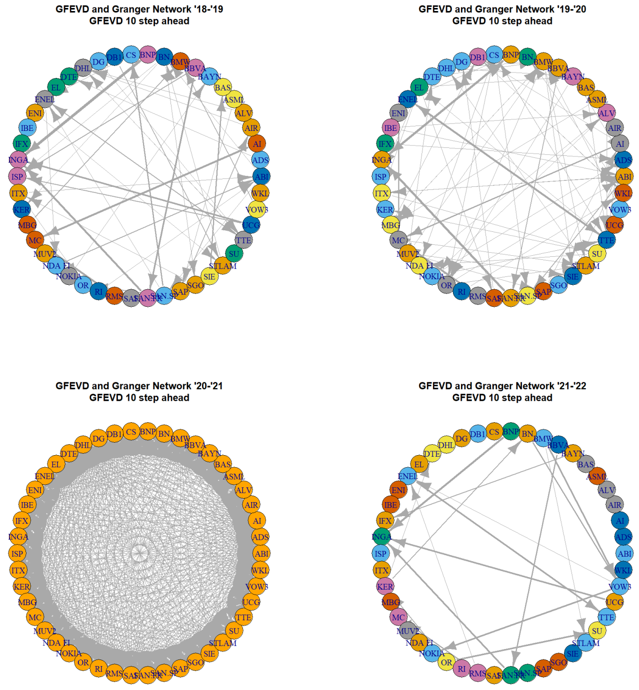

Figure 2.3: Factor-adjusted idiosyncratic Granger Causality volatility networks during the Covid-19 period, 2018–2022.

Notes: Networks are estimated from the idiosyncratic residuals obtained after GDFM extraction. Directed edges are identified via sparse VAR Granger causality and weighted using 10-step-ahead GFEVD contributions. Node colours indicate communities detected using the Walktrap algorithm. Each panel corresponds to a rolling one-year estimation window.

The key finding is the sharp collapse of modularity during both crises. In economically stable periods, modularity exceeds 0.50, indicating that stocks cluster into well-defined sectoral communities with limited cross-sector spillovers. During the GFC and the Covid-19 crisis, modularity drops to near zero, signalling that the community structure dissolves as shocks propagate indiscriminately across the network.

Table 2.4: Network Topology (NT) measures for one-year estimation windows

Years	*NT Measure*	*NT Results*
2006-2007	**Diameter**	2.00
	Average Path Length	1.32
	Modularity (Q)	0.60
2007-2008	**Diameter**	5.00
	Average Path Length	2.05
	Modularity (Q)	0.51
2008-2009	**Diameter**	2.00
	Average Path Length	1.21
	Modularity (Q)	0.01
2009-2010	**Diameter**	8.00
	Average Path Length	3.59
	Modularity (Q)	0.49
2018-2019	**Diameter**	3.00
	Average Path Length	1.40
	Modularity (Q)	0.64
2019-2020	**Diameter**	2.00
	Average Path Length	1.04
	Modularity (Q)	0.35
2020-2021	**Diameter**	2.00
	Average Path Length	1.20
	Modularity (Q)	0.00
2021-2022	**Diameter**	4.00
	Average Path Length	1.98
	Modularity (Q)	0.71

Notes: Network measures computed on the GFEVD-weighted idiosyncratic Granger causality network (Equation 2.23) estimated over one-year (256 trading days) windows. Modularity is maximised using the Louvain algorithm. Diameter and average path length are computed on the directed version of the network.

This pattern is broadly consistent with theoretical mechanisms in which systemic stress increases cross-sectional dependence across assets and weakens sectoral segmentation. In particular, models of liquidity spirals and fire-sale externalities (Brunnermeier and Pedersen, 2009; Cont and Schaanning, 2017; Shleifer and Vishny, 2011) imply that periods of market-wide stress may generate increasingly broad-based shock propagation across firms and sectors. While the present analysis does not identify these mechanisms directly, the sharp decline in modularity during crisis periods is qualitatively consistent with a temporary weakening of sector-specific clustering under systemic stress. The lowest modularity values are observed during the core crisis windows (2008–2009 and 2020–2021), whereas higher modularity values are typically observed during relatively stable market periods.

Dynamic Network Analysis: Rolling-Window Modularity and Spillover

To track the evolution of network structure at higher frequency, I compute the modularity and the common-component spillover index on daily rolling windows of $W = 256$ observations over both the GFC and *Covid-19* periods.

Figure 2.4 plots the rolling-window idiosyncratic network modularity and the common-component spillover index (computed following Greenwood-Nimmo et al., 2021) over the full sample.

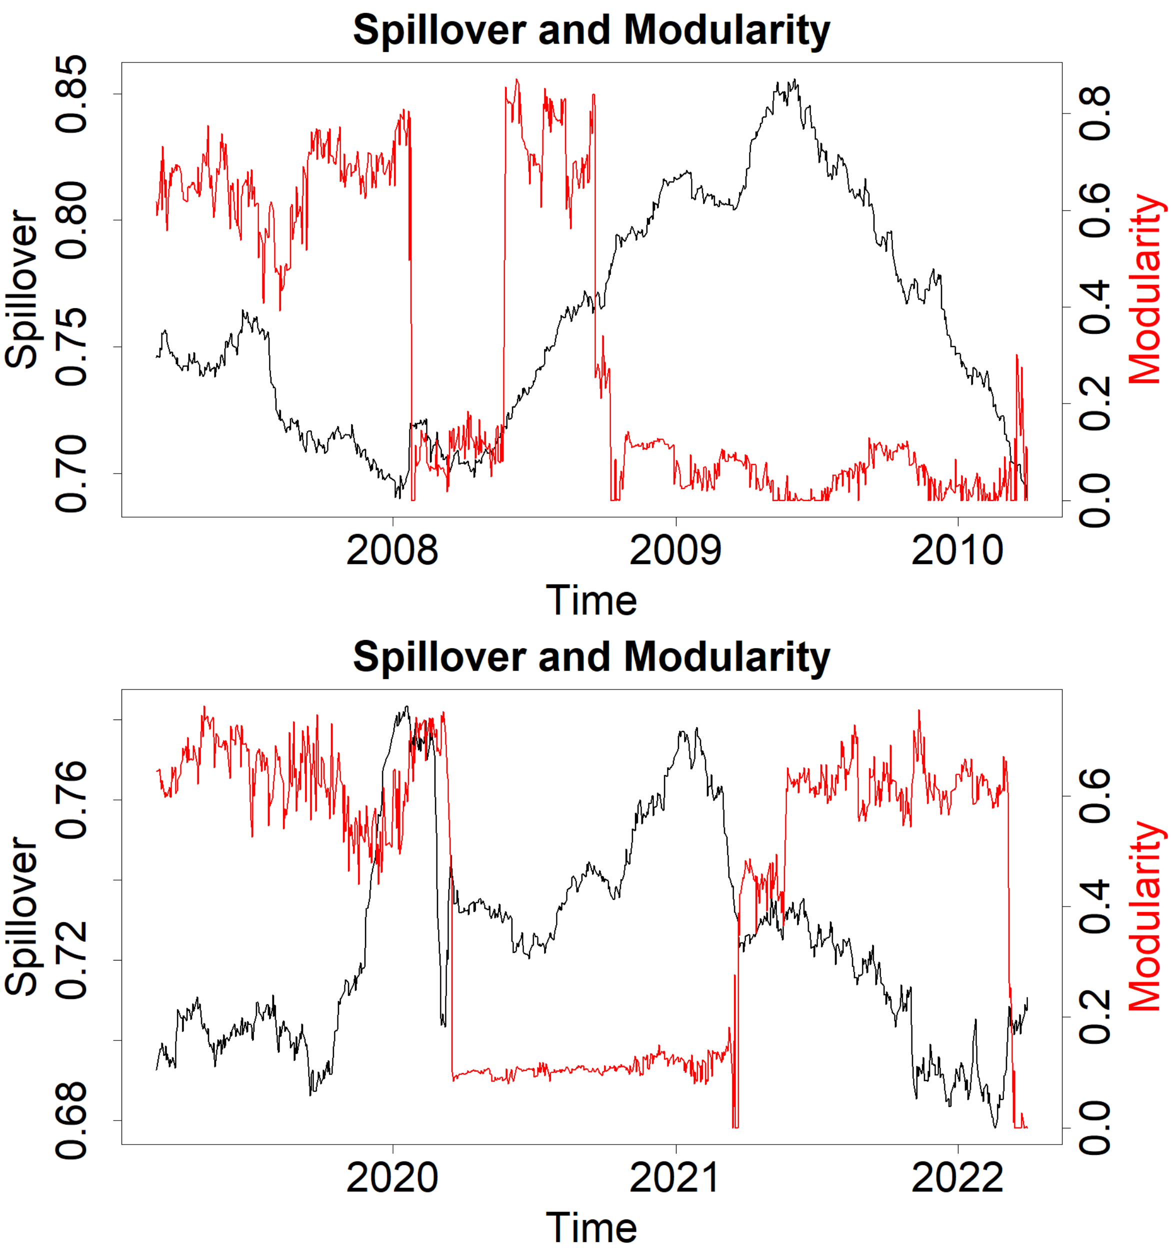

Figure 2.4: Common-component spillover index and idiosyncratic network modularity during systemic crises. The top panel reports the Global Financial Crisis (2007–2010), while the bottom panel reports the Covid-19 crisis (2020–2022). The black line shows the spillover index computed on the common component of volatility, and the red line shows the modularity of the idiosyncratic Granger Causality Networks.

Notes: The spillover index is computed from the common component extracted via the Generalized Dynamic Factor Model (GDFM), following the methodology of Greenwood-Nimmo et al. (2021). Idiosyncratic modularity is computed on the Granger Causality Network estimated from the idiosyncratic residuals $\hat{\xi}_t$, using rolling windows of $W = 256$ daily observations. Network edges are weighted using 10-step-ahead Generalized Forecast Error Variance Decomposition (GFEVD) contributions, as defined in Eq. (2.23). Modularity is obtained via the Walktrap community-detection algorithm. Both series are standardised for comparability. The figure illustrates the collapse of modularity during systemic stress episodes and its negative co-movement with the common-component spillover index.

The correlation between the two series is examined both in levels and in first differences. In levels, both series display substantial persistence — the autocorrelation at lag 1 exceeds 0.99 for both — which inflates the standard correlation coefficient and makes inference unreliable. I therefore also report the correlation in first differences, which removes the persistence and provides a cleaner measure of co-movement at the daily frequency. The first-difference specification is not motivated by the presence of a deterministic trend — neither series exhibits trending behaviour — but rather by the need to address the near-unit-root persistence that characterises rolling-window estimates computed with overlapping data. The correlation between ΔQ_t and ΔS_t is significantly negative, robust to Newey-West correction for serial dependence induced by the overlapping windows. Regressions estimated with the Newey-West standard errors yield a statistically significant negative coefficient during the GFC ($\hat{\beta} = -1.614$, $p = 0.036$) and a marginally significant one, but still negative, during Covid-19 ($\hat{\beta} = -1.270$, $p = 0.078$).

These results should be interpreted primarily as an internal consistency check of the GDFM-FNETS decomposition rather than evidence of a structural relationship: the two measures are derived from complementary components of the same model, so their alignment confirms that the factor decomposition is internally coherent. When the common-component spillover index rises — indicating that market-wide shocks account for a larger share of forecast error variance — the idiosyncratic network modularity falls, meaning that the residual interactions lose their clustering. The results suggest that the common and idiosyncratic components reflect complementary dimensions of the same underlying systemic dynamics.

Economically, the results suggest that systemic crises are characterised by a shift from sector-specific volatility propagation toward increasingly pervasive market-wide dynamics. During stable periods, idiosyncratic interactions remain relatively segmented across communities of firms. During crisis periods, however, common shocks dominate the volatility structure, reducing the importance of sectoral boundaries and increasing cross-sectional synchronisation across firms.

2.6 Conclusion

This chapter applies a GDFM–FNETS framework to daily Parkinson volatility series for 46 EuroSTOXX 50 constituents over 2002–2022. The empirical evidence points to two main conclusions. First, factor-based volatility forecasts improve upon standard univariate benchmarks at the portfolio level, with gains concentrated among firms and sectors whose volatility contains a larger common component. This suggests that cross-sectional information is most useful when volatility is driven by pervasive market-wide forces rather than by purely firm-specific persistence.

Second, the factor-adjusted idiosyncratic networks reveal that the structure of residual volatility propagation changes substantially during periods of systemic stress. Network modularity declines during the Global Financial Crisis and the Covid-19 episode, indicating that residual linkages become less clearly organised along sectoral communities when common volatility shocks dominate the market. The forecasting and network results therefore, provide complementary evidence: the GDFM captures the predictive content of common volatility, while the idiosyncratic network describes the residual propagation structure after common shocks have been filtered out.

The analysis is deliberately limited to statistical forecast accuracy and network diagnostics.

Assessing the economic value of these forecasts requires a separate portfolio exercise. Future research could evaluate whether the GDFM-based forecasts improve portfolio allocation, hedging performance, or Value-at-Risk forecasts relative to standard volatility models. Exploring alternative rolling-window lengths (see Feng and Zhang, 2025) may help clarify the trade-off between forecast accuracy and sensitivity to structural change bias and variance. Finally, applying the framework to other equity markets would help assess the external validity of the results.

Appendix

A. Data Descriptive Statistics and Details

Table A.1: Descriptive statistics for European blue-cips

	Volatilities Summary				
	ABI.BR	ADS.DE	AI.PA	AIR.PA	ALV.DE
Min.	-3.246	-3.008	-2.868	-2.964	-2.460
1st Qu.	-0.678	-0.697	-0.714	-0.717	-0.693
Median	-0.089	-0.044	-0.074	-0.072	-0.114
Mean	0.000	0.000	0.000	0.000	0.000
3rd Qu.	0.585	0.638	0.612	0.651	0.572
Max.	4.917	4.965	4.354	4.398	3.780
	ASML.AS	BAS.DE	BAYN.DE	BBVA.MC	BMW.DE
Min.	-3.037	-2.763	-2.945	-3.591	-6.973
1st Qu.	-0.699	-0.709	-0.683	-0.690	-0.690
Median	-0.072	-0.061	-0.054	-0.042	-0.044
Mean	0.000	0.000	0.000	0.000	0.000
3rd Qu.	0.610	0.639	0.598	0.651	0.632
Max.	4.174	4.655	5.452	3.477	4.116
	BN.PA	BNP.PA	CS.PA	DB1.DE	DG.PA
Min.	-3.073	-2.957	-2.851	-3.035	-3.031
1st Qu.	-0.683	-0.694	-0.683	-0.691	-0.698
Median	-0.055	-0.080	0.000	-0.072	-0.046
Mean	0.000	0.000	0.000	0.000	0.000
3rd Qu.	0.609	0.606	0.567	0.606	0.626
Max.	4.488	4.100	3.579	4.069	4.661
	DPW.DE	DSY.PA	ENGI.PA	ENI.MI	EOAN.DE
Min.	-3.187	-2.875	-2.890	-3.485	-3.412
1st Qu.	-0.681	-0.684	-0.683	-0.678	-0.687
Median	-0.059	-0.053	-0.051	-0.049	-0.045
Mean	0.000	0.000	0.000	0.000	0.000
3rd Qu.	0.607	0.612	0.605	0.611	0.603
Max.	4.372	3.690	4.021	4.398	3.874
	EXO.PA	FR.PA	G.MI	HEI.DE	HEN3.DE
Min.	-3.342	-2.990	-3.274	-2.942	-2.890
1st Qu.	-0.695	-0.689	-0.702	-0.677	-0.693
Median	-0.067	-0.053	-0.069	-0.043	-0.059
Mean	0.000	0.000	0.000	0.000	0.000
3rd Qu.	0.621	0.601	0.598	0.610	0.599
Max.	4.457	3.914	4.214	4.318	3.891
	IFX.DE	ITX.MC	LHA.DE	MC.PA	OR.PA
Min.	-3.019	-2.973	-3.206	-3.085	-2.991
1st Qu.	-0.697	-0.685	-0.691	-0.694	-0.680
Median	-0.061	-0.062	-0.057	-0.051	-0.063
Mean	0.000	0.000	0.000	0.000	0.000
3rd Qu.	0.601	0.613	0.606	0.602	0.608
Max.	4.318	4.274	4.002	3.979	4.126
	ORA.PA	PHIA.AS	RNO.PA	SAF.PA	SAN.MC
Min.	-3.017	-2.905	-3.134	-3.111	-3.028
1st Qu.	-0.687	-0.692	-0.699	-0.690	-0.691
Median	-0.048	-0.056	-0.068	-0.058	-0.064
Mean	0.000	0.000	0.000	0.000	0.000
3rd Qu.	0.603	0.617	0.601	0.607	0.612
Max.	4.215	4.064	4.354	3.920	4.412
	UL.PA				
Min.	-3.150				
1st Qu.	-0.689				
Median	-0.059				
Mean	0.000				
3rd Qu.	0.610				
Max.	4.230				

Notes: Descriptive statistics of daily Parkinson range-based volatility (Equation 2.1) for 46 EuroSTOXX 50 constituents, 2 January 2002 to 30 December 2022. *Source*: Yahoo Finance. Personal evaluation.

Table A.2: Table of stocks in EuroSTOXX 50 and related information

Ticker	Sector	Location	Exchange	Currency
ABI	Consumer Staples	Belgium	Nyse Euronext - Euronext Brussels	EUR
ADS	Consumer Discretionary	Germany	Xetra	EUR
AI	Materials	France	Nyse Euronext - Euronext Paris	EUR
AIR	Industrials	France	Nyse Euronext - Euronext Paris	EUR
ALV	Financials	Germany	Xetra	EUR
ASML	Information Technology	Netherlands	Euronext Amsterdam	EUR
BAS	Materials	Germany	Xetra	EUR
BAYN	Health Care	Germany	Xetra	EUR
BBVA	Financials	Spain	Bolsa De Madrid	EUR
BMW	Consumer Discretionary	Germany	Xetra	EUR
BN	Consumer Staples	France	Nyse Euronext - Euronext Paris	EUR
BNP	Financials	France	Nyse Euronext - Euronext Paris	EUR
CS	Financials	France	Nyse Euronext - Euronext Paris	EUR
DB1	Financials	Germany	Xetra	EUR
DG	Industrials	France	Nyse Euronext - Euronext Paris	EUR
DHL	Industrials	Germany	Xetra	EUR
DTE	Communication	Germany	Xetra	EUR
EL	Health Care	France	Nyse Euronext - Euronext Paris	EUR
ENEL	Utilities	Italy	Borsa Italiana	EUR
ENI	Energy	Italy	Borsa Italiana	EUR
IBE	Utilities	Spain	Bolsa De Madrid	EUR
IFX	Information Technology	Germany	Xetra	EUR
INGA	Financials	Netherlands	Euronext Amsterdam	EUR
ISP	Financials	Italy	Borsa Italiana	EUR
ITX	Consumer Discretionary	Spain	Bolsa De Madrid	EUR
KER	Consumer Discretionary	France	Nyse Euronext - Euronext Paris	EUR
MBG	Consumer Discretionary	Germany	Xetra	EUR
MC	Consumer Discretionary	France	Nyse Euronext - Euronext Paris	EUR
MUV2	Financials	Germany	Xetra	EUR
NDA	Financials	Finland	Nasdaq Omx Helsinki Ltd.	EUR
NOKIA	Information Technology	Finland	Nasdaq Omx Helsinki Ltd.	EUR
OR	Consumer Staples	France	Nyse Euronext - Euronext Paris	EUR
RI	Consumer Staples	France	Nyse Euronext - Euronext Paris	EUR
RMS	Consumer Discretionary	France	Nyse Euronext - Euronext Paris	EUR
SAF	Industrials	France	Nyse Euronext - Euronext Paris	EUR
SAN.FR	Health Care	France	Nyse Euronext - Euronext Paris	EUR
SAN.SP	Financials	Spain	Bolsa De Madrid	EUR
SAP	Information Technology	Germany	Xetra	EUR
SGO	Industrials	France	Nyse Euronext - Euronext Paris	EUR
SIE	Industrials	Germany	Xetra	EUR
STLAM	Consumer Discretionary	Italy	Borsa Italiana	EUR
SU	Industrials	France	Nyse Euronext - Euronext Paris	EUR
TTE	Energy	France	Nyse Euronext - Euronext Paris	EUR
UCG	Financials	Italy	Borsa Italiana	EUR
VOW3	Consumer Discretionary	Germany	Xetra	EUR
WKL	Industrials	Netherlands	Euronext Amsterdam	EUR

Notes: List of stocks listed on the EuroSTOXX 50 from January 2, 2002, to December 30, 2022. Source: Yahoo Finance and BlackRock.

3. Global and local effects of Federal Reserve communication: uncertainty and real activity

We document a new channel for the international transmission of US monetary policy: communication. We identify a Fed communication shock using a novel method based on a textual sentiment extracted from Fed speeches, and frequency domain restrictions in a structural dynamic factor model. Unlike previous studies, we consider both domestic US effects and international spillovers. A positive communication shock reduces various uncertainty measures and temporarily increases the real activity. These conclusions apply to the US and extend across borders with no qualitative difference between advanced economies and emerging markets.[1]

[1]This chapter is based on joint research. Accordingly, the first-person plural is used throughout the chapter.

3.1 Introduction and Literature

Uncertainty is an important driver of the economy as it affects economic behaviour such as investment decisions, household expenditures, and financial markets (Jurado et al., 2015; Baker et al., 2016). At the same time, central banks communicate the likely future course of monetary policy to inform individuals and firms ahead of their economic decisions. In the last few years, with major economies stuck at the zero lower bound, this communication has become a central policy tool, especially through forward guidance (Del Negro et al., 2012; Filardo and Hofmann, 2014; McKay et al., 2016). As a result, central bank communication (CBC) attracted the interest of an increasing number of scholars (Cieslak et al., 2023; Ahrens et al., 2024; Barry et al., 2025; Hansen et al., 2019; Blinder et al., 2024; Hanifi et al., 2022; Gorodnichenko et al., 2023; Aruoba and Drechsel, 2024; Bordo and Istrefi, 2023; Istrefi and Piloiu, 2020; Istrefi et al., 2023; Fadda et al., 2025).

It is widely agreed that CBC can be very effective. For example, the following statement by Mario Draghi (ECB president at the time) during the EU sovereign debt crisis in 2012 has been widely credited as beneficial for the stability of the Euro area: "Within our mandate, the ECB is ready to do 'whatever it takes' to preserve the euro. And believe me, it will be enough."

Does CBC reduce uncertainty? As of today, there is mixed evidence in the literature. Several studies show that initial market uncertainty decreases over time if central bank communications are properly understood (Rosa and Verga, 2008) and that the impact of the CBC on market volatility is most important during periods of high uncertainty, affecting interest rates, risk premia, and financial market volatility (Gertler and Horvath, 2018; Hansen et al., 2019). There is strong evidence in favour of the uncertainty transmission channel as an important factor in explaining why long-term interest rates react to the release of the Inflation Report (Hansen et al., 2019). The tone of communications and the consistency of topics covered by the Central Bank are influenced by the language of the CBC (Tillmann and Walter, 2019). The CBC can increase or decrease uncertainty with significant effects on stock prices and market volatilities (Bauer et al., 2021). Moreover, Forward Guidance shocks manage to have a positive effect on future rate uncertainty but a negative effect on stock market volatility (Gorodnichenko et al., 2023). In contrast, individual central bank members' speeches do not seem to resolve the dynamics of intraregional uncertainty calculated as realised volatility and tail risk (Ahrens et al., 2024).

The recent literature on central bank communication points to a broad set of measurable informational channels. Fadda et al. (2025) show that central banks also communicate uncertainty directly, while Istrefi et al. (2023) document that Fed speeches contain information about financial stability concerns relevant for monetary policy decisions. Istrefi and Piloiu (2020) emphasize that central bank communication is shaped not only by policy content but also by the informational environment in which the public forms opinions about central banks, and Bordo and Istrefi (2023) show that perceptions of FOMC members can be organized along hawkish, dovish, and swing-oriented dimensions. Read together, these contributions reinforce a central point for this thesis: the language of central banks is an economically meaningful object, and its effects should be studied as a source of information that moves expectations, uncertainty, and the interpretation of policy.

Forward guidance refers to the use of central bank communication to influence expectations about the future path of monetary policy. Rather than operating only through current changes in the policy rate, forward guidance works by shaping beliefs about future policy decisions, macroeconomic conditions, and the likely reaction function of the central bank. In this sense, Federal Reserve communication is not merely descriptive: speeches by policymakers can convey information about the expected policy path, the assessment of risks, and the central bank's interpretation of the economic outlook.

Compared to other studies, this chapter focuses on Federal Reserve speeches and treats them as a forward-guidance-related source of macro-financial information. When a Fed member speaks, market participants, firms, and households may revise their expectations because the speech can provide signals about the future stance of monetary policy and the broader economic outlook. We therefore extract sentiment from speeches of the Federal Reserve Board of Governors and use it to identify a Fed communication shock. The shock is not intended to replicate conventional monetary policy shocks identified from changes in policy rates or high-frequency surprises. Instead, it is designed to capture a communication-based innovation associated with the low-frequency co-movement between Fed sentiment and the Federal Funds Rate (FFR) within a Structural Dynamic Factor Model (SDFM) (Stock and Watson, 2005; Forni et al., 2009, 2025). This interpretation is consistent with the logic of forward guidance: communication can affect expectations immediately, while the policy instrument may adjust only gradually over time.

Empirical studies of central bank communication typically transform speeches, official documents, newspapers, or social media content into quantitative textual indicators (Baker et al., 2016; Shapiro et al., 2022; Gorodnichenko et al., 2024). These indicators are then used to study how communication shapes expectations, financial markets, and perceptions of the future policy stance. We focus on the central bank board's speeches because the frequency of such speeches is higher than that of institutional documents and because the speeches of the board and central bank presidents turn out to be more important than Federal Open Market Committee announcements for stock prices, treasury yields, and all but the shortest maturity interest rate futures (Swanson and Jayawickrema, 2023). Furthermore, as discussed by Moschella and Pinto (2019), speeches are addressed to multiple audiences and their format is much less constrained than official central bank documents.

In this chapter, we study the domestic and international effects of the Fed communication shock. At the domestic level, we examine its transmission to uncertainty, expectations, real activity, inflation, labour-market conditions, and financial variables. At the international level, we estimate its propagation to real and financial activity across an unbalanced panel of advanced economies and emerging markets. The empirical analysis proceeds in three steps:

1. *Fed sentiment extraction.* We construct sentiment indicators from speeches delivered by members of the Federal Reserve Board of Governors. Textual sentiment is extracted using both lexicon-based and transformer-based NLP methods, allowing us to compare tradi-

tional dictionary approaches with more context-sensitive language models. The resulting sentiment series provides the textual input used in the subsequent structural analysis.

