MENTAL MODELS AND FINANCIAL FORECASTS *

Francesca Bastianello Paul H. Décaire Marius Guenzel†

November 18, 2024

Abstract

We uncover financial professionals' mental models—the narratives they use to explain their subjective beliefs. Using 82,000 equity reports, we prompt large language models (LLMs) to extract 3.5 million narratives, each combining a topic, valuation channel, sentiment, and time outlook. To validate the reliability of our output, we introduce a multi-step LLM-based approach and new diagnostic tools. We establish three sets of findings. First, narratives are centered around a limited set of topics, primarily focused on top-line items, with variation in topic focus over time and across industries. Narratives are mostly forward-looking, with three times as many arguments focusing on the future as on the past. Second, differences in topic focus, sentiment, and time outlook across forecasters strongly predict the extent of disagreement in their subjective quantitative forecasts. Lastly, time-series variation in the average narrative's sentiment and in the average narrative's focus on top-line items closely track Shiller's CAPE ratio ($\rho = 0.84$ and $\rho = 0.42$), and the cross-sectional variation in narratives predicts key asset pricing patterns. Narratives associated with growth stocks are more optimistic and forward-looking than those for value stocks, consistent with forecasters (mis)perceiving growth stocks as having above-average growth potential. Overall, this paper helps bridge the gap between 'what forecasters believe' and 'why they believe it.'

Keywords: Mental Models, Professional Forecasters, Large Language Models, Beliefs

*First Version: October 30, 2024. This Version: November 18, 2024. All errors and omissions are our own.

†Bastianello is at the University of Chicago, Booth School of Business. Décaire is at the W. P. Carey School of Business at Arizona State University. Guenzel is at the Wharton School of the University of Pennsylvania. We thank conference participants at the 2023 ASU-UA Junior Conference and the 1st Workshop on LLMs and Generative AI for Finance. Bastianello gratefully acknowledges support from the Liew Fama Junior Faculty Fellowship, the Fama Faculty Fellowship, and the Fama-Miller Center at the University of Chicago Booth School of Business. Décaire gratefully acknowledges financial support from the Center for Responsible Investing at Arizona State University. Guenzel gratefully acknowledges financial support from the Jacobs Levy Equity Management Center for Quantitative Financial Research. Contact information: francesca.bastianello@chicagobooth.edu, paul.decaire@asu.edu, and mguenzel@wharton.upenn.edu.

1 Introduction

How do people form expectations of financial variables? According to the rational expectations benchmark, individuals use all available information and are able to produce statistically optimal forecasts. However, a growing body of research has challenged this hypothesis, with the benefit of increasingly better data. While early evidence was mostly indirect and relied heavily on observed outcomes and anomalies, recent efforts using survey data and lab experiments offer ways to validate and enrich belief formation theories by directly measuring individuals' expectations.[1] Despite these advances, we still have limited direct evidence of the underlying reasoning behind these beliefs.

This paper fills this gap by measuring professional forecasters' mental models— the reasoning and narratives they use to explain their beliefs—using a large sample of 82 thousand equity reports written by financial analysts, and linking the narratives to forecasters' quantitative expectations within the same reports. This setting is well-suited for studying mental models, as analysts use their reports to connect, add nuance, and emphasize the key factors they believe will shape a firm's future performance. These narratives open a window into the mental models forecasters use to make sense of the world.

To systematically elicit mental models, we prompt large language models (LLMs) to extract more than 3.5 million separate narratives from the text portions of the sell-side analysts' equity reports. Our dataset spans 11,308 stocks over a 24-year period and closely aligns with commonly used commercial databases such as the Institutional Brokers' Estimate System (IBES) in firm coverage and characteristics. The extracted narratives identify forecasters' mental models as a combination of attention allocation to topics (e.g., about customer demand or market share), each with its associated valuation channel (e.g., impacting sales, costs, margins, or discount rates), sentiment (positive, neutral, or negative) and time outlook (past, present, near future, and distant future).

Our results are organized around one methodological contribution and three sets of

[1] See De Bondt and Thaler (1990), Vissing-Jorgensen (2003), Bacchetta et al. (2009), Malmendier and Nagel (2011), Case et al. (2012), Amromin and Sharpe (2014), Greenwood and Shleifer (2014), Bordalo et al. (2020b), Bordalo et al. (2020a), Giglio et al. (2021), Nagel and Xu (2022) among others, and Adam and Nagel (2023) for a review.

analyses. (1) We introduce a multi-step, LLM-based approach as well as new diagnostic tools to help researchers evaluate the validity of LLM-based outputs in the context of economic analysis. (2) We show that narratives are derived from a simplified model of the world, in which forecasters subjectively consider a limited subset of topics to be of first-order importance. This leads to a sparse mental representation: when two analysts evaluate the same firm, only 32% of the topics they discuss overlap between their reports. (3) Differences across forecasters' individual narratives significantly predict the heterogeneity in their subjective quantitative beliefs. (4) Forecasters' average narratives significantly predict key asset pricing patterns. In the time series, fluctuations in forecasters' narrative sentiment, time-outlook, and focus on key valuation channels (i.e., sales) predict time-series variation in valuation ratios. In the cross-section, the spread in forecasters' narrative sentiment between value and growth stocks or small and large stocks replicates the patterns documented in the literature when using returns.

Methodological contribution. To tackle the challenges associated with extracting information from large bodies of text, we make two methodological contributions. First, we introduce a new approach which maximizes the richness, informativeness, and stability of the extracted information. Second, we introduce a new diagnostic tool to ensure the robustness of our LLM-based text analysis and synthesis.

We first show that, when working with large bodies of text, naive methods that prompt LLMs to extract narratives in a single step result in outputs that are both unstable and incomplete. Part of the challenge lies in working with large bodies of text, and LLMs' greater focus on the beginning of the text provided. To overcome these challenges, we fine-tune our LLM-based mental model extraction methods. Instead of a naive, single-step prompt, we iteratively provide the artificial intelligence (AI) with sub-segments of each report. Once completed, we provide the AI with all extracted narratives and the full documents, and ask it to validate the selected topics. This step allows the AI to ensure that our "chunking method" does not alter the context associated with each narrative, even in cases where a narrative spans more than one sub-segment.

We demonstrate that not only does this multi-step strategy identify more narratives than the naive one-step elicitation approach, but it also better captures the nuances and complexity

of forecasters' theses—crucial features for studying the intricacies of mental models. For example, an analyst might discuss the same topic (e.g., inflation) and attribute it to two valuation channels (e.g., sales and costs). In these instances, single-step approaches are less likely to pick up more than one narrative associated with a given topic. Moreover, we show that the additional narratives extracted through the multi-step strategy are important, as they are strongly predictive of forecasters' quantitative beliefs. Our employed methodology can be readily applied to many other contexts involving AI-based text processing of large volumes of text.

To validate the robustness and stability of our LLM-based text analysis and synthesis, we also introduce a new diagnostic tool. The increased processing speeds and cost reductions associated with LLMs enable us to prompt the AI to extract topics from the same equity report *multiple times*. We then calculate the likelihood of each topic being included across attempts, effectively quantifying the consistency of the AI's output. We find a very high level of stability associated with the output of our multi-step prompt, with substantial improvements over the single-step prompt.

Attention Allocation. Forecasters predominantly focus their attention on a limited subset of topics. The average report discusses 18 distinct topics, out of 139 total topics ever discussed across all reports. Longer reports tend to encompass a broader range of topics, and provide more detailed explanations.

To assess whether topics omitted from textual discussions are outside forecasters' mental models, rather than excluded due to space constraints in reports, we conduct two exercises. First, we show that topics not included in a report—but utilized by other analysts forecasting the same firm in the same year—have no explanatory power for a forecaster's subjective beliefs. This approach enables us to concentrate on topics that are likely relevant for evaluating the firm, while filtering out irrelevant ones, thereby reducing noise in our measurement. Second, building on this line of reasoning, we focus on the subset of reports that fall within the top 10% of length in our sample. For pairs of analysts evaluating the same firm-year, we examine whether the number of overlapping topics approaches 100% as space constraints in the reports diminish. Despite the reduced constraints, we still observe a significant gap in the topics covered between these pairs of reports, indicating that analysts do not converge on the same

set of topics even when given more space.

We also show that forecasters' mental models, while topically sparse, involve nuanced "within-topic" narratives that often consider the effect of a given topic (e.g., inflation) on more than one valuation channel (e.g., sales and cost) and over multiple time horizons. On average, a topic is articulated in 2.43 individual narratives. It is associated with more than one valuation channel in 15% of cases, and forecasters express mixed sentiment about a topic's effect on valuation in 23% of cases.

Across forecasters' reports and mental models, two distinct topic categories emerge: persistent topics associated with the key dynamics driving firms' operations in each respective industry, and transitory topics that quickly capture the spotlight but disappear quickly from narratives. Persistent topics include customer demand, market expansion, and market share, whereas transient topics include pandemics and public health issues and inflation. Moreover, forecasters' mental models vary in an intuitive manner not just over time and through the business cycle, but also across industries. For example, compared to the average firm in the sample, forecasters' narratives for utility firms skew towards regulation and compliance topics, whereas those for mining and oil firms tilt towards production capacity (i.e., reserves) topics. Likewise, attention allocation to transient topics increases when the underlying economic variables associated with these topics deviate from heuristic benchmarks or historical trends (e.g., inflation exceeding 2%), and when there are significant news surprises.

Next, we analyze the valuation channels associated with the discussed topics: (i) sales (accounting for 45% of discussed topics), (ii) costs (15%), (iii) profitability margins (10%), (iv) earnings/cash flows (25%), and (v) discount rates (5%). Forecasters' attention allocation to the different valuation channels is time varying, but unconditionally, sales-related topics receive most of the focus. In the aggregate, attention allocation to sales variables peaks in the run-up to financial crises, whereas costs and earnings variables gain prominence in their aftermath. This shift in focus plausibly underscores a new dimension of market euphoria, wherein financial professionals transition from a narrative-driven approach to valuation—premised on the expectation that sales will eventually translate into free cash flows—to a more grounded perspective where actual cash flows serve as the fundamental basis for determining value. The greater focus on top-line as opposed to bottom-line items also echoes

survey evidence from CEOs and CFOs (Graham, 2022).[2]

We then link topics to their time outlook, and show that in the average report only 14% of narratives focus on past events, while near- and distant-future arguments make up for 53% and 4.5% of narratives, respectively.[3] This analysis highlights the forward-looking nature of analysts' mental models. This is consistent with analysts using a bottom-up approach in valuation, with forward-looking inputs playing a role in their valuation methodologies.

Finally, we show that the perception of specific topics is likely shaped by characteristics unique to the forecasters. For example, individuals from different countries tend to allocate attention to different topics, and are less likely to attribute the same effect to the same topics, suggesting a role for local inference. For a concrete example, forecasters located in regions with high inflation are not only more likely to discuss the impact of inflation on firms located in other countries than other analysts evaluating the same firm, but they also use a more pessimistic tone when discussing the effects of inflation.

Linking narratives to the quantitative forecasts. To deepen our analysis, we then link the mental model components—topics, sentiment, and time outlook—with the corresponding subjective quantitative forecasts within the same report, creating a tight link between forecasters' reasoning and resulting beliefs.

We first find that analysts' focus on specific drivers of earnings or cash flows corresponds to their numerical valuations and beliefs being more sensitive to fluctuations in those factors, when comparing forecasters of the same firm and during the same time period. Studying two key building blocks of earnings—sales and costs—we show that when analysts allocate more attention to sales-related topics in their discussions, their quantitative forecasts are more sensitive to changes in sales. Conversely, increased focus on cost-related variables in narratives is associated with valuations reacting more aggressively to fluctuations in costs.

We next find that narrative sentiment is positively associated with the deviation between analysts' forecasted price target and the market price at the time of the report. This dynamic is predominantly driven by narrative sentiment about the present and future, while sentiment

[2] While it may be tempting to interpret these magnitudes as unambiguous evidence that professionals overlook discount rates in valuation, one aspect to consider is that discount rates are often derived from simple (and thus easier-to-discuss) models such as the CAPM, albeit with subjective adjustments (Decaire, 2023). By contrast, no such streamlined framework exists for cash flow forecasts.

[3] Statements about the present or those that were unclear to the AI account for the remaining share.

from past narratives has a limited and statistically less robust effect on the quantitative forecasts.

Finally, we directly examine the impact of narrative and mental model heterogeneity on belief disagreement by leveraging the granularity of our data and analyzing (randomly) paired forecasters evaluating the same firm at the same time. We find substantial heterogeneity in the topical focus of narratives across paired forecasters, with only 32% of all topics included in at least one equity report overlapping between both. This is consistent with the evidence on sparse mental models discussed above, and indicates a substantial subjective component in forecasters' mental models and narratives. Further, the narrative interpretation of overlapping topics between paired forecasters mediates the extent of their narrative disagreement. Specifically, analyst pairs discussing the same topic agree on its directional impact on firm valuation in 47% of cases. However, when forecaster pairs also agree on both the topic's valuation channel and time horizon, the likelihood of agreeing on the topic's directional effect increases to 74%.

Importantly, the extent of narrative (mis)alignment between forecaster pairs evaluating the same firm at the same time strongly contributes to the extent of their quantitative belief (dis)agreement. For example, relative to the average belief disagreement within forecaster pairs about the firm's twelve-month price target, corresponding to 9% of earnings per share, a one standard deviation reduction in narrative topic similarity increases belief disagreement by 7%.

Taken together, these results provide evidence that forecasters' narratives, and the subjective reasoning they use in the process, directly and systematically affect their resulting quantitative expectations about firms, contributing to (bias-driven) disagreement in quantitative forecasts.

Asset pricing implications: time series. With respect to the relation between narratives and asset pricing implications, one set of analyses focuses on the sentiment of forecasters' narratives. We construct a sentiment measure by equally weighting the sentiment of all extracted individual narratives. Over the full sample period, the correlation between Shiller's cyclically adjusted price-to-earnings (CAPE) ratio and the average narrative share of topics about top-line drivers for U.S. firms is 0.42. Further, CAPE strongly co-moves with our

narrative sentiment measure ($\rho = 0.84$), with most volatility being associated with narratives that discuss near-future outcomes.

The richness of our data then allows us to elucidate the underlying drivers of sentiment. We start by examining its temporal dimensions. When we break down sentiment based on the horizon of each associated narrative—past, present, near-future (1–3 years), distant-future (> 3 years)—we find drastically different time-series patterns. The increased volatility of sentiment associated with near-future narratives compared to that of narratives about the past, present, or distant future, is most pronounced around economic downturns and subsequent recoveries. Further, distant-future sentiment is substantially more persistent than sentiment at other horizons while also being more positive on average.

Asset pricing implications: cross-section. A second set of analyses relates forecasters' narratives to well-known cross-sectional asset pricing characteristics and anomalies. Narrative sentiment and time outlook strongly, and nearly monotonically, predict whether the forecasted firm is a growth or value stock. Consistent with the view that forecasters (mis)perceive growth stocks as having above-average growth potential, narrative sentiment for growth stocks is substantially higher, and the narrative time outlook for growth stocks is significantly more forward-looking. We also find, among other things, a monotone relation between narrative sentiment and momentum stocks. This is consistent with a mechanism where strong past returns lead to increasingly positive narrative sentiment from analysts and forecasters, reinforcing subsequent return continuation.

Literature. To our knowledge, this is the first paper to use LLMs to extract professional forecasters' mental models from text data on a large scale, and simultaneously relate the extracted narratives to their quantitative beliefs in field data ('IBES' data), providing insight into how professional market participants reason and discuss their belief formation. We specifically contribute to the literature in four ways by documenting: (i) mental models as the combination of attention allocation to topics with each their associated sentiment and time outlook; (ii) how mental models vary across industries, regions, and periods; (iii) what individual characteristics shape mental models, and by (iv) drawing economic implications from these novel facts.

Our paper contributes to the growing literature on the properties and heterogeneity of

professional forecasters' valuation practices and forecasting (Andrei et al., 2019; Cookson and Niessner, 2020; Meeuwis et al., 2022; Decaire, 2023; Gormsen and Huber, 2023; Ben-David and Chinco, 2024; Decaire and Graham, 2024; Decaire et al., 2024; Decaire and Guenzel, 2024). We advance this research by extracting the narratives underlying quantitative forecasts, thereby bridging the gap between 'what forecasters believe' and 'why they believe what they do.'

To do so, we build on prior work that seeks to measure people's mental models: the reasoning behind their expectations. A growing body of research has used open-ended and structured survey questions to capture the thought process in people's reasoning in a variety of different fields, including macroeconomics (Andre et al., 2022, Andre et al., 2024a, Binetti et al., 2024, Stantcheva, 2024), and finance (Bailey et al., 2019, Chinco et al., 2022, Chopra and Haaland, 2023, Andre et al., 2024b, Bauer et al., 2024, Laudenbach et al., 2024).[4] We build on these papers, and consider a setting where we are able to link the reasoning of financial professionals to their forecasts, without having to rely on eliciting reasoning and expectations via additional surveys.

Additionally, our paper contributes to the emerging literature at the intersection of AI, machine learning, and finance (Bybee, 2023, Gabaix et al., 2023, Van Binsbergen et al., 2023, Bybee et al., 2024, Charles and Sui, 2024, Chen, 2024, Lopez-Lira and Tang, 2024). We add to this line of research by analyzing the text data in sell-side analysts' equity reports, exploring the heterogeneity therein, and demonstrating that the extracted narratives closely align with their subjective quantitative beliefs.

Related to forecaster disagreement, Decaire and Guenzel (2024) and Decaire et al. (2024) explore various complementary aspects. employ a Latent Dirichlet Allocation (LDA) model to relate topics estimated from analyst reports to their very long-term cash flow growth forecasts. Using this data, they show that disagreement about long-term growth forecasts between analyst pairs is driven more by the differences in the intensity with which each analyst covers overlapping topics (for example, one analyst focusing more on sales while the other emphasizes costs) than by whether both analysts cover the same topics at all. We build on these results, showing that not only does the intensity of topic coverage matter, but

[4] See Haaland et al. (2024) for a recent survey of this literature.

also how these topics are integrated into forecasters' mental models. Decaire et al. (2024) show that disagreement in analysts' expected valuations stems not only from differences in cash flow forecasts but also from the use of different discount rates due to variations in fundamental model implementations (e.g., CAPM). Our evidence offers a different perspective, showing that differences in narrative focus on both discount rate and cash flow-related topics contribute to explaining disagreement in beliefs and expected valuations.

Finally, we add to the ongoing debate about how and why professionals allocate attention to different sources of information. By showing that analysts tend to focus on a limited subset of topics, out of a broad menu of options, our results align closely with theories of sparse mental representation (Gabaix, 2014, Fan et al., 2021, Ba et al., 2022, Bordalo et al., 2023, Bordalo et al., 2024), and those suggesting attention allocation should shift over business cycles (e.g., Kacperczyk et al., 2016). The patterns we document suggest fundamental drivers of value creation, such as customer demand, product development, capital expenditures almost always receiving coverage, whereas others topics are included in analysts' mental models on the basis of salience and surprise, e.g., inflation, pandemics, and wars.

Overall, our paper offers new insights into the mental models and narratives behind financial market participants' quantitative beliefs, thereby shedding light on the inner workings of investors' belief formation process.

Organization of the paper. The rest of the paper proceeds as follows. Section 2 introduces the data and reviews different (ML-based) text analysis methods. Section 3 discusses our LLM-based mental model and narrative extraction method. Section 4 presents aggregate time-series and cross-sectional results. Section 5 presents disaggregate results at the forecaster–report level. Section 6 concludes.

2 Data Collection and (ML-Based) Text Analysis Methods

2.1 Data Collection

We initially downloaded 157,549 equity reports from professional forecasters[5] employed by 46 major equity research firms. We restricted the time window to reports published in the first

[5] The equity report data is taken directly from Decaire and Graham (2024).

quarter of the calendar year (January 1^{st} to April 1^{st}) from 2000 to 2023, to ensure that our data items were systematically measured at a similar time point during the year.

Decaire and Graham (2024) use the same set of equity reports as a starting point for their analysis and provide an in-depth review of the sample properties, including geographic, brokerage house, and industry coverage, as well as firm-specific characteristics. They note minor differences compared to the content of other mainstream analyst databases, such as IBES, but demonstrate that the composition of the present data is comprehensive, representative, and comparable to its commercial counterparts. We refer the reader to Decaire and Graham (2024) for more details.

We start by extracting the text from each report using a series of Python routines. Specifically, in this step, we *exclude* information contained in figures, tables, and the legal boilerplate disclaimers at the end of reports. The objective is instead to capture the valuation thesis that professionals articulate to support their numerical analysis. At last, we do a final pass on the extracted text to address extraction issues that might have been missed by the Python routine by using Meta's LLM, Llama-2-70b.[6] Lastly, for an equity report to be included in the sample, it must contain at least 200 words and no more than 5,000. In cases where an analyst published more than one report on the same firm during those months, we randomly select one publication of that calendar year to avoid duplicates for a given analyst-firm-year pair. This procedure resulted in 81,757 reports, covering 11,308 firms located in 92 countries.

2.2 LLMs Vs. Other Text Analysis Methods: Topic Interpretability Versus Output Robustness

Multiple strategies have been employed in the economic and finance literature to extract parts of individual's mental models, including (i) standard machine learning techniques, particularly Latent Dirichlet Allocation (LDA) models (Decaire and Guenzel, 2024; Bybee et al., 2020), and (ii) human workers (e.g., Andre et al., 2024a). Both of these approaches have benefits and limitations. Employing large language models (LLMs), such as Claude from

[6] For example, in some cases, the extracted text does not contain spaces between words. LLMs are efficient at breaking sentences into individual words.

Anthropic, helps overcome two of the key challenges affecting these previous methodologies, while retaining their key respective benefits.

While LDA models are efficient in identifying topics within documents, the way those models process information often leads to topics that are not easily interpretable, limiting the type of questions researchers can explore with these tools. For instance, in the context of equity reports, the topics produced by LDA models tend to be thematics, broad categories grouping themes related to particular industries, institutions or regions together. For example, Decaire and Guenzel (2024) extract topics from equity reports using a LDA model and provide a detailed list of the topics obtained with this approach as well as the main keywords associated with each topic. Although this type of output is suitable for addressing many research questions, such as those in Decaire and Guenzel (2024), it is less ideal in scenarios where the research question requires intuitive and interpretable topics or when there is a need to establish a clear and intuitive causal relation between a cause and its effect as discussed in the text. However, LDA models offer the advantage of scalability, allowing researchers to perform analyses on a large set of documents, and multiple tests are available to assess LDAs' output quality, allowing researchers to employ statistical tools to benchmark their results.

In contrast, human workers can process texts and apply human interpretation to extract topics that are both intuitive and contextually meaningful within an equity report. Human judgement is also well suited to evaluate and document cause-to-effect relations from textual data, expanding the set of questions that researchers can explore (e.g., Andre et al., 2024a). However, concerns arise regarding the consistency of the output–the likelihood that two workers will identify all and the same topics-especially when the documents under evaluation span dozens of pages and contain complex discussions. Employing human workers generally limits researchers ability to process large quantities of data, typically restricting analyses to smaller subsets of a few hundred documents. Moreover, due to the expense and difficulty of recruiting a large number of human workers for such analyses, it is often not feasible to evaluate the same documents twice, raising concerns about the robustness of the results and the potential for noisy and inconsistent output. The difficulty to easily validate the output of human workers at scale becomes increasingly concerning in cases where the task at hand is complex and involves a great quantity of text to process.

LLMs enable us to generate clear and interpretable topics, akin to those produced by human workers, and to evaluate documents that could be highly challenging even for the most skilled human workers given their length and complexity (i.e., equity reports length ranges from 200 to 5,000 words). Furthermore, LLMs allow for repeated prompting at a relatively low cost, enabling us to perform the analysis at scale (i.e., 82,000 documents in our case), as well as perform additional robustness and validity checks, similar to what is feasible with other machine learning techniques such as LDA models.

3 LLM-Based Extraction of Mental Models and Narratives

3.1 Employed Approach

The ultimate goal of this exercise is to obtain sets of causal chains that combine topics with their associated sentiment and time horizon. Our data collection allows for the possibility that multiple arguments can be used to reference the same topic in different ways. To extract the mental model from the text of equity reports, we employ Claude 3.5 Sonnet from Anthropic,[7] and proceed in three steps.

Steps 1 and 2. In steps 1 and 2, our goal is to extract three categories of variables from equity reports: (i) the mental model variables needed for the analysis, (ii) the categorization of these mental model variables, (iii) a set of words directly taken from the report to perform verification of the output. The AI is tasked with identifying all arguments related to topics discussed in the report that are used to motivate the valuation exercise. The sentiment variable can take four possible values, positive, neutral, negative, or unclear based on the tone and sentiment associated with the topic's argument in the report. Unclear is used in cases where the AI cannot clearly determine a sentiment associated with the context of the report. The time outlook takes five possible values, past, present, near-future (1-3 years), distant-future (3 years and more), and unclear. To assist the AI in determining the time outlook, we also provide the date on which the report was prepared.

Second, we require the AI to report whether the topic's argument is firm-related, industry-related, or macroeconomy-related given the context from the report, as well as whether

[7] We use the *claude-3-5-sonnet-20240620* version.

the topic is associated with top-line items (revenues or sales), costs, profitability margins, bottom-line items (Ebitda, Cash flows, earnings), discount rates, or if it is associated with a relative valuation statement toward the firm peers or historical average. This allows us to gather additional structure to evaluate the mental model of the analyst that produced the report.

Finally, we ask the AI to extract two sets of words directly taken from the text. We first require 1 or 2 keywords that best characterize the sentiment associated with the topic. This allows us to evaluate the associated sentiment with its main reference point. Last, we obtain the first 5 words related to the portion of the text that relates to the topic. This last step allows us to gather the exact position of each topic in the overall document, but to also cross-verify that the topics collected by the AI are plausible given the context of the report.

The distinction between steps 1 and 2 pertains to the quantity and type of information we provide the AI to accomplish its task. In step 1, we break down the equity reports in subsets of 200 words,[8] and we then iterate on each 200-words subsets using the prompt from step 1 (see Appendix A for additional detail). Breaking down large documents into smaller subsets drastically improves the AI's ability to extract a richer and diversified set of topics and improve the consistency of the output, which we discuss in more detail below in Section 3.3. However, some topics might be split between two subsets, introducing the risk that the AI might fail to fully collect them.

Step 2 addresses this concern. Once step 1 is finished, we aggregate all topics that have been identified and ask the AI to review the results, to ensure that step 1 did not miss any relevant topics. At this stage, the AI is provided with the full equity report as well as the full set of results obtained from step 1 (see again Appendix A for additional detail). This ensures that the analysis is also performed on the full report at once. If the AI identifies topics that were missing in the first pass, we add them to our list of results. A benefit of our approach in steps 1 and 2 regarding the selection of topics is that as the AI searches for topics, we do not impose any structure on the type of topics that are collected by the AI, nor do we require the AI to label topics with predetermined standardized labels. However, this lack of

[8] The decision to use 200 words as our cut-off is motivated by the fact that the shortest documents included in our analysis have at least 200 words.

structure makes it difficult to perform systematic analysis and identify patterns in the data.

Step 3. In step 3, we group the topics into 139 standardized labels. Our choice of baseline standardized labels is based on two pieces of information: (i) a set of 200 word-to-vec clustering of the topics, and (ii) financial intuition to ensure that the topics are labeled with terms that are intuitive and economically meaningful in the context of financial valuation (see again Appendix A for additional detail). The complete list of standardized labels is available in Appendix B. Our list of topics features a two-tier design, where we first ask the AI to label the raw topics using a refined list of standardized labels (see Appendix Table B1), and if the AI fails to find a clear and direct match within this list, we then offer a second coarser list (see Appendix Table B2). The intuition for this strategy stems from how analysts structure their discussion in the reports. For example, in some instances an analyst might say "Inflation pressure with increase the cost of good sold," effectively implying a causal link between inflation and costs, but in other cases the analyst might simply state "The cost of goods sold will increase." Thus, in an effort to be as systematic as possible in our data collection, the design of the final prompt step is sufficiently flexible to favor obtaining the more refined statements, while also ensuring that coarser statements are captured as well.

3.2 Report Anatomy: Location of Topics, Sentiment, and Time Outlook

Consistent with the general intuition that a report's introduction and conclusion have a higher content density, we find that a greater number of collected information and topics are collected at the very beginning or the very end of a report (Figure I, Panel A; Figure II, Panel A). At the same time, when studying the narrative arc of the reports, our results indicate that the nature of arguments made at different points of the report progression tend to vary systematically (Figure I, Panels B, C, and D). Broadly, it is possible to break down the average report progression into three subsections: (i) executive summary, (ii) investment thesis, and (iii) key risks.

Professional forecasters tend to first review the previous period's results by highlighting past statements, and provide a broad executive summary. The early part of a report generally emphasizes factors related to the firm itself with a distinct focus on sales and earnings. As the report progresses into the second part, the discussion becomes more forward-looking,

with sales continuing to capture the greatest share of the discussion. This is consistent with survey evidence presented in Graham (2022), finding that managers focus mainly on ensuring high-quality sales forecasts when making investment decisions. In the last part of an average report, forecasters tend to offer potential scenarios that could realize in the near or distant further. These discussions tend to focus on possible downside risks, reflected in the large drop in sentiment toward the end of reports. At that stage, industry and macroeconomic factors become more prevalent as well.

Broadly, these patterns suggest that the arguments associated with a given topic can have a different sentiment and time outlook when discussed at different points in a report. Considering that we are interested in studying professional forecasters' mental models and narratives at a granular level, this highlights the need to ensure that the data collection methodology is able to capture a high level of granularity. The next subsection presents corroborating evidence on this.

3.3 Diagnostic Tools

The recent significant reduction in the cost of performing our analysis using AI, combined with its scalability, allow us to introduce a novel method for evaluating the robustness of the mental model outputs, and establish a baseline for the degree of AI-induced "randomness" in our results. The strategy is executed in two steps. First, we prompt the AI to extract the mental model from a given equity report multiple times (ten attempts), collecting the results of each attempt. In the second step, we calculate the Jaccard similarity score of the extracted topics between each possible pair of the ten attempts for a given report. The Jaccard similarity score is defined as:

$$\text{Jaccard similarity} = \frac{|A \cap B|}{|A \cup B|} \tag{1}$$

where A and B represent the sets of topics extracted in attempts A and B of a given report, $|A \cap B|$ is the number of topics that overlap in both attempts, and $|A \cup B|$ denotes the total number of distinct topics discussed in at least one of the attempts. Essentially, the Jaccard similarity quantifies the overlap between the sets of topics included in both attempts, ranging

from 0 (no overlap) to 1 (complete overlap).

This allows us to assess how sensitive our results are to random perturbations in the LLM. This exercise is akin to asking the same human workers to extract the mental model from the same report multiple times and evaluate the robustness of the output across different attempts. In an ideal scenario, the AI, or the human worker, would systematically output the same mental model.

Figure II presents the results of this test in Panel E, for a stratified random subsample of 240 reports.[9] As the figure reveals, the Jaccard score is close to 1, consistently through the spectrum of different report lengths. This is reassuring, as it confirms that the LLM output is consistent across different attempts.

This exercise also helps highlight a key benefit of breaking down the equity reports into 200-word subsets in the first step of our data collection process. To illustrate how breaking the text in 200-word subset helps improve the reliability of data collection, we repeat the mental model extraction while using the full report in one step. When we provide with the full text at once, instead of using our multi-step approach, the AI tends to behave "*lazily*," where it focuses the data collection on the early part of reports almost exclusively (Figure II, Panel B), compared to our strategy where data is collected persistently throughout the report (Figure II, Panel A). The naive one-step patterns are concerning, considering that the nature of the information discussed at various points in the report tends to be different, both in terms of topics, sentiment, and time outlook, as discussed above. Appendix Table D.1 presents results to confirm the incremental informativeness of the narratives collected at various point throughout reports. Sentiment collected on each quintile of reports help explain analysts price-to-earnings ratio, and when combined, narrative from each subsequent part of reports contain incremental information on the previously parts of reports as shown by the monotonic increase in Within R^2 from Columns 6 to 9.

The resulting output using a naive one-step strategy tends to have lower breadth (i.e., number of topic collected), and lower depth (i.e., number of argument per topics), as shown in Panel D of Figure II. Indeed, when provided with the full text, the AI tends to collect only

[9] We stratify the random subsample based on the length of the reports, by increments of 100 words. We then randomly select 5 reports in each strata. This allows us to study the robustness property of the data extraction approach over the entire range of report lengths available using a well-populated subsample.

one argument per topic. This pattern holds true for documents of all lengths. In contrast, our approach allows us to extract the additional amount of detail discussed in longer reports.

Additionally, the output obtained when directly providing the full text at once tends to be substantially less reliable along two key dimensions. As the number of identified arguments increases, the likelihood that two attempts will generate the same output declines rapidly, to the point of being almost completely random when exceeding 10 arguments (Figure II, Panel F). By contrast, our methodology obtains a consistent output for any number of identified arguments. Further, we find that our multi-step strategy nearly always performs better, in terms of output robustness and Jaccard topic similarity, than directly providing the full text for any level of report length (Figure II, Panel E).

Broadly, this evidence suggests that our data collection strategy generates consistent and robust output, even when the task becomes increasingly complex, as captured by the number of arguments contained in a report or its length.

3.4 Descriptive Versus Predictive Information Processing Using LLMs

LLMs have been used in various ways to extract and produce information from unstructured text. With respect to data generation, it is possible to break down the type of assignments into two categories: (i) descriptive, and (ii) predictive text extraction.

Our exercise is of *descriptive* nature, in that we ask the system to describe the information we present it, instead of requiring it to make use of that information to perform a *prediction*. We also restrict the LLM's temperature to 0 and the top_p to 0, which further limits the AI's ability to provide results that are beyond the information that we provide in the task, or to output results for which the model has lower levels of confidence. At the same time, by requiring the AI to validate its answers by providing direct excerpts from the text, we can ensure that the topics, sentiment, and outlook variables that are extracted from the text that is truly contained in the reports, and not extrapolated by the AI itself based on its training sample. For a concrete example, Appendix C presents one excerpt from our sample provided to the AI, alongside the associated output.

The above considerations and checks are important, considering that the AI training set–all documents used to construct the AI model– might induce look-ahead bias, which can

be a cause for concern when performing prediction (Sarkar and Vafa, 2024). Sarkar and Vafa (2024) design a novel test to investigate whether pre-training data can contaminate output when asking AI to perform predictions, and find instances where such leakage is likely affecting the results of existing studies performing prediction with LLMs. The risk is that the AI might combine the information in the report with the context of its training data, blurring the line between the actual driver of its prediction.

To further mitigate concerns that the topics extracted by the AI were not the result from leakages of its training dataset, and instead truly contained in the documents, we also perform a concrete empirical test. Specifically, we show that 100.00%[10] of the 5-word snippets extracted by the AI are truly contained in the original documents. This indicates that the AI indeed used the context of the provided documents to extract the topics.

4 Aggregate Patterns

In terms of aggregate patterns, we investigate the average characteristics of forecasters' mental models and narratives, i.e., we focus on how narrative topic focus, sentiment, and time outlook averaged across forecasters varies by cross-sectional characteristics, industry, and over the business cycle. Section 5 will then explore these dynamics more granularly at the disaggregated, forecaster–report level.

4.1 Texts and Topics

The average (median) document contains 1,333 (1,026) words, 43 (37) arguments, and 18 (18) distinct topics (Table I). Longer texts tend to contain richer discussions, and include a broader set of topics, than shorter ones (Figure II Panel D). While these features are expected, they also highlight the need to account for this aspect of our data when performing analyses that can be sensitive to these dimensions of the data.

Combining the topics into valuation model inputs, (i) sales command 45% of discussed

[10] In rare instances (1.4% of all snippets), the AI fixed typos, or extracted two non-contiguous pieces of a given sentence to maximize the meaningfulness of the 5-word snippet. For example, in a report for the firm RAND.AS from RBC Capital Market on January 13^{th}, 2011, the AI extracted "*expect further recovery from later-cycle*," to reference the snippet from the text "*expect further recovery from* ***[the]*** *later-cycle*." We view these idiosyncrasies as reasonable in the context of our exercise.

topics, (ii) costs (15%), (iii) profitability margins (10%), (iv) earnings/cash flows (25%), and (v) discount rates (5%). Further, we find that the average (median) report allocates the largest share of the discussion to topics related to top-line items (sales). However, the focus on top-line items varies substantially throughout the sample period, with a low of (44.23%) in 2005 and a high of (57.64%) in 2021 (Figure III, Panel F). In the aggregate, we find that the overall focus on sales for US firms strongly correlates with key valuation metrics such as Shiller's CAPE variable, exhibiting a correlation of 0.42 (Figure IV).