2. *Identification of the Fed communication shock.* We identify a Fed communication shock within a structural dynamic factor model. The identification exploits the co-movement between Fed sentiment and the FFR over the frequency band relevant for policy transmission. This frequency-domain restriction is motivated by the idea that forward guidance operates through expectations about the future path of policy rather than only through contemporaneous changes in the policy instrument. The shock is therefore interpreted as a forward-guidance-related communication innovation, rather than as a conventional monetary policy shock.

3. *International propagation.* We then project non-U.S. variables onto the identified Fed communication shock using local projections (Jordà, 2005) to estimate its international transmission across an unbalanced panel of countries. We adopt this sequential strategy because the Fed communication shock is defined within the U.S. monetary policy system. Identifying the shock using only U.S. information preserves its domestic interpretation, while the international propagation is assessed empirically in a second stage.

We obtain three main results. First, at the domestic level, a positive Fed communication shock reduces uncertainty across a broad set of indicators. The evidence is not limited to a single proxy: the shock lowers disagreement among professional forecasters across several macroeconomic variables, reduces financial uncertainty as measured by the VIX, and is associated with declines in broader uncertainty indicators such as the WUI. This extensive set of responses suggests that Fed communication affects uncertainty through multiple margins, including financial markets, expectations, and macroeconomic perceptions.

Second, the Fed communication shock has expansionary effects on the U.S. macro-financial system. Real activity increases, unemployment declines, and inflation responds mildly but positively. These responses are consistent with the interpretation of clearer and more positive Fed communication as operating through an aggregate-demand channel. Third, the shock propagates internationally. Positive Fed communication reduces uncertainty and stimulates real and financial activity abroad. The qualitative pattern of the responses is similar across advanced economies and emerging markets, although the magnitude of the effects is stronger in advanced economies.

A key result of the analysis is that the identified Fed communication shock is empirically distinct from existing shocks in the literature. We compare it with standard monetary policy shocks, high-frequency monetary policy surprises, forward-guidance and LSAP shocks, fiscal shocks, and news shocks (Christiano et al., 1999; Romer and Romer, 2004; Ramey, 2016; Swanson, 2021). The Fed communication shock is not significantly correlated with these existing measures. This evidence supports the interpretation that the shock captures a distinct driver of economic fluctuations rather than a relabelling of conventional monetary policy, fiscal, or news shocks.

These results relate to several strands of the literature. On the domestic side, they are consistent with Rosa and Verga (2008), who show that central bank communication has a measurable impact on financial markets, and with Gertler and Horvath (2018), Hansen et al. (2019), and Bauer et al. (2021), who document effects of CBC on financial volatility and risk premia. Unlike Ahrens et al. (2024), who find that individual speeches do not resolve intraregional uncertainty, our identification strategy — which aggregates sentiment across all Board members — uncovers a robust uncertainty-reducing effect. On the international side, Georgiadis (2016) shows that domestic policies can make economies more resilient to US monetary policy spillovers, while Georgiadis and Jarocinski (2023) highlights that different Fed instruments — interest rates, forward guidance, asset purchases — create distinct trade-offs for emerging market central banks. Our finding that the magnitude of spillovers is stronger in advanced economies than in emerging markets contrasts with Chen et al. (2016) and Dedola et al. (2017), who report larger effects on EMs through QE and tightening channels, respectively. In our case, the qualitative pattern of the responses is similar across country groups, but the size of the effects is larger in advanced economies. This suggests that communication operates as a distinct spillover channel, with a transmission mechanism that differs from those associated with conventional and unconventional monetary policy tools.

The rest of the chapter is structured as follows. In Section 3.2, we present the dataset used for the domestic and international analysis. In Section 3.3, we describe the empirical methodology: Fed sentiment extraction, SDFM and the shock identification strategy, and local projections for the analysis of international spillovers. In Section 3.4, we present the domestic results and the evidence on the distinctiveness of the identified Fed communication shock. In Section 3.5, we present the results on international spillovers. Section 3.6 concludes.

3.2 Data

3.2.1 Data for Domestic Analysis

We work with a high-dimensional panel of time series.

- The dataset comprises 112 variables from the FRED-QD, spanning from the first quarter of 1986 to the second quarter of 2021 that come from the dataset of Forni et al. (2025) (henceforth FGGSS). The transformation of the variables was done following the procedure performed in FGGSS. To the previous variables, we aggregated 23 disagreement variables from the Survey of Professional Forecasters (SPF) of the Federal Reserve Bank of Philadelphia, using both Cross-Sectional dispersion measures and Mean Growth (Expectation) measures for quarterly forecasts. We used current forecast variables for the estimation.

- We consider other measures of uncertainty, such as the EPU Index and the Smoothed WUI Index for the United States, the spread between bullish and bearish investor sentiment (in percentage points — available from the first quarter of 1988), and the CBOE-VIX (available from the first quarter of 1986). All variables were constructed at quarterly frequency, considering both the cumulative sum and the period average. We use the period average in the estimation.

- We also aggregated sentiment data extracted from informal speeches of central bank board members — presidents included. The speeches are taken from the CBS Dataset constructed by Campiglio et al. (2025), which includes and expands the Bank of International Settlements (BIS) dataset. We used three textual sentiment indicators to identify the communication shock. One of those derives from financial lexicon developed by Loughran and McDonald (2011) and the other two are based on finBERT, a language model based on BERT (Devlin et al., 2019), to tackle Natural Language Processing (NLP) tasks in the financial domain: the first is the ratio of the difference between positive and negative sentences to their sum, and the second includes neutral sentences in the denominator (Araci, 2019; Yang et al., 2020).

For the domestic analysis, we also include additional measures of uncertainty, including the bull–bear investor sentiment spread from the American Association of Individual Investors (AAII) Investor Sentiment Survey (American Association of Individual Investors, 2026). The indicator is computed as the difference between the percentage of respondents reporting a bullish outlook and the percentage reporting a bearish outlook for the stock market. Since this series is available from the first quarter of 1988 and represents the latest starting date among the uncertainty measures used in the domestic specification, our baseline sample begins in the first quarter of 1988. Excluding this indicator allows the domestic analysis to start in the first quarter of 1986.

3.2.2 Data for International Spillover Analysis

In the spirit of Fernández-Villaverde et al. (2024), the dataset comprises 17 variables: 10 variables come from the IMF - IFS (International Financial Statistics), 2 variables are our Fed communication shocks, 3 variables are WUI, WSI (World Sentiment Index) and EPU, and the last 2 variables are log volatility of countries' stock prices and countries' disagreement on GDP forecast. The variables are categorised into global and country components. All the variables are at the quarterly level.

Global: Our Fed communication shock (both finBERT and LM), the CBOE-VIX, the log of the S&P 500 index, the log of the WTI price of oil, the yield on two-year US Treasuries, the Chicago Federal Reserve National Financial Conditions Index (NFCI), and the log of world real GDP.

Country: The log of a country's Stock Price Index, the log of Industrial Production Index, the log of Fixed Investment, the log of per capita GDP, the Smoothed WUI at the country level, the WSI at the country level, the EPU (quarterly) at the country level, the log Stock Volatility, and the disagreement on GDP forecast.

3.3 Methodology

3.3.1 Fed Sentiment Extraction

Sentiment extraction from the Federal Reserve Board of Governors' speeches is done in two ways. The first approach is lexicon-based and follows the spirit of Aruoba and Drechsel (2024), which is a conventional tool as it is based on word-by-word analysis of the text and n-grams, which reduce some researcher discretion but remain sensitive to context and preprocessing choices (Loughran and McDonald, 2011).[2] For example, if we consider a word such as ”Growth” — in the lexicon-based analysis — this can create confusion; in fact, if it is preceded by ”GDP”, it may have a very positive meaning, whereas if it is preceded by ”Inflation”, this may have an entirely negative meaning.

Lexicon-based analysis removes stopwords from the text, ensuring accurate word-by-word analysis but making it more difficult to capture the overall meaning of the discourse. Moreover, the lexicon introduced by Loughran and McDonald (2011) contains indications of positivity or negativity in words and criteria such as litigious, uncertain, constraining, or superfluous.

The lexicon-based sentiment measure is defined as:

$$Fed\ Sent_{i,t}^{Lexicon} = \frac{\#\text{positive words}_{i,t} - \#\text{negative words}_{i,t}}{\#\text{positive words}_{i,t} + \#\text{negative words}_{i,t}} \tag{3.1}$$

The second approach relies on finBERT, a BERT-based language model pre-trained on financial text (Devlin et al., 2019; Araci, 2019; Yang et al., 2020). Unlike dictionary-based methods, finBERT exploits contextual information at the sentence level and is therefore better suited to capturing the polarity of financial statements.

A model such as finBERT, rather than cleaning the text of stopwords, prefers a context-focused approach, analysing speech sentences directly to better understand their context. Unlike the financial lexicon developed by Loughran and McDonald (2011), finBERT classifies analysed sentences into negative, positive, or neutral polarity categories.

The sentiment measures based on finBERT are defined as:

$$Fed\ Sent_{i,t}^{finBERT} = \frac{\#\text{positive sentences}_{i,t} - \#\text{negative sentences}_{i,t}}{\#\text{positive sentences}_{i,t} + \#\text{negative sentences}_{i,t}} \tag{3.2}$$

and,

[2] An *n-gram* is a sub-sequence of n elements of a given sequence. Depending on the application, these elements may be phonemes, syllables, letters, words, etc. An *n-gram* of length 1 is called a *unigram*, of length 2 a *digram*, of length 3 a *trigram* and from length 4 onwards an *n-gram*.

$$Fed\ Sent_{i,t}^{finBERT\ Neu} = \frac{\#\text{positive sentences}_{i,t} - \#\text{negative sentences}_{i,t}}{\#\text{positive sentences}_{i,t} + \#\text{neutral sentences}_{i,t} + \#\text{negative sentences}_{i,t}} \tag{3.3}$$

where:

- i represents the speaker who gave the speech at that given date t;
- t refers to the speeches on a specific date.

In the speech data, it is possible to find some days with more than a single speech, delivered by different members of the Board of Governors. Firstly, we consider the sentiment for every speaker i for each day in the data t. Second, to consider the sentiment of the Federal Reserve Board of Governors on that specific day, and not just the sentiment of every single speaker, we consider the total number of positive words (sentences) and negative words (sentences) for that specific day and then proceed as explained above, simply considering the time index t.

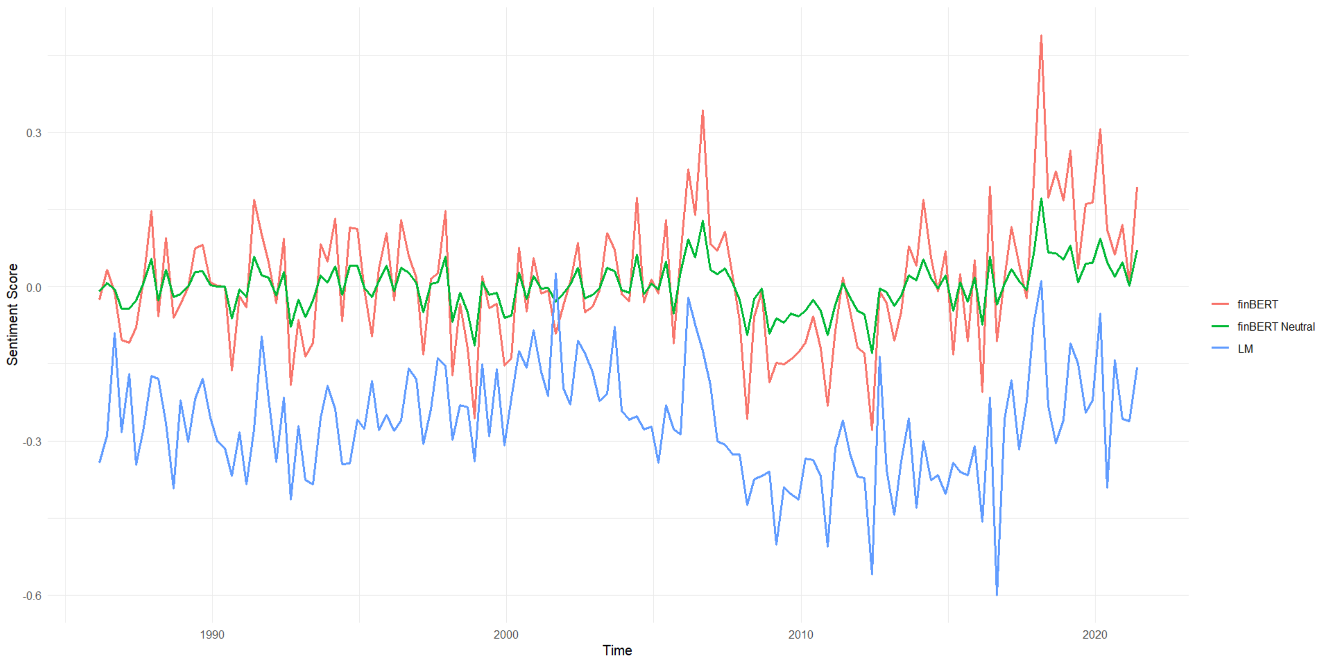

Figure 3.1: Sentiment extracted from Federal Reserve Board of Governors' speeches, quarterly frequency, 1986Q1–2021Q2. The figure compares sentiment time series obtained using finBERT, finBERT-Neutral, and Loughran–McDonald (LM).

Notes: Sentiment is computed from the text of Board of Governors speeches using three NLP models. finBERT and finBERT-Neutral rely on transformer-based contextual embeddings; LM uses a financial dictionary. Values are standardised.

Central bank speeches have a frequency that is not daily. On some days, it is possible that no speeches will be delivered. For these reasons, we decided to aggregate sentiments extracted for a specific day, at the weekly, monthly, and quarterly levels. Given the frequency of FRED data at the quarterly level, we define the quarterly generic Fed's sentiment for both *lexicon* and *finBERT* as s_t.

Figure 3.1 compares the three sentiment series visually and shows their evolution over time. Figure 3.2 summarises their distributions.

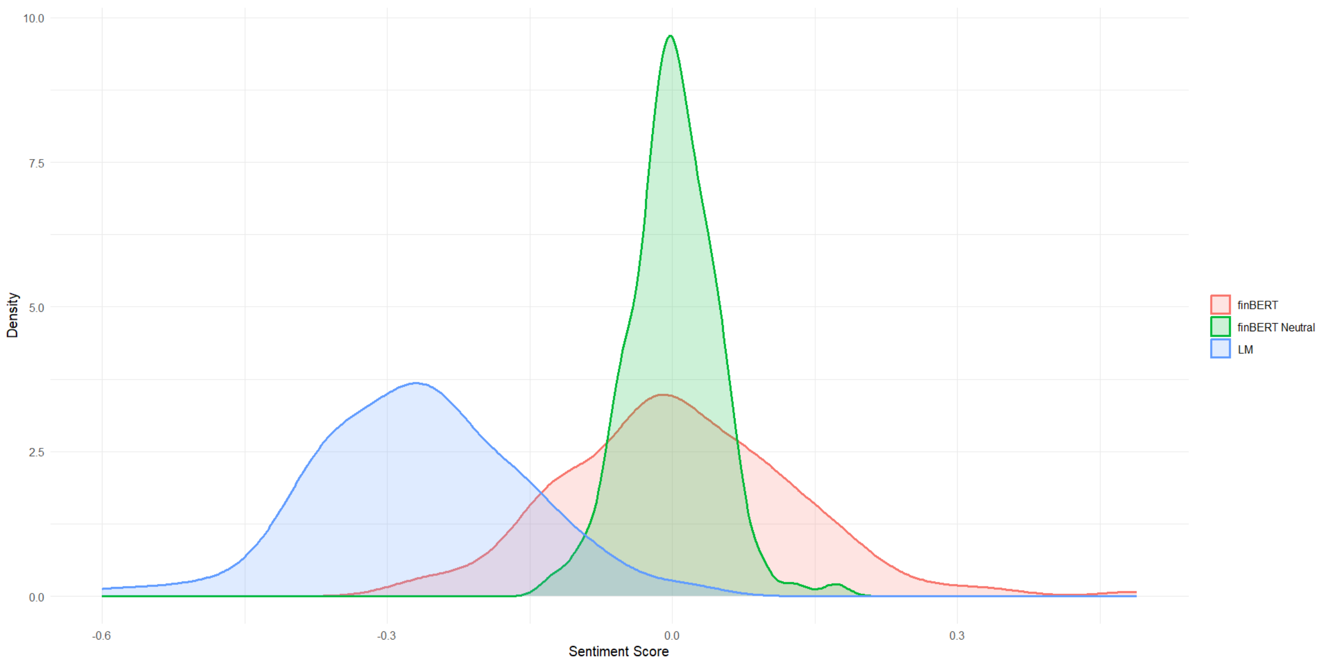

Figure 3.2: Kernel density estimates of the distributions of sentiment scores extracted from Federal Reserve Board of Governors' speeches using three text-based sentiment measures: finBERT, finBERT with neutral sentences, and the Loughran–McDonald (LM) dictionary, over the full sample.

Notes: Sentiment scores are computed from Board of Governors speeches using three alternative NLP methods: (i) finBERT sentiment probabilities; (ii) finBERT-Neutral, where neutral classifications are retained; (iii) LM dictionary-based. All series are standardised to a zero mean and unit variance. Kernel densities use a Gaussian bandwidth.

We refer to this measure as *net positive sentiment*, as it reflects the balance between positive and negative language in the statements of the Federal Reserve's Board of Governors. This terminology distinguishes the indicator from other communication-based measures, such as hawkish/dovish tone, uncertainty tone, or topic-specific sentiment.

The finBERT-based *net positive sentiment* indicator is computed in two versions. The first normalises the difference between positive and negative sentences by the sum of positive and negative sentences only. The second includes neutral sentences in the denominator, thereby scaling the positive-negative balance by the total number of classified sentences. The former captures the polarity of non-neutral language, while the latter captures the overall intensity of net positive sentiment relative to the full speech content.

In the baseline empirical analysis, the Fed communication shock is identified using the finBERT sentiment measure that includes neutral sentences in the denominator, denoted as *Fed Sent finBERT neu*. This choice is motivated by the fact that the measure scales the positive-minus-negative balance by the total number of classified sentences, thereby capturing the overall net positive tone of each speech relative to its full textual content. The alternative finBERT measure, *Fed Sent finBERT*, excludes neutral sentences from the denominator and therefore captures the polarity of non-neutral sentences only. The Loughran–McDonald measure, *Fed Sent Lexicon*, is used as a dictionary-based robustness check directly into the main analysis.

Unless otherwise stated, all impulse response functions in the main empirical section are based on the communication shock identified using *Fed Sent finBERT neu* and *Fed Sent Lexicon*.

For further information about speakers, the number of speeches, the sentiment expressed during the period in which they have a position on the Board of Governors, and other details, see Appendix A.

3.3.2 Fed Communication Shocks

Structural Dynamic Factor Model (SDFM)

This subsection outlines the Structural Dynamic Factor Model (SDFM) proposed by Stock and Watson (2005) and Forni et al. (2009). Let x_t be a n-dimensional, zero-mean stationary vector of high-dimensional observable economic variables, we can consider the following representation of SDFM:

$$x_t = \chi_t + \xi_t = \Lambda F_t + \xi_t \tag{3.4a}$$

$$C(L)F_t = \epsilon_t \tag{3.4b}$$

$$\epsilon_t = Ru_t \tag{3.4c}$$

where Λ is a $n \times r$ matrix of coefficients, F_t is a $r \times 1$ vector of unobserved factors and ξ_t a $n \times 1$ vector of idiosyncratic components. In this model, it is assumed that the idiosyncratic component and the common component are orthogonal at all leads and lags. The vector $\epsilon_t \sim WN(0, \Sigma_\epsilon)$ where $C(L)$ is an $r \times r$ stable polynomial matrix, u_t is a $q \times 1$ vector of orthonormal structural shocks (with $q \leq r$), and R is an $r \times q$ matrix of coefficients with rank q.

We focus on the common component $\chi_t = \Lambda F_t$ because economic shocks, forward guidance included, are pervasive, so they live in the space of common components, which coincides with the space of factors. By inverting the stable polynomial matrix $C(L)$ we can derive $F_t = C(L)^{-1}\epsilon_t = C(L)^{-1}Ru_t$; thus the dynamic relationship between u_t and the common component is:

$$\chi_t = \Lambda C(L)^{-1}Ru_t = B(L)u_t \tag{3.5}$$

Then, we can consider the structural dynamic representation:

$$x_t = B(L)u_t + \xi_t \tag{3.6}$$

where macroeconomic variables are affected by some structural shocks found within the $B(L)$ matrix, plus measurement error. However, at this stage, the impulse responses are not identified because the representation of Eq.(3.6) is not unique. In this case, the identification is achieved up to orthogonal rotations, as in a Structural VAR model.

We define $R = SH$, where S is an $r \times q$ matrix such that $SS^{'} = \Sigma_\epsilon$, and H is a $q \times q$ orthonormal matrix (a matrix such that $H^{-1} = H^{'}$) which imposes the identifying restrictions. Thus, the structural representation for χ_t is

$$\chi_t = \Lambda C(L)^{-1} SHu_t = D(L)Hu_t = B(L)u_t \tag{3.7}$$

and the structural shocks are given by $u_t = R^{-1}\epsilon_t = H^{'}S^{-1}\epsilon_t = H^{'}\eta_t$, where $\eta_t \sim WN(0, I)$. Assuming that, without loss of generality, the Fed communication shock is u_{1t}, then we need the first column h_1 such that $h_1^{\prime}\eta_t = u_{1t}$.

Identification of the Fed Communication Shock

The identification strategy is closely related to the frequency-domain approach of FGGSS, but it is adapted to the present setting. In their application, the target variable is an observed macroeconomic variable. Here, by contrast, the target is a constructed textual indicator: Federal Reserve sentiment extracted from speeches of the Board of Governors. The identification, therefore, aims to recover the structural innovation that maximises the low-frequency covariance between this targeted sentiment measure and the FFR.

Let x_t denote the vector of observable variables, ordered such that the first two entries correspond to (i) the communication sentiment index s_t and (ii) the monetary policy instrument m_t (the FFR for the US). The frequency–band covariance matrix is:

$$V(\underline{\theta}, \overline{\theta}) = \int_{\underline{\theta}}^{\overline{\theta}} \Re\left(D(e^{-i\theta})D(e^{i\theta})^{'}\right)\, d\theta \tag{3.8}$$

where the interval $[\underline{\theta}, \overline{\theta}]$ is a band of frequencies such that $0 \leq \underline{\theta} \leq \overline{\theta} \leq \pi$, $\Re(.)$ represent the real part of (.) and the matrix $V(\underline{\theta}, \overline{\theta})$ is that matrix that captures the entire frequency band covariance of the variables that stay in a range $[\underline{\theta}, \overline{\theta}]$ — defined as the *frequency band covariance matrix.*[3]

The identification strategy consists of imposing restrictions on the contribution of a candidate structural shock to the elements of the *frequency band covariance matrix.* The contribution of a

[3]This $\Re(.)$ is useful because we are interested in identifying the shock in the common component of the model because the imaginary components of $V(.)$ vanish in the time domain when all frequencies are aggregated and, for this reason, do not contribute to covariance.

generic shock to the band covariance is:

$$\Psi(\underline{\theta}, \overline{\theta}) = \int_{\underline{\theta}}^{\overline{\theta}} \Re\left(D(e^{-i\theta})hh^{'}D(e^{i\theta})^{'}\right) d\theta \tag{3.9}$$

where $\Psi(\underline{\theta}, \overline{\theta})$ represent the variance (or covariance) contribution of any generic shock $h^{'}\eta_t$, where h is such that $h^{'}h = 1$, to $V(\underline{\theta}, \overline{\theta})$.

Let $\Psi_{lk}(\underline{\theta}, \overline{\theta})$ be the l, k element of $\Psi(\underline{\theta}, \overline{\theta})$. In the single target case, identification is obtained by selecting the shock that maximises the relevant element of this matrix over the chosen frequency band. Let $(M_1, N_1), (M_2, N_2), \ldots, (M_m, N_m)$ be the m entries of interest, it's possible to target a weighted sum of such entries $\sum_{k=1}^{m} \omega_k \Psi_{M_k N_k}(\underline{\theta}, \overline{\theta})$, where ω are weights assigned to each element of the sum. It's possible to rewrite the sum as:

$$\sum_{k=1}^{m} \omega_k \Psi_{M_k N_k}(\underline{\theta}, \overline{\theta}) = h^{'}O_{MN}(\underline{\theta}, \overline{\theta})h \tag{3.10}$$

where $O_{MN}(\underline{\theta}, \overline{\theta}) = \int_{\underline{\theta}}^{\overline{\theta}} R\left(D(e^{i\theta})^{'}P_M^{'}\Omega P_N D(e^{-i\theta})\right) d\theta$.[4]

In this application, we are interested in the first column of the rotation matrix H, denoted by h_1. We identify the Fed communication shock as the structural innovation that maximises the low-frequency covariance between the communication sentiment index s_t and the monetary policy instrument m_t. Formally given by:

$$h_1 = \arg\max_{h \in \mathbb{R}^q} h^{'}O_{MN}(\underline{\theta}, \overline{\theta})h \quad \text{s.t.} \quad h^{'}h = 1. \tag{3.11}$$

We have $m = 1$ and the frequency band $[\underline{\theta}, \overline{\theta}] = [0, 2\pi/32]$. With quarterly data, this band captures low-frequency movements corresponding to cycles of eight years or longer. This choice is intended to isolate the persistent component of the covariance between communication sentiment and the policy rate, consistent with the idea that forward guidance affects expectations and monetary policy dynamics at horizons that may be distant from the original communication event.

Since the communication sentiment index and the monetary policy instrument are ordered as

[4]The contribution of the shock $h^{'}\eta_t$ to the weighted sum $\sum_{k=1}^{m} \omega_k \Psi_{M_k N_k}(\underline{\theta}, \overline{\theta})$ is given by

$$h^{'}\left[\int_{\underline{\theta}}^{\overline{\theta}} R\left(D(e^{-i\theta})^{'}\sum_{k=1}^{m} \omega_k \mathcal{E}_{M_k}^{'}\mathcal{E}_{N_k} D(e^{-i\theta})\right) d\theta\right] h$$

where ω_k are the weights (a reasonable choice for the weights is to take the reciprocals of the cyclical variances of the variables). The weighted sum $\sum_{k=1}^{m} \omega_k \mathcal{E}_{M_k}^{'}\mathcal{E}_{N_k} = P_M^{'}\Omega P_N$, where $P_M = \left(\mathcal{E}_{M_1}^{'}, \mathcal{E}_{M_2}^{'}, \ldots, \mathcal{E}_{M_m}^{'}\right)^{'}$ and $P_N = \left(\mathcal{E}_{N_1}^{'}, \mathcal{E}_{N_2}^{'}, \ldots, \mathcal{E}_{N_m}^{'}\right)^{'}$ are $m \times n$ matrices, and $\Omega = \text{diag}(\omega_1, \omega_2, \ldots, \omega_m)$ is an $m \times m$ matrix. In the case of multiple targets, this is the objective function for the identification problem. Of course, this objective function reduces to the single target objective function in the case $m = 1$.

the first two variables in x_t, we set $M = 1$ and $N = 2$. The targeted entry of the frequency-band covariance matrix is therefore the covariance between Federal Reserve sentiment and the FFR over the selected low-frequency band. The economic intuition is that speeches by members of the Federal Reserve Board of Governors may convey forward-guidance signals by shaping expectations about the future path of monetary policy. The identification strategy, therefore, targets the persistent co-movement between Fed sentiment and the FFR rather than high-frequency surprises around individual announcements.