Our data allows also us to evaluate the type of economic entities that are mostly discussed by analysts in their reports, (i) firms (81.69%), (ii) industry (13.99%), or (iii) macroeconomy (5.28%), as shown in Panel E of Figure III, which also displays the variation over time. Most topics discussed in reports are directly related to the firm itself, with industry receiving the second most amount of focus. While the relative importance of each entities is relatively constant throughout our sample, attention to industry and macroeconomic topics command a greater share of discussion following economic and financial crises (e.g., 2008 and 2020). Also, when macroeconomic factors are discussed, this is more frequently to address topics related to discount rates or sales than topics associated with firm's operating performance (Table I, Panel B).

Topics that are related to fundamental drivers of cash flow generation tend to receive a large and persistent share of attention. These include capital expenditures, cost structure and efficiency, customer demand, leverage, dividends, market shares, and mergers and acquisitions (see Figure V, measuring the average amount of attention that the average report in our sample allocates to each topic in a given year). At the same time, we find that some topics capture the discussion at specific moments in time, most strikingly for public health issues (pandemics), but also for topics such as taxation and exchange rates.

Finally, to assess whether topics omitted from textual discussions are outside forecasters' mental models rather than excluded due to space constraints in reports for instance, we conduct two exercises. First, we show that topics not included in a report–but utilized by other analysts evaluating the same firm in the same year–have no explanatory power on a forecaster's subjective valuation (D.2). This strategy allows us to focus on topics that are likely to be relevant to evaluate the firms while avoiding the irrelevant ones, effectively

mitigating the effect of noise in our measurement. Second, pursuing on this line of thought, we focus on the subset of reports that are both in the top 10% in terms of length within our sample. We examine, between pairs of analysts evaluating the same firm-year, whether the number of overlapping topics approaches 100% as the space constraint in reports diminishes. We still find a significant gap between the topics covered in these pairs of reports, suggesting that, even with limited constraints, analysts do not end up focusing on the same set of topics (Table I, Panel C).

Industry-specific patterns. Breaking down attention allocation further across major industries, we highlight additional strikingly different patterns. For industries that are highly regulated, and mostly mature, such as utilities or the finance and insurance industry, forecasters tend to allocate a large share of their attention to the (i) regulation environment, (ii) payoffs, (iii) capital structure, (iv) cost management, (v) pricing, and merger & acquisitions (Appendix Figures D.1 and D.2). This highlights the likely key drivers of value creation for shareholders when firms face limited growth potential and low competition. Dynamic and innovative industries, such as the information sector or professional and scientific services, tend to focus more aggressively on (i) capital investment, (ii) customer acquisition, (iii) market share, (iv) product and service mix, and (v) clinical trials (Appendix Figures D.3 and D.4). When discussing resource-dependent sectors, such as the oil & gas and mining industries, forecasters focus on (i) commodity and raw material prices, (ii) oil & gas prices, and (iii) production capacity, which in this context as to do with a firm's ability to secure natural resource reserves (Appendix Figure D.5). Firms evolving in industries in which customer acquisition is crucial, and firms facing physical local and geographic constraints to sustain their growth, either due to production challenges or consumer demand, such as manufacturing and retail trades, receive substantial emphasis on their ability to generate sales in their (i) multiple locations, expand their footprint, (ii) maintain their market shares, (iii) introduce new product and services, and (iv) manage customers demand (Appendix Figures D.6 and D.7). The general pattern of low margin and high volume strategy observed with most major retailers prompt forecasters to pay particular attention to (v) cost structure and efficiency (Appendix Figure D.7). For industry where supply is relatively fixed in the short and medium run, and where most of fluctuation in sales comes from variation in consumer

demand or pricing, such as the transportation industry, firms' ability to (i) stimulate demand, (ii) efficiently control their costs and prices, as well as (iii) optimizing their capital structure are key factors forecasters focus on (Appendix Figure D.8). Here, it also appears reasonable that forecasters pay significant attention to regulation and compliance issues, considering that airlines are particularly exposed to regulation changes in their activities.

Additionally, while the above patterns are industry-specific, some major events, such as 2008 financial crisis or the 2020 pandemic, tend to capture attention and affect narratives across all industries, as shown by the the economic recession, economic recovery, and pandemic and public heath issue patterns we document in all of those figures.

Firm, industry, and macro topics. Next, we explore the top 20 topics discussed when forecasters refer to the three entities evaluated in our sample–firm, industry, and macroeconomy (Appendix Figure D.9). The breakdown of topics associated with each entity is intuitive. Topics that are generally associated with macroeconomic factors, such as commodity prices or economic recession, tend to be clustered with the macroeconomy entity, whereas topics that are related to product market competition and supply chain are mainly discussed with the industry entity. Finally, topics related to the firm entity tend to be related to topics that are under a firm's control, such as dividend policy and capital expenditure. Broadly, these topic clusters are reassuring as it suggests that our data collection strategy associates topics to their logical entity.

This cut of the data also allows us to evaluate the key topics within each entity, and the way topics are discussed in reports (i.e., whether the key dynamic associated with each topic is at the firm, industry, or macro level). For topics articulated around firm actions and dynamics, we find that (i) capital expenditure, (ii) cost structure, (iii) discount rate, (iv) dividend policies, (v) merger & acquisition, and (vi) product and service mix appear to be the main areas that forecasters believe to be mainly under firms' control. These are among the main pillars that practitioners, business school curricula, and researchers emphasize when thinking about firms' ability to create value. In terms of transient topics, discussions about restructuring increase after the burst of the dot.com bubble, discussions about cash and liquidity gain traction in times of market distress, and pricing power receives more focus when inflation becomes more salient, such as during 2022–2024.

At the industry level, (i) customer demand, (ii) market size, (iii) new entrants, (iv) pricing pressure, and (v) regulatory environment are the key factors that receive systematic attention from forecasters. This selection of topics squares with how industry threat and opportunity analysis is taught in business schools. For the transient topics, while supply chain discussion are generally overlooked, following the 2020 pandemic, forecasters have increasingly discussed possible implications of supply chain disruption.

At last, macroeconomic topics are mainly focused on the role of (i) foreign exchanges and currencies (ii) interest rates, (iii) gross domestic product, (iv) commodities, and (v) oil & gas prices. When it comes to transient topics, the patterns are striking. For some of the topic that generally receive limited attention, we note significant spikes in attention at key points in times, including: (i) the 2018 Trump tax cut (taxation), (ii) the 2020 pandemic, (iii) the significant inflation that started in 2020, (iv) the 2008 economic recession, and (v) the economic recovery that followed.

4.2 Sentiment and Time Outlook

We next turn to the relationship between narrative sentiment and time outlook, i.e., we examine the temporal drivers of sentiment. As shown in Figure VI, near-future narratives are associated with the largest volatility in sentiment over time, whereas distant-future sentiment is both more persistent and more positive than sentiment at any other time horizon. The differences in the degree of volatility in narrative sentiment across the different horizons are most pronounced in the lead-up to and aftermath of crises, such as the Great Recession.

The average report sentiment strongly covaries with Shiller's CAPE measure, with a correlation of 0.84 (Panel B of Figure III). Refining the analysis of this pattern, Appendix Figure D.10 overlays Shiller's CAPE ratio onto the time-series plots in Figure VI to breakdown the sentiment associated with different levels of time outlook at their relation with the Shiller's CAPE. CAPE most strongly covaries with near-future narrative sentiment, though there is a clearly visible and meaningful positive correlation between CAPE and narrative sentiment across all time horizons. Thus, combined with the evidence from Figure IV, our evidence suggests that key aspects of forecasters' narratives evolve over time in ways that are mirrored in valuation ratios.

Sentiment and entities, industries, geographies, and topics. To further examine the drivers of narrative sentiment, we next explore the variation by discussed entities and across industries and geographies, as well as the association between sentiment and underlying topics. Appendix Figure D.11 shows that narrative sentiment over time is remarkably similar irrespective of whether focusing on firm, industry, and macro topics. In other words, when sentiment is high, it tends to be high for both firm-specific and general economic aspects, and the same holds true when sentiment is low. One possible exception is the end of the sample period, where narrative firm and industry sentiment have recovered after the COVID pandemic, whereas narrative macro sentiment has further deteriorated.

Appendix Figure D.12 shows variation in narrative sentiment by industry. In line with priors, some industries are particularly pro-cyclical sentiment-wise, such as the transportation and finance industries (especially during the Great Recession), whereas others are less cyclical, such as Utilities.

Narrative sentiment is similar over time across geographies, as shown in Appendix Figure D.13. There are clear troughs in narrative sentiment during the Great Recession across the globe. Interestingly, during the COVID crisis, it appears that narrative sentiment in the U.S. and Canada was somewhat more stable than in other parts of the world.

Finally, Appendix Figure D.14 reveals which topics are typically associated the positive or negative sentiment, respectively. The most negative-sentiment inducing topics include wars and recessions, whereas the most positive-sentiment inducing topics include market expansion and automation. The latter is interesting in and of itself, showing that the advent of AI and automation enters the average financial professional's narratives and mental model positively, rather than negatively (e.g., due to concerns about automation leading to job obsolescence).

Outlook and industries, geographies, and topics. Additionally, we break down the average time outlook of narratives by industry, geography, and topic. Appendix Figure D.15 reveals non-trivial variation in trends in narrative time outlook across industries. Sectors such as Transportation, Professional Services, and Manufacturing have become more forward-looking over time, in terms of forecasters' narrative discussions, whereas there is little or nor change in terms of the past-versus-future narrative focus in Finance and Mining.

In Appendix Figure D.16, we further find that the average narrative has become more

forward-looking over the sample period in the U.S. and Canada, and the same is true to a lesser degree in Europe. By contrast, the average narrative time outlook has remained relatively steady in Asia and Oceania, especially at lower frequencies.

In terms of underlying topics that drive the variation in narrative outlook, natural disasters enter discussions mostly through narratives about the past, whereas energy transition and demographic changes are predominantly included as narratives about the future (Appendix Figure D.17). These results provide novel insights into which broad economic trends and risks market participants attend to in their mental models retroactively and only after they occur, and which risks they account for proactively.

4.3 Sentiment and Time Outlook in the Cross-Section

We now turn to the narrative implications for key cross-sectional asset pricing patterns.

Sentiment. Narrative sentiment is roughly U-shaped in firms' market beta (Figure VII, Panel A). This pattern is consistent with firms at both low and high ends of market risk exposure eliciting more positive emotional responses from forecasters. For example, forecasters may highlight the low-risk nature of utility firms while simultaneously emphasizing the upside return potential of high-risk technology firms.

In Panel B, we find a strikingly monotone link between narrative sentiment and book-to-market ratio. Narrative sentiment is substantially higher for firms with low book-to-market ratios, i.e., for growth or "glamour" stocks, consistent with the idea that investors (mis)perceive these stocks as having above-average growth potential.

Similarly, in Panel C, we find a monotone relation between narrative sentiment and momentum stocks. This is consistent with a dynamic where strong past returns trigger increasingly positive narrative sentiment, driving subsequent return continuation.

Finally, in Panels D and E, we find consistent strong relationships between narrative sentiment and firm size and age, respectively. Narrative sentiment about younger (and thus likely smaller) firms is systematically more positive than that about more established firms, consistent with an average mental model across forecasters that places greater weight on younger firms offering more attractive growth opportunities, despite their higher risks.

Time Outlook. Figure VIII shows the associations between the same set of cross-sectional characteristics from Figure VII and time outlook. Cross-sectional characteristics linked to a stronger focus on narratives about the future include low or high market beta stocks, low book-to-market stocks, large stocks, and young firms. Conversely, characteristics linked to a stronger focus on narratives about the past include medium beta stocks, high book-to-market stocks, small stock, and old firms.[11] Broadly, this is consistent with an average mental model that emphasizes growth potential and innovation for younger, riskier firms, while viewing mature firms as being increasingly tied to historical performance.

5 Disaggregate Patterns

To refine and extend the broad conclusions derived from the analysis of aggregate patterns, we conduct our next set of analyses at the report level.

5.1 Linking Narrative Sentiment and Time Outlook to Quantitative Forecasts

Table II examines how narrative sentiment and time outlook relate to subjective quantitative beliefs, namely the subjective expectation about stock's price appreciation over the next twelve months. The analysis is disaggregated, with the unit of observation at the level of forecasting firm (i.e., contributor), forecasted firm, and year. Panel A shows, in columns 1 to 6, that both sentiment and time outlook positively predict subjective quantitative beliefs, i.e., that more positive and future-oriented discussions are associated with higher level of subjectively expected price appreciation. This holds both in baseline regressions with sentiment and outlook included individually, when adding both variables jointly, and in regressions that exploit within firm-by-year variation only. In columns 7 and 8, we investigate different sentiment time horizons, i.e., the interaction of sentiment and time outlook. We note the final two columns show normalized coefficients for ease of comparison. As these columns reveal, present-based, near-future, and distant-future sentiment all strongly and positively predict higher subjective beliefs, whereas the evidence for past-oriented sentiment

[11] The association between time outlook and momentum is more scattered, possibly reflecting that momentum captures both recent past performance and beliefs about future trends, making its relationship with time outlook less straightforward.

is mixed, both economically and statistically speaking.

Panel B examines the correlations of sentiment at different horizons. There is a relatively strong correlation between present, near-term, and distant future sentiment. At the same time, sentiment about the past shows a markedly weaker connection to sentiment at other horizons, especially with distant-future sentiment. This suggests that beliefs about the future are not merely derived from past events in a straightforward, mechanical way.

Overall, the results in Table II suggest that subjective sentiment and time outlook in forecasters' narratives play a significant role in shaping their subsequent quantitative beliefs.

5.2 Linking Valuation Channels to Quantitative Forecasts

We then investigate whether attention allocation to specific valuation channels impacts how sensitive quantitative expectations are to changes in the underlying channel. This exercise allows us to directly link the narratives we measure with the LLM from reports discussions with the quantitative prediction generated by the analysts. Precisely, we evaluate the following relation:

$$\text{Price target}_0^{12-month} = \frac{E_1}{(1+r)} + \frac{E_2}{(1+r)^2} + ... \tag{2}$$

$$\text{Price targe}_0^{12-month} = E_0 \left[\frac{(1+g_1)}{(1+r)} + \frac{(1+g_1)(1+g_2)}{(1+r)^2} + ... \right. \tag{3}$$

$$\frac{\text{Price targe}_0^{12-month}}{E_{-1}} = \frac{E_0}{E_{-1}} * \left[\frac{(1+g_1)}{(1+r)} + \frac{(1+g_1)(1+g_2)}{(1+r)^2} + ... \right. \tag{4}$$

$$\ln\left(\frac{\text{Price targe}_0^{12-month}}{E_{-1}}\right) = \ln\left(\frac{E_0}{E_{-1}}\right) + fct(g, r) \tag{5}$$

where, $\ln\left(\frac{E_0}{E_{-1}}\right)$ denotes the change in earnings that the analyst just observed over the previous period. We then focus our analysis on two potential drivers of recent earnings growth: sales growth and cost growth. This allows us to break down the drivers of recent earnings growth into their primitives, while also offering us a way to directly link the nature of attention allocation in reports. The underlying assumption of this strategy is that professional forecasters who focus more heavily on costs, instead of sales, might ultimately adjust their valuation more aggressively in response to changes in a firm's costs than in its sales. Focusing

on these two valuation drivers offers two benefits. First, sales growth and cost growth over the previous period are readily available at the time of producing the reports, which ensures that analysts have access to this data when determining their quantitative beliefs. Moreover, these quantities are also observable to the econometrician, therefore facilitating measurement. Second, the *direct* effect of sales and cost growth is intuitive: positive growth in sales should be positively associated with earnings growth and the price target, whereas cost increases should be negatively associated with earnings and the price target, all else equal. Broadly, this offers a simple setting to investigate how changes in analysts' focus on these key drivers are associated with how sensitive their subsequent valuation exercise becomes to changes in those variables.

Panels A and B of Table III report the results of this test for both the sales and cost channels, respectively. We measure attention allocation to a valuation channel as the share of a report that relates to that specific channel, with 32% of the average report being related to sales and 10% to costs (Table I, Panel B). First, in Panels A and B of Table III, we find that recent sales growth is positively associated with the price target, whereas recent cost increases have a negative relation, consistent with economic intuition. Then, focusing on attention allocation to sales, we find that greater attention allocation to the sales channel is associated with analysts' price targets being increasingly sensitive to variation in sales (i.e., the relation becomes increasingly positive). Columns 1 to 4 gradually introduce increasingly granular fixed effects. In our favored specification, Column 4, we include contributor interacted with firm fixed effects. This allows us to study how changes in the intensity of attention allocation for a given analyst evaluating a specific firm are associated with their price target sensitivity over time. In Panel B of Table III, we repeat the exercise for attention allocation to the cost channel and find a similar interpretation: greater attention allocation to the cost channel is associated with analysts' price targets being increasingly sensitive to variation in costs (i.e., the relation becomes increasingly negative).

Broadly, the results presented in this table suggest a direct link between the measure of attention allocation obtained from our data collection and how professional forecasters form their quantitative beliefs. A greater focus on particular valuation drivers is associated with their valuation outcomes being more sensitive to fluctuations in those factors.

5.3 Subjective Narrative Attention Allocation and Forecast (Dis-)Agreement

Topics and Subjective Directional Effects. Next, we investigate how differences in attention allocation impact the likelihood that two forecasters agree on the directional effect that a topic will have on valuations. Forecasting and valuation are inherently complex, multi-dimensional problems, where a given topic can have multiple potential effects on valuation outcomes. For example, inflation can negatively impact a firm's costs, by increasing its procurement costs, but also improve its sales, by allowing it to sell its product at higher price. Topics can also have different effects on firm valuation at different horizons. For instance, firms might not be able to contain their cost in the short-run efficiently, but might be able to adjust their sales in the longer horizon, effectively mitigating the effect of inflation. The allocation of attention to different aspects of these topics can significantly shape how forecasters interpret their effects on firm valuation. Panel A of Appendix Figure D.18 illustrate this intuition, and shows how inflation, depending on the channel analysts pinpoint, can be associated with different directional effects on average[12]. Further, Table IV confirms this pattern at the within-report level, such that in the average report, 15% of topics are discussed in association with more than one valuation channel (i.e., sales, costs, etc.), and that in more than 23% of cases, a topic is related to more than one sentiment (e.g., inflation will *positively* impact valuation through a firm's sales and *negatively* through its costs). As a result, while analysts might discuss the same topic, their interpretation of its impact may vary depending on which dimension of the issue they perceive as most critical, reflecting differences in their underlying mental models.

The effect of topic on valuation can be decomposed into a (i) direction, such as positive or negative (i.e., sentiment), and (ii) a magnitude. We focus our analysis on the directional effect to maximize emphasis on forecasters' "fundamental thinking" about the topic, examining analyst pairs that forecast the same firm at the same firm, and restricting the sample to cases where both analysts discuss the same topic (e.g., Inflation and CPI).

Table V presents the results. The likelihood that forecasters agree on the directional effect of a topic on the valuation outcome is strongly related to whether professionals focus on the

[12] Panels B-E of Appendix Figure D.18 present word cloud of the keywords used by analysts when discussing inflation for each drivers of cash-flows (earnings).

same valuation channel (e.g., both discuss the effect of inflation on sales), and whether they both evaluate the effect of the topic at the same time horizon (e.g., both discuss the effect over the near-future). Precisely, focusing on the same outlook (valuation channel) increases the likelihood of interpreting the directional effect of the topic in the same way by 6% (11%), relative to a sample average of 47 percentage points (columns 1 and 2 of Panel A).

To refine the analysis, and to consider the fact that most topics in our sample are broad–two analysts discussing inflation can focus on different aspect of inflation–we redo the analysis focusing on instances where both analysts discuss the same argument. In this case, focusing on the same outlook (valuation channel) increases the likelihood to share the same directional sentiment about the effect of the argument on the valuation outcome increases by 7% (6%), relative to a sample average of 71 percentage points (columns 3 and 4 of Panel A). As a final exercise, and to ensure that these patterns are not driven by differences in report lengths, we repeat the analysis with documents that are of similar length (Panel B). Our results are unchanged.

Broadly, these findings suggest that the individual ways in which forecasters incorporate information and new signals into their valuation exercises play an important role. This indicates that attention allocation and how information is integrated into mental models are significant factors in explaining levels of belief agreement or disagreement.

Quantitative Belief Disagreement. Finally, we investigate whether the topics discussed in reports, and the heterogeneity in narratives across forecasters, are related to forecasters' subjective quantitative beliefs and belief disagreement, performing two tests.

The first test reveals that individuals who allocate attention to different topics tend to form substantially different quantitative beliefs and, as a result, significantly disagree about expected stock price levels. Column 1 from Panel A of Table VI reports the results of this test with no additional controls and fixed effects. Intuitively, as the share of topics discussed in two reports increases, the gap between the price targets issued by the respective analysts, i.e., their belief disagreement, declines. Broadly, a one-standard deviation increase in the similarity between two reports, the Jaccard similarity, is associated with a 7%[13] reduction in

[13] We obtain this statistic by multiplying the coefficient in Column 1 (−0.06) with the Jaccard similarity standard deviation (0.11), and dividing the value by the average valuation disagreement (0.09).

the difference between both reports price target.

To investigate the robustness of this result, columns 2 to 4 gradually include firm, year, and contributor$_A$*contributor$_B$ fixed effects. The firm fixed effects allow for a within-firm analysis, accounting for different characteristics such as harder-to-value firms, ensuring that firm-specific time-invariant characteristics do not drive out results. The year fixed effect enables us to control for the fact that in some years, global shocks might affect how analysts allocate to topics. Finally, the contributor$_A$*contributor$_B$ allows us to track how changes in topic overlap between pairs of analysts is associated with changes in their valuation disagreement. We find economically and statistically robust and corroborating results throughout.

The second variable of interest, *Same Argument* enables us to refine the analysis further and investigate whether similarities in the particular way information is incorporated in analysts' mental models also plays a role in their likelihood to converge on their quantitative beliefs. Precisely, *Same Argument* corresponds to the proportion of overlapping topics for which analysts share the same sentiment-outlook-valuation channel. A value of 0 indicates that no pair of overlapping topics share those three characteristics, and a value of 1 means that all overlapping topics are discussed in the exact same manner in the reports. Columns 3 and 4 add this variable of interest to the regression. The results reveal that the way information in incorporated in reports indeed plays an important incremental role in explaining the extent to which analysts converge on their quantitative beliefs.

As a final robustness exercise, columns 5 to 8 redo the analysis only for the subset of report pairs for which the number of words in both reports is within a 20% margin.Our conclusions are again unaffected.

In Panel B of Table VI, we explore how attention allocation to topics, and the way information is incorporated through mental models in the quantitative beliefs, explains co-movement in expected price appreciation for pairs of analysts. Our results show that, as analysts' discussed attention allocation becomes more similar, the co-movement between their subjective expectations increases. In other words, they increasingly respond to the same signals. Column 1 present the results with no fixed effects and controls, and column 2 introduces firm, year, and contributor pair fixed effects. Our conclusions are similar in both cases. Columns 3 and 4 also include the *Same argument* variable and show that the

way overlapping topics are incorporated in the analysis is an additional important driver in explaining the extent to which quantitative beliefs comove. Columns 5 to 8 again repeat the analysis only for the subset of report pairs for which the number of words in both report is within a 20% margin. As above, our conclusions remain the same.

Broadly, the results from Table VI suggest that differential attention to topics, as well as the way topics are subjectively processed and incorporated into forecasters' mental models, play an important role in shaping professionals quantitative beliefs and the co-movement of beliefs across forecasters.

Drivers of Attention Allocation. To understand why analysts allocate attention to different topics and ultimately derive different valuations, we first explore the factors that drive analyst attention and then study whether analyst-specific characteristics help explain differences in focus.

In our first test, we examines the drivers of attention allocation by focusing on inflation, an ex-ante (and ex-post) crucial economic variable that has experienced varying salience over time. Naturally, to understand why different forecasters attend differentially to the same signals, and how signals are differentially and subjectively processed, it is important to understand the broad mechanisms that lead forecasters to incorporate economic developments into their discussions and narratives. As a first step in this direction, Panel A of Table VII shows that as inflation increases, forecasters allocate more narrative attention to inflation-related topics. This effect is concentrated in cases when inflation rises above the 2% target, and vanishes when inflation is below that target. In other words, when inflation "exogenously" becomes a more salient economic concern, forecasters react to this and adjust their narratives accordingly. Broadly, this highlights the responsiveness of forecasters' narratives to economic signals.

We then study how analyst-specific characteristics, such as their location, impact how they allocate attention and process information. Under the framework of local inference, an individual's own experience should influence the features that are top of mind. For example, an individual experiencing episodes of high inflation might be more sensitive to the effect of inflation on a firm's operations than someone without such experiences. To formally test this hypothesis, we begin broadly by showing that a pair of analysts located in the same country

are more likely to allocate attention to the same topics and share the same sentiment about how overlapping topics impact valuations (Table VIII).

To deepen our understanding of how local inference impacts information processing, we refine our analysis of attention allocation to inflation by incorporating information about the inflation situation in the analyst's country (Panel A of Table VIII). In this test (Columns 1 to 3), we focus on cases where analysts are located in a different country than the firm's headquarters to minimize risks of collinearity between the two variables of interest, $\text{Inflation}^{\text{Analyst Country}}c,t$ and $\text{Inflation}^{\text{Firm HQ Country}}c,t$. We first show that attention allocation to inflation in reports increases with inflation levels in both the analyst's and the firm's countries. To evaluate how heightened focus on inflation translates into how inflation is incorporated into valuation models, we then study the relationship between inflation levels and the sentiment associated with inflation when it is discussed in reports (Columns 4 and 5). First, we find that analysts do not appear to directly use the inflation level from their own location to evaluate how inflation impacts firm valuation. This suggests that they do not mistakenly use information that is likely unrelated to the firm's operations. Instead, we find that an analyst's inflation experience influences how they evaluate the inflation affecting the firm. Analysts who experience above-target inflation levels tend to be more pessimistic about the effect of inflation on firm valuation.

Broadly, the evidence presented in this section indicates that analysts' attention is shaped by factors that are important to the firm but also by factors related to their own experiences. We also find that the way their own experiences shape their beliefs does not appear to involve confusing which pieces of information should be included in their valuation models. Instead, their experiences influence how they evaluate information related to the firm.

6 Conclusion

This paper measures the reasoning behind professional forecasters' quantitative expectations. We do so by extracting the chain of thoughts that analysts use to justify their valuations from more than 82,000 equity analyst reports.

Our first contribution is methodological, and allows us to overcome the challenges of

gathering information from large bodies of text data. Specifically, while recent advancements in LLMs have made it possible to extract information from text, the output is often incomplete and unstable. To overcome these issues, we depart from a naive single-step prompt. Instead, we use a multi-step approach, and develop a diagnostic tool akin to bootstrapping to ensure a more comprehensive and stable extraction of analysts' narratives.

We then turn to four sets of analysis. First, in terms of aggregate patterns, we find that while attention to different valuation channels is time-varying, forecasters predominantly focus on sales-related topics and near-future outcomes. Second, turning to the time series implications of financial professionals' reasonings, we construct a measure of narrative sentiment, and find that it strongly correlates with Shiller's price-to-earnings ratio, with most volatility being associated with narratives that discuss near-future outcomes. Third, we explore the cross-sectional asset pricing implications. Among other things, we find that the reasonings used in the valuation of growth stocks have a higher sentiment and are associated with more forward-looking arguments than the reasonings used in the valuation of value stocks. Similarly, the reasonings that characterize the valuation of small stocks and momentum stocks are also associated with higher sentiment than those for large and contrarian stocks. Finally, we show that heterogeneity in reasoning systematically predicts the heterogeneity in professionals' quantitative forecasts.

Our findings provide direct evidence from the field that a seemingly homogeneous set of financial market participants, with access to similar detailed public information and sophisticated tools, not only report different quantitative beliefs about firms and financial markets, but also use differing reasonings to justify these beliefs. As such, our findings signal the importance of asset pricing and macro finance models that incorporate heterogeneous subjective views of how the (financial) world operates.

References

Adam, K., and S. Nagel. 2023. Expectations data in asset pricing. In *Handbook of economic expectations*, pp. 477–506. Elsevier.

Amromin, G., and S. A. Sharpe. 2014. From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science* 60:845–866.

Andre, P., I. Haaland, C. Roth, M. Wiederholt, and J. Wohlfart. 2024a. Narratives about the Macroeconomy. Tech. rep., SAFE Working Paper.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart. 2022. Subjective models of the macroeconomy: Evidence from experts and representative samples. *The Review of Economic Studies* 89:2958–2991.

Andre, P., P. Schirmer, and J. Wohlfart. 2024b. Mental Models of the Stock Market. Tech. rep., Working Paper.

Andrei, D., B. Carlin, and M. Hasler. 2019. Asset Pricing with Disagreement and Uncertainty about the Length of Business Cycles. *Management Science* 65:2900–2923.

Ba, C., J. A. Bohren, and A. Imas. 2022. Over-and underreaction to information. *Available at SSRN 4274617* .

Bacchetta, P., E. Mertens, and E. Van Wincoop. 2009. Predictability in financial markets: What do survey expectations tell us? *Journal of International Money and Finance* 28:406–426.

Bailey, M., E. Dávila, T. Kuchler, and J. Stroebel. 2019. House price beliefs and mortgage leverage choice. *The Review of Economic Studies* 86:2403–2452.

Bauer, R., K. Gödker, P. Smeets, and F. Zimmermann. 2024. Mental Models in Financial Markets: How Do Experts Reason about the Pricing of Climate Risk? *European Corporate Governance Institute–Finance Working Paper* .

Ben-David, I., and A. Chinco. 2024. Expected EPS× Trailing P/E. Tech. rep., National Bureau of Economic Research.

Binetti, A., F. Nuzzi, and S. Stantcheva. 2024. People's Understanding of Inflation. Tech. rep., National Bureau of Economic Research.

Bordalo, P., J. J. Conlon, N. Gennaioli, S. Y. Kwon, and A. Shleifer. 2023. How people use statistics. Tech. rep., National Bureau of Economic Research.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer. 2020a. *Expectations of fundamentals and stock market puzzles*. National Bureau of Economic Research.

Bordalo, P., N. Gennaioli, G. Lanzani, and A. Shleifer. 2024. A Cognitive Theory of Reasoning and Choice .

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer. 2020b. Overreaction in macroeconomic expectations. *American Economic Review* 110:2748–2782.

Bybee, J. L. 2023. The ghost in the machine: Generating beliefs with large language models. *Working Paper* .

Bybee, L., B. Kelly, A. Manela, and D. Xiu. 2024. Business news and business cycles. *The Journal of Finance* 79:3105–3147.

Bybee, L., B. T. Kelly, A. Manela, and D. Xiu. 2020. The structure of economic news. Tech. rep., National Bureau of Economic Research.

Case, K. E., R. J. Shiller, and A. Thompson. 2012. What have they been thinking? Home buyer behavior in hot and cold markets. Tech. rep., National Bureau of Economic Research.

Charles, C., and P. Sui. 2024. Marketwide Memory. *Working paper* .

Chen, Z. 2024. Memory and Beliefs in Financial Markets: A Machine Learning Approach. *Working Paper* .

Chinco, A., S. M. Hartzmark, and A. B. Sussman. 2022. A new test of risk factor relevance. *The Journal of Finance* 77:2183–2238.

Chopra, F., and I. Haaland. 2023. Conducting qualitative interviews with AI .

Cookson, J. A., and M. Niessner. 2020. Why Don't We Agree? Evidence from a Social Network of Investors. *American Economic Review* 111:173–228.

De Bondt, W. F., and R. H. Thaler. 1990. Do security analysts overreact? *The American economic review* pp. 52–57.

Decaire, P., and M. Guenzel. 2024. What Drives Very Long-Run Cash Flow Expectations? *Working Paper* .

Decaire, P. H. 2023. Discount Rate Uncertainty and Capital Investment. *Working paper* .

Decaire, P. H., and J. Graham. 2024. Valuation Fundamentals. *Working paper* .

Decaire, P. H., D. Sosyura, and M. Wittry. 2024. Resolving Estimation Ambiguity. *Working paper* .

Fan, T., Y. Liang, and C. Peng. 2021. The inference-forecast gap in belief updating. *Unpubished Manuscript* .

Gabaix, X. 2014. A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics* 129:1661–1710.

Gabaix, X., R. S. Koijen, R. Richmond, and M. Yogo. 2023. Asset embeddings. *Available at SSRN 4507511* .

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus. 2021. Five facts about beliefs and portfolios. *American Economic Review* 111:1481–1522.

Gormsen, N. J., and K. Huber. 2023. Corporate discount rates. Tech. rep., National Bureau of Economic Research.

Graham, J. R. 2022. Presidential address: Corporate finance and reality. *The Journal of Finance* 77:1975–2049.

Greenwood, R., and A. Shleifer. 2014. Expectations of returns and expected returns. *The Review of Financial Studies* 27:714–746.

Haaland, I. K., C. Roth, S. Stantcheva, and J. Wohlfart. 2024. Measuring what is top of mind. Tech. rep., National Bureau of Economic Research.

Kacperczyk, M. T., S. van Nieuwerburgh, and L. Veldkamp. 2016. A Rational Theory of Mutual Funds' Attention Allocation. *Econometrica* 84:571–626.

Laudenbach, C., A. Weber, J. Wohlfart, and R. Weber. 2024. Beliefs about the Stock Market and Investment Choices: Evidence from a Field Experiment. *Review of Financial Studies* .

Lopez-Lira, A., and Y. Tang. 2024. Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models. *Working paper* .

Malmendier, U., and S. Nagel. 2011. Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics* 126:373–416.

Meeuwis, M., J. A. Parker, A. Schoar, and D. I. Simester. 2022. Belief Disagreement and Portfolio Choice. *Journal of Finance* 77:3191–3247.

Nagel, S., and Z. Xu. 2022. Asset pricing with fading memory. *The Review of Financial Studies* 35:2190–2245.

Sarkar, S. K., and K. Vafa. 2024. Lookahead Bias in Pretrained Language Models. Tech. rep., Harvard Business School, Working Paper.

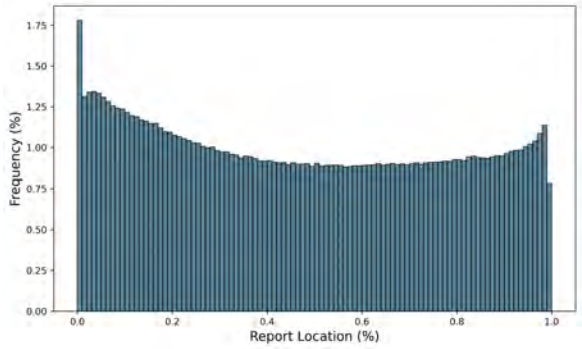
Stantcheva, S. 2024. Why do we dislike inflation? Tech. rep., National Bureau of Economic Research.

Van Binsbergen, J. H., X. Han, and A. Lopez-Lira. 2023. Man versus machine learning: The term structure of earnings expectations and conditional biases. *The Review of Financial Studies* 36:2361–2396.

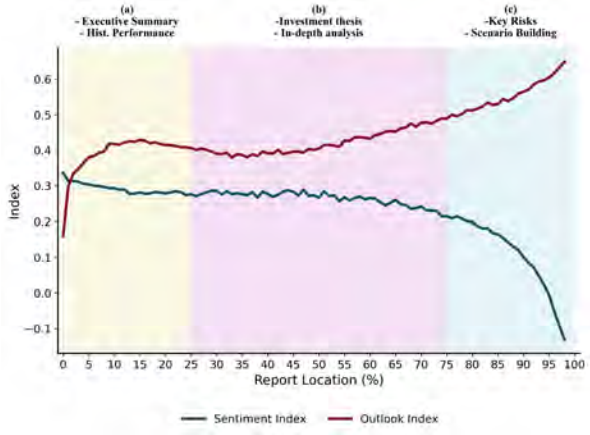
Vissing-Jorgensen, A. 2003. Perspectives on behavioral finance: Does" irrationality" disappear with wealth? Evidence from expectations and actions. *NBER macroeconomics annual* 18:139–194.