The resulting shock is interpreted as a Fed communication shock operating through the forward-guidance channel. It captures the component of Federal Reserve communication sentiment that is systematically associated with medium- to long-run monetary policy dynamics. This interpretation is assessed empirically below by comparing the identified shock with standard monetary policy, forward guidance, LSAP, fiscal, and news shocks.

The interpretation of the identified communication shock is also related to the literature on central bank information effects. Nakamura and Steinsson (2018) show that monetary policy announcements may reveal information about the central bank's assessment of economic fundamentals, while Jarociński and Karadi (2020) distinguish between pure monetary policy shocks and central bank information shocks using the high-frequency comovement of interest rates and stock prices around policy announcements. This distinction is important for the present analysis because the shock identified here is not intended to capture a conventional monetary policy surprise. Rather, it captures innovations in the positive tone of Federal Reserve communication that are systematically associated with persistent monetary-policy-related movements. In this sense, the shock is closer to an information or communication shock than to a pure policy-rate shock, although it is identified from speeches and low-frequency macro-financial comovement rather than from high-frequency asset-price surprises around FOMC announcements.

Relation with Forni et al. 2025 The identification strategy adopted in this chapter is closely related to the frequency-domain structural dynamic factor approach of Forni et al. (2025), who identify macroeconomic structural shocks by exploiting low-frequency comovement between observed macroeconomic variables and the common component of a large information set. The present application departs from that framework in two main respects.

First, the target variable is not a standard observed macroeconomic aggregate, but a constructed textual indicator extracted from the speeches of members of the Federal Reserve Board of Governors. Second, the structural innovation is selected to maximise the low-frequency covariance between this targeted communication measure and the FFR, thereby isolating the component of Federal Reserve communication that is most closely associated with persistent monetary-policy-related movements. Hence, while the econometric machinery follows the same frequency-domain logic, the object of identification is different: the shock is designed to capture a communication innovation embedded in central bank language rather than a conventional macroeconomic disturbance.

3.3.3 International Spillover - Unbalanced Panel Local Projections

To study the international spillover effect of the Fed communication shock coming from the Fed and identified using the FGGSS methodology, we implemented an unbalanced panel Local Projection (henceforth LP). To avoid problems with data size, we applied the LP for each variable in the country block (see the section 3.2.2).

We implement an LP by estimating the following:

$$y_{i,t+h} - y_{i,t-1} = \beta^h s_t + \sum_{l=1}^{L} \alpha_l^h \, \Delta y_{i,t-l} + \sum_{l=1}^{L} \gamma_l^h \, s_{t-l} + \delta^h X_{i,t} + \mu_i^h + \varepsilon_{i,t}^h \tag{3.12}$$

for $h = 0, 1, 2, \ldots$, where $y_{i,t+h}$ represents the outcome variable in country i at time $t+h$, i.e., $y_{i,t+h} = \{ln(GDP_{it+h}), IP_{it+h}, ln(I_{it+h}), ln(SP_{it+h}), WUI_{it+h}, WSI_{it+h}, EPU_{it+h}, ln(SV_{it+h}), GDP\ disagreement_{it+h}\}$, and s_t is the Fed communication shock identified using the SDFM and the identification strategy from FGGSS. Following Fernández-Villaverde et al. (2024), we include lagged outcome and explanatory variables, $\Delta y_{i,t-l}$ and s_{t-l}, to address serial correlation, choosing a lag length of two.

The motivation for using local projections is that the Fed communication shock is first identified within the US economy and then used as an external measure of the underlying communication disturbance in the international analysis. Since this estimated shock is an imperfect proxy for the true structural innovation, the local projection framework provides a direct way to estimate average dynamic responses across countries while conditioning on global and country-specific controls.

The regressor vector $X_{i,t}$ includes global and country-specific controls: the first and second lags of a country's GDP growth rate per capita, the Industrial Production growth rate, the Fixed Investment growth rate, the Smoothed WUI, the WSI, the EPU and the global variables — the VIX, the S&P 500 index, the WTI oil price, the yield on two-year US Treasuries, the NFCI, and world GDP. WTI oil price and world GDP are in log-difference form. μ_i^h represents country FEs, and $\varepsilon_{i,t}^h$ is the error term. Standard errors are clustered by time. The β^h represents the IRFs from LP estimation for $h = 0, 1, 2, \ldots$. The inclusion of both global and country-specific controls aims to reduce the risk that the estimated responses capture broader global macro-financial disturbances unrelated to Fed communication.

We run the regression separately for each country variable. The number of observations is different across variables and depend on data availability: 4,987 for GDP per capita, 6,145 for Industrial Production Index, 4,498 for Fixed Investment, 3,691 for Stock Prices Index, 28,897 for Smoothed WUI, 10,739 for WSI, 1,954 for EPU, 6,421 for Stock Volatility, and 1,879 for GDP forecast disagreement in the longest horizon estimation ($h = 20$). Because data availability differs substantially across variables and countries, the resulting panel is highly unbalanced. Consequently, the estimated spillover effects should be interpreted as average responses conditional on data availability rather than as effects estimated on a perfectly homogeneous cross-country sample.

In the table 3.1, we show how many countries for each variable we consider in the analysis.

Table 3.1: Count of countries with non-missing data by variable and group.

Variable	# of countries for LP	# of EM	# of AE	Total of countries in the dataset
GDP per capita	67	26	41	152
Industrial Production	52	13	39	152
Fixed Investment	62	21	41	152
Stock Prices	35	14	21	152
World Uncertainty Index	142	98	44	152
World Sentiment Index	142	98	44	152
Economic Policy Uncertainty	20	5	15	152
Stock Volatilities	56	25	31	152
GDP forecast disagreement	22	7	15	152

Notes: EM stands for Emerging Markets and AE stands for Advanced Economies.

For more information on the sample countries and descriptive statistics of the countries and variables involved in the analysis of international spillover effects, see Appendix F.

3.4 Domestic Analysis Results

3.4.1 Identified Shock Visualization

Figure 3.3 plots the estimated Federal Reserve communication shock obtained from the SDFM. The shock captures innovations in the positive tone of Federal Reserve speeches that are associated with persistent movements in the FFR at low frequencies. Positive realisations of the shock, therefore, correspond to episodes in which Fed communication becomes more favourable, confident, or reassuring in a way that is systematically related to the monetary policy stance. Importantly, the shock is not simply the raw sentiment series. It is the structural innovation extracted from the common component of the large macro-financial information set and targeted to the low-frequency comovement between the textual sentiment indicator and the policy rate.

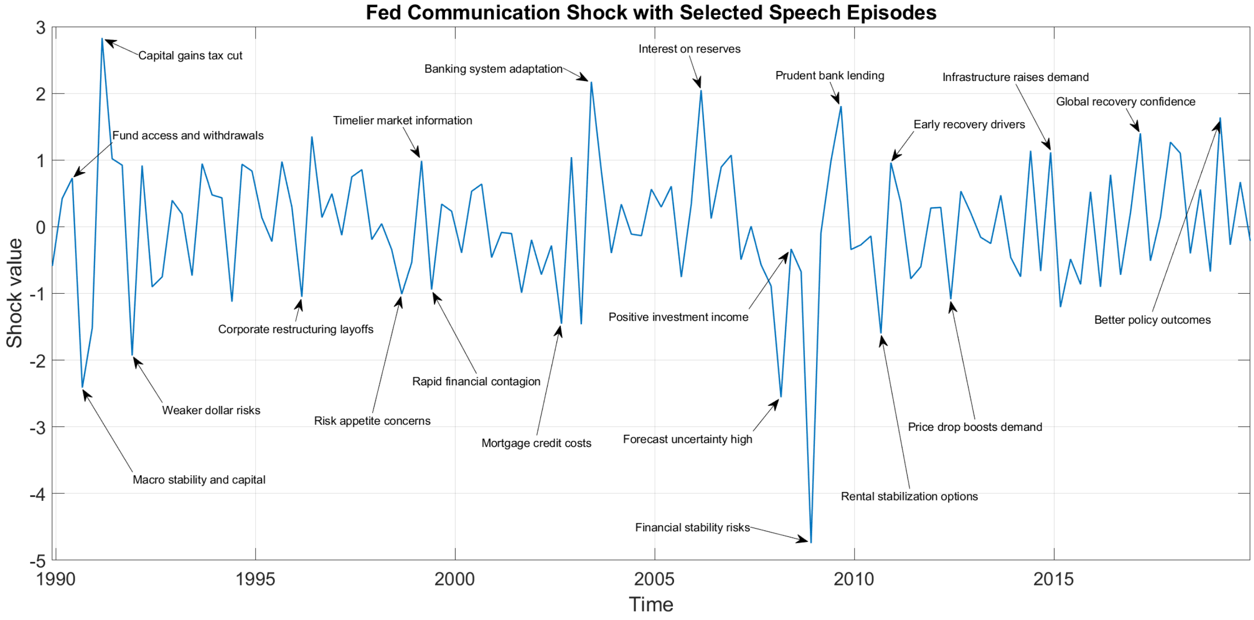

Figure 3.3: Identified Fed communication shock sentiment and selected high-impact speech episodes

Notes: The figure reports the standardised Fed communication shock identified using finBERT sentiment. Labels indicate selected high-impact episodes and are intentionally shortened for readability. The communication shock is identified using the procedure in Section 3.3.2. The speech excerpts correspond to dates of large positive or negative shocks and illustrate the underlying narrative content.

The visual inspection of the shock is useful for two reasons. First, it provides a transparent assessment of the timing and magnitude of the identified communication innovations. Second, it allows the estimated shock to be compared with major monetary policy and macro-financial episodes, thereby helping to interpret whether the extracted signal behaves consistently with the narrative evidence on Fed communication.

3.4.2 Validation and Robustness of the Fed Communication Shock

Correlation Results

A first validation exercise compares the identified Fed communication shock with existing shock series used in the literature. This comparison is central to the interpretation of the shock: if the identified series were strongly correlated with standard monetary policy, forward guidance, LSAP, fiscal, news shocks, or Central Bank Information shocks, it would be difficult to interpret it as a distinct communication driver of macroeconomic fluctuations. We therefore compute correlations between the Fed communication shock and the benchmark shock series at contemporaneous and lagged horizons from $k = 0$ to $k = 4$ quarters.

For each quarterly lag $k = 0, 1, \ldots, 4$, we test the null of zero correlation:

$$H_0 : \rho_k = 0 \qquad \text{vs.} \qquad H_1 : \rho_k \neq 0.$$

Table 3.2: Correlation test between Fed communication shock — identified using finBERT and Loughran-McDonald (LM) — and other shocks in literature

	k = 0		k = 1		k = 2		k = 3		k = 4	
	finBERT	LM	finBERT	LM	finBERT	LM	finBERT	LM	finBERT	LM
FFR (HFI)	0.08	0.08	0.07	0.07	0.07	0.08	0.07	0.08	0.07	0.08
FG (HFI)	0.05	0.04	0.09	0.06	0.08	0.05	0.08	0.04	0.08	0.05
LSAP (HFI)	-0.17	-0.09	-0.18	-0.08	-0.15	-0.08	-0.16	-0.09	-0.16	-0.08
Ramey Gvt. Spend.	-0.03	-0.01	-0.04	-0.04	-0.04	-0.04	-0.03	-0.01	-0.04	-0.04
BZP Gvt. Spend.	0.04	0.06	0.00	0.01	0.00	0.00	-0.02	0.06	0.00	0.01
BP Gvt. Spend.	0.13	0.16	0.06	-0.09	0.06	-0.09	0.05	0.16	0.06	-0.09
CEE mp	0.07	0.08	0.08	0.09	0.09	0.10	0.06	0.08	0.06	0.09
CEE mp (full)	0.08	0.07	0.09	0.09	0.09	0.09	0.06	0.07	0.07	0.08
CEE mp (later)	0.09	0.08	0.10	0.09	0.11	0.10	0.07	0.08	0.08	0.10
R&R mp	0.15	0.19	0.15	0.19	0.16	0.20	0.16	0.19	0.16	0.20
Ramey Tax	0.06	-0.02	0.04	-0.04	0.04	-0.05	0.03	-0.02	0.03	-0.07
Jarocinski–Karadi MP	0.08	0.09	0.10	0.10	0.09	0.10	0.09	0.10	0.09	0.10
Jarocinski–Karadi CBI	-0.02	0.00	-0.02	0.00	-0.01	0.01	-0.01	0.01	0.00	0.01
Miranda Agrippino–Ricco	-0.07	-0.01	-0.06	0.00	-0.05	0.00	-0.05	0.00	0.01	0.05

Notes: Factors used for Structural shock identification in DFM are $r = 14$. (Christiano et al., 1999) — CEE; (Romer and Romer, 2004; Ramey, 2016; Swanson, 2021) — FFR (Federal Funds Rate), FG (Forward Guidance), LSAP (Large-Scale Asset Purchase), HFI (High-Frequency Identification); *Ramey Gvt. Spend.* — Ramey news shock; *BZP Gvt. Spend.*: Ben-Zeev–Pappa shock; *BP*: Blanchard-Perotti shock (the shocks from *Ramey Gvt. Spend.* to *Ramey Tax*, come from Ramey (2016)), *Jarocinski–Karadi MP and CBI* come from Jarociński and Karadi (2020); *Miranda Agrippino–Ricco* is the shock estimated in Miranda-Agrippino and Ricco (2021). Results shown for each quarterly lag $k = 0, 1, \ldots, 4$. P-value: 0.001 (***), 0.01 (**), 0.05 (*).

Ljung-Box Test and One-Sample t-test

The *Ljung-Box Test for Autocorrelation* shows us that there is no evidence of autocorrelation in the shock, using 24 lags, for both finBERT and Loughran-McDonald identified shock — see Table 3.3.

Table 3.3: Ljung-Box test for Autocorrelation

Test	X^2	df	p-value
finBERT shock	21.18	24	0.628
LM shock	35.15	24	0.066

Notes: The Ljung-Box test verifies the null hypothesis of no autocorrelation up to a lag of 24. In both cases, there is no statistical evidence of autocorrelation. P-value: 0.001 (***), 0.01 (**), 0.05 (*).

The *One Sample t-test for Zero Mean* shows us that the mean of the shock is not significantly different from zero, for both finBERT and Loughran-McDonald identified shock — see Table 3.4.

Table 3.4: One-Sample t-test for Zero Mean

Test	t	df	p-value	95% CI
finBERT shock	4.3×10^{-15}	120	1	[-0.18, 0.18]
LM shock	-5.7×10^{-15}	120	1	[-0.18, 0.18]

Notes: The table reports t-tests for the null hypothesis that the mean of each identified shock is equal to zero. In both cases, the null cannot be rejected. P-value: 0.001 (***), 0.01 (**), 0.05 (*).

Communication shocks appear to be free of autocorrelation, have zero mean and are uncorrelated with other shocks in the literature — up to a lag of 4 quarters *—, consistent with the exogenous shock hypothesis.*

Robustness to Alternative Constructions of the Communication Shock

We then assess whether the domestic impulse responses are robust to alternative constructions of the Fed communication shock. In the baseline specification, the shock is identified using both Loughran–McDonald and finBERT with neutral sentences sentiment over the communication band $[0, 2\pi/32]$, by targeting the low-frequency co-movement between the textual sentiment indicator and the FFR. We compare this specification with shocks obtained using the Loughran–McDonald and finBERT sentiment measures over alternative frequency bands, and with an extended domestic sample that excludes the AAII bull–bear spread. These exercises do not test the exogeneity of the shock. Rather, they evaluate whether the main domestic responses depend on specific modelling choices used to construct the communication shock. We report these results in Appendix B., Appendix C., and Appendix D.

In addition, we address the concern that the Federal Funds Rate may be an imperfect policy indicator during the zero lower bound period. During this period, the policy rate was constrained for several quarters, while Federal Reserve communication continued to convey information about the expected path of monetary policy, the policy reaction function, and macroeconomic conditions. To account for this issue, we repeat the identification by replacing the FFR with the Wu–Xia shadow rate (Wu and Xia, 2016). The shadow rate provides a broader measure of the monetary policy stance, especially when the zero lower bound constrains conventional short-term rates. Since the Wu–Xia shadow rate starts in the first quarter of 1990, this robustness exercise is estimated over the $1990Q1$–$2021Q2$ sample, ending in the same quarter as the baseline domestic analysis but starting two years later.

We implement this robustness exercise for both the finBERT sentiment measure with neutral sentences in the denominator and the Loughran–McDonald sentiment measure. The resulting impulse responses are reported in Appendix E. Overall, the main qualitative patterns are preserved. A positive communication shock identified using the Wu–Xia shadow rate is followed by an expansion in real activity, an increase in short-term and longer-term interest rates, and a decline in several uncertainty indicators. SPF disagreement also falls on impact for most forecast variables, while mean growth expectations generally display responses that are consistent with the baseline evidence, although some variables exhibit more heterogeneous medium-run dynamics. These results suggest that the main domestic findings are not mechanically driven by using the FFR as the policy variable in the identification step.

3.4.3 Impulse Response Functions for Identified Shocks in the US

For the baseline results, we use the sample starting in the first quarter of 1988, which allows us to include the bull–bear sentiment spread among investors in the information set. The impulse responses trace the effects of a one-standard-deviation positive Fed communication shock on disagreement, expectations, macroeconomic variables, financial conditions, and uncertainty indicators. Throughout the section, the shock should be interpreted as the communication-related macro-financial innovation identified in Section 3.3.2, rather than as a purely exogenous communication disturbance. This distinction is important because Fed communication is partly endogenous to the macroeconomic and financial environment in which it occurs.

IRFs on Disagreement

The Figure 3.4 and Figure 3.5 show the IRFs for the Disagreement of SPFs from a standard deviation Fed communication shock (positive). The disagreement is a measure of dispersion in the SPF and, for this reason, is used as a proxy for the uncertainty perceived by economic operators. As expected, a positive Fed communication shock generates an initial robust reduction in disagreement among professional forecasters for all the variables.

To quantify the magnitude of the effects, we express the peak impulse responses relative to the standard deviation of each disagreement measure. A one-standard-deviation Fed communication shock reduces Real GDP disagreement by about 0.06 standard deviations, and disagreement in Nominal GDP and GDP price forecasts by 0.09–0.12 standard deviations. Real Consumption disagreement falls by 0.08 standard deviations, while Industrial Production disagreement shows a somewhat larger response of 0.28 standard deviations. Two series with different underlying scales (Net Profit and Housing disagreement) display larger standardised effects, around 0.7 and 1.9 standard deviations, respectively. These estimates suggest that the Fed communication shock has economically meaningful effects in reducing forecast dispersion. The confidence bands, in turn, are informative about statistical uncertainty around the point estimates rather than the economic size of the responses.

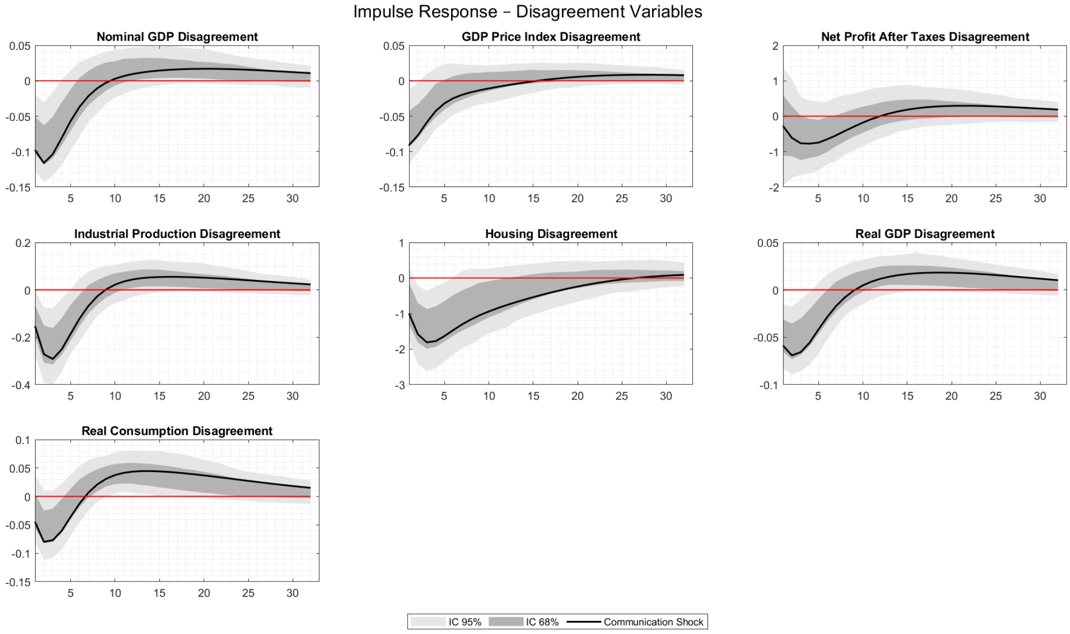

Figure 3.4: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are computed from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is a structural innovation identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different SPF disagreement measure (cross-sectional dispersion of forecasts) over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, with fourteen factors. The variable used for identification is the Sentiment extracted using finBERT with neutral sentences.

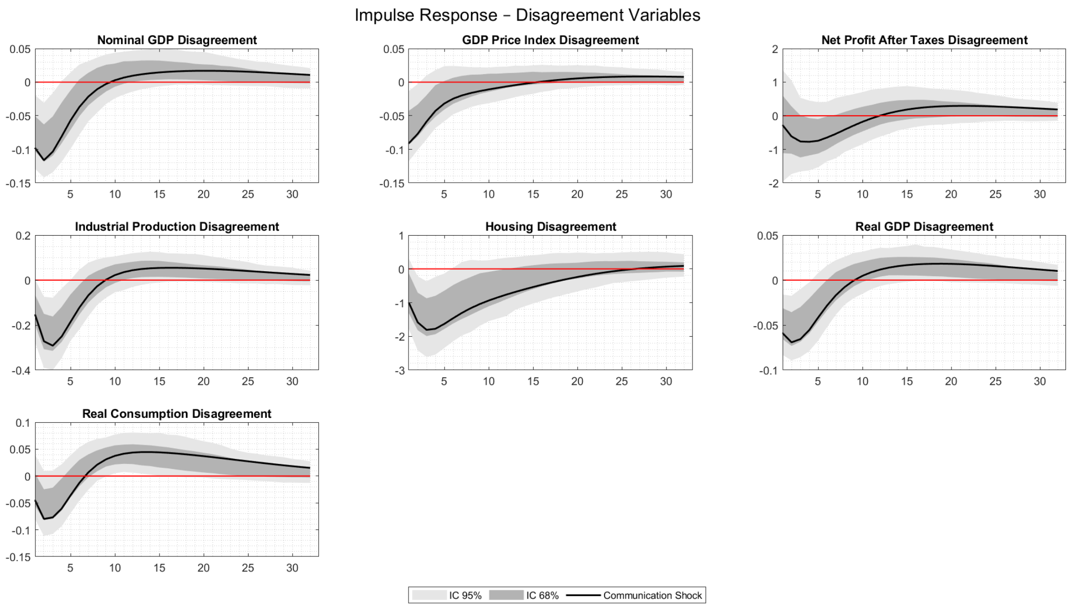

Figure 3.5: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Loughran-McDonald Fed communication shock. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are computed from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is a structural innovation identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different SPF disagreement measure (cross-sectional dispersion of forecasts) over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, with fourteen factors. The variable used for identification is the Sentiment extracted using the Loughran-McDonald Vocabulary.

Taken together, the results suggest that a positive Fed communication shock is associated with a decline in forecast dispersion across professional forecasters. This pattern is consistent with the view that central bank communication can coordinate expectations by reducing disagreement about the future path of macroeconomic conditions.

IRFs on Expectation

Figures 3.6 and 3.7 report the IRFs for SPF mean-growth expectations under the same communication shock. A positive Fed communication shock generates an initial increase in expectations for nominal and real growth, output, profits, and consumption. In the medium term, roughly after 5–10 quarters, most series revert toward slightly negative values.

The most pronounced and persistent negative effects, emerging after 5–10 quarters, are observed for housing and production expectations, suggesting a trade-off between the short-run boost in confidence induced by communication and its medium-run persistence.

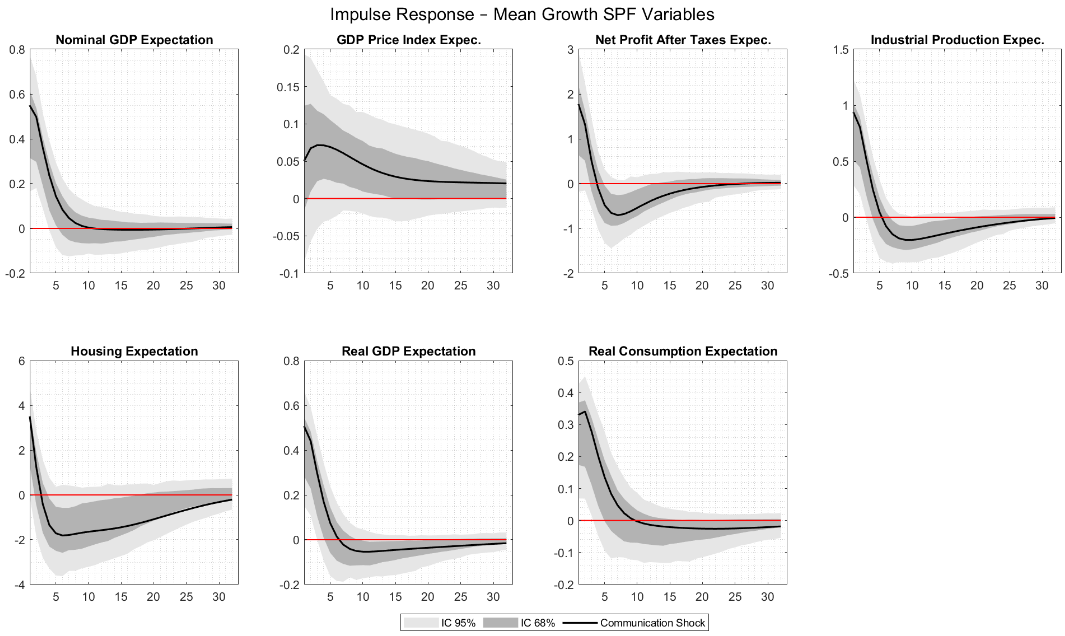

Figure 3.6: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are computed from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is a structural innovation identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different SPF mean growth expectation measures over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, with fourteen factors. The variable used for identification is the Sentiment extracted using finBERT with neutral sentences.