Figures and Tables

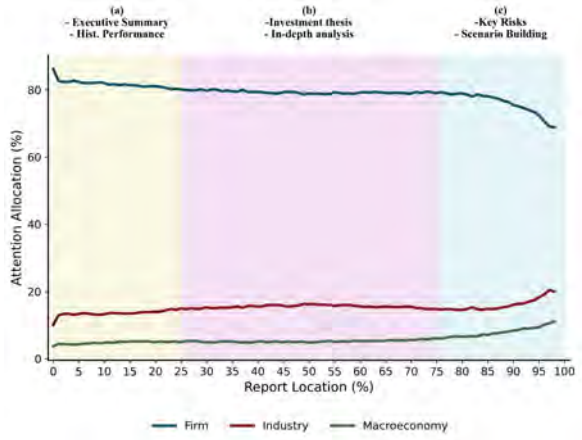
(A) Where Does the AI Collect the Information?.



(B) Narrative Arc - Sentiment and Outlook.



(C) Narrative Arc - Entities.



(D) Narrative Arc - Valuation Drivers.

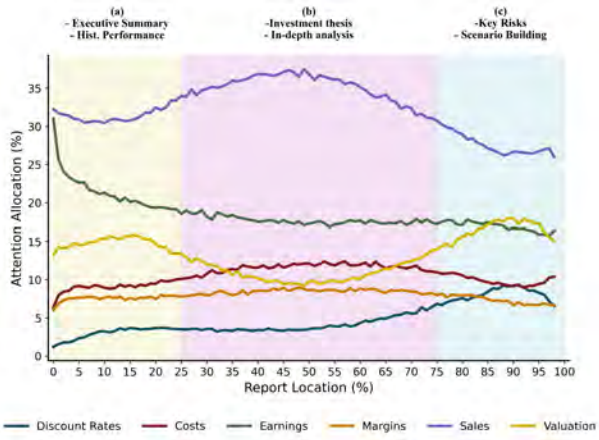
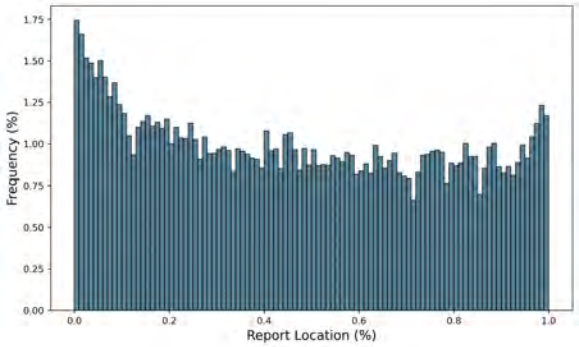


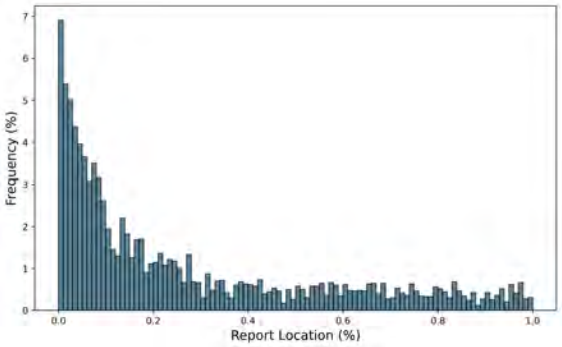
Figure I:
Report Narrative Arc

This figure shows where the information is collected in the average report, as well as the type of information that is collected in each part. The x-axis for all 4 panels denotes the part of the report, in percentiles of its progression from the start to the end. For example, a value of 1% indicates the measure for the topics discussed in the first 1% of a report. The sample period is 2000–2023, and includes all the firms in our sample. Panel A shows the proportion of the topics that are collected at each point of a report's progress. Panel B plots the average sentiment and average time outlook throughout the average report. Panel C plots the attention allocation to entities (i.e., firm, industry, and the macroeconomy) throughout the average report. Panel D plots the attention allocation to valuation drivers (i.e., sales, costs, earnings/cash flows, margins, discount rates, and relative valuation statements) throughout the average report. For Panels B, C, and D, the y-axis indicates the values for the two indices plotted on the graph, sentiment and time outlook. For the *Sentiment Index*, the index denotes the sentiment associated with the portion of the text located in the X% of the average report (the average is measured using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports). For the *Time Outlook Index*, the index denotes the time outlook associated with the portion of the text located in the X% of the average report (the average is measured using values of -1, 0, 1, 2 for past, present, near-future, and distant-future outlook in the reports). The three shaded areas indicate three generally occurring sections to provide an intuitive progression over the x-axis.

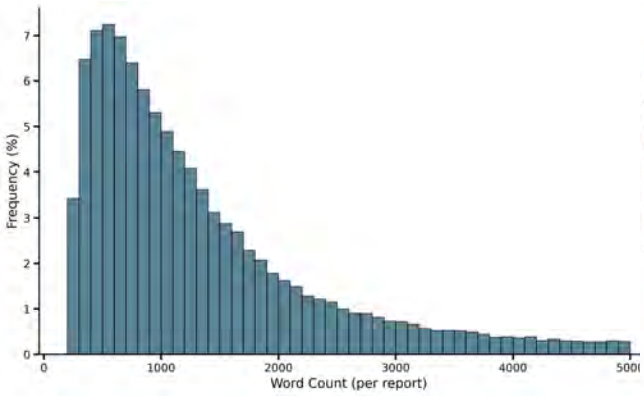
(A) Topic Location - Chunking Method.



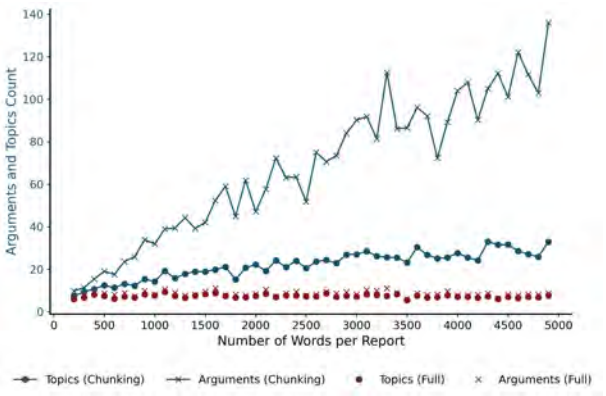
(B) Topic Location - Full Report Method.



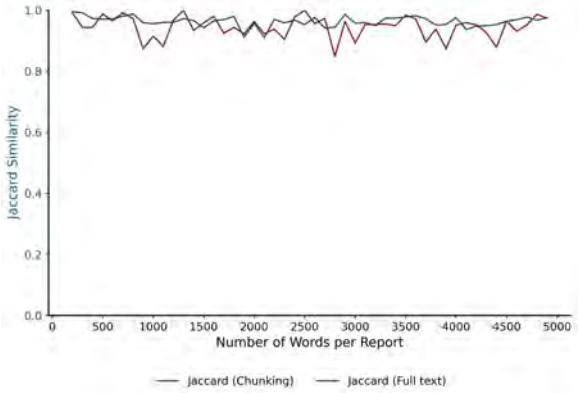
(C) Words per Report (Full Sample).



(D) Arguments and Topics per Report.



(E) Output Robustness (Text length).



(F) Output Robustness (Depth).

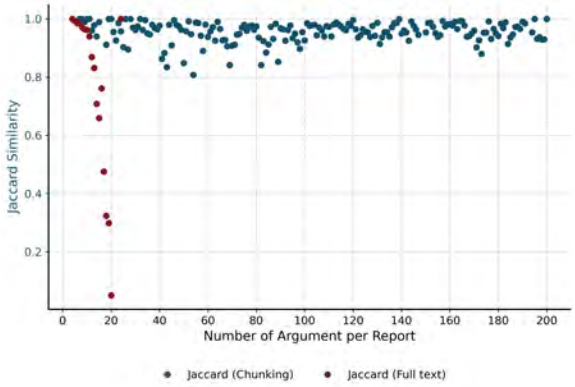


Figure II:
Diagnostic Tool

This figure plots the results from our robustness exercise. With the exception of Panel C, the sample is generated from performing the mental model extraction 10 times on a subsample of 240 reports randomly selected from our full sample. Panel C uses the full sample. Panels A and B show where the topics are collected on average in the reports for our main data collection method (Panel A), and the alternative method that extracts topics from the full text at once (Panel B). Panel C plots the histogram of the number of words per report. Panel D shows the number of topics and narratives as a function of report length, for both our data collection technique (in blue), and the alternative collection strategy (in red). Panel E plots the similarity score measured between the 10 attempts to extract topics of each of the 240 reports as a function of report length. In blue is the Jaccard similarity for our data collection strategy, and in red the alternative collection strategy. Panel F plots the Jaccard similarity for both our data collection strategy (in blue), and the alternative data collection strategy (in red).

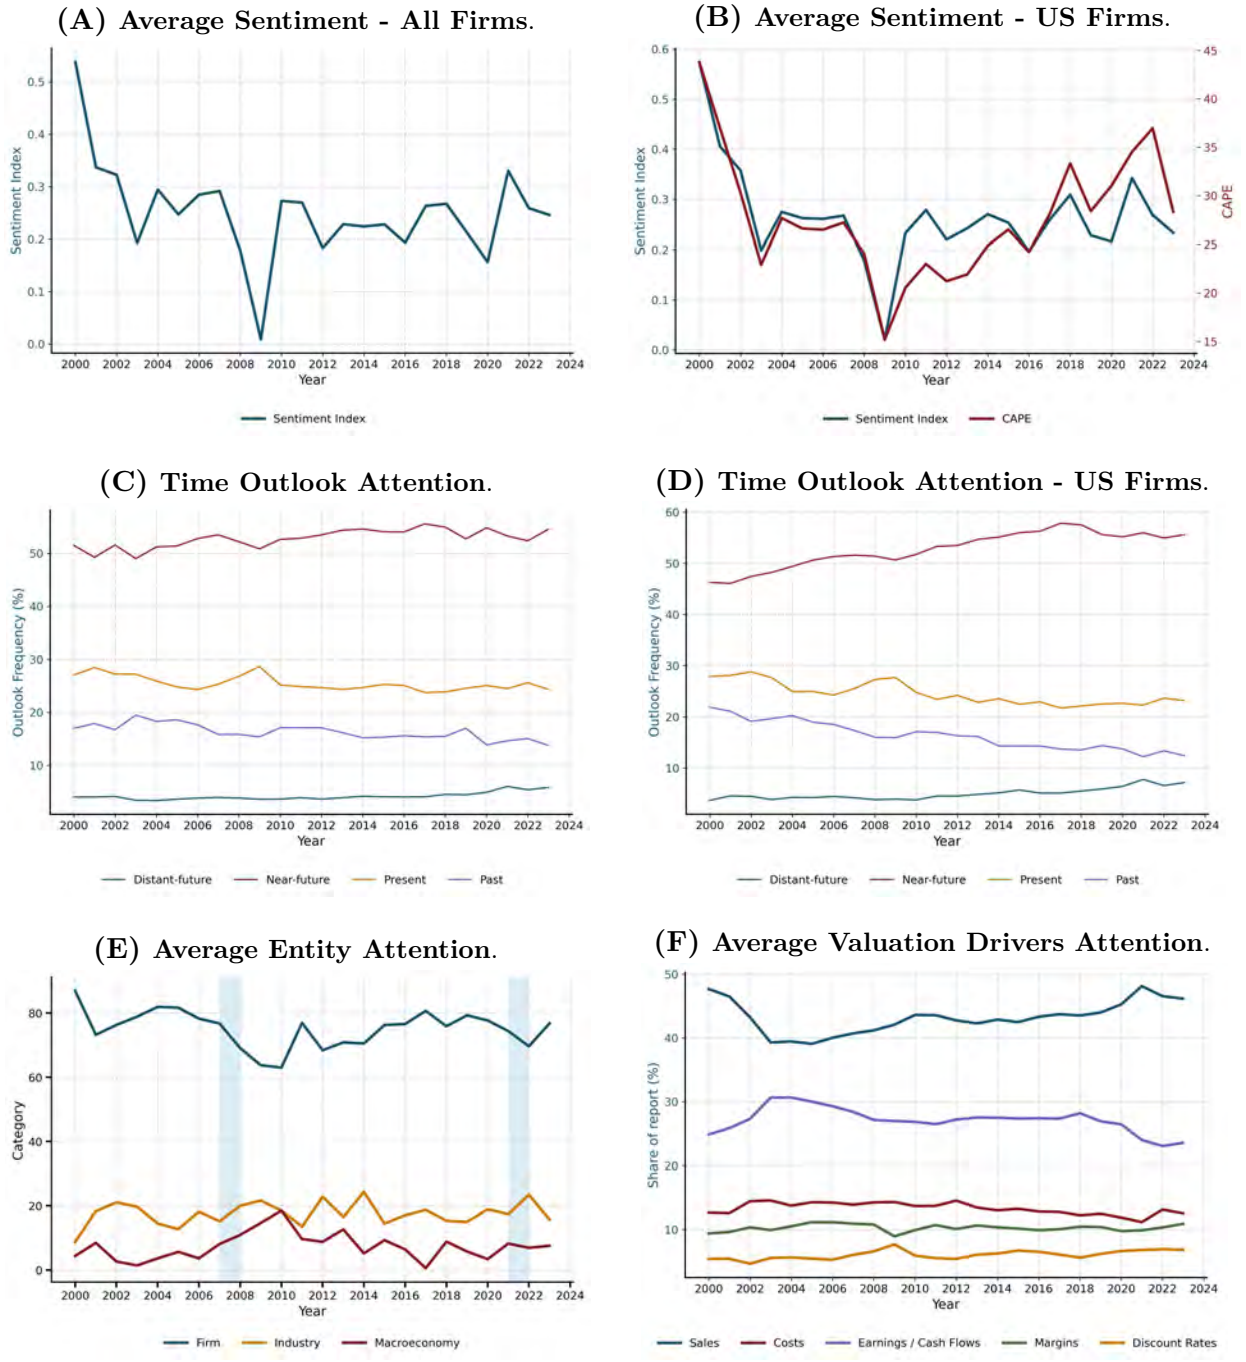


Figure III:
Sentiment, Outlook, Entities, and Valuation Drivers

This figure plots the time trends for the key variables collected in our sample, 2000–2023. The x-axis is expressed in years for all panels. For Panels A–B, the y-axis denotes the sentiment associated with the average report written in a given year (the average is measured using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports). The rightmost y-axis of Panel B indicates the scale for Shiller's CAPE index. For Panels C–D, the y-axis denotes the frequency, in percentage, at which the average report uses statements discussing topics about the past, present, near-future (<3 years), and distant-future (>3 years). Panel E denotes the frequency, in percentage, at which the average report uses statements discussing the three entities–firm, industry, and macroeconomy. Panel F denotes the frequency, in percentage, at which the average report uses statements discussing the time outlook associated with the different valuation drivers–sales, costs, margins, earnings, and discount rates.

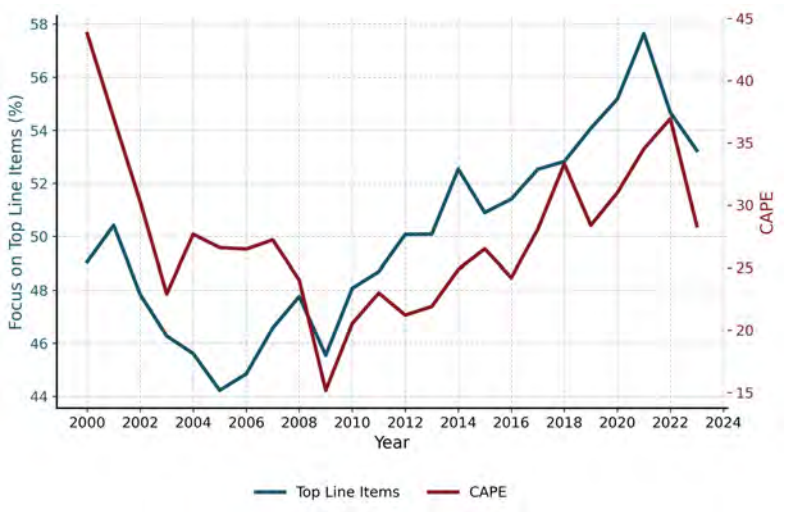


Figure IV:
Attention to Topline Items and Market Valuations - US Firms

This figure plots the average report attention allocation to top-line items–sales/revenues–for all US firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the leftmost y-axis denotes the proportion of the average report that is allocated to discuss statements related to sales. The rightmost y-axis denotes Shiller's CAPE index. The correlation between Shiller's CAPE and attention allocation to sales/revenues is 0.42.

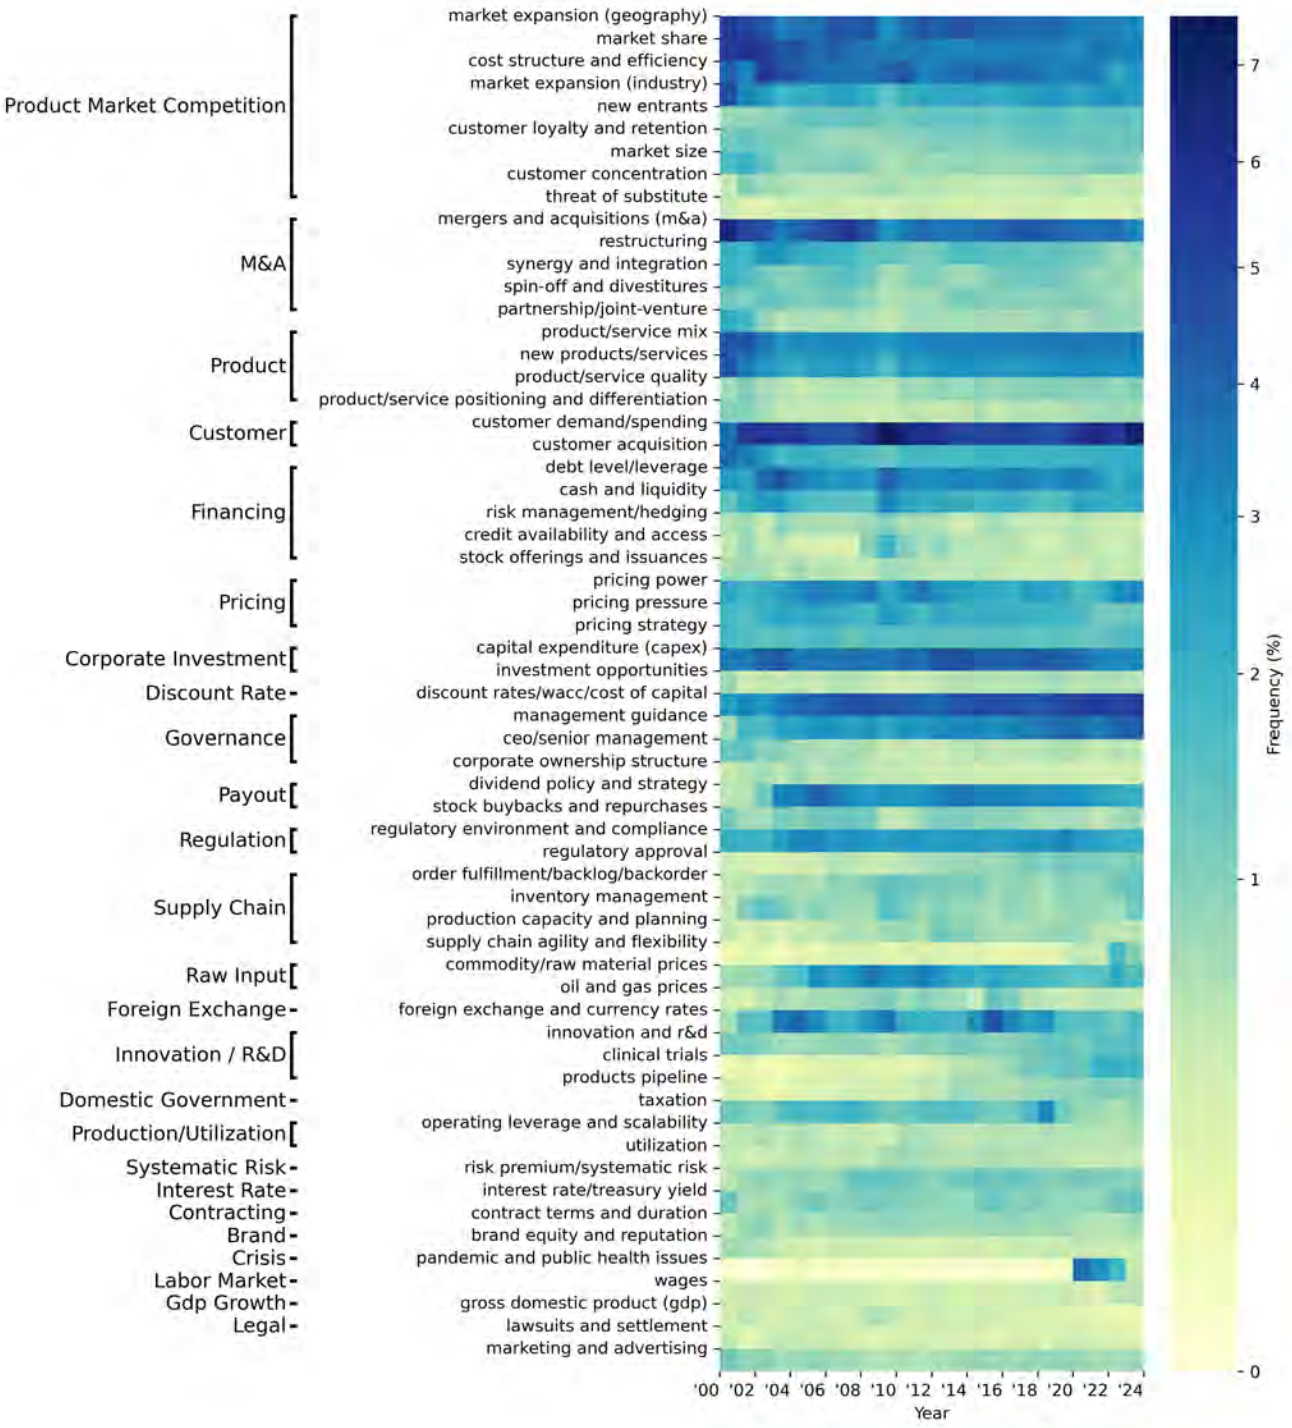


Figure V:
Attention Allocation Over Time – Full sample

This figure plots the attention allocation to the 60 most important topics for all the firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the right, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

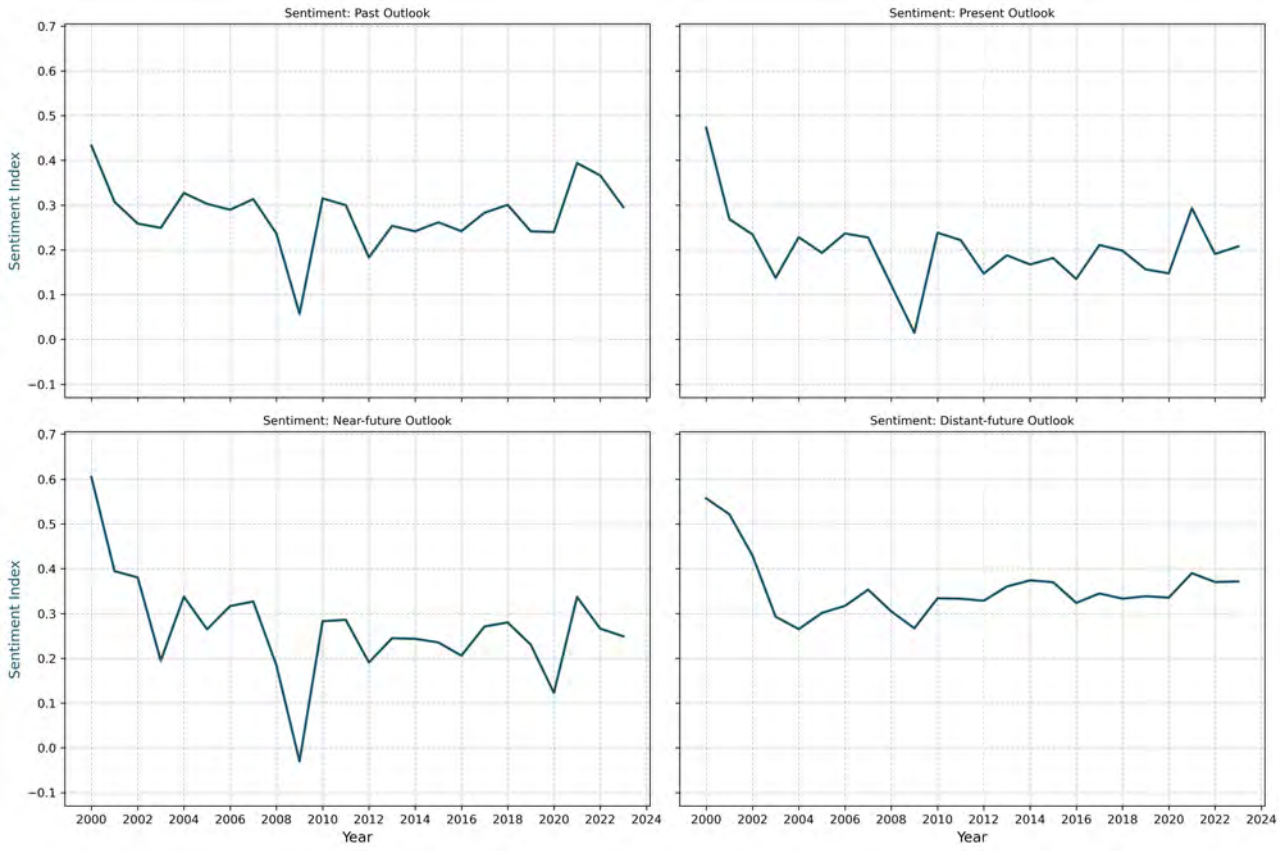


Figure VI:
Average Sentiment by Time Outlook

This figure plots the average sentiment over time for each time outlook in our sample: past, present, near-future (1–3 years), distant-future (>3 years) for the period 2000–2023. The x-axis is expressed in years. The y-axis denotes the sentiment associated with the average report written in a given year (the average is measured using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports).

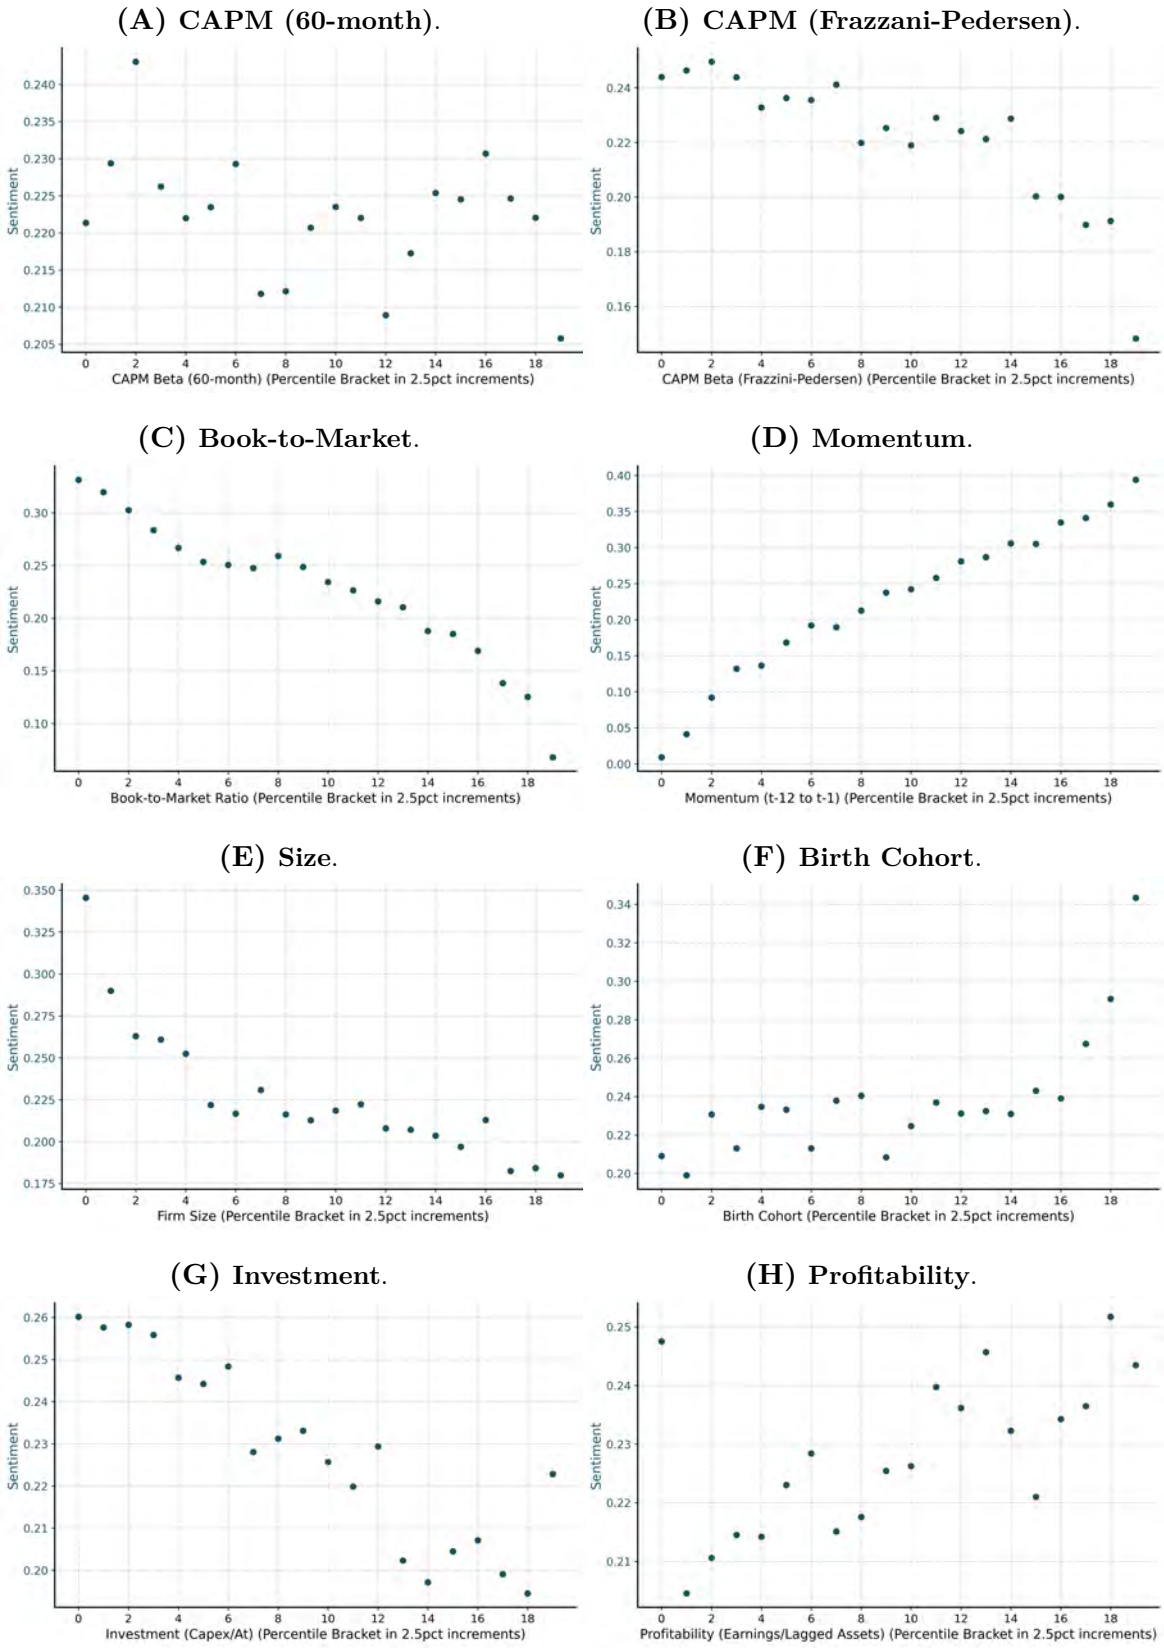


Figure VII:
Sentiment in the Cross-Section

This figure plots the cross-sectional patterns of sentiment for various key variables: CAPM beta, book-to-market, momentum, firm size, and birth cohort, across all the firms in our sample for which we have data on these characteristics. The x-axis corresponds to percentile brackets in 5% increments. For example, a value of 1 indicates the first percentile bracket, which includes observations where the x-variable falls within the 0 and 5 percentile range in each of their respective year. In all panels, the y-axis represents the average report sentiment within each percentile bracket. Report sentiment is calculated as the average sentiment of all topics discussed in the report, with values assigned as follows: -1 for negative sentiment, 0 for neutral sentiment, and 1 for positive sentiment. Larger values indicate more positive statements on average for the report.