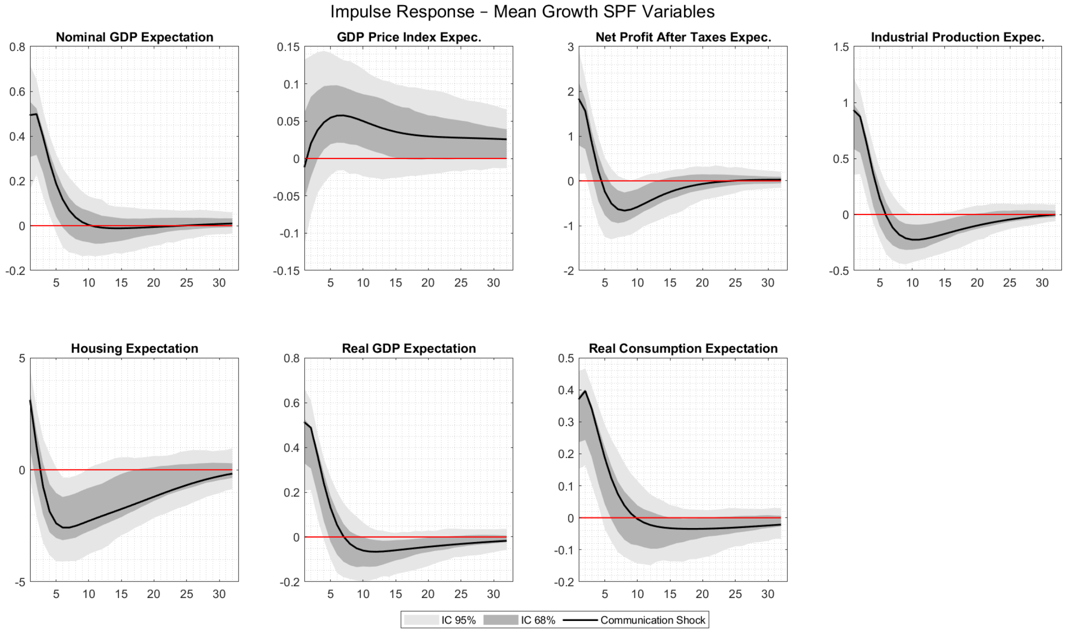

Figure 3.7: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Loughran-McDonald Fed communication shock. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are computed from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is a structural innovation identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different SPF mean growth expectation measures over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, with fourteen factors. The variable used for identification is the Sentiment extracted using the Loughran-McDonald Vocabulary.

To assess the magnitude of the effects on expectations, we express the peak IRFs in units of each SPF mean-growth series' standard deviation. A one-standard-deviation Fed communication shock increases real GDP expectations by about 0.55 standard deviations and real consumption expectations by 0.36 standard deviations. Nominal GDP expectations rise by 0.58 standard deviations, while inflation expectations respond more modestly, by around 0.08 standard deviations. Industrial production expectations display a stronger reaction of roughly one standard deviation. Net profit and housing expectations show larger standardised responses, at about 2.1 and 4.2 standard deviations, respectively.

IRFs on Macroeconomic Variables

Figures 3.8 and 3.9 report the responses of US macroeconomic and financial variables to the same positive Fed communication shock. Since the responses are standardised by the unconditional standard deviation of each variable, comparisons across series should be interpreted as relative magnitudes rather than as direct comparisons of economic units. The responses are broadly consistent with a confidence-enhancing or accommodative interpretation of the shock. Uncertainty indicators decline, financial conditions improve, and some real activity variables respond positively in the short run.

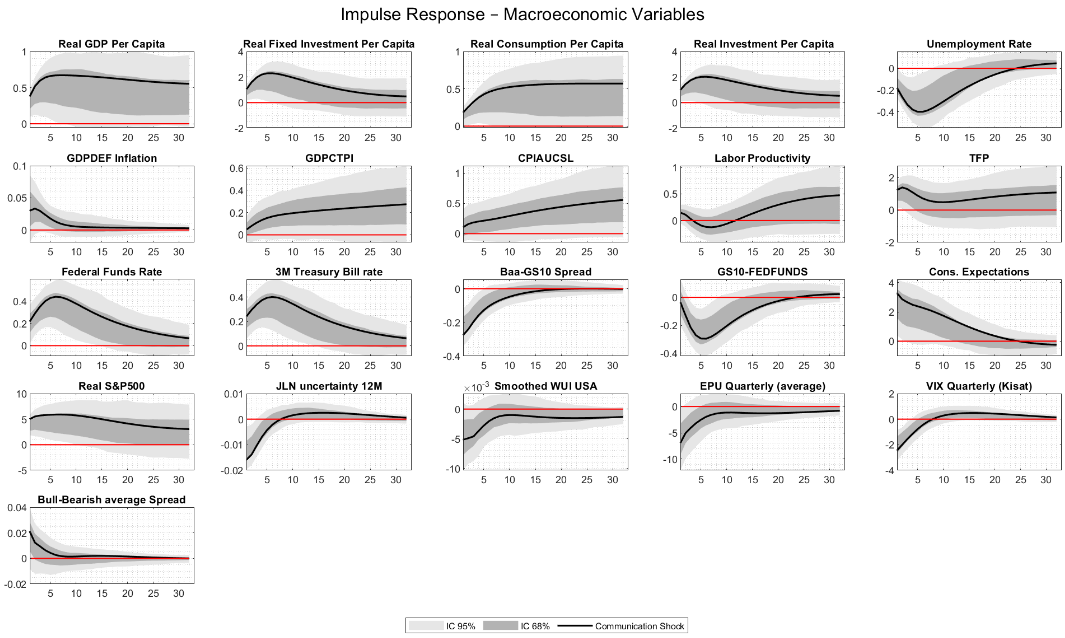

Figure 3.8: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. The figure reports responses for real activity, inflation, labour productivity, financial conditions, and uncertainty indicators over a 35-quarter horizon.

Notes: Impulse responses are computed from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is a structural innovation identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different macroeconomic and financial US variables over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, with fourteen factors. The variable used for identification is the Sentiment extracted using finBERT with neutral sentences.

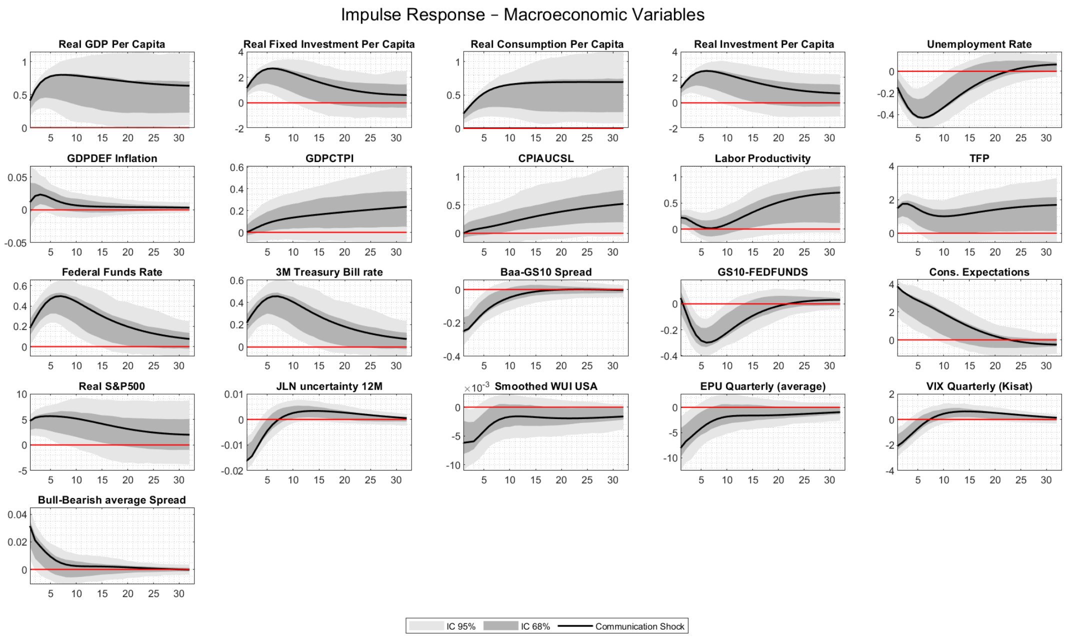

Figure 3.9: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Loughran-McDonald Fed communication shock. The figure reports responses for real activity, inflation, labour productivity, financial conditions, and uncertainty indicators over a 35-quarter horizon.

Notes: Impulse responses are computed from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is a structural innovation identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different macroeconomic and financial US variables over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, with fourteen factors. The variable used for identification is the Sentiment extracted using the Loughran-McDonald Vocabulary.

To quantify the macroeconomic effects, we express the peak IRFs in units of each variable's standard deviation. A one-standard-deviation Fed communication shock raises real GDP per capita and real consumption per capita by about 0.8 standard deviations. Investment responds more strongly, with increases of 2.3–2.6 standard deviations across measures. The unemployment rate declines by around 0.42 standard deviations, while inflation responses are smaller, ranging between 0.03 and 0.43 standard deviations depending on the measure.

The FFR and short-term Treasury rates increase by approximately 0.47–0.52 standard deviations, while uncertainty measures generally decline. These magnitudes indicate that the identified Fed communication shock is associated with economically relevant movements across real, nominal, financial, and uncertainty dimensions. Since the responses are standardised by the unconditional standard deviation of each variable, comparisons across series should be interpreted as relative magnitudes rather than as direct comparisons of economic units.

3.4.4 Summary of Domestic Transmission

Overall, the domestic responses are consistent with the recent literature emphasising the informational role of central bank communication. Fadda et al. (2025) show that central banks communicate uncertainty directly; in the present framework, a positive Fed communication shock is followed by lower uncertainty indicators and reduced SPF disagreement, suggesting that the information embedded in communication is perceived as reassuring. The decline in credit spreads and the increase in equity valuations are also consistent with Istrefi et al. (2023), who document that Federal Reserve speeches contain signals about financial stability concerns.

The reduction in disagreement is in line with Istrefi and Piloiu (2020), who emphasise that communication affects the informational environment in which agents form expectations. Finally, the positive responses of expectations and some real activity variables resonate with Bordo and Istrefi (2023), who show that communication is interpreted through perceived hawkish or dovish orientations. In this chapter, the positive Fed communication shock appears closer to a confidence-enhancing or relatively accommodative signal. However, given the identification strategy, this interpretation should be understood as suggestive rather than as evidence of a pure exogenous communication shock.

Appendix C. reports impulse responses estimated using alternative frequency bands for the identification process and the same sentiment extracted from finBERT and Loughran–McDonald, which confirms the robustness of the baseline results. Appendix D. provides additional robustness checks based on alternative sample windows with the same frequency band as the baseline results.

3.5 International Spillover Results

This section studies whether the Fed communication shock identified from Federal Reserve Board of Governors' speeches is transmitted internationally. Following the logic of Fernández-Villaverde et al. (2024), we estimate unbalanced panel local projections for a broad set of real, financial, uncertainty, and expectation variables. The unbalanced structure of the panel allows us to exploit the widest possible country coverage for each outcome variable, but it also implies that the estimated responses should be interpreted as average effects conditional on data availability.

Appendix E. lists the countries included in the international spillover analysis and reports descriptive statistics.

3.5.1 Impact Across all Countries

Figures 3.10 and 3.11 report the responses of real, financial, uncertainty, and expectation variables across the full international panel. The results suggest that the Fed communication shock is transmitted beyond the United States. Uncertainty-related variables, including WUI, EPU, stock price volatility, and GDP forecast disagreement, generally decline after the shock, although the magnitude and persistence of the responses differ across indicators.

Real and financial variables display a broadly positive short-run response. GDP per capita, industrial production, fixed investment, and stock prices tend to increase after the shock, both in the aggregate international panel and in the domestic US responses shown by the green line. This pattern is consistent with the interpretation that positive Fed communication can operate as a global confidence-enhancing signal, improving expectations and reducing uncertainty abroad.

The World Sentiment Index also increases in the short run, with a stronger response in the United States than in the international panel. This suggests that the informational content of Fed communication is most directly absorbed domestically, but also propagates to foreign economies through financial and expectation channels.

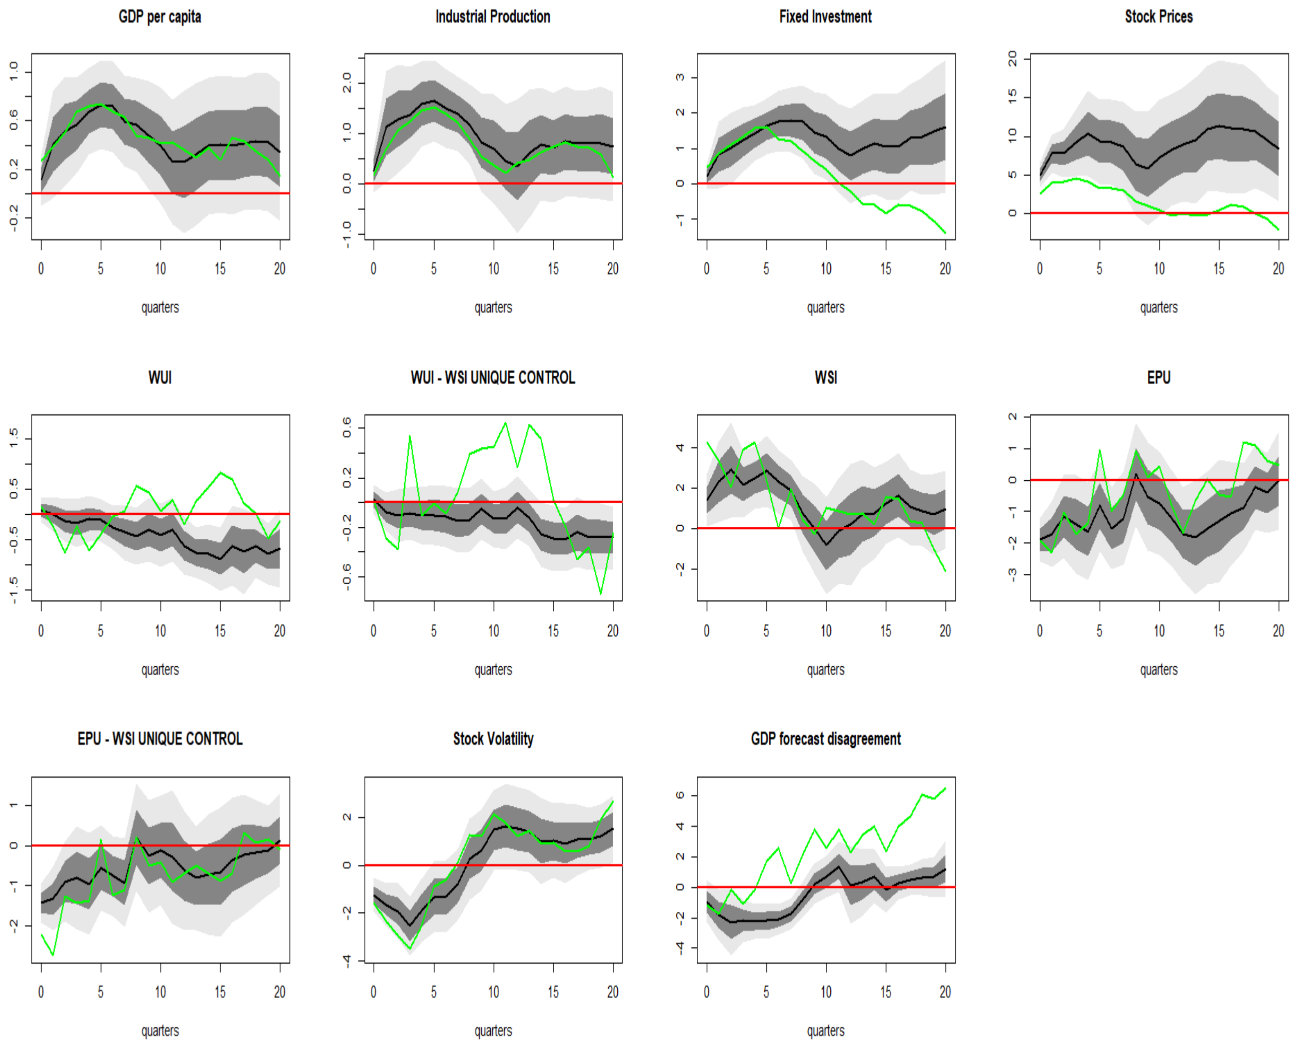

Figure 3.10: International spillover effects of a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the average response across all countries in the sample.

Notes: IRFs are estimated using the unbalanced panel local projection model in Eq. (3.12). Variables include GDP per capita, industrial production, investment, stock prices, WUI, WSI, EPU, stock volatility, and SPF disagreement. The green line represents Domestic effects. Shaded areas correspond to 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$(cycles from 8+ years), with fourteen factors. The variable used for identification is the Sentiment extracted using finBERT with neutral sentences.

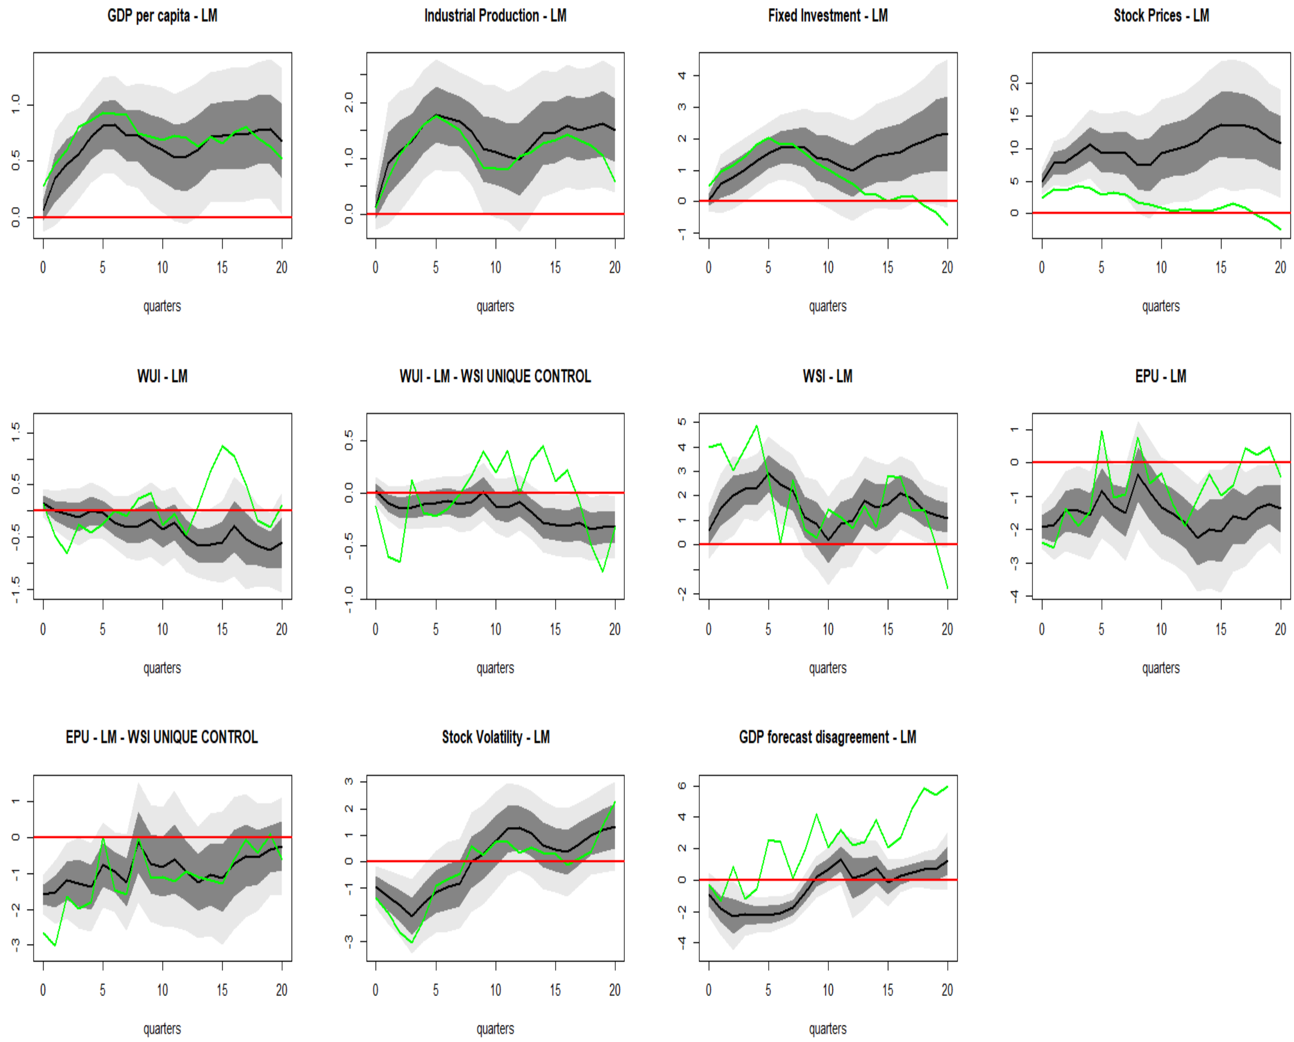

Figure 3.11: International spillover effects of a one-standard-deviation Fed communication shock using Loughran-McDonald sentiment. Each panel reports the average response across all countries in the sample.

Notes: IRFs are estimated using the unbalanced panel local projection model in Eq. (3.12). Variables include GDP per capita, industrial production, investment, stock prices, WUI, WSI, EPU, stock volatility, and SPF disagreement. The green line represents Domestic effects. Shaded areas correspond to 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$(cycles from 8+ years), with fourteen factors. The variable used for identification is the Sentiment extracted using the Loughran-McDonald vocabulary.

It can be observed in figures 3.10 and 3.11, that the effect of the Fed communication shock on the main real and financial variables (GDP per capita, industrial production, fixed investment and stock prices) is always positive in the short term, suggesting a boost to production, investment and financial markets indexes, both at the aggregate level and at the US level (green line). These results indicate that the spillover effect of CBC exists and is statistically significant for other global economies as well.

Regarding the World Sentiment Index (WSI), we observe that the Fed communication shock is positive at the aggregate level and also for the US economy in the short term. In particular, in the United States, the initial increase in the WSI is significantly more intense and larger. Once again, all these results suggest that there is a spillover effect from the communications of the

Federal Reserve Board of Governors.

3.5.2 Impact across Advanced Economies (AE) and Emerging Markets (EM)

To better understand global dynamics, we have decided to divide countries into two groups: Advanced Economies (AE) and Emerging Markets (EM). This allows us to see the difference in magnitude of the impact of the Fed communication shock.

Figures 3.12 e 3.13 show the impulse response ($0-20$ quarters) of real, financial and expectation variables to a Fed communication shock from the Fed Board of Governors using finBERT and Lexicon Approach Loughran-McDonald vocabulary for Advanced Economies (AE) and Emerging Markets (EM). As before, the green line represents the effect of the Fed communication shock on the US economy.

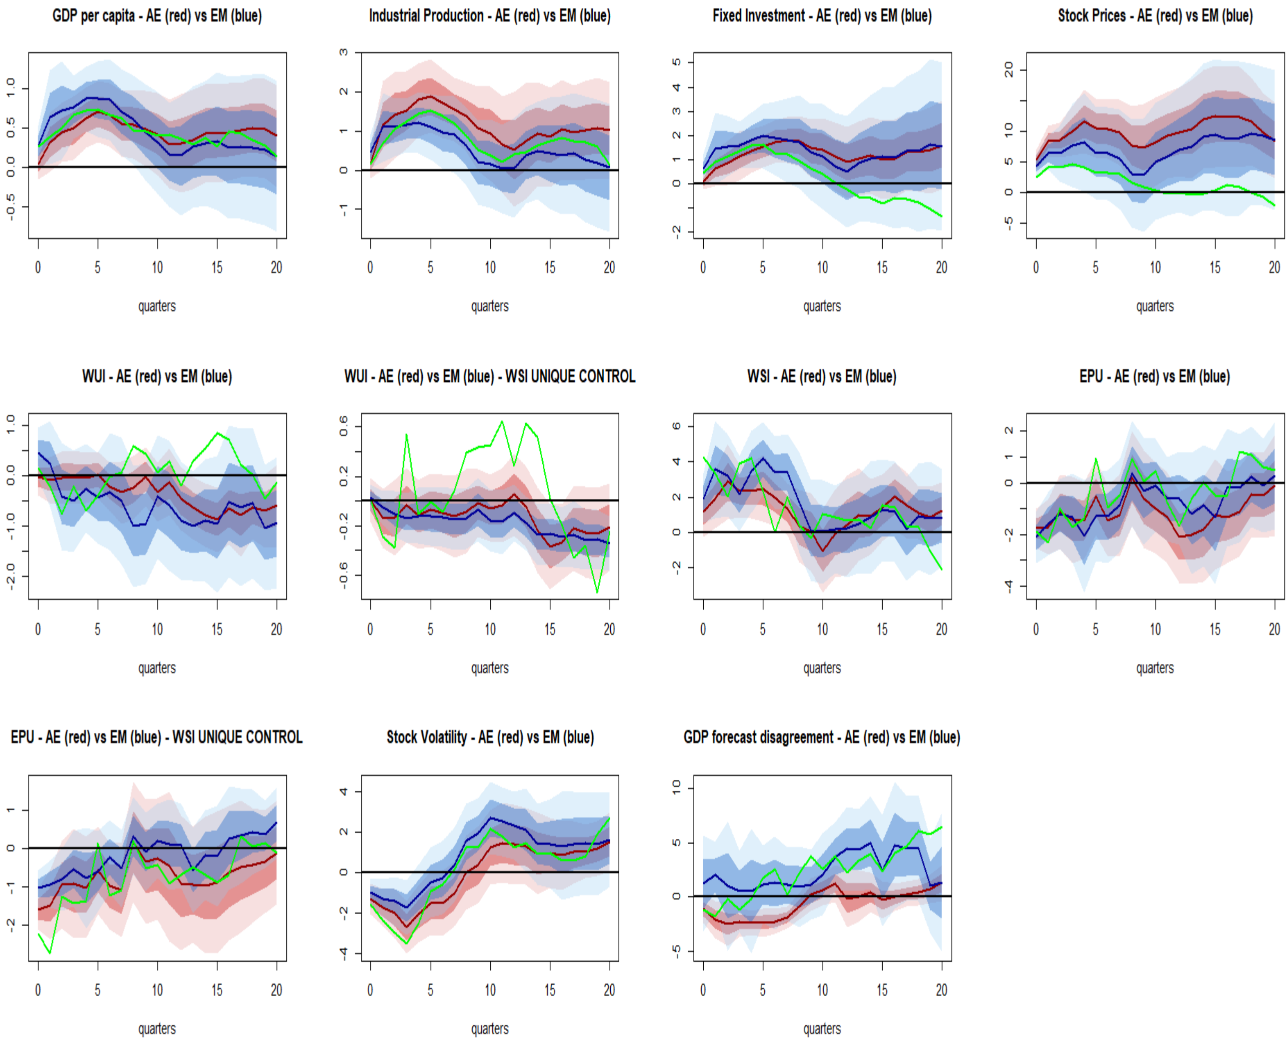

Figure 3.12: International spillover effects of a communication shock: Advanced Economies (AE) vs Emerging Markets (EM). Each panel reports the impulse response of a macroeconomic or uncertainty variable over 20 quarters from the finBERT Fed communication shock with neutral sentences in the denominator.

Notes: IRFs are estimated using the unbalanced panel local projection model in Eq. (3.12). Variables include GDP per capita, industrial production, investment, stock prices, WUI, WSI, EPU, stock volatility, and SPF disagreement. The green line represents Domestic effects. The red area and the red line represent AE, and the blue area and the blue line represent EM. Shaded areas correspond to 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$(cycles from 8+ years), with fourteen factors. The variable used for identification is the Sentiment extracted using finBERT with neutral sentences.

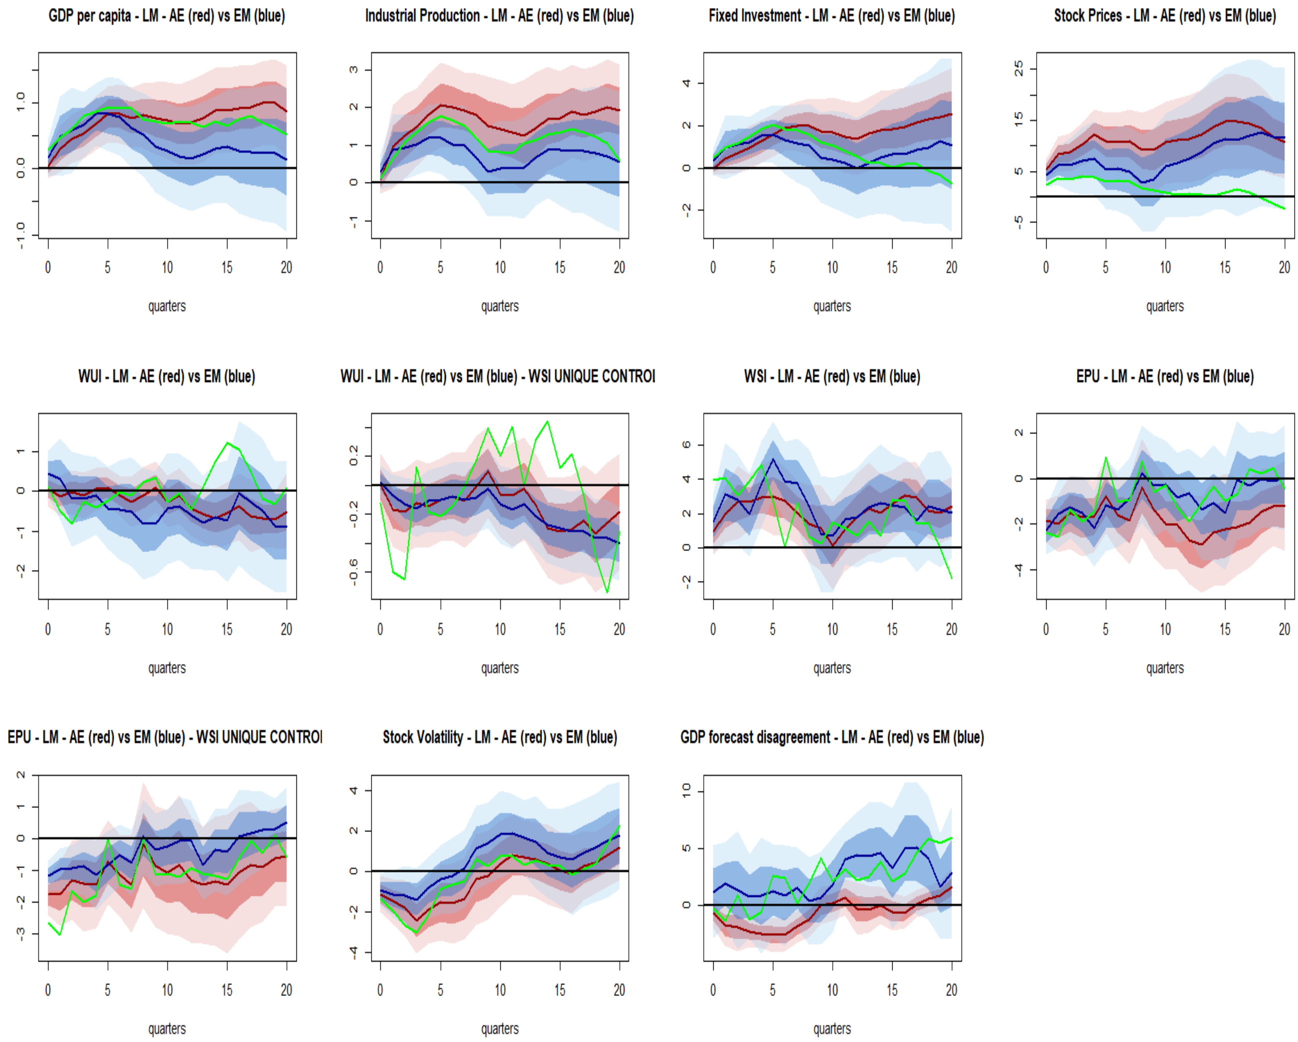

Figure 3.13: International spillover effects of a communication shock: Advanced Economies (AE) vs Emerging Markets (EM). Each panel reports the impulse response of a macroeconomic or uncertainty variable over 20 quarters from the Loughran-McDonald Fed communication shock.

Notes: IRFs are estimated using the unbalanced panel local projection model in Eq. (3.12). Variables include GDP per capita, industrial production, investment, stock prices, WUI, WSI, EPU, stock volatility, and SPF disagreement. The green line represents Domestic effects. The red area and the red line represent AE, and the blue area and the blue line represent EM. Shaded areas correspond to 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$(cycles from 8+ years), with fourteen factors. The variable used for identification is the Sentiment extracted using the Loughran-McDonald vocabulary.

Figures 3.12 and 3.13 distinguish between Advanced Economies (AE) and Emerging Markets (EM). The responses reveal substantial heterogeneity across country groups. Real and financial variables generally respond more strongly and more persistently in advanced economies, while the responses in emerging markets are weaker and often less precisely estimated.

For the World Sentiment Index, the Fed communication shock generates a positive short-run response in both groups, although the domestic US response remains larger. Uncertainty indicators such as WUI, EPU, and stock price volatility tend to decline more clearly in advanced economies than in emerging markets. This suggests that Fed communication may be transmitted more effectively to economies with deeper financial markets, stronger institutional linkages, and information environments more closely integrated with US monetary policy news.

GDP forecast disagreement displays an important difference across groups. In advanced economies and in the United States, disagreement declines initially and then gradually reverts. In emerging markets, by contrast, disagreement tends to increase over the horizon. This pattern suggests that Fed communication may reduce uncertainty in economies more closely integrated with US financial conditions, while generating more ambiguous signals in emerging markets, where the same shock may be interpreted through exchange-rate, capital-flow, or external-financing risks.

These findings suggest that the international transmission of Fed communication depends on the degree to which each economy is integrated with US financial markets and exposed to Federal Reserve policy signals. The stronger responses in advanced economies are consistent with deeper financial linkages, greater sensitivity to US monetary conditions, and faster incorporation of Fed communication into domestic expectations and asset prices. These results differ from Chen et al. (2016) and Dedola et al. (2017), who find stronger effects of US monetary policy shocks in emerging markets. The difference is not necessarily contradictory: those papers focus on conventional monetary policy, quantitative easing, or tightening-related volatility, whereas the shock identified here is a forward guidance-related Fed communication shock.

3.6 Conclusion

Central bank communication plays an important role in shaping expectations, financial conditions, and the interpretation of macroeconomic risks. The evidence in this chapter suggests that Fed communication contains economically relevant information that is associated with sizeable macro-financial responses.

The results show that a positive Fed communication shock is followed by a decline in disagreement among professional forecasters and by a reduction in several measures of macro-financial uncertainty. At the same time, mean growth expectations improve, while real activity responds positively over the short to medium run. These findings are consistent with the view that central bank communication can clarify the policy outlook, reduce uncertainty, and strengthen confidence among economic agents.

Overall, the results reinforce the view emerging from recent work by Fadda et al. (2025); Istrefi et al. (2023); Istrefi and Piloiu (2020); Bordo and Istrefi (2023) that central bank communication is an economically meaningful source of information. In our setting, a positive communication shock reduces uncertainty, lowers disagreement, improves expectations, and is associated with a positive response of real activity, consistent with the idea that communication shapes how agents interpret risks, policy intentions, and the broader economic outlook.

Beyond its impact on uncertainty, the response of real GDP per capita, TFP, and the unemployment rate suggests that the identified shock is associated with an expansionary macroeconomic response. Moreover, the positive reaction of the spread between bullish and bearish investor sentiment is consistent with a confidence or credit-channel interpretation. This suggests that favourable Fed communication may reduce perceived risk and support more accommodative financial conditions.

The international evidence indicates that Fed communication shocks also transmit beyond the United States. The responses are generally more persistent at the domestic level, but the spillover results suggest that Fed communication is associated with improvements in real and financial conditions and with reductions in uncertainty abroad. The effects are stronger and more stable for AEs than for EMs, consistent with differences in financial integration, policy regimes, institutional structures, and exposure to the global financial cycle.

These findings should be interpreted with appropriate caution. The shock identified in this chapter is not a directly observed communication disturbance, but a text-based and model-dependent object extracted from Federal Reserve Board of Governors' speeches and identified within a structural dynamic factor framework. The results, therefore, support the view that central bank communication is an important component of monetary policy transmission, while also highlighting the need for careful interpretation of communication shocks constructed from textual indicators.

Appendix

A. Details About Speakers

Table A.1: Summary of speeches by Board of Governors of the Federal Reserve members (Loughran–McDonald Vocabulary)

Speaker	N. Speeches	First speech	Last speech	Period (days)	Posit.	Negat.	Difference P-N	Average Sentiment	Ratio Pos.
Alan Greenspan	446	1987-10-05	2005-12-02	6633	17326	32768	-15442	-0.284	0.346
Ben S Bernanke	221	2002-10-15	2014-01-03	4098	9215	18371	-9156	-0.284	0.334
Roger W Ferguson	132	1998-03-04	2006-04-17	2966	5772	8679	-2907	-0.141	0.399
John Patrick LaWare	98	1988-09-08	1995-02-14	2350	3068	6319	-3251	-0.317	0.327
Lael Brainard	96	2014-12-03	2023-01-19	2969	3476	6419	-2943	-0.245	0.351
Laurence H Meyer	94	1996-09-08	2002-12-18	2292	4492	8672	-4180	-0.282	0.341
Donald L Kohn	86	2002-11-22	2010-05-13	2729	3017	6812	-3795	-0.372	0.307
Susan Schmidt Bies	82	2002-02-28	2007-02-26	1824	3168	4614	-1446	-0.151	0.407
Jerome H Powell	76	2013-02-22	2023-12-01	3934	2641	5561	-2920	-0.331	0.322
Daniel K Tarullo	75	2009-03-19	2017-04-04	2938	2835	7598	-4763	-0.406	0.272
Janet L Yellen	63	1996-03-13	2017-10-20	7891	2797	6035	-3238	-0.314	0.317
Edward M Gramlich	60	1998-02-27	2005-06-03	2653	1910	3641	-1731	-0.258	0.344
Randal K Quarles	56	2017-11-30	2021-12-02	1463	2262	3581	-1319	-0.203	0.387
Susan M Phillips	53	1992-02-13	1998-03-26	2233	2056	3062	-1006	-0.182	0.402
Michelle W Bowman	52	2019-02-11	2023-12-08	1761	2262	2955	-693	-0.087	0.434
Mark W Olson	51	2002-02-07	2006-06-13	1587	1753	2178	-425	-0.072	0.446
Randall S Kroszner	48	2006-04-03	2008-12-08	980	2021	3414	-1393	-0.227	0.372
Edward W Kelley	47	1987-07-09	2000-03-30	4648	1627	2649	-1022	-0.159	0.380
Elizabeth A Duke	46	2009-02-11	2013-05-09	1548	1826	4478	-2652	-0.383	0.290
Wayne D Angell	42	1986-04-22	1993-10-27	2745	1778	2027	-249	-0.049	0.467
Manuel H Johnson	41	1986-04-26	1990-05-02	1467	1510	2600	-1090	-0.246	0.367
Stanley Fischer	41	2014-07-10	2017-09-28	1176	1268	2871	-1603	-0.362	0.306
H Robert Heller	34	1986-12-29	1989-05-02	855	1348	1976	-628	-0.166	0.406
Richard H Clarida	33	2018-10-25	2021-11-30	1132	796	1333	-537	-0.212	0.374
Lawrence Lindsey	32	1992-03-12	1995-12-14	1372	995	2057	-1062	-0.315	0.326
Frederic S Mishkin	29	2006-10-12	2008-07-28	655	1315	2927	-1612	-0.377	0.310
Christopher J Waller	28	2021-03-29	2023-11-28	974	725	1500	-775	-0.316	0.326
Paul A Volcker	23	1986-01-29	1987-07-21	538	1461	2704	-1243	-0.226	0.351
Alice M Rivlin	19	1996-12-19	1999-06-01	894	774	1200	-426	-0.138	0.392
Kevin M Warsh	18	2006-07-18	2010-11-08	1574	727	1496	-769	-0.324	0.327
Sarah Bloom Raskin	18	2010-11-12	2013-07-17	978	609	1849	-1240	-0.498	0.248
Jeremy C Stein	15	2012-10-11	2014-05-06	572	297	830	-533	-0.451	0.264
Martha Romayne Seger	15	1986-03-06	1989-12-07	1372	364	1003	-639	-0.431	0.266
Michael S Barr	15	2022-09-07	2023-12-01	450	539	1093	-554	-0.275	0.330
Lisa D Cook	14	2022-10-06	2023-11-16	406	320	523	-203	-0.111	0.380
David W Mullins	13	1991-09-04	1993-07-28	693	389	893	-504	-0.383	0.303
Philip N Jefferson	11	2022-10-04	2023-11-24	416	234	569	-335	-0.433	0.291
Alan S Blinder	9	1995-03-09	1996-01-10	307	232	670	-438	-0.473	0.257
Lawrence B Lindsey	7	1996-02-19	1996-10-11	235	206	480	-274	-0.344	0.300
Preston Martin	5	1986-02-05	1986-04-09	63	181	469	-288	-0.437	0.278
Emmett John Rice	4	1986-01-28	1986-10-30	275	71	207	-136	-0.515	0.255
Brian F Madigan	1	2009-08-21	2009-08-21	0	40	264	-224	-0.737	0.132

Notes: List of all members of the Federal Reserve Board of Governors, including the chairpersons. The list — compiled using the Loughran-McDonald vocabulary — includes the total number of speeches per board member, the dates of their first and last speeches, the duration of their terms of office (in days), the number of positive and negative words used by each member, the net difference of positive and negative words used by each member, and the average sentiment.

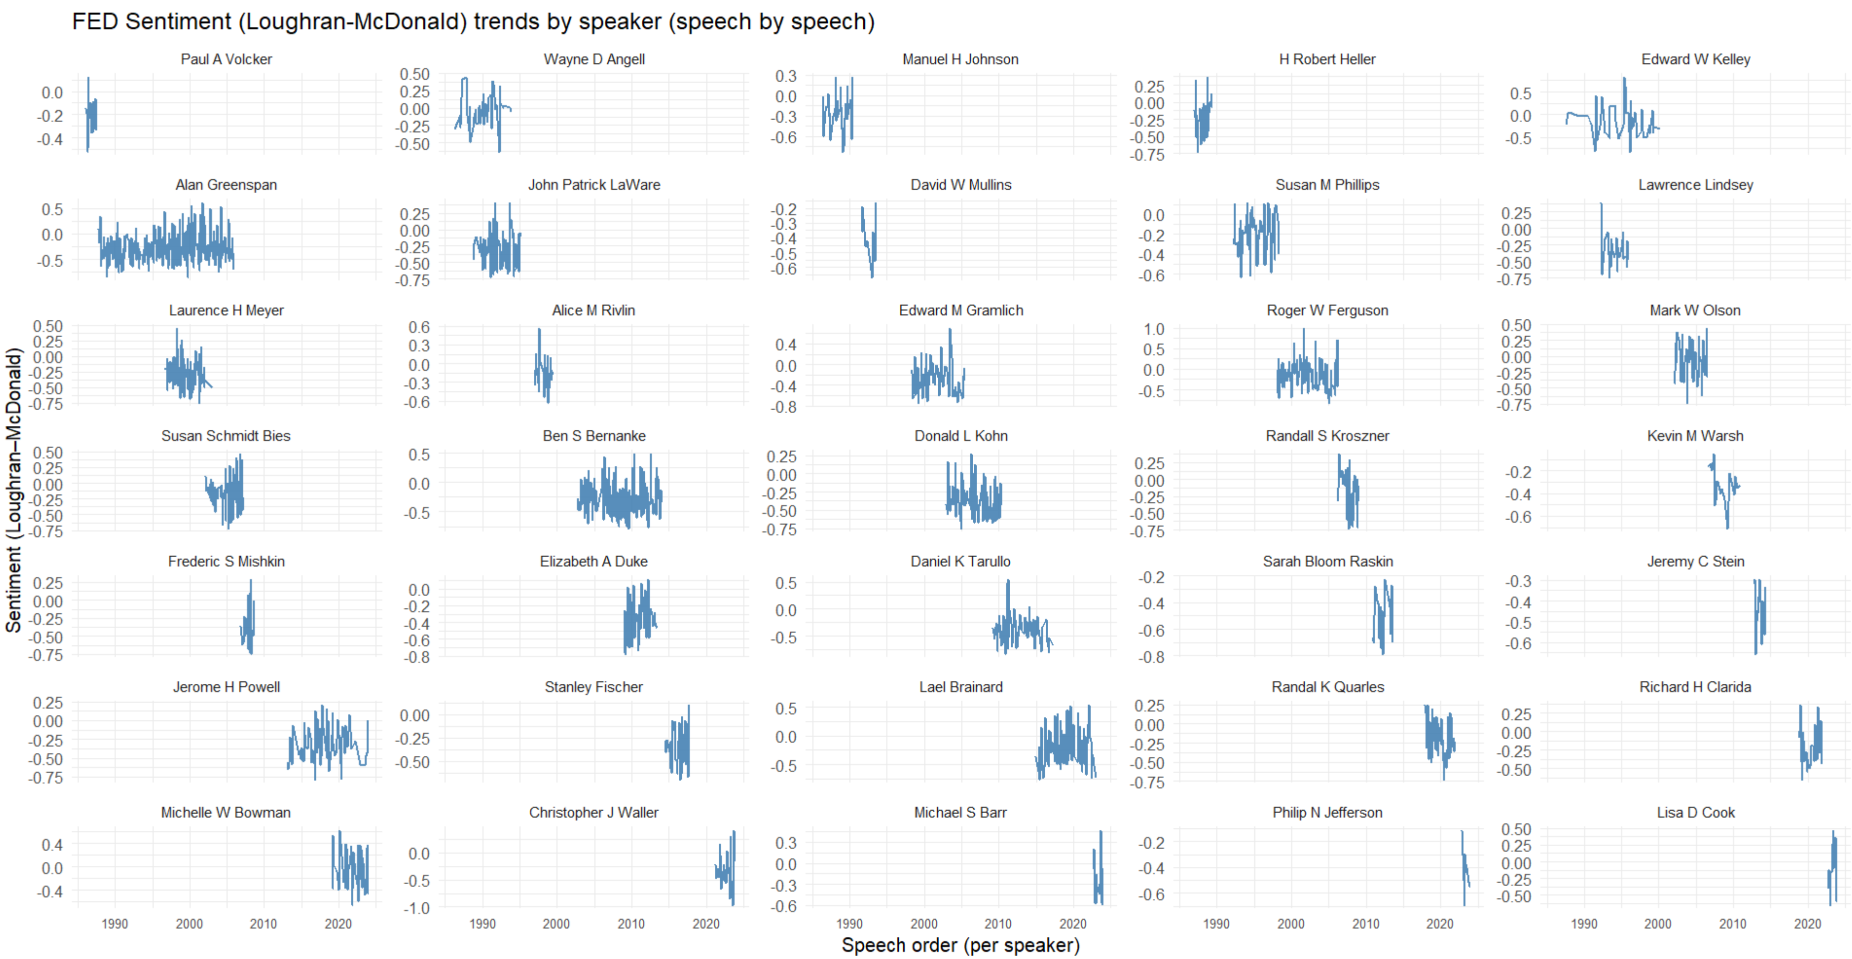

Figure A.1: Fed Sentiment using Lexicon Approach - Loughran-McDonald vocabulary, for each speaker.

Notes: The figure shows the daily sentiment of each member of the Federal Reserve Board of Governors. The sentiment is calculated using the Loughran-McDonald vocabulary. The period of the daily sentiment is the same as the baseline analysis, 1989–2021.

B. Impulse Response Functions for Identified Shocks on Alternative Frequency Bands - finBERT

IRFs on Disagreement

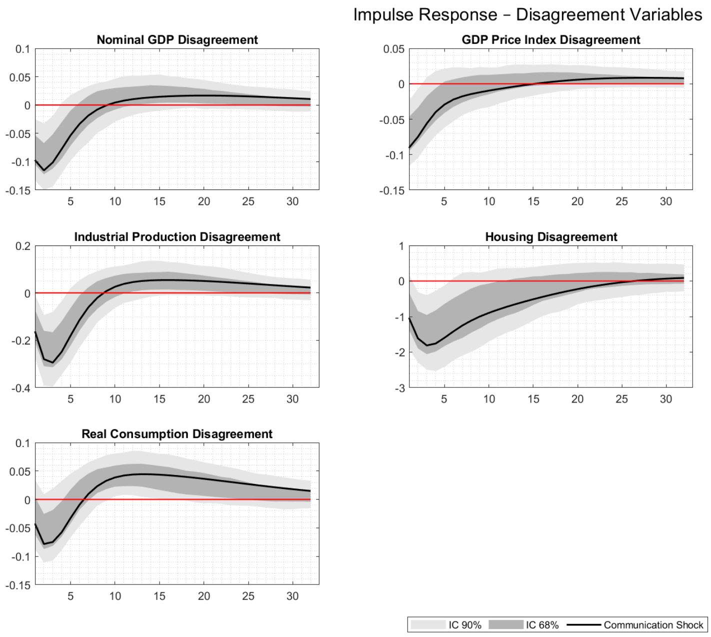

Figure B.1: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/20]$ (cycles longer than five years), using fourteen factors. The identification variable is the finBERT with neutral sentences.

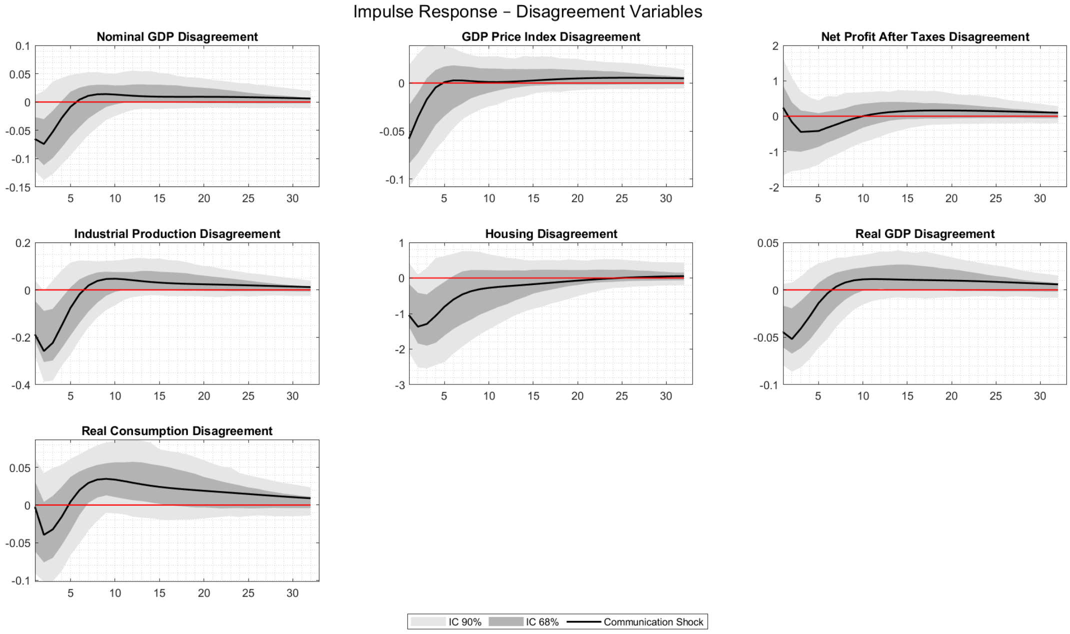

Figure B.2: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[2\pi/40, 2\pi/20]$ (cycles between five and ten years), using fourteen factors. The identification variable is the finBERT with neutral sentences.

IRFs on Expectation

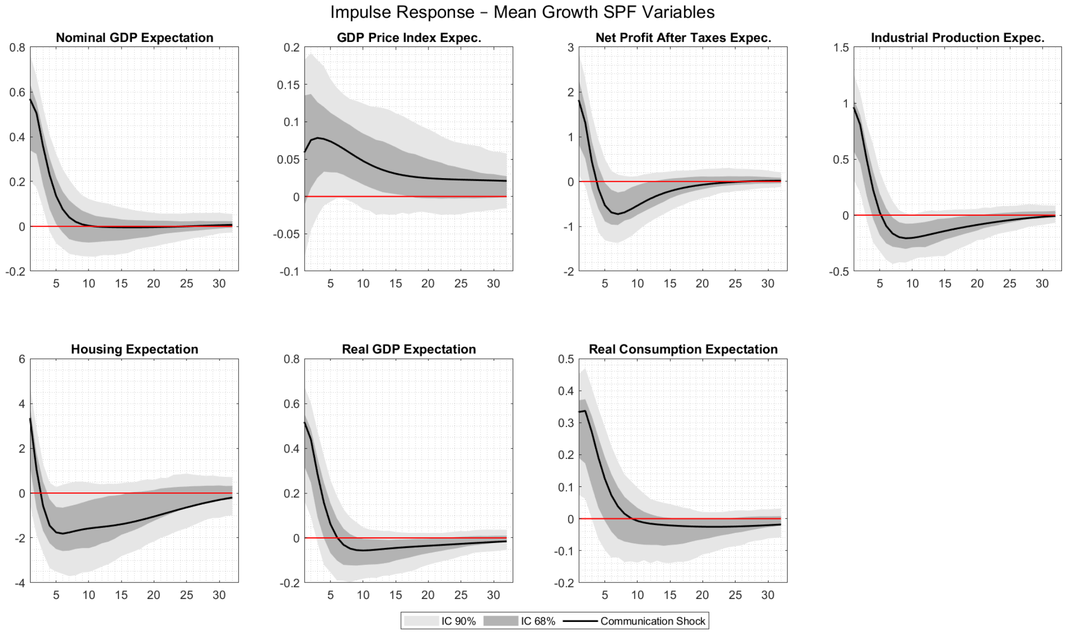

Figure B.3: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/20]$ (cycles longer than five years), using fourteen factors. The identification variable is the finBERT with neutral sentences.