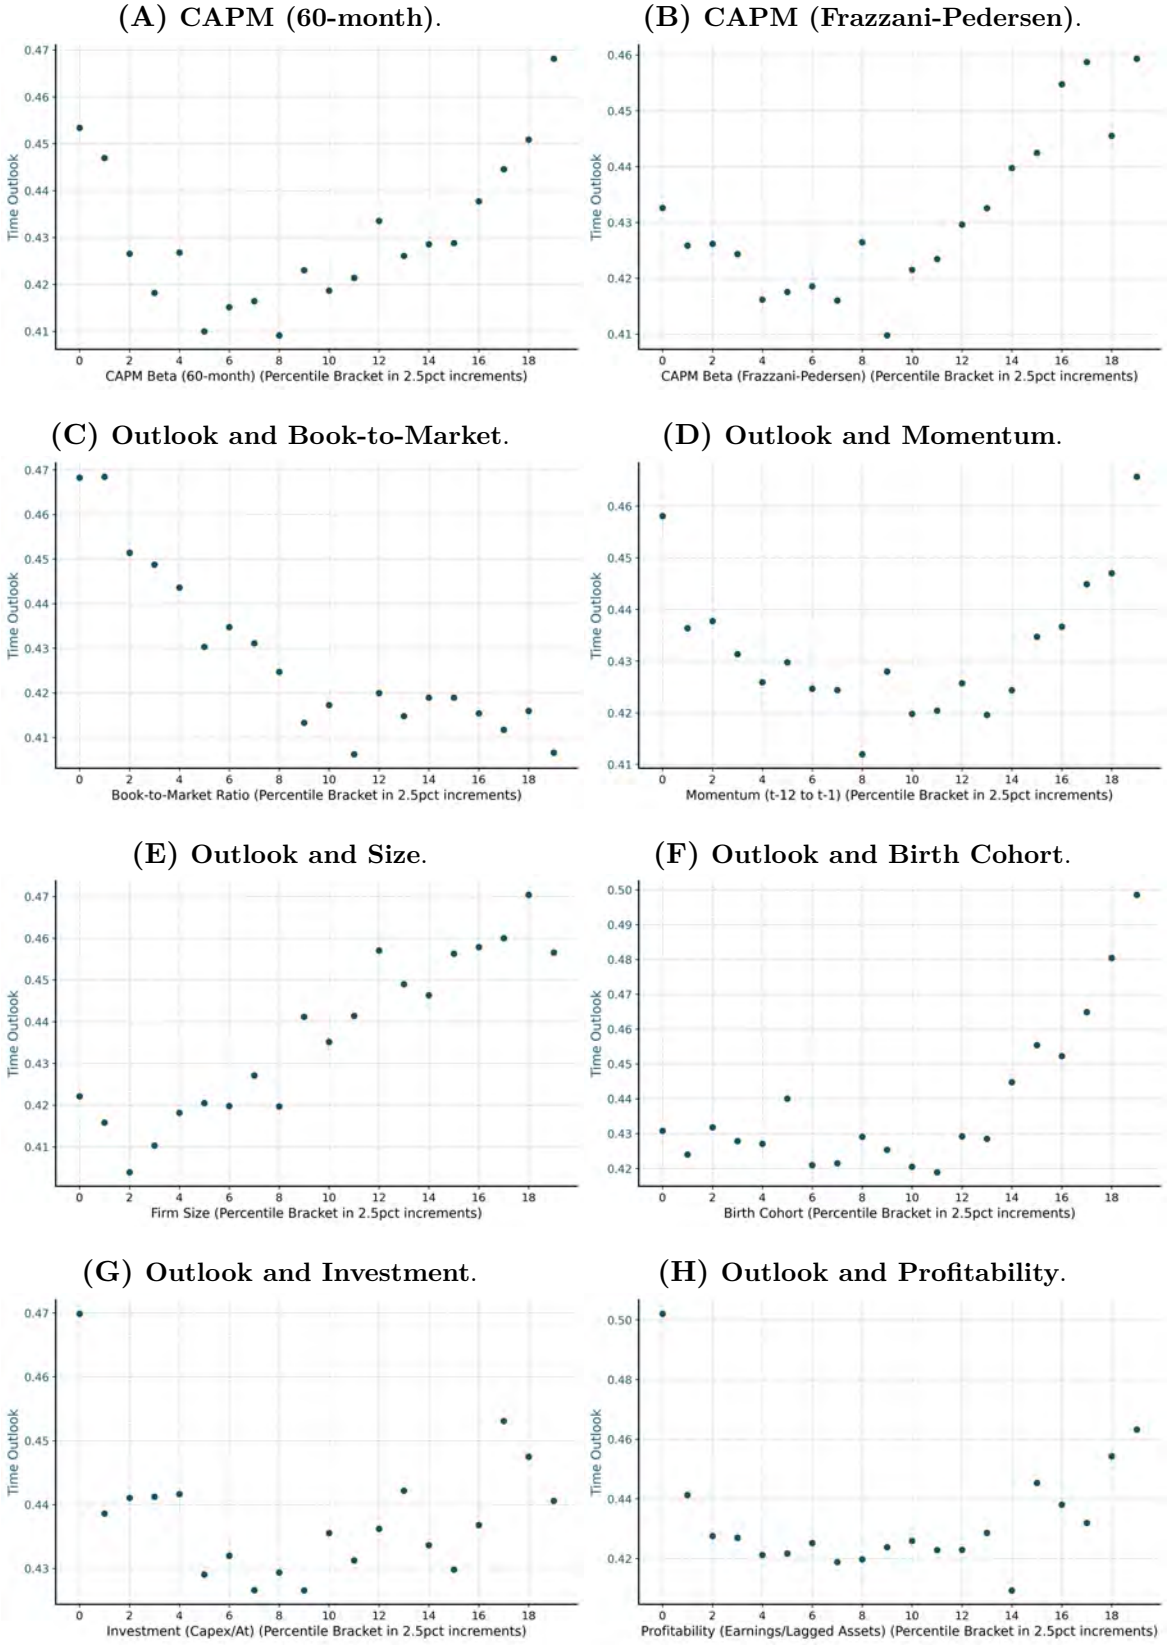


Figure VIII:
Time Outlook in the Cross-Section

This figure plots the cross-sectional patterns of time outlook for various key variables: CAPM beta, book-to-market, momentum, firm size, and birth cohort, across all the firms in our sample for which we have data on these characteristics. The x-axis corresponds to percentile brackets in 5% increments. For example, a value of 1 indicates the first percentile bracket, which includes observations where the x-variable falls within the 0 and 5 percentile range in each of their respective year. In all panels, the y-axis represents the average report time outlook in each percentile bracket. Report outlook is calculated as the average outlook of all topics discussed in the report, with values assigned as follows: -1 for past outlook, 0 for present outlook, 1 for near-future outlook, and 2 for distant-future outlook. Larger values indicate more forward-looking statements for the report.

Table I:
Summary Statistics

This table presents summary statistics. The sample period is 2000–2023. Panel A describes the sample coverage. Panel B focuses on variables associated with forecasters' mental models, and Panel C examines measures of narrative similarity.

Panel A: Structure	Avg. Obs. per Category-Year	Avg. Obs per Category	Nb. Unique Category
Category:			
Equity report			81,757
Year		3405.29	24
Brokerage house	97.29	1,776.67	46
Firm	1.61	7.23	11,308
Headquarters Country	56.88	882.30	92
Industry	179.53	4270.95	19

Panel B: Mental Model	Mean	25^{th} pct.	Median	75^{th} pct.	Standard Deviation	Nb. Observation
Topics:						
Nb. Unique topic per report$_{i,j,t}$	18.36	14.00	18.00	22.00	6.62	81,757
Nb. Narrative per report$_{i,j,t}$	43.23	24.00	37.00	55.00	25.89	81,727
Nb. Word per report$_{i,j,t}$	1,333.25	611.00	1,026.00	1,722.00	994.29	81,727
Nb. Narrative per Topic (avg. report) $_{i,j,t}$	2.43	1.68	2.11	2.67	3.32	81,727
Valuation Drivers:						
Sales: Share of report (%)$_{i,j,t}$	32.02	20.97	31.23	41.94	15.05	81,727
Costs: Share of report (%)$_{i,j,t}$	9.85	4.97	8.45	13.21	6.58	81,727
Margins: Share of report (%)$_{i,j,t}$	7.44	3.56	6.23	10.06	5.12	81,727
Earnings: Share of report (%)$_{i,j,t}$	18.84	11.15	17.05	24.67	10.53	81,727
Discount rates: Share of report (%)$_{i,j,t}$	4.58	2.01	3.35	6.02	3.75	81,727
Valuation: Share of report (%)$_{i,j,t}$	14.09	8.20	12.44	18.05	8.30	81,727
Entities:						
Firm: Share of report (%)$_{i,j,t}$	80.69	73.01	83.26	91.03	14.04	81,727
• Sales share of Firm (%)$_{i,j,t}$	37.16	23.69	35.54	49.08	18.32	47,819
• Discount rates share of Firm (%)$_{i,j,t}$	6.95	3.01	5.13	9.04	6.09	81,556
Industry: Share of report (%)$_{i,j,t}$	13.99	6.54	11.61	18.72	10.06	81,727
• Sales share of Industry (%)$_{i,j,t}$	73.18	59.12	80.47	91.97	21.91	47,819
• Discount rates share of Industry (%)$_{i,j,t}$	0.83	0.27	0.57	1.13	0.77	68,030
Macroeconomy: Share of report (%)$_{i,j,t}$	5.28	2.39	4.12	6.92	4.16	81,727
• Sales share of Macro (%)$_{i,j,t}$	43.02	31.26	41.68	62.52	18.51	47,819
• Discount rates share of Macro (%)$_{i,j,t}$	22.50	12.00	20.57	36.00	11.33	47,819
Sentiment:						
Sentiment$_{i,j,t}$	0.23	-0.03	0.26	0.52	0.38	81,726
Sentiment$^{past}_{i,j,t}$	0.27	-0.14	0.35	0.91	0.65	69,701
Sentiment$^{present}_{i,j,t}$	0.18	-0.14	0.22	0.54	0.47	81,247
Sentiment$^{near-future}_{i,j,t}$	0.24	-0.07	0.29	0.60	0.46	81,554
Sentiment$^{distant-future}_{i,j,t}$	0.34	0.00	0.40	1.00	0.59	54,278
Outlook:						
Outlook$_{i,j,t}$	0.44	0.26	0.47	0.64	0.29	81,727
Past: Share of report (%)$_{i,j,t}$	14.41	5.69	11.37	20.31	11.22	81,727
Present: Share of report $_{i,j,t}$	31.04	22.44	30.17	38.52	12.25	81,727
Near-future: Share of report $_{i,j,t}$	48.81	38.62	48.91	58.95	14.95	81,727
Distant-future: Share of report $_{i,j,t}$	4.51	2.08	3.42	5.80	3.62	81,727

Panel C: Similarity measures	Mean	25^{th} pct.	Median	75^{th} pct.	Standard Deviation	Nb. Observation
Between analysts:						
Jaccard$_{A,B,j,t}$	31.81	25.00	31.58	38.71	10.71	51,723
Jaccard$^{45\text{th pct} < \text{Report length} < 55\text{th pct}}_{A,B,j,t}$	33.92	26.92	33.33	40.74	10.58	23,280
Jaccard$^{90\text{th pct} < \text{Report length}}_{A,B,j,t}$	40.38	33.33	39.58	45.95	10.04	777
jaccard$^{Top5}_{A,B,j,t}$	21.88	11.11	25.00	25.00	15.30	51,723
Sentiment agreement (same topic)						
Unconditional	0.47	0.00	0.00	1.00	0.50	2,006,167
Different argument & outlook & channel	0.43	0.00	0.00	1.00	0.50	459,701
Same argument & outlook & channel	0.74	0.00	1.00	1.00	0.44	29,738
Within analyst:						
Jaccard$_{A,j,t,t-1}$	38.04	29.17	37.93	46.67	12.85	34,643
Jaccard$^{Top5}_{A,j,t,t-1}$	25.22	11.11	25.00	42.86	16.59	34,643
Within report - Outlook:						
jaccard$^{past,present}_{i,j,t}$	10.90	0.00	10.00	16.67	10.66	68,501
jaccard$^{past,near-future}_{i,j,t}$	16.96	8.33	15.38	25.00	12.70	68,703
jaccard$^{past,distant-future}_{i,j,t}$	8.41	0.00	0.00	14.29	14.56	45,638
jaccard$^{present,near-future}_{i,j,t}$	16.63	7.69	15.38	25.00	11.76	80,380
jaccard$^{present,distant-future}_{i,j,t}$	8.08	0.00	7.14	12.50	8.81	53,333
jaccard$^{near-future,distant-future}_{i,j,t}$	9.79	0.00	8.33	14.29	9.52	53,463

Table II:
Sentiment, Time Outlook, and 12-Month Expected Price Growth

This table presents results related to the role of sentiment and time outlook on the 12-month price target appreciation of forecasters. In Panel A, the dependent variable, *12-month price target appreciation*, corresponds to the analysts' 12-month price target divided by the closing price at the time of the forecast, multiplied by 100. The variables of interest are $Sentimenti, j, t$, which is equal to the average sentiment measured in the report of the contributor i, for firm j, in year t, and $Outlooki, j, t$, which is equal to the average outlook measured in the report of the contributor i, for firm j, in year t. Positive values for Sentiment and Outlook indicate positive sentiment and forward-looking outlook, respectively. In both Panels A and B, $Sentiment^{Distant-future}i, j, t$ is equal to the average sentiment measured in the report of the contributor i, for firm j, in year t only for the statements about the distant future. $Sentiment^{Near-future}i, j, t$ is equal to the average sentiment measured in the report of the contributor i, for firm j, in year t only for the statements about the near future. $Sentiment^{Present}i, j, t$ is equal to the average sentiment measured in the report of the contributor i, for firm j, in year t only for the statements about the present. $Sentiment^{Past}i, j, t$ is equal to the average sentiment measured in the report of the contributor i, for firm j, in year t only for the statements about the past. The regressions are estimated using ordinary least squares. Standard errors (in parentheses) are estimated using heteroskedasticity-consistent estimators clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: Sentiment and Expectations	$\ln(\frac{\text{Price target}_{i,j,t}}{EPS_{j,t-1}})$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(β_1) Sentiment$_{i,j,t}$	0.43***	0.40***			0.43***	0.40***		
	(0.03)	(0.01)			(0.03)	(0.01)		
Outlook$_{i,j,t}$			0.32***	0.06***	0.33***	0.06***		
			(0.04)	(0.02)	(0.04)	(0.02)		
Sentiment$^{Distant-future}_{i,j,t}$							0.03**	0.01
							(0.01)	(0.01)
Sentiment$^{Near-future}_{i,j,t}$							0.10***	0.11***
							(0.01)	(0.01)
Sentiment$^{Present}_{i,j,t}$							0.05***	0.05***
							(0.01)	(0.01)
Sentiment$^{Past}_{i,j,t}$							0.07***	0.03***
							(0.01)	(0.01)
Firm*Contributor FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	49,313	40,074	49,313	40,074	49,313	40,074	27,787	19,664
F Statistics	235.55	894.14	69.94	12.65	175.18	454.55	41.13	112.98
R^2	0.01	0.87	0.00	0.86	0.01	0.87	0.01	0.89
H_0^A: $\beta_3 = \beta_4$ (P-value)							0.00	0.01
H_0^B: $\beta_5 = \beta_4$ (P-value)							0.02	0.01
H_0^C: $\beta_6 = \beta_4$ (P-value)							0.05	0.00
Observations	61,426	35,227	61,427	35,228	61,426	35,227	34,658	15,081
F Statistics	456.08	1022.64	9.27	12.53	232.73	517.09	117.97	153.23
R^2	0.01	0.81	0.00	0.80	0.01	0.81	0.02	0.82

Panel B: Sentiment and Outlook	Sentiment$^{Near-future}_{i,j,t}$				Sentiment$^{Distant-future}_{i,j,t}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sentiment$^{Past}_{i,j,t}$	0.31***	0.28***	0.13***	0.14***	0.10***	0.07***	0.01	0.00
	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Sentiment$^{Present}_{i,j,t}$			0.49***	0.45***			0.16***	0.13***
			(0.00)	(0.01)			(0.01)	(0.01)
Sentiment$^{Near-future}_{i,j,t}$							0.15***	0.10***
							(0.01)	(0.01)
Firm*Contributor FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	69,563	55,441	69,257	55,150	46,451	33,440	46,332	33,332
F Statistics	5710.73	2836.40	11511.32	5396.32	435.88	111.39	850.84	240.18
R^2	0.10	0.42	0.31	0.54	0.01	0.46	0.07	0.48
Within R^2	0.10	0.08	0.31	0.27	0.01	0.01	0.07	0.04

Table III:
Attention Allocation and Valuation Adjustment

This table presents results exploring how analysts attention allocation to sales impact the sensitivity of their valuation to sales. The analysis is conducted at the pair of forecasters, A and B, topic k, firm j and year t level. The dependent variable, *Same sentiment*, takes the value of 1 if both analysts share the same sentiment for the overlapping topic k, and 0 otherwise. The first variable of interest, Same outlook, takes the value of 1 if both analysts share the same outlook when discussing topic k, and 0 otherwise. The second variable of interest, Same valuation channel, takes the value of 1 if both analysts assign the same valuation channel when discussing topic k, and 0 otherwise. In Panel A, the topics over which the pair of analysts are matched refer to the standardized labels. Panel B, the topics over which the pair of analysts are matched refer to the raw topics (the individual narratives). On average, analysts employ various individual narratives to organize their discussion of topics. The analysis performed in columns 5 to 8 is conducted on a subsample for which the number of words included in the reports of analysts A and B is within a 20% margin. The regressions are estimated using ordinary least squares. Standard errors (in parentheses) are estimated using heteroskedastic-consistent estimators and clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: Sales Focus	$\ln(\frac{\text{Price target}_{i,j,t}}{EPS_{j,t-1}})$			
	(1)	(2)	(3)	(4)
$\ln(\text{Costs growth})_{i,j,t}$	-0.23***	-0.21***	-0.21***	-0.22***
	(0.05)	(0.05)	(0.05)	(0.05)
$\ln(\text{Sales growth})_{i,j,t}$	0.44***	0.42***	0.42***	0.45***
	(0.11)	(0.11)	(0.11)	(0.12)
$\ln(\text{Sales growth})_{i,j,t}$*Attention Allocation$^{\text{Sales Channel}}_{i,j,t}$	1.13***	1.17***	1.17***	1.03***
	(0.23)	(0.22)	(0.22)	(0.26)
Attention Allocation$^{\text{Sales Channel}}_{i,j,t}$	0.06	0.08**	0.07*	0.09**
	(0.04)	(0.04)	(0.04)	(0.04)
Firm* FE	Yes	Yes	Yes	No
Year FE	No	Yes	Yes	No
Year FE	No	No	Yes	No
Firm*Contributor FE	No	No	No	Yes
Contributor*Year FE	No	No	No	Yes
Observations	40,178	40,178	40,178	33,778
F Statistics	38.33	39.72	39.28	28.65
R^2	0.82	0.83	0.83	0.87
Panel B: Costs Focus	$\ln(\frac{\text{Price target}_{i,j,t}}{EPS_{j,t-1}})$			
	(1)	(2)	(3)	(4)
$\ln(\text{Costs growth})_{i,j,t}$	-0.16***	-0.15***	-0.15***	-0.17***
	(0.05)	(0.05)	(0.05)	(0.05)
$\ln(\text{Sales growth})_{i,j,t}$	0.77***	0.76***	0.76***	0.75***
	(0.09)	(0.09)	(0.09)	(0.10)
$\ln(\text{Costs growth})_{i,j,t}$*Attention Allocation$^{\text{Costs Channel}}_{i,j,t}$	-0.75**	-0.64*	-0.64*	-0.58*
	(0.36)	(0.34)	(0.34)	(0.34)
Attention Allocation$^{\text{Costs Channel}}_{i,j,t}$	-0.35***	-0.33***	-0.33***	-0.32***
	(0.07)	(0.07)	(0.07)	(0.08)
Firm* FE	Yes	Yes	Yes	No
Year FE	No	Yes	Yes	No
Year FE	No	No	Yes	No
Firm*Contributor FE	No	No	No	Yes
Contributor*Year FE	No	No	No	Yes
Observations	40,178	40,178	40,178	33,778
F Statistics	36.60	34.29	34.11	25.88
R^2	0.82	0.83	0.83	0.87
Within R^2	0.04	0.04	0.04	0.04

Table IV:
Topic Interpretation and Complexity

This table presents results exploring the role of topic interpretation in reports. The statistics reports focus on the narratives related to cash-flow channels: sales, costs, earnings, or margins. Panel A reports the proportion of topic-report that are associated with 1, 2, 3, or 4 cash-flow channels. For example, in 86.85% of cases, a topic discuss in a report will be associated with only one valuation channel. Panel B reports the proportion of topic-report that are associated with 1, 2, or 3 sentiment for the cash-flow. For example, in 76.8% of cases, a topic discuss in a report will be associated with only one sentiment/effect on the valuation outcome.