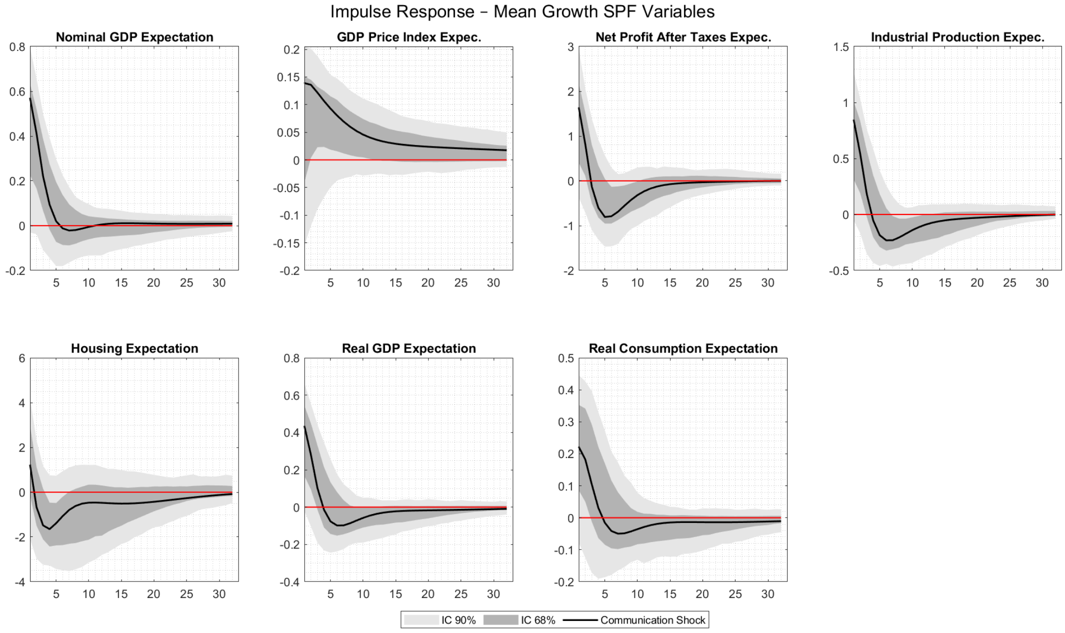

Figure B.4: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[2\pi/40, 2\pi/20]$ (cycles between five and ten years), using fourteen factors. The identification variable is the finBERT with neutral sentences.

IRFs on Macroeconomic Variables

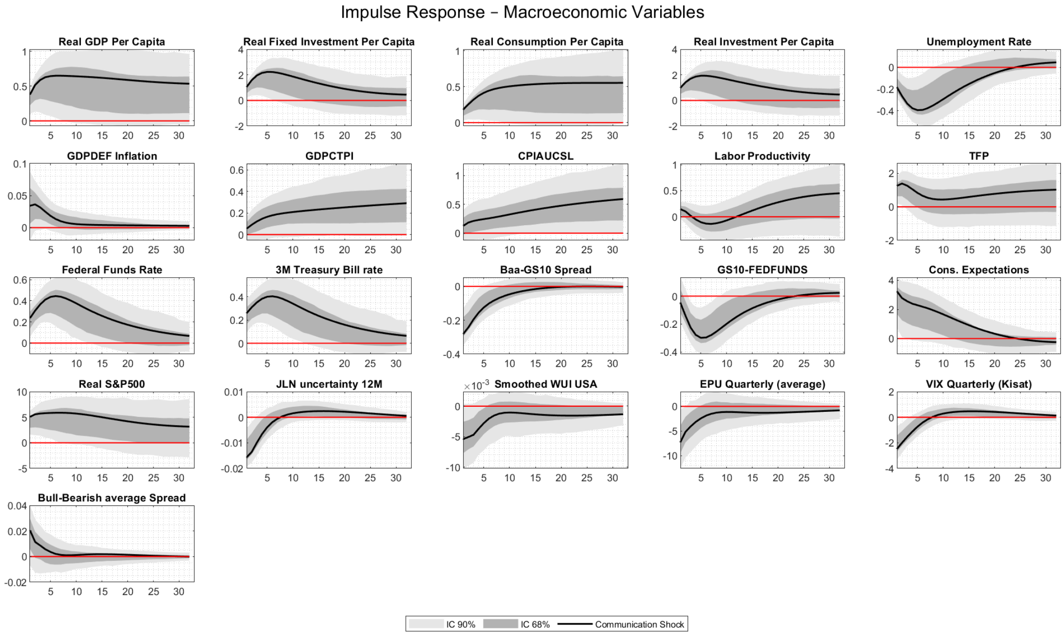

Figure B.5: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/20]$ (cycles longer than five years), using fourteen factors. The identification variable is the finBERT with neutral sentences.

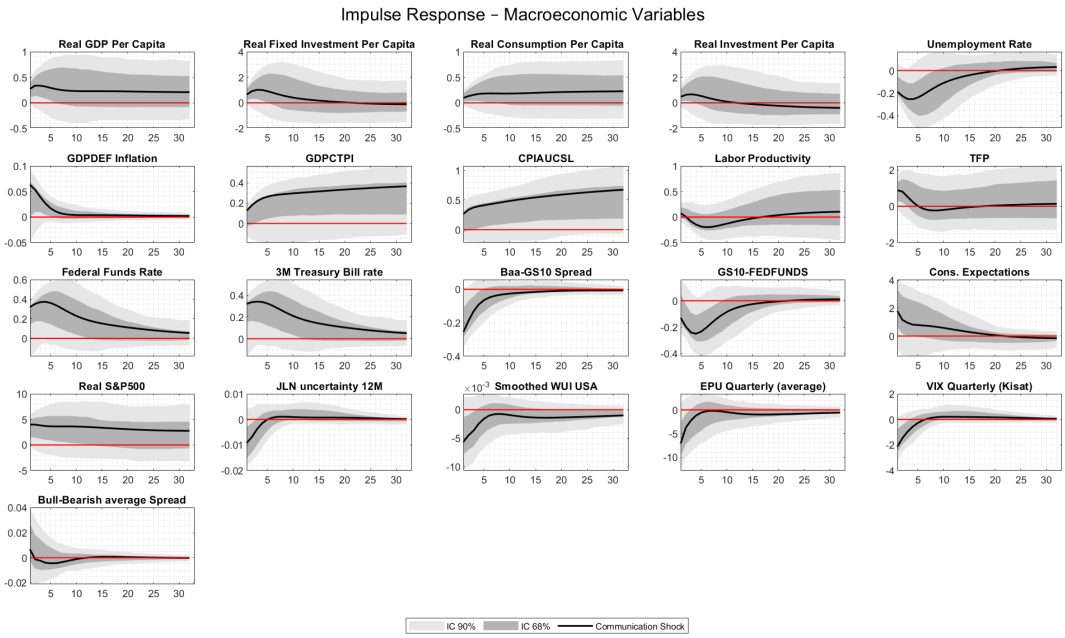

Figure B.6: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[2\pi/40, 2\pi/20]$ (cycles between five and ten years), using fourteen factors. The identification variable is the finBERT with neutral sentences.

C. Impulse Response Functions for Identified shocks on Alternative Frequency Bands - Loughran-McDonald

IRFs on Disagreement

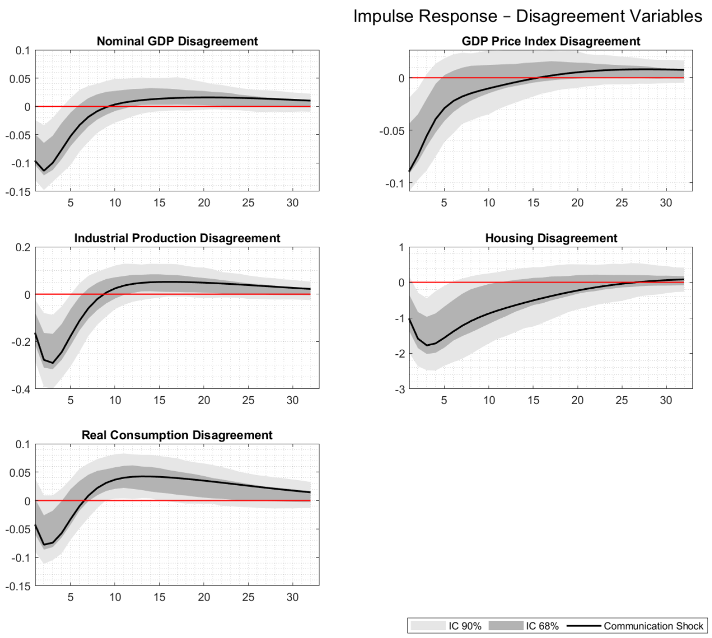

Figure C.1: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/20]$ (cycles longer than five years), using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

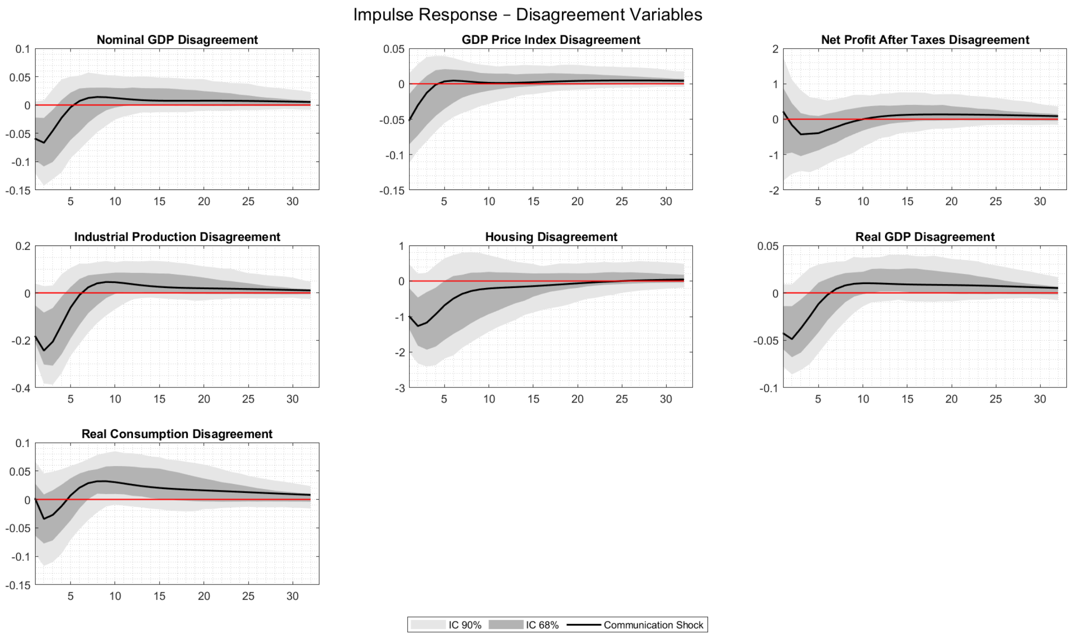

Figure C.2: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[2\pi/40, 2\pi/20]$ (cycles between five and ten years), using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

IRFs on Expectations

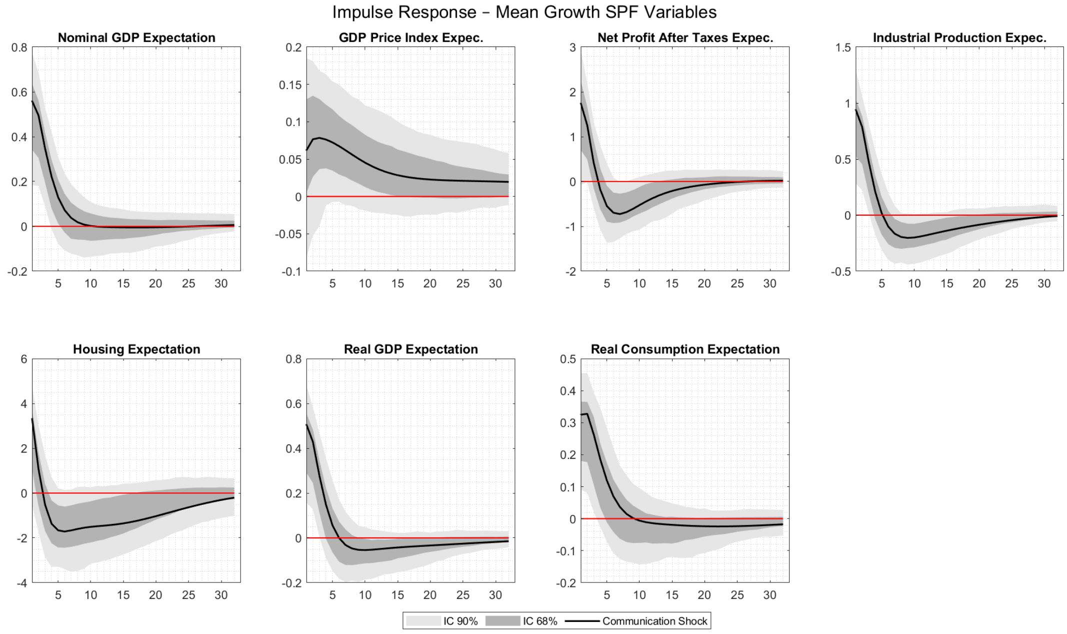

Figure C.3: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/20]$ (cycles longer than five years), using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

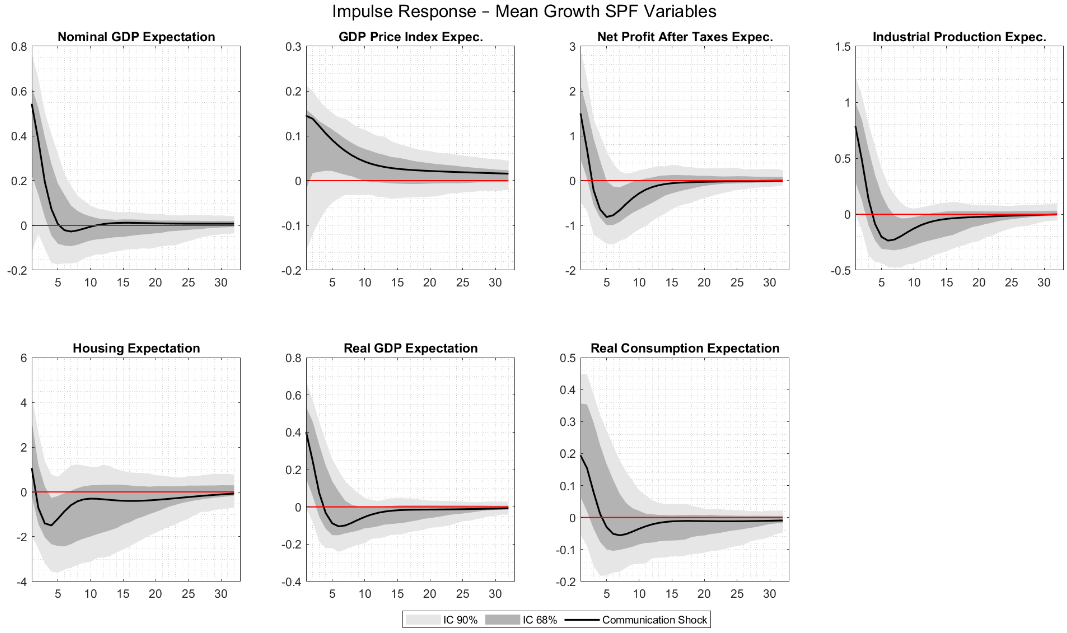

Figure C.4: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[2\pi/40, 2\pi/20]$ (cycles between five and ten years), using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

IRFs on Macroeconomic Variables

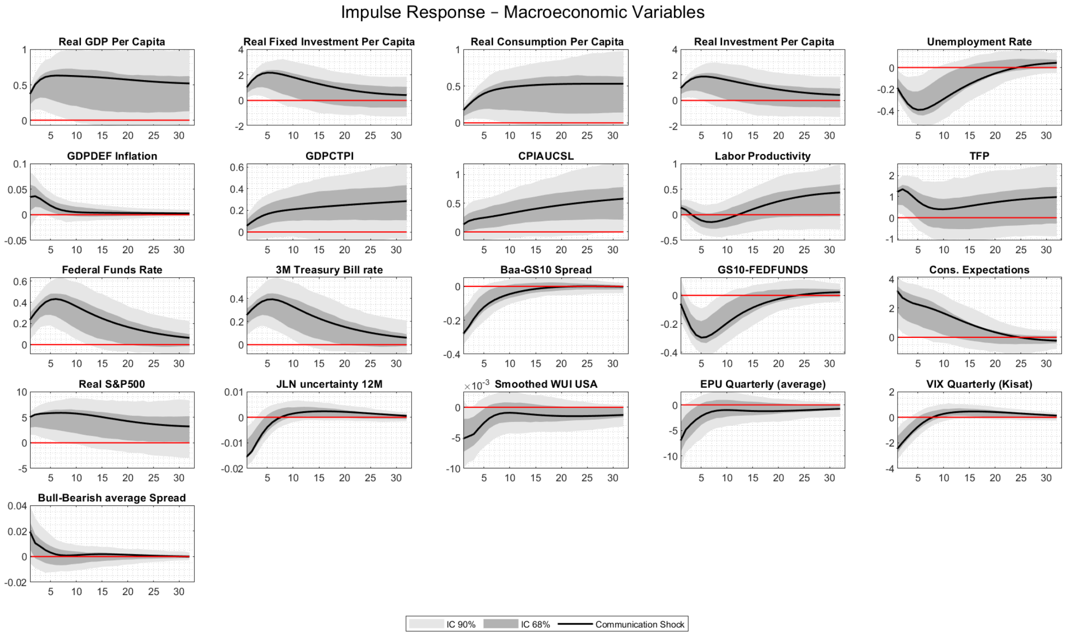

Figure C.5: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/20]$ (cycles longer than five years), using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

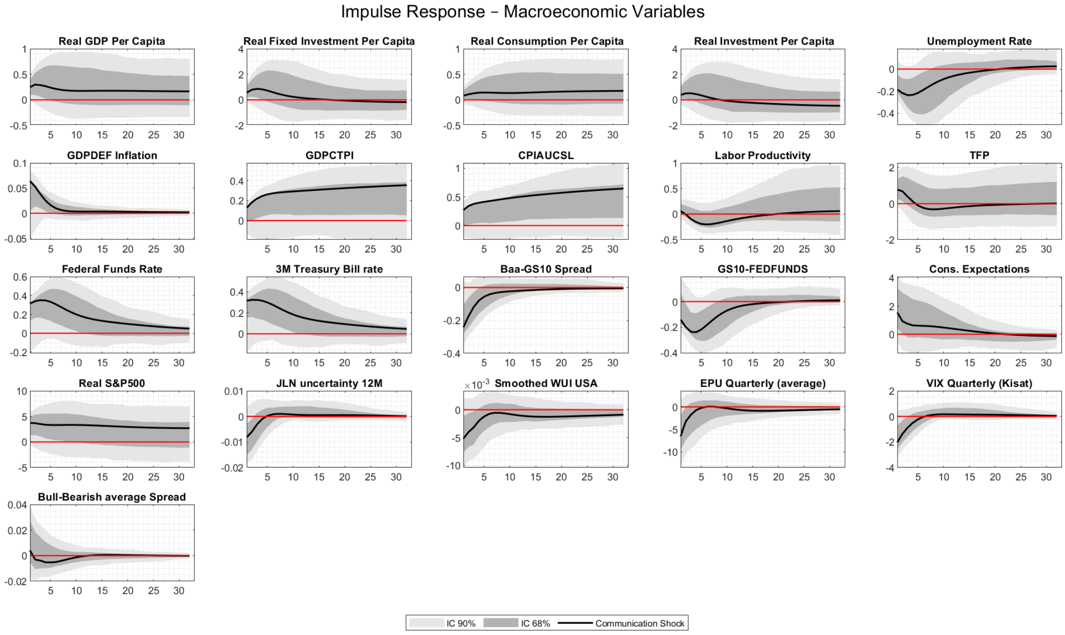

Figure C.6: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[2\pi/40, 2\pi/20]$ (cycles between five and ten years), using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

D. Impulse Response Functions for Identified Shocks on the Extended Domestic Sample Data

Compared to the previous analysis, here we have two more years of observations, but we lose the information relating to the spread difference between Bullish and Bearish sentiment investors — in percentage points. The IRFs in this section show the same effect as before for all the variables.

IRFs on Disagreement

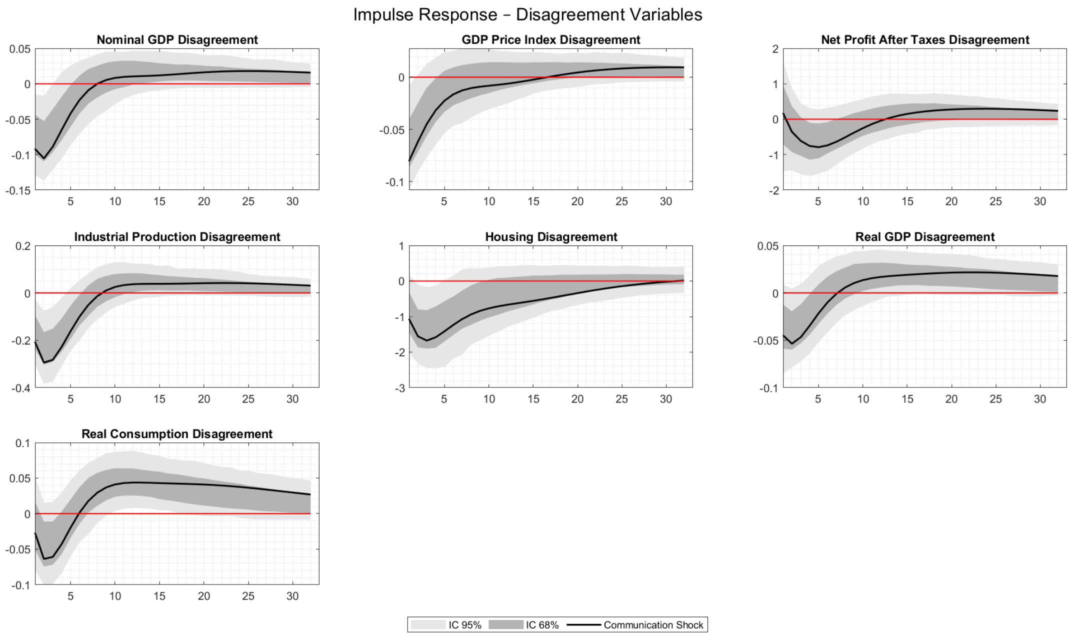

Figure D.1: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the finBERT with neutral sentences.

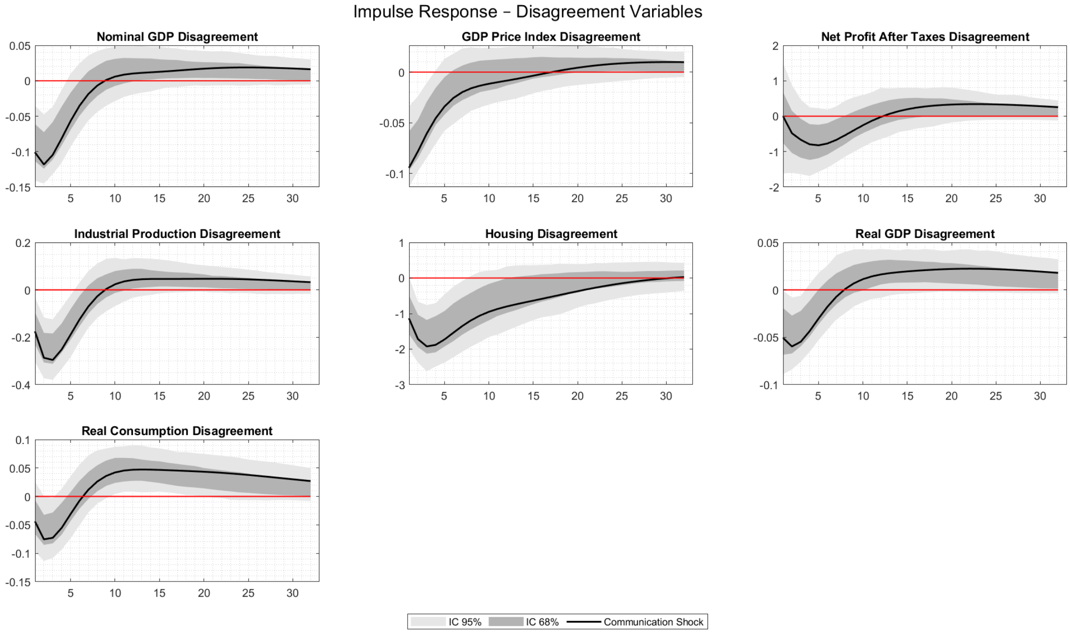

Figure D.2: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

IRFs on Expectation

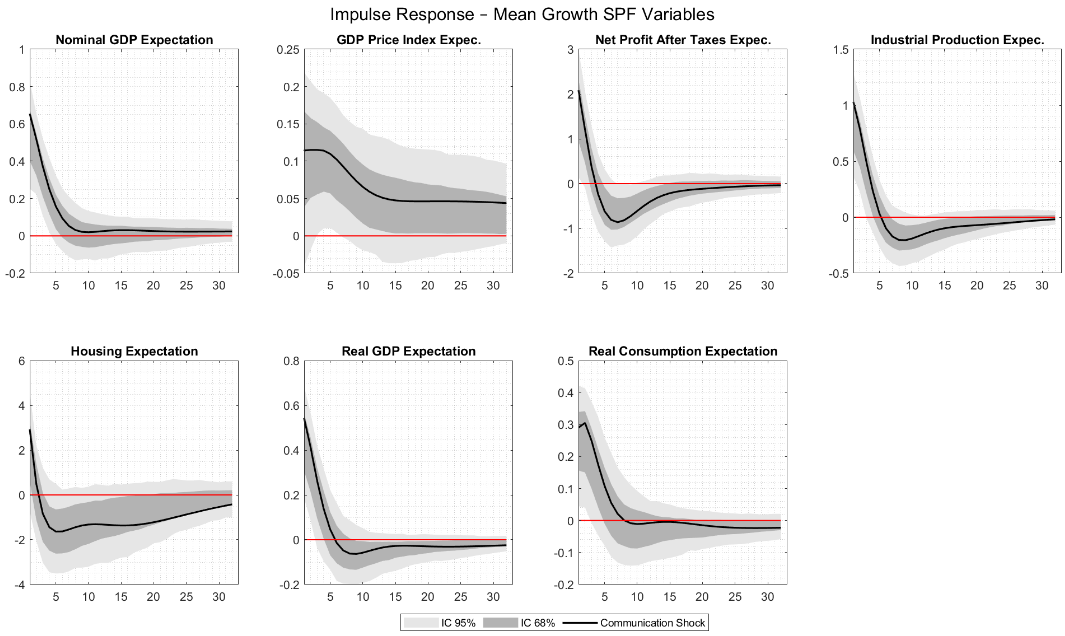

Figure D.3: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the finBERT with neutral sentences.

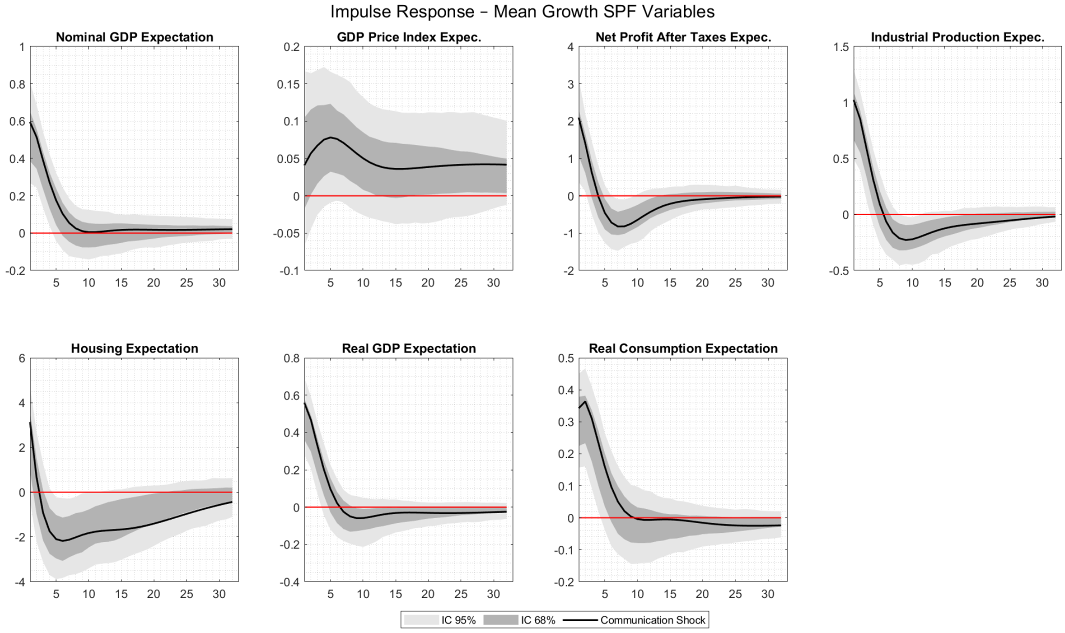

Figure D.4: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

IRFs on Macroeconomic Variables

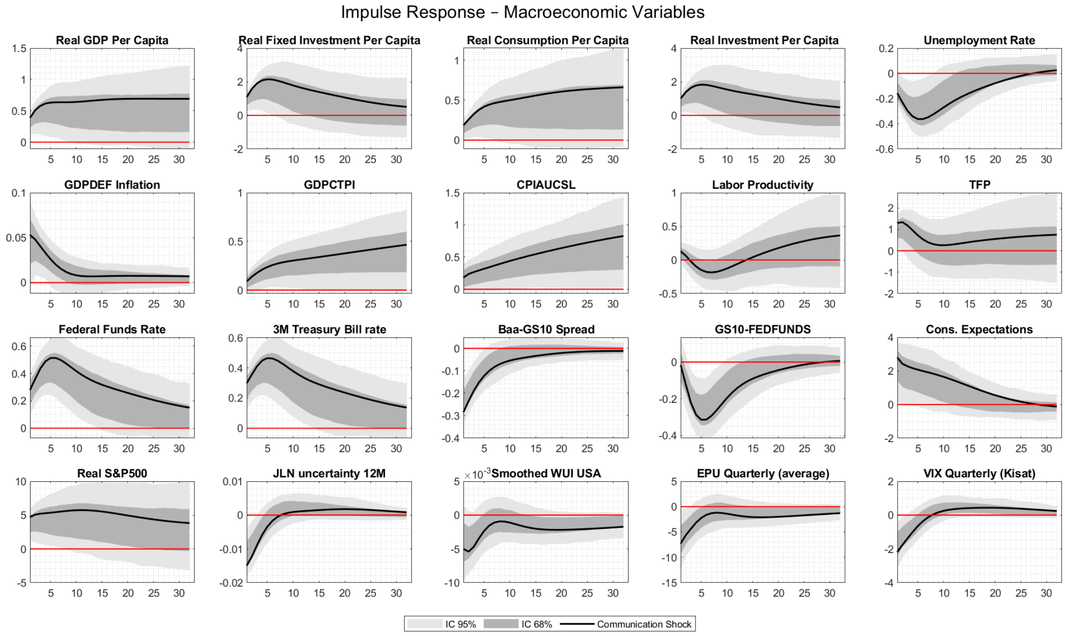

Figure D.5: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. The figure reports responses for real activity, inflation, labour productivity, financial conditions, and uncertainty indicators over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the finBERT with neutral sentences.

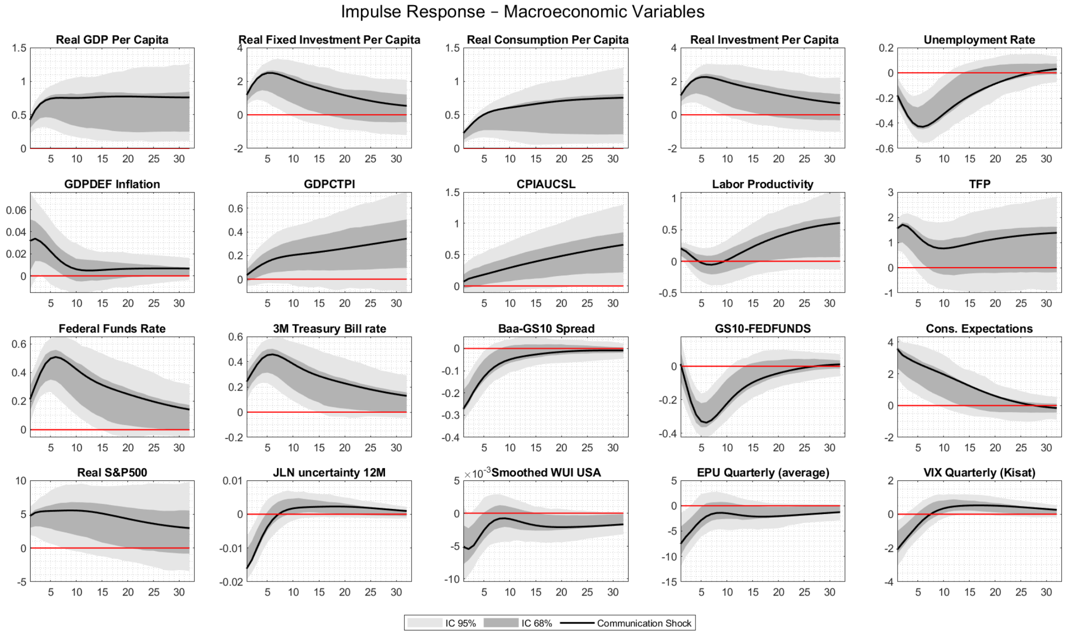

Figure D.6: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the Loughran–McDonald dictionary.

E. Impulse Response Functions for Identified Shock with the Wu–Xia Shadow Rate

This appendix reports an additional robustness exercise in which the Fed communication shock is identified by replacing the FFR with the Wu–Xia shadow rate (Wu and Xia, 2016). The purpose of this exercise is to assess whether the domestic impulse responses are sensitive to the use of the FFR as the interest-rate variable in the identification step, especially given that the policy rate was constrained during the zero lower bound period. The Wu–Xia shadow rate provides a broader measure of the monetary policy stance when conventional short-term interest rates are close to their lower bound, and is therefore particularly useful for evaluating whether the baseline findings are driven by the limited variation in the FFR during those episodes.

The identification follows the baseline analysis setup. The shock is identified on the communication band $[0, 2\pi/32]$ using fourteen factors, but the low-frequency co-movement is targeted between the net positive sentiment indicator and the Wu–Xia shadow rate rather than the FFR. Since the Wu–Xia shadow rate starts in the first quarter of 1990, the estimation sample runs from the first quarter of 1990 to the second quarter of 2021, ending in the same quarter as the baseline domestic analysis but starting two years later. The exercise is implemented for both the finBERT sentiment measure with neutral sentences in the denominator and the Loughran–McDonald dictionary-based sentiment measure. The figures below report the resulting impulse responses for SPF disagreement, SPF mean growth expectations, and US macroeconomic and financial variables.

IRFs on Disagreement

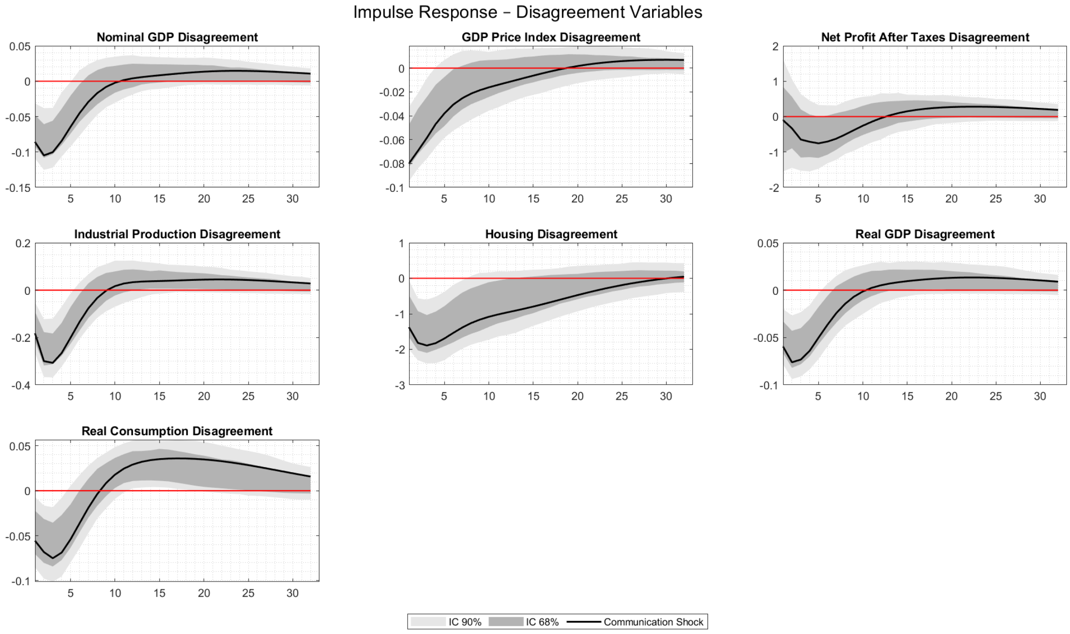

Figure E.1: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the finBERT with neutral sentences. The interest rate used for the identification is the Wu–Xia Shadow Rate (Wu and Xia, 2016).

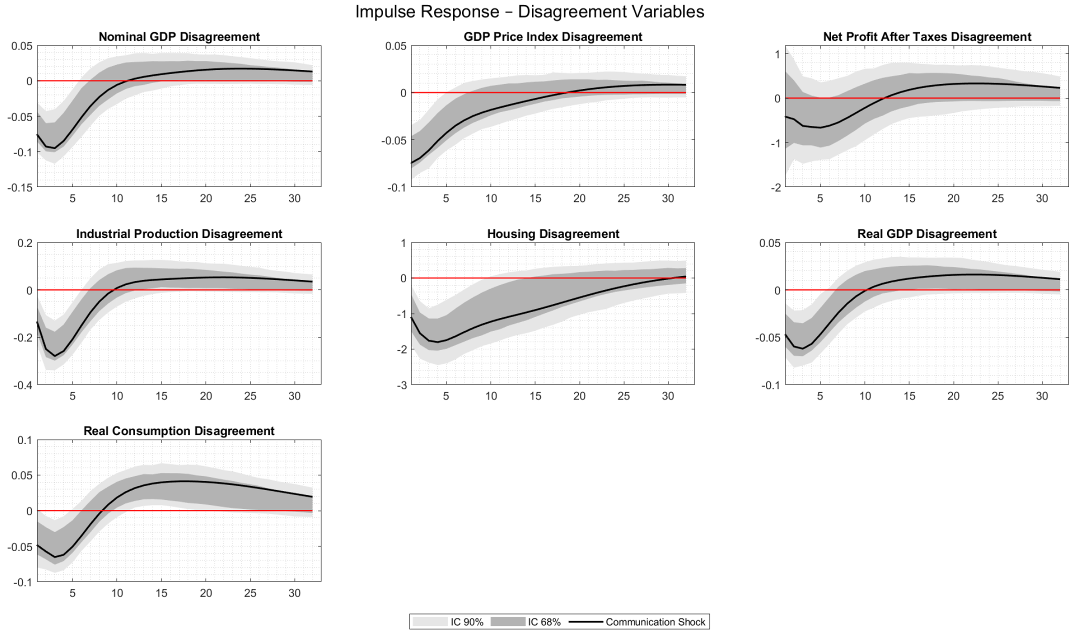

Figure E.2: Impulse responses of Survey of Professional Forecasters (SPF) disagreement variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the Loughran–McDonald dictionary. The interest rate used for the identification is the Wu–Xia Shadow Rate (Wu and Xia, 2016).

IRFs on Expectation

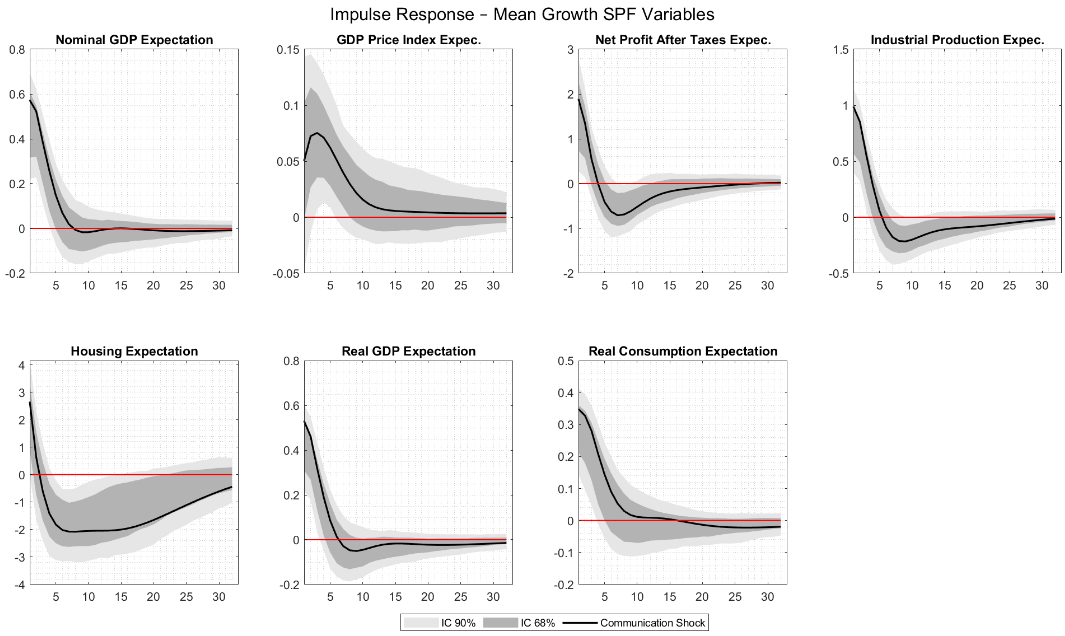

Figure E.3: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. Each panel reports the response of a different disagreement measure over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the finBERT with neutral sentences. The interest rate used for the identification is the Wu–Xia Shadow Rate (Wu and Xia, 2016).

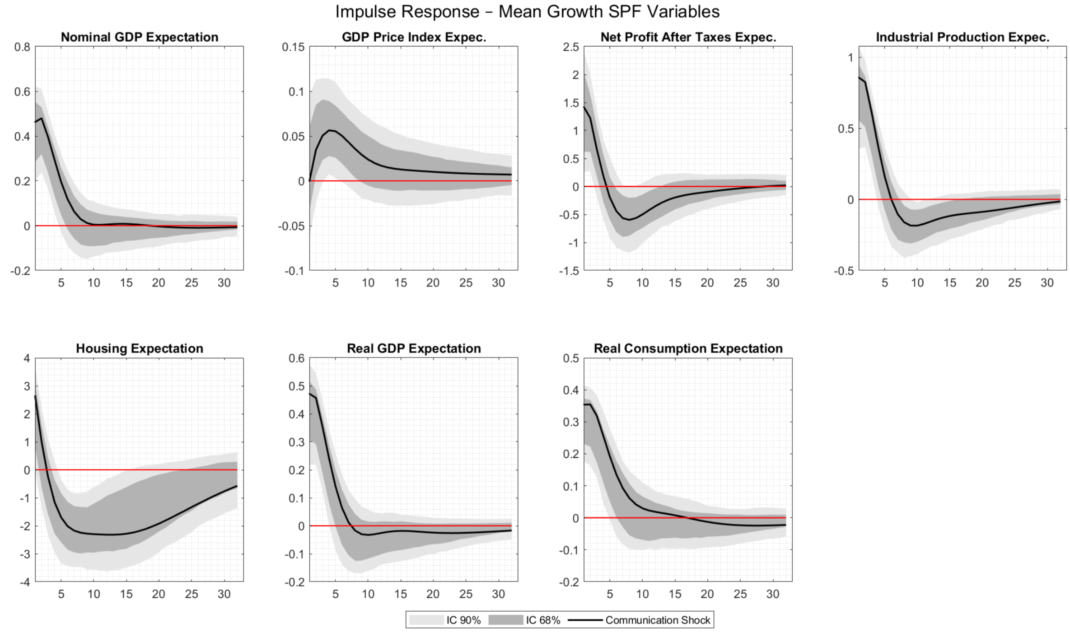

Figure E.4: Impulse responses of Survey of Professional Forecasters (SPF) mean growth expectation variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the Loughran–McDonald dictionary. The interest rate used for the identification is the Wu–Xia Shadow Rate (Wu and Xia, 2016).

IRFs on Macroeconomic Variables

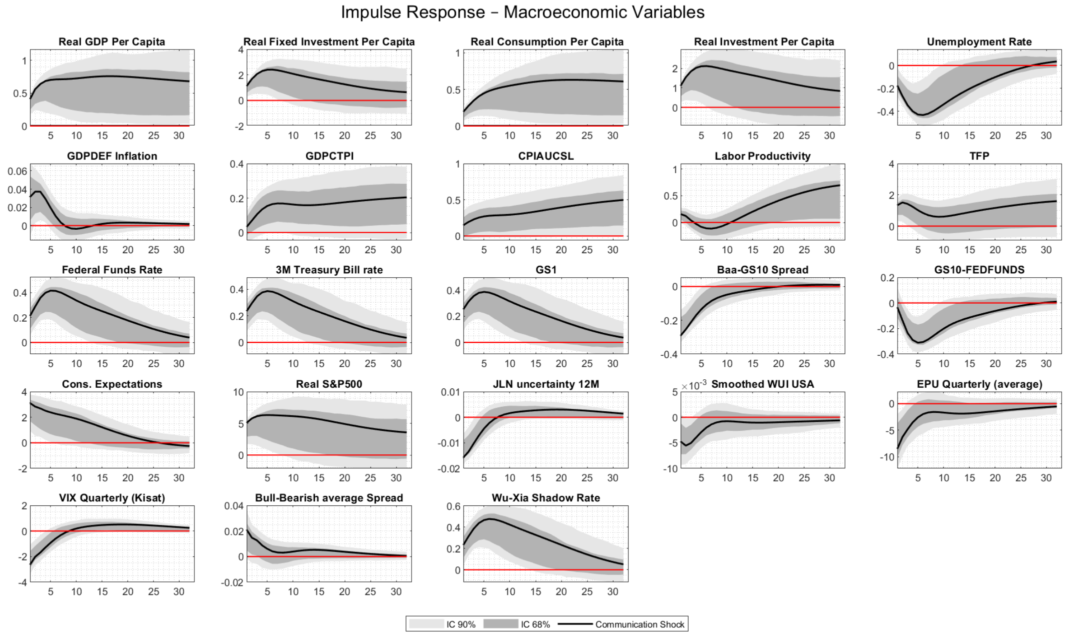

Figure E.5: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Fed communication shock identified using finBERT sentiment with neutral sentences in the denominator. The figure reports responses for real activity, inflation, labour productivity, financial conditions, and uncertainty indicators over a 35-quarter horizon.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the finBERT with neutral sentences. The interest rate used for the identification is the Wu–Xia Shadow Rate (Wu and Xia, 2016).

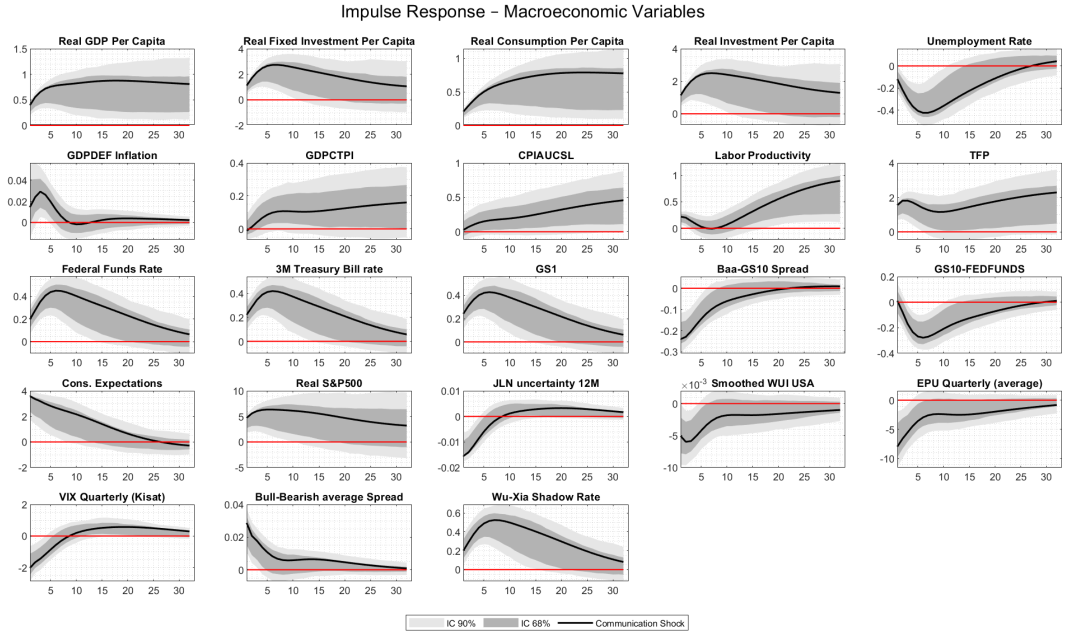

Figure E.6: Impulse responses of US macroeconomic and financial variables to a one-standard-deviation Loughran-McDonald Fed communication shock.

Notes: Impulse responses are estimated from the Structural Dynamic Factor Model (SDFM) described in Section 3.3.2. The communication shock is identified within the factor model using the restrictions discussed in Section 3.3.2. Each panel reports the response of a different variable over a 35-quarter horizon. Shaded areas denote 68% and 95% confidence intervals. The shock is identified on the communication band $[0, 2\pi/32]$, using fourteen factors. The identification variable is the Loughran–McDonald dictionary. The interest rate used for the identification is the Wu–Xia Shadow Rate (Wu and Xia, 2016).

F. Sample Countries and Descriptive Statistics for International Spillover Analysis

Table F.1: Countries included in the international spillover analysis

Group	Countries
Advanced Economies (50)	Australia, Austria, Belgium, Bulgaria, Canada, Chile, Costa Rica, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, Ireland, Israel, Italy, Japan, South Korea, Kuwait, Latvia, Lithuania, Luxembourg, Malta, Netherlands, New Zealand, Norway, Oman, Panama, Poland, Portugal, Qatar, Romania, Russian Federation, Saudi Arabia, Singapore, Slovakia, Slovenia, Spain, Sweden, Switzerland, United Arab Emirates, United Kingdom, United States, Uruguay.
Emerging Markets (102)	Afghanistan, Albania, Algeria, Angola, Argentina, Armenia, Azerbaijan, Bangladesh, Belarus, Benin, Bolivia, Bosnia and Herzegovina, Botswana, Brazil, Burkina Faso, Burundi, Cambodia, Cameroon, Central African Republic, Chad, China, Colombia, Côte d'Ivoire, Democratic Republic of the Congo, Dominican Republic, Ecuador, Egypt, El Salvador, Eritrea, Ethiopia, Fiji, Gabon, Gambia, Georgia, Ghana, Guatemala, Guinea, Guinea-Bissau, Haiti, Honduras, India, Indonesia, Iran, Iraq, Jamaica, Jordan, Kazakhstan, Kenya, Kyrgyzstan, Laos, Lebanon, Lesotho, Liberia, Libya, Madagascar, Malawi, Malaysia, Maldives, Mali, Mauritania, Mauritius, Mexico, Moldova, Mongolia, Montenegro, Morocco, Mozambique, Myanmar, Namibia, Nepal, Nicaragua, Niger, Nigeria, North Macedonia, Pakistan, Papua New Guinea, Paraguay, Peru, Philippines, Rwanda, Senegal, Serbia, Sierra Leone, South Africa, Sri Lanka, Sudan, Taiwan, Tajikistan, Tanzania, Thailand, Togo, Tunisia, Turkey, Turkmenistan, Uganda, Ukraine, Uzbekistan, Venezuela, Vietnam, Yemen, Zambia, Zimbabwe.

Notes: The classification follows the country grouping adopted in the empirical analysis. Advanced economies (AEs) and emerging markets (EMs) are defined according to the IMF classification.

Table F.2: Descriptive statistics for the international spillover analysis

Variable	**Obs.**	**Mean**	**Median**	**S.D.**	**P10 / P90**
All Countries					
GDP per capita (ln)	6383	-3.89	-4.61	2.21	-6.09 / -0.01
Industrial Production (ln)	7258	4.24	4.49	0.74	3.46 / 4.74
Fixed Investment (ln)	5809	24.78	24.84	2.94	21.13 / 28.24
Stock Prices (ln)	4441	3.79	4.18	1.47	2.01 / 5.03
WUI	31879	0.04	0.03	0.05	0.00 / 0.11
WSI	13722	0.82	0.79	0.99	-0.32 / 2.02
EPU	2416	4.61	4.59	0.50	4.03 / 5.24
Stock volatility (ln)	7618	-4.51	-4.53	0.49	-5.11 / -3.86
SPF GDP disagreement	2346	0.44	0.38	0.27	0.22 / 0.72
Advanced Economies					
GDP per capita (ln)	4510	-3.89	-4.72	2.11	-5.84 / 0.04
Industrial Production (ln)	6004	4.18	4.45	0.79	3.29 / 4.73
Fixed Investment (ln)	4252	24.38	24.48	2.74	21.03 / 27.90
Stock Prices (ln)	2826	3.87	4.20	1.23	2.23 / 5.03
WUI	10262	0.04	0.03	0.05	0.00 / 0.10
WSI	4224	0.96	0.98	1.19	-0.47 / 2.38
EPU	1652	4.64	4.61	0.47	4.08 / 5.25
Stock volatility (ln)	4607	-4.55	-4.61	0.47	-5.12 / -3.95
SPF GDP disagreement	1606	0.40	0.35	0.21	0.21 / 0.63
Emerging Markets					
GDP per capita (ln)	1873	-3.90	-3.78	2.43	-6.71 / -1.70
Industrial Production (ln)	1141	4.49	4.58	0.29	4.06 / 4.79
Fixed Investment (ln)	1457	25.79	25.36	3.29	21.70 / 30.04
Stock Prices (ln)	1519	3.59	4.15	1.86	1.85 / 5.05
WUI	21617	0.04	0.02	0.06	0.00 / 0.12
WSI	9498	0.76	0.73	0.89	-0.25 / 1.83
EPU	580	4.49	4.44	0.61	3.80 / 5.21
Stock volatility (ln)	3011	-4.44	-4.41	0.51	-5.11 / -3.78
SPF GDP disagreement	740	0.52	0.43	0.35	0.25 / 0.85

Notes: The table reports descriptive statistics for the variables used in the international spillover analysis. P10 and P90 denote the 10th and 90th percentiles, respectively. All variables are expressed in logarithms where indicated.

4. General Conclusion

The central argument developed in this thesis is that economic uncertainty and financial volatility are best understood as complex, information-intensive phenomena that require equally sophisticated empirical tools. Across the three parts of the work, a common theme emerges: the relevant signals are rarely observed directly, and they must therefore be recovered from large, heterogeneous datasets through disciplined statistical and econometric methods. This is true in the measurement of uncertainty, in the forecasting of volatility within the EuroSTOXX 50 universe, and in the analysis of Federal Reserve communication. Although the substantive questions differ, the underlying challenge is the same: how to identify the economically meaningful component of a noisy and high-dimensional environment.

The survey chapter establishes the conceptual background. It shows that uncertainty is not a single, unambiguous measure, but rather a family of related concepts that have been operationalised in different ways. News-based indices, textual measures, forecast-based definitions, and market proxies each capture distinct aspects of the economic environment. Their coexistence is not a defect of the literature, but a reflection of the fact that uncertainty itself is multidimensional. The literature reviewed in the first part also makes clear that the use of text in economics is no longer peripheral. Textual analysis has become a standard route to measuring expectations, sentiment, policy tone, and uncertainty, yet it remains subject to important choices regarding source selection, classification, and interpretation. The survey, therefore, provides the conceptual and methodological basis for the rest of the thesis.

The second part of the thesis shows how factor models can be used to address a different but related problem: how to predict volatility in a large and interconnected financial system. The Generalised Dynamic Factor Model is particularly appropriate in this context because it is designed to recover common dynamic forces from a large panel of time series. In the EuroSTOXX 50 setting, where assets are exposed to both common European shocks and idiosyncratic movements, this framework offers a natural way of organising the data and extracting predictive structure. The empirical contribution of this part lies in demonstrating that volatility forecasting can benefit from a factor-based approach that is both flexible and economically interpretable. More broadly, it underlines the value of working with high-dimensional methods when the object of interest is not an isolated series but a system of mutually dependent financial variables.

The third part of the thesis places communication at the centre of the analysis. The communication of the Federal Reserve matters because policy is not transmitted solely through interest-rate decisions or balance-sheet operations. It is also transmitted through words: through statements, press conferences, minutes, speeches, and forward guidance. These forms of communication shape expectations, alter market beliefs, and can move asset prices even in the absence of an immediate policy action. The analysis developed in this chapter, therefore, reinforces a broader point made throughout the thesis: language is not merely descriptive, but economically consequential. Insofar as policymakers use language strategically, communication becomes part of the policy instrument set and must be studied as such.

Read together, the three parts of the thesis suggest a coherent way of thinking about uncertainty, volatility, and policy communication. First, uncertainty must be measured with tools that are sensitive to the informational content of text and to the structure of large datasets. Second, volatility in financial markets is best modelled as a dynamic and partly common phenomenon, for which factor methods offer a suitable representation. Third, communication by central banks affects markets precisely because it changes the informational environment in which agents form expectations and update beliefs. The link between these themes is not accidental. All three concern how information is produced, aggregated, and transmitted in an economy characterised by complexity and incomplete observability.

The thesis also points to the practical value of maintaining this integrated perspective. For researchers, it suggests that empirical work on uncertainty and communication should not remain confined to a single methodological tradition. Text-based measures, factor models, and event-based analyses can be combined in ways that improve both measurement and interpretation. For policy analysis, it implies that language and information deserve the same analytical seriousness as more conventional quantitative variables. For financial econometrics, it reinforces the usefulness of multivariate and dynamic approaches when the objective is to forecast, detect, or interpret volatility in interconnected markets.

At the same time, the thesis makes clear that methodological sophistication does not remove the need for careful interpretation. Textual indicators depend on how language is selected and coded; factor models depend on how latent structure is specified and estimated; communication analyses depend on how events are timed, classified, and linked to market responses. None of these methods is self-validating. Their value lies in the discipline with which they are applied and in the transparency with which their assumptions are stated. The contribution of this thesis is therefore not to present a final answer, but to show how a set of complementary methods can be used to study related economic phenomena in a coherent way.

The broader implication is that uncertainty, volatility, and communication should be approached as part of a common empirical architecture. They are different objects, but they are connected by the same underlying problem: the economy is shaped by information that is incomplete, unevenly distributed, and constantly revised. The survey of uncertainty measurement, the GDFM analysis of European equity volatility, and the study of Federal Reserve communication each illuminate one part of that architecture. Together, they show that the combination of textual analysis and high-dimensional econometric modelling is not only feasible, but also fruitful for understanding modern financial and macroeconomic dynamics.

This thesis closes with that claim, but not with closure in the stronger sense. The methods reviewed and applied here open several lines of further work, especially where textual information, factor structure, and market response intersect. What the present study offers is a framework:

a way of organising the empirical problem so that future research can build on it with greater precision. In that sense, the main contribution of the thesis is to clarify how uncertainty and communication can be studied together, and why doing so requires a careful balance between interpretability, dimensionality reduction, and institutional context.

Bibliography

Ahir, H., N. Bloom, and D. Furceri (2022). The world uncertainty index. National bureau of economic research working paper, n. 29763.

Ahrens, M., D. Erdemlioglu, M. McMahon, C. J. Neely, and X. Yang (2024). Mind your language: Market responses to central bank speeches. *Journal of Econometrics*, 105 – 121.

Alessi, L., M. Barigozzi, and M. Capasso (2009). Forecasting large datasets with conditionally heteroskedastic dynamic common factors. ECB working paper no. 1115.

Alizadeh, S., M. W. Brandt, and F. X. Diebold (2002). Range-based estimation of stochastic volatility models. *the Journal of Finance 57*(3), 1047–1091.

Alonso Alvarez, I., J. J. Pérez, and M. Diakonova (2025). Rethinking gpr: The sources of geopolitical risk. Banco de España working paper n°2522, Banco de España.

Altissimo, F., R. Cristadoro, M. Forni, M. Lippi, and G. Veronese (2010). New eurocoin: Tracking economic growth in real time. *Review of Economics and Statistics 92*(4), 1024–1034.

American Association of Individual Investors (2026). AAII Investor Sentiment Survey. Historical survey data on bullish, neutral, and bearish investor sentiment.

Andrews, D. W. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica: Journal of the Econometric Society*, 817–858.

Araci, D. (2019). Finbert: Financial sentiment analysis with pre-trained language models. arXiv:1908.10063.

Aruoba, S. B., F. X. Diebold, and C. Scotti (2009). Real-time measurement of business conditions. *Journal of Business & Economic Statistics 27*(4), 417–427.

Aruoba, S. B. and T. Drechsel (2024). Identifying monetary policy shocks: A natural language approach. NBER wp 32417.

Ash, E. and S. Hansen (2023). Text algorithms in economics. *Annual Review of Economics 15*, 659–688.

Aslam, F., P. Ferreira, K. S. Mughal, and B. Bashir (2021). Intraday volatility spillovers among European financial markets during COVID-19. *International Journal of Financial Studies 9*(1), 5.

Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica 71*(1), 135–171.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bai, J. and P. Wang (2015). Identification and bayesian estimation of dynamic factor models. *Journal of Business & Economic Statistics 33*(2), 221–240.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Barabási, A.-L. and M. Pósfai (2016). *Network science*. Cambridge: Cambridge University Press.

Barhoumi, K., O. Darné, and L. Ferrara (2013). Dynamic factor models: A review of the literature. *Journal of Business Cycle Measurement and Analysis 2*, 73–107.

Barigozzi, M. and C. Brownlees (2019). Nets: Network estimation for time series. *Journal of Applied Econometrics 34*(3), 347–364.

Barigozzi, M., H. Cho, and D. Owens (2024). Fnets: Factor-adjusted network estimation and forecasting for high-dimensional time series. *Journal of Business & Economic Statistics 42*(3), 890–902.

Barigozzi, M. and M. Hallin (2016). Generalized dynamic factor models and volatilities: Recovering the market volatility shocks. *The Econometrics Journal 19*(1), C33–C60.

Barigozzi, M. and M. Hallin (2017). A network analysis of the volatility of high dimensional financial series. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 66*(3), 581–605.

Barigozzi, M. and M. Hallin (2020). Generalized dynamic factor models and volatilities: Consistency, rates, and prediction intervals. *Journal of Econometrics 216*(1), 244–272.

Barigozzi, M., M. Hallin, and S. Soccorsi (2019). Identification of global and local shocks in international financial markets via general dynamic factor models. *Journal of Financial Econometrics 17*(3), 462–494.

Barigozzi, M., M. Hallin, S. Soccorsi, and R. von Sachs (2021). Time-varying general dynamic factor models and the measurement of financial connectedness. *Journal of Econometrics 222*(1), 324–343.

Barry, M.-L., B. J. Bruns, S. Kandemir, J. Klose, V. Smirnov, and P. Tillmann (2025). The emotions of monetary policy. Ssrn 5099686.

Bauer, M. D., A. Lakdawala, and P. Mueller (2021). Market-based monetary policy uncertainty. *The Economic Journal 132*(644), 1290–1308.

Bekaert, G., M. Bergbrant, and H. Kassa (2025). Expected idiosyncratic volatility. *Journal of Financial Economics 167*, 104023.

Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological) 57*(1), 289–300.

Bernanke, B. S., J. Boivin, and P. Eliasz (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *Quarterly Journal of Economics 120*(1), 387–422.

Beyhum, J. and J. Striaukas (2024). Testing for sparse idiosyncratic components in factor-augmented regression models. *Journal of Econometrics 244*(1), 105845.

Bianchi, F., S. C. Ludvigson, and S. Ma (2022). Belief distortions and macroeconomic fluctuations. *American Economic Review 112*(7), 2269–2315.

Birru, J. and T. Young (2022). Sentiment and uncertainty. *Journal of Financial Economics 146*(3), 1148–1169.

Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research 3*, 993–1022.

Blinder, A. S., M. Ehrmann, J. de Haan, and D.-J. Jansen (2024). Central Bank Communication with the General Public: Promise or False Hope? *Journal of Economic Literature 62*(2), 425–457.

Bloom, N. (2009). The impact of uncertainty shocks. *econometrica 77*(3), 623–685.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics 31*(3), 307–327.

Bordo, M. and K. Istrefi (2023). Perceived fomc: The making of hawks, doves and swingers. *Journal of Monetary Economics 136*, 125–143.

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *The review of financial studies 22*(6), 2201–2238.

Bunčić, D. and K. I. M. Gisler (2016). Global equity market volatility spillovers: A broader role for the United States. *International Journal of Forecasting 32*(4), 1317–1339.

Bybee, L., B. T. Kelly, A. Manela, and D. Xiu (2020). The structure of economic news. NBER working paper n°26648.

Caldara, D. and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review 112*(4), 1194–1225.

Campiglio, E., J. Deyris, D. Romelli, and G. Scalisi (2025). Warning words in a warming world: Central bank communication and climate change. Grantham research institute on climate change and the environment working paper no 418.

Carriero, A., T. E. Clark, and M. Marcellino (2018). Measuring uncertainty and its impact on the economy. *Review of Economics and Statistics 100*(5), 799–815.

Cascaldi-Garcia, D., C. Sarisoy, J. M. Londoño, J. Rogers, B. Sun, D. Datta, T. R. T. Ferreira, O. V. Grishchenko, M. R. Jahan-Parvar, F. Loria, S. Ma, M. Rodriguez, and I. Zer (2023). What is certain about uncertainty? *Journal of Economic Literature 61*(2), 624–654.

Chen, J., D. Li, Y.-N. Li, and O. Linton (2025). Estimating time-varying networks for high-dimensional time series. *Journal of Econometrics*, 105941.

Chen, Q., A. Filardo, D. He, and F. Zhu (2016). Financial crisis, US unconventional monetary policy and international spillovers. *Journal of International Money and Finance 67*, 62–81.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of macroeconomics 1*, 65–148.

Cieslak, A., S. Hansen, M. McMahon, and S. Xiao (2023). Policymakers' uncertainty. NBER working paper n°31849.

Cong, L. W., T. Liang, X. Zhang, and W. Zhu (2025). Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information. *Management Science 71*(12), 10727–10739.

Cont, R. and E. Schaanning (2017). Fire sales, indirect contagion and systemic stress testing. *Indirect Contagion and Systemic Stress Testing (June 13, 2017)*.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics 7*(2), 174–196.

Davis, S. J. (2016). An index of global economic policy uncertainty. Technical report, National Bureau of Economic Research.

Dedola, L., G. Rivolta, and L. Stracca (2017). If the fed sneezes, who catches a cold? *Journal of International Economics 108*, S23–S41.

Deev, O. and v. Lyócsa (2020). Connectedness of financial institutions in Europe: A network approach across quantiles. *Physica A: Statistical Mechanics and its Applications 550*, 124035.

Del Negro, M., M. Giannoni, C. Patterson, et al. (2012). The forward guidance puzzle. Federal reserve bank of new york new york.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171–4186.

Diebold, F. X. and K. Yılmaz (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of econometrics 182*(1), 119–134.

Doz, C., D. Giannone, and L. Reichlin (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics 164*(1), 188–205.

Doz, C., D. Giannone, and L. Reichlin (2012). A quasi–maximum likelihood approach for large, approximate dynamic factor models. *Review of Economics and Statistics 94*(4), 1014–1024.

Ehrmann, M. and J. Talmi (2020). Starting from a blank page? semantic similarity in central bank communication and market volatility. *Journal of Monetary Economics 111*, 48–62.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, 987–1007.

Fadda, P., R. Hanifi, K. Istrefi, and A. Penalver (2025, None). Central bank communication of uncertainty. *Journal of International Money and Finance 157*(C), None.

Feng, Y. and Y. Zhang (2025). Forecasting realized volatility: The choice of window size. *Journal of forecasting 44*(2), 692–705.

Fernández-Villaverde, J., T. Mineyama, and D. Song (2024). Are we fragmented yet? measuring geopolitical fragmentation and its causal effect. NBER working paper n°32638.

Filardo, A. J. and B. Hofmann (2014). Forward guidance at the zero lower bound. *BIS Quarterly Review March*.

Foglia, M. and E. Angelini (2020). From me to you: Measuring connectedness between Eurozone financial institutions. *Research in International Business and Finance 54*, 101238.

Forni, M. and L. Gambetti (2010). The dynamic effects of monetary policy: A structural factor model approach. *Journal of Monetary Economics 57*(2), 203–216.

Forni, M. and L. Gambetti (2014). Sufficient information in structural vars. *Journal of Monetary Economics 66*, 124–136.

Forni, M., L. Gambetti, A. Granese, L. Sala, and S. Soccorsi (2025). An american macroeconomic picture: Supply and demand shocks in the frequency domain. *American Economic Journal: Macroeconomics 17*(3), 311–341.

Forni, M., D. Giannone, M. Lippi, and L. Reichlin (2009). Opening the black box: Structural factor models with large cross sections. *Econometric Theory 25*(5), 1319–1347.

Forni, M., A. Giovannelli, M. Lippi, and S. Soccorsi (2018). Dynamic factor model with infinite-dimensional factor space: Forecasting. *Journal of Applied Econometrics 33*(5), 625–642.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2000). The generalized dynamic-factor model: Identification and estimation. *Review of Economics and Statistics 82*(4), 540–554.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2005). The generalized dynamic factor model: One-sided estimation and forecasting. *Journal of the American Statistical Association 100*(471), 830–840.

Forni, M., M. Hallin, M. Lippi, and P. Zaffaroni (2017). Dynamic factor models with infinite-dimensional factor space: Asymptotic analysis. *Journal of Econometrics 199*(1), 74–92.

Forni, M. and M. Lippi (2001). The generalized dynamic factor model: representation theory. *Econometric theory 17*(6), 1113–1141.

Gambacorta, L., B. Kwon, T. Park, P. Patelli, and S. Zhu (2024). Cb-lms: Language models for central banking. BIS Working Paper 1215, Bank for International Settlements.

Gambetti, L., D. Korobilis, J. Tsoukalas, and F. Zanetti (2023). Agreed and disagreed uncertainty. Technical report.

Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature 57*(3), 535–574.

Georgiadis, G. (2016). Determinants of global spillovers from US monetary policy. *Journal of International Money and Finance 67*, 41–61.

Georgiadis, G. and M. Jarocinski (2023). Global spillovers from multi-dimensional US monetary policy. ECB working paper n°2881.

Gertler, P. and R. Horvath (2018). Central bank communication and financial markets: New high-frequency evidence. *Journal of Financial Stability 36*, 336–345.

Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica 74*(6), 1545–1578.

Giannone, D., L. Reichlin, and D. Small (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics 55*(4), 665–676.

Giglio, S. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy 129*(7), 1947–1990.

Giovannelli, A., D. Massacci, and S. Soccorsi (2021). Forecasting stock returns with large dimensional factor models. *Journal of Empirical Finance 63*, 252–269.

Gorodnichenko, Y., T. Pham, and O. Talavera (2023). The voice of monetary policy. *American Economic Review 113*(2), 548–584.

Gorodnichenko, Y., T. Pham, and O. Talavera (2024). Central bank communication on social media: What, to whom, and how? *Journal of Econometrics*. Forthcoming.

Greenwood-Nimmo, M., V. H. Nguyen, and Y. Shin (2021). Measuring the connectedness of the global economy. *International Journal of Forecasting 37*(2), 899–919.

Grilli, R., G. Tedeschi, and M. Gallegati (2015). Markets connectivity and financial contagion. *Journal of Economic Interaction and Coordination 10*, 287–304.

Haeran Cho, Hyeyoung Maeng, I. A. E. and P. Fearnhead (2024). High-dimensional time series segmentation via factor-adjusted vector autoregressive modeling. *Journal of the American Statistical Association 119*(547), 2038–2050.

Hallin, M. and M. Lippi (2013). Factor models in high-dimensional time series—a time-domain approach. *Stochastic processes and their applications 123*(7), 2678–2695.

Hallin, M. and R. Liška (2007). Determining the number of factors in the general dynamic factor model. *Journal of the American Statistical Association 102*(478), 603–617.

Hanifi, R., K. Istrefi, and A. Penalver (2022). Central bank communication of uncertainty. Banque de france wp898.

Hansen, S. and M. McMahon (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics 99*, S114–S133.

Hansen, S., M. McMahon, and M. Tong (2019). The long-run information effect of central bank communication. *Journal of Monetary Economics 108*, 185–202.

Hassan, T. A., S. Hollander, A. Kalyani, L. van Lent, M. Schwedeler, and A. Tahoun (2025). Text as data in economic analysis. *Journal of Economic Perspectives 39*(3), 193–220.

Huang, A. H., H. Wang, and Y. Yang (2023). Finbert: A large language model for extracting information from financial text. *Contemporary Accounting Research 40*(2), 806–841.

Huynh, T. L. D., M. Foglia, and J. A. Doukas (2021). COVID-19 and tail-event driven network risk in the Eurozone. *Finance Research Letters 44*, 102070.

Istrefi, K., F. Odendahl, and G. Sestieri (2023). Fed communication on financial stability concerns and monetary policy decisions: Revelations from speeches. *Journal of Banking & Finance 151*, 106820.

Istrefi, K. and A. Piloiu (2020). Public opinion on central banks when economic policy is uncertain. *Revue d'économie politique 130*(2), 283–306.

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

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review 95*(1), 161–182.

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring Uncertainty. *American Economic Review 105*(3), 1177–1216.

Katsouris, C. (2024). Robust estimation in network vector autoregression with nonstationary regressors. arXiv:2401.04050.

Knight, F. H. (1921). *Risk, uncertainty and profit*, Volume 31. Houghton Mifflin.

Lakdawala, A., T. Moreland, and M. Schaffer (2021). The international spillover effects of US monetary policy uncertainty. *Journal of International Economics 133*, 103525.

Legerstee, R. and P. H. Franses (2015). Does disagreement amongst forecasters have predictive value? *Journal of Forecasting 34*(4), 290–302.

Loughran, T. and B. McDonald (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-ks. *The Journal of finance 66*(1), 35–65.

Loughran, T. and B. McDonald (2016). Textual analysis in accounting and finance: A survey. *Journal of accounting research 54*(4), 1187–1230.

Lucchetti, L. and D. O. Cajueiro (2026). Language as data: An overview of natural language processing in economics and finance. *Journal of Economic Surveys*.

Ludvigson, S. C., S. Ma, and S. Ng (2021). Uncertainty and business cycles: exogenous impulse or endogenous response? *American Economic Journal: Macroeconomics 13*(4), 369–410.

Lütkepohl, H. (2005). *New introduction to multiple time series analysis*. Springer.

Magkonis, G. and A. Tsopanakis (2020). The financial connectedness between eurozone core and periphery: A disaggregated view. *Macroeconomic Dynamics 24*(7), 1674–1699.

Martin, B., F. S. Passino, M. Cucuringu, and A. Luati (2024). Nirvar: Network informed restricted vector autoregression.

McKay, A., E. Nakamura, and J. Steinsson (2016). The power of forward guidance revisited. *American Economic Review 106*(10), 3133–3158.

Mensi, W., F. Z. Boubaker, K. H. Al-Yahyaee, and S. H. Kang (2018). Dynamic volatility spillovers and connectedness between global, regional, and GIPSI stock markets. *Finance Research Letters 25*, 230–238.

Miranda-Agrippino, S. and H. Rey (2020). US monetary policy and the global financial cycle. *The Review of Economic Studies 87*(6), 2754–2776.

Miranda-Agrippino, S. and G. Ricco (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics 13*(3), 74–107.

Moschella, M. and L. Pinto (2019). Central banks' communication as reputation management: How the fed talks under uncertainty. *Public administration 97*(3), 513–529.

Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics 133*(3), 1283–1330.

Newey, W. K. and K. D. West (1986). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix.

Owens, D., H. Cho, and M. Barigozzi (2023). fnets: An r package for network estimation and forecasting via factor-adjusted var modelling. *The R Journal 15*, 214–239.

Paltalidis, N., D. Gounopoulos, R. Kizys, and Y. Koutelidakis (2015). Transmission channels of systemic risk and contagion in the european financial network. *Journal of Banking & Finance 61*, S36–S52.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of business*, 61–65.

Petropoulos, A. and V. Siakoulis (2021). Can central bank speeches predict financial market turbulence? evidence from an adaptive nlp sentiment index analysis using xgboost machine learning technique. *Central Bank Review 21*(4), 141–153.

Poncela, P., E. Ruiz, and K. Miranda (2021). Factor extraction using kalman filter and smoothing: This is not just another survey. *International Journal of Forecasting 37*(4), 1399–1425.

Priola, M. P., P. Lorenzini, G. Tizzanini, et al. (2021). Measuring central banks' sentiment and its spillover effects with a network approach. Ssrn wp: https://ssrn.com/abstract=3764004.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. *Handbook of macroeconomics 2*, 71–162.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American economic review 94*(4), 1055–1084.

Rosa, C. and G. Verga (2008). The impact of central bank announcements on asset prices in real time. *International Journal of Central Banking 4*(2), 175–217.

Shapiro, A. H., M. Sudhof, and D. J. Wilson (2022). Measuring news sentiment. *Journal of Econometrics 228*(2), 221–243.

Shleifer, A. and R. Vishny (2011). Fire sales in finance and macroeconomics. *Journal of economic perspectives 25*(1), 29–48.

Sill, K. (2012). Measuring economic uncertainty using the survey of professional forecasters. *Business Review* (Q4), 16–27.

Sill, K. (2014). Forecast disagreement in the survey of professional forecasters. *Business Review (Federal Reserve Bank of Philadelphia) 97*(2).

Silva, T. C., K. Moriya, and R. M. Veyrune (2025). From text to quantified insights: A large-scale LLM analysis of central bank communication. IMF Working Paper 2025/109, International Monetary Fund.

Stock, J. H. and M. W. Watson (2002a). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association 97*(460), 1167–1179.

Stock, J. H. and M. W. Watson (2002b). Macroeconomic forecasting using diffusion indexes. *Journal of business & economic statistics 20*(2), 147–162.

Stock, J. H. and M. W. Watson (2005). Implications of dynamic factor models for var analysis. National bureau of economic research cambridge, mass., usa.

Stock, J. H. and M. W. Watson (2016). Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, pp. 415–525. Elsevier.

Swanson, E. T. (2021). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics 118*, 32–53.

Swanson, E. T. and V. Jayawickrema (2023). Speeches by the fed chair are more important than FOMC announcements: An improved high-frequency measure of US monetary policy shocks. Unpublished manuscript.

Symitsi, E., L. Symeonidis, A. Kourtis, and R. N. Markellos (2018). Covariance forecasting in equity markets. *Journal of Banking & Finance 96*, 153–168.

Tang, Y., F. Ma, M. I. M. Wahab, and Y. Wei (2022). Does the US stock market information matter for European equity market volatility: A multivariate perspective? *Applied Economics 54*(49), 5703–5718.

Tillmann, P. and A. Walter (2019). The effect of diverging communication: The case of the ECB and the bundesbank. *Economics Letters 176*, 68–74.

Wu, J. C. and F. D. Xia (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking 48*(2-3), 253–291.

Yang, Y., M. C. S. Uy, and A. Huang (2020). Finbert: A pretrained language model for financial communications. arXiv:2006.08097.

Yu, Z., W. Liu, and F. Yang (2023). A central bankers' sentiment index of global financial cycle. *Finance Research Letters 57*. Forthcoming.

Zeileis, A. (2004). Econometric computing with hc and hac covariance matrix estimators. *Journal of statistical software 11*, 1–17.