Panel A: Number of Cash-flow Channels per Topic-Report	**Share of Sample**
1	86.85%
2	11.78%
3	1.26%
4	0.11%
Panel B: Number of Sentiment per Topic-Report	**Share of Sample**
1	76.8%
2	19.2%
3	3.9%

Table V:
Topic, Direction, and Agreement

This table presents regression results for the drivers agreement on the directional effect of a topic on the valuation outcome–the 12-month price target. The analysis is conducted at the pair of forecasters, A and B, topic k, firm j and year t level. The dependent variable, *Same sentiment*, takes the value of 1 if both analysts share the same sentiment for the overlapping topic k, and 0 otherwise. The first variable of interest, Same outlook, takes the value of 1 if both analysts share the same outlook when discussing topic k, and 0 otherwise. The second variable of interest, Same valuation channel, takes the value of 1 if both analysts assign the same valuation channel when discussing topic k, and 0 otherwise. In Panel A, the topics over which the pair of analysts are matched refer to the standardized labels. Panel B, the topics over which the pair of analysts are matched refer to the raw topics (the individual narratives). On average, analysts employ various individual narratives to organize their discussion of topics. The analysis performed in columns 5 to 8 is conducted on a subsample for which the number of words included in the reports of analysts A and B is within a 20% margin. The regressions are estimated using ordinary least squares. Standard errors (in parentheses) are estimated using heteroskedastic-consistent estimators and clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: Merge on Topic	Same sentiment$_{A,b,k,j,t}$			
Sample (Filter 1):	Full Sample		Both Reports are within a 20% Word Count Margin	
	(1)	(2)	(3)	(4)
Same outlook$_{A,B,k,j,t}$	0.03***	0.03***	0.04***	0.03***
	(0.00)	(0.00)	(0.00)	(0.00)
Same valuation channel$_{A,B,k,j,t}$	0.04***	0.03***	0.04***	0.03***
	(0.00)	(0.00)	(0.00)	(0.00)
Firm FE	No	Yes	No	Yes
Year FE	No	Yes	No	Yes
Contibutor$_A$*Contibutor$_B$ FE	No	Yes	No	Yes
Observations	2,006,167	2,006,165	375,729	375,728
F Statistics	953.98	884.82	388.62	309.04
R^2	0.00	0.02	0.00	0.03
Panel B: Merge on Narratives	Same sentiment$_{A,b,k,j,t}$			
Sample (Filter 1):	Full Sample		Both Reports are within a 20% Word Count Margin	
	(1)	(2)	(3)	(4)
Same outlook$_{A,B,k,j,t}$	0.05***	0.03***	0.08***	0.04***
	(0.00)	(0.00)	(0.01)	(0.01)
Same valuation channel$_{A,B,k,j,t}$	0.04***	0.02***	0.06***	0.03***
	(0.01)	(0.00)	(0.01)	(0.01)
Firm FE	No	Yes	No	Yes
Year FE	No	Yes	No	Yes
Contibutor$_A$*Contibutor$_B$ FE	No	Yes	No	Yes
Observations	106,056	105,552	23,388	23,050
F Statistics	93.94	55.28	61.29	18.19
R^2	0.00	0.09	0.01	0.19

Table VI:
Attention Allocation Alignment and Quantitative Beliefs Agreement

This table presents results exploring the role of topic focus alignment on valuation disagreement. In Panel A, the dependent variable, *Valuation Disagreement*, corresponds to the absolute difference between forecaster A and B's 12-month price targets divided by the sum of forecaster A and B's 12-month price targets, for firm j on year t. The calculation is performed for pairs of analysts that evaluate the same firm at the same time. The variables of interest are $Jaccard A, B, j, t$, which is equal to the number of topics that overlap in analyst A and B's reports, divided by the total number of distinct topics included in at least one report. $Same\ Argument A, B, j, t$ measures the proportion of the overlapping topics in both reports that share the same time outlook, sentiment, and valuation channel, multiplied by 100. In Panel B, the dependent variable, $\ln(Price\text{-}to\text{-}earnings\ Ratio)_A$, corresponds to analyst A's 12-month price target divided by the earnings-per-share measured at the time of the report. $\ln(Price\text{-}to\text{-}earnings\ Ratio)_B$, corresponds to analyst B's 12-month price target divided by the earnings-per-share measured at the time of the report. For both Panels A and B, the analysis performed in columns 5 to 8 is conducted on a subsample for which the number of words included in the reports of analysts A and B is within a 20% margin. The regressions are estimated using ordinary least squares. Standard errors (in parentheses) are estimated using heteroskedasticity-consistent estimators clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Dependent variable:	Valuation Disagreement$_{A,B,j,t}$							
Sample:	Full Sample				Both Reports are within a 20% Word Count Margin			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Jaccard$_{A,B,j,t}$	-0.06***	-0.04***	-0.05***	-0.04***	-0.06***	-0.05***	-0.04***	-0.04***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Same Argument (%)$_{A,B,j,t}$			-0.06***	-0.05***			-0.08***	-0.06***
			(0.01)	(0.01)			(0.01)	(0.01)
Contibutor$_A$*Contibutor$_B$ FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	40,627	39,309	30,822	29,779	8,859	7,805	6,705	5,840
F Statistics	97.78	60.10	80.82	55.02	40.04	16.30	31.84	12.50
R^2	0.00	0.25	0.01	0.26	0.01	0.35	0.01	0.38
Dependent variable:	ln(Price-to-Earnings Ratio)$_{A,j,t}$							
Sample:	Full Sample				Both Reports are within a 20% Word Count Margin			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(Price-to-Earnings Ratio)$_{B,j,t}$	0.96***	0.96***	0.96***	0.96***	0.96***	0.96***	0.94***	0.95***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Jaccard$_{A,B,j,t}$*ln(Price-to-Earnings Ratio)$_{B,j,t}$	0.06***	0.06***	0.06***	0.06***	0.07***	0.05**	0.07***	0.06**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Jaccard$_{A,B,j,t}$	-0.22***	-0.21***	-0.25***	-0.23***	-0.27***	-0.23**	-0.32***	-0.27**
	(0.08)	(0.08)	(0.09)	(0.09)	(0.10)	(0.11)	(0.11)	(0.12)
Same Arg. (%)$_{A,B,j,t}$*ln(Price-to-Earnings Ratio)$_{B,j,t}$			0.03	0.04*			0.06*	0.06*
			(0.02)	(0.02)			(0.03)	(0.04)
Same Argument (%)$_{A,B,j,t}$			-0.11	-0.13			-0.19	-0.24*
			(0.08)	(0.08)			(0.13)	(0.14)
Contibutor$_A$*Contibutor$_B$ FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	29,261	29,177	23,662	23,575	5,096	4,972	4,096	4,012
F Statistics	82174.33	75025.23	39833.21	37077.92	25595.62	21145.44	14064.59	12671.11
R^2	0.97	0.97	0.97	0.97	0.97	0.97	0.97	0.97

Table VII:
Drivers of Attention Allocation – The Case of Inflation

This table presents results on the drivers of attention allocation. The sample period is 2000–2023. The first dependent variable, *Attention to Inflation*, corresponds to the share of the equity report allocated to discussing the topic "Inflation and CPI" for forecaster i, firm j, and year t. The second dependent variable, *Sentiment about inflation*, denotes a report's average sentiment toward inflation, when the report includes "Inflation and CPI" as a topic. $\text{Inflation}^{FirmHQCountry}$ and $\text{Inflation}^{AnalystCountry}$ respectively correspond to the inflation measured in the firm HQ country or analyst country on year t. $1^{\text{Inflation} > 2\%}$ is a binary variable equal to 1 if the current inflation is above 2%, and 0 otherwise. The regressions are estimated using ordinary least squares. Standard errors (in parentheses) are estimated using heteroskedastic-consistent estimators and clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: Firm HQ Country	Attention to $\text{Inflation}_{i,j,t}$				
Sample:	Full Sample			$\text{Inflation}^{\text{Firm}} \leq 2\%$	$\text{Inflation}^{\text{Firm}} > 2\%$
	(1)	(2)	(3)	(4)	(5)
$\text{Inflation}^{\text{Firm HQ Country}}_{c,t}$	0.04***	0.04***	0.02***	0.00	0.02***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Observations	73,887	70,778	70,778	33,224	35,367
F Statistics	124.27	95.87	11.69	0.01	7.50
R^2	0.01	0.21	0.22	0.25	0.26
Panel B: Analyst Country	Attention to $\text{Inflation}_{i,j,t}$			Sentiment about $\text{Inflation}_{i,j,t}$	
	(1)	(2)	(3)	(4)	(5)
$\text{Inflation}^{\text{Analyst Country}}_{c,t}$	0.04***	0.04***	0.02*	-0.00	
	(0.01)	(0.01)	(0.01)	(0.02)	
$\text{Inflation}^{\text{Firm HQ Country}}_{c,t}$	0.02***	0.02***	0.02*	-0.04*	-0.04*
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
$1_{Inflation^{\text{Analyst Country}}_{c,t} > 2\%}$					-0.10
					(0.15)
$1_{Inflation^{\text{Analyst Country}}_{c,t} > 2\%} \times \text{Inflation}^{\text{Firm HQ Country}}_{c,t}$					-0.05*
					(0.02)
Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Observations	15,141	13,792	13,792	461	461
F Statistics	45.01	40.21	3.52	1.87	3.14
R^2	0.02	0.25	0.26	0.47	0.48

Table VIII:
Analyst Personal Experience and Attention Allocation

This table presents results exploring the relation between analysts' country of location and overall attention allocation, and agreement in information processing. The sample period is 2000–2023. The first dependent variable, Jaccard, corresponds to the Jaccard similarity score between the reports of analysts A and B, about firm j, in year t. The second dependent variable, Same Argument, measures the proportion of topics that share the same sentiment among the topics that overlap between the reports written by analysts A and B, about firm j, in year t. The variable of interest, $1^{Analysts\ Location\ in\ the\ same\ Country}$, is a binary variable equal to 1 if both analysts are located in the same country, and zero otherwise. The regressions are estimated using ordinary least squares. Standard errors (in parentheses) are estimated using heteroskedastic-consistent estimators and clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Dependent variable:	$\text{Jaccard}_{A,B,j,t}$		$\text{Same Argument}_{A,B,j,t}$	
	(1)	(2)	(3)	(4)
$1^{\text{Analysts Located in the Same Country}}_{A,B,j,t}$	0.88***	0.78***	0.41**	0.51*
	(0.22)	(0.27)	(0.19)	(0.28)
Firm*Year FE	No	Yes	No	Yes
Observations	19,722	13,543	15,035	9,932
F Statistics	15.93	8.65	4.66	3.21
R^2	0.00	0.44	0.00	0.38

Appendix A AI Prompt

Step 1: Identify key drivers of valuation exercise

Prompt input: Equity report.

Prompt output: Main mental model.

1. Main Task Analyze the provided equity report excerpt to identify all topics that clearly and directly explain the analyst's price target. For each topic:

- Categorize it into one of the following "category": "firm," "industry," or "macroeconomy."
- Categorize it into one of the following "channel": "earnings/cash flow: sales", "earnings/cash flow: costs", "earnings/cash flow: margins", "earnings/cash flow: earnings", "discount rate", "valuation", or "other".
- Identify its associated sentiment: "positive", "negative", "neutral", or "unclear."
- Identify the best keyword that captures the sentiment associated with the topic.
- Determine the time outlook: "past", "present", "near-future", "distant-future", or "unclear."
- Identify a 5 word snippet from the text to capture where the topic is first discussed.

2. Equity Report

```
<equityreport>
{EquityReport}
</equityreport>
```

3. Topic Identification

- Focus exclusively on topics directly tied to the justification of the price target.
- Each identified topic should be specific and clearly related to a single factor influencing the price target.

- Disregard general statements or any information not directly relevant to the price target justification.

4. Category Identification

- Assign each topic to the most appropriate category: "firm" (company-specific factors), "industry" (sector or industry trends), or "macroeconomy" (broader economic conditions).
- If a topic could reasonably belong to more than one category, select the category that best fits the context provided in the report.

5. Channel Identification Assign each topic to one of the following channels based on its context:

- "earnings/cash flow: sales": If related to sales or revenues.
- "earnings/cash flow: costs": If related to costs or operating expenses.
- "earnings/cash flow: margins": If related to profitability margins.
- "earnings/cash flow: earnings": If related to earnings, EBITDA, EBIT, or cash flows.
- "discount rate": If related to terms like "discount rate", "weighted average cost of capital/wacc", "equity beta", "idiosyncratic risk", "credit rating", "systematic risk", "equity risk premium (erp)", "market risk premium (mrp)", "interest rate", "risk-free rate (rf)", or "treasury yield".
- "valuation": If related to valuation terms such as "undervalued," "overvalued," or "fair-valued."
- "other": If none of the above apply.

6. Sentiment Identification Determine the sentiment associated with each topic, selecting one of the following:

- "positive": Indicates optimism or favorable conditions such as "beat expectations", "surpass expectations", "confidence", "bullish", "optimistic", ...
- "negative": Indicates pessimism or unfavorable conditions such as "uncertainty", "challenges", "caution", "concerns", "bearish", "underperform expectations", ...

- "neutral": Indicates neutrality or conditions in line with expectations such as "in line with", "meet consensus", ...

7. Keyword Identification

- Identify one or two keywords used by the analyst that best characterize the "Sentiment" variable.
- If no keywords can clearly characterize the sentiment, use "".

8. Time Outlook Identification Assign the time outlook based on the report's content. Remember that the report was written in the year {Year}. Possible values:

- "past": Events that have already occurred.
- "present": Events happening now.
- "near-future": Events expected within 1-3 years or short-term guidance.
- "distant-future": Events expected in more than 3 years or long-term guidance.
- "unclear": When the time frame is not clear.

9. Snippet Identification

- Identify a 5 word snippet from the text to capture where the topic is first discussed.

10. Output Format For each topic, provide your answer in the following format:

`{Topic}//{Category}//{Channel}//{Sentiment}//{Keyword}//{Timeoutlook}//{Snippet}`

- `{Topic}` is the specific topic.
- `{Category}`: "firm," "industry," or "macroeconomy."
- `{Channel}`: "earnings/cash flow: sales", "earnings/cash flow: costs", "earnings/cash flow: margins", "earnings/cash flow: earnings", "discount rate", "valuation", or "other".
- `{Sentiment}`:"positive", "negative", "neutral", or "unclear."
- `{Keyword}`: the best keyword that captures the sentiment associated with the topic.
- `{Timeoutlook}`:"past", "present", "near-future", "distant-future", or "unclear."

- `{Snippet}`: 5 word snippet.

11. Example Here is an example of a text with the associated output to help you: "Investment summary: Inflation will raise the cost of inputs while also reducing demand. We have a negative outlook. Despite lagging behind its peers in the EV market, we remain optimistic about the firm's ability to increase its market presence over the next 5 years, as reflected in our hold recommendation. Pricing pressure in the company's main segment has compressed the firm's margins." In this example, the output should be:

```
inflation//macroeconomy//earnings/cash flow: costs//negative//negative
↪ outlook//near-future//Inflation will raise the cost
market share//firm//earnings/cash flow: sales//positive//remain
↪ optimistic//distant-future//Despite lagging behind its peers
pricing pressure//industry//earnings/cash flow:
↪ margins//negative//pressure//past//Pricing pressure in the company
```

This should be the only output you provide.

Step 2: Review the selection of valuation drivers and identify any missing relevant drivers

Prompt input: Equity report and previous analysis.

Prompt output: Omitted mental model.

1. Main Task Review the previous equity report analysis and identify all topics that clearly and directly explain the analyst's price target and that were omitted in the previous analysis, based on the provided equity report. For each omitted topic:

- Categorize it into one of the following "category": "firm," "industry," or "macroeconomy."
- Categorize it into one of the following "channel": "earnings/cash flow: sales", "earnings/cash flow: costs", "earnings/cash flow: margins", "earnings/cash flow: earnings", "discount rate", "valuation", or "other".
- Identify its associated sentiment: "positive", "negative", "neutral", or "unclear."
- Identify the best keyword that captures the sentiment associated with the topic.
- Determine the time outlook: "past", "present", "near-future", "distant-future", or "unclear."
- Identify a 5 word snippet from the text to capture where the topic is first discussed.

Importantly, do not output topics already included in the previous analysis.

2. Equity Report

```
<equityreport>
{EquityReport}
</equityreport>
```

3. Previous Analysis

```
<Previous Analysis>
{Initial Topics}
</Previous Analysis>
```

Previous analysis output structure is:

`{Topic}//{Category}//{Channel}//{Sentiment}//{Keyword}//{Timeoutlook}//{Snippet}`

Where

- `{Topic}` is a topic.
- `{Category}`: "firm," "industry," or "macroeconomy."
- `{Channel}`: "earnings/cash flow: sales", "earnings/cash flow: costs", "earnings/cash flow: margins", "earnings/cash flow: earnings", "discount rate", "valuation", or "other".
- `{Sentiment}`:"positive", "negative", "neutral", or "unclear."
- `{Keyword}`: the best keyword that captured the sentiment associated with the topic.
- `{Timeoutlook}`:"past", "present", "near-future", "distant-future", or "unclear."
- `{Snippet}`: 5 word snippet.

4. Omitted Topic Identification

- Focus exclusively on omitted topics directly tied to the justification of the price target.
- Each identified omitted topic should be specific and clearly related to a single factor influencing the price target.
- Disregard general statements or any information not directly relevant to the price target justification.

5. Category Identification

- Assign each omitted topic to the most appropriate category: "firm" (company-specific factors), "industry" (sector or industry trends), or "macroeconomy" (broader economic conditions).
- If an omitted topic could reasonably belong to more than one category, select the category that best fits the context provided in the report.

6. Channel Identification Assign each omitted topic to one of the following channels based on its context:

- "earnings/cash flow: sales": If related to sales or revenues.
- "earnings/cash flow: costs": If related to costs or operating expenses.
- "earnings/cash flow: margins": If related to profitability margins.
- "earnings/cash flow: earnings": If related to earnings, EBITDA, EBIT, or cash flows.
- "discount rate": If related to terms like "discount rate", "weighted average cost of capital/wacc", "equity beta", "idiosyncratic risk", "credit rating", "systematic risk", "equity risk premium (erp)", "market risk premium (mrp)", "interest rate", "risk-free rate (rf)", or "treasury yield".
- "valuation": If related to valuation terms such as "undervalued," "overvalued," or "fair-valued."
- "other": If none of the above apply.

7. Sentiment Identification Determine the sentiment associated with each omitted topic, selecting one of the following:

- "positive": Indicates optimism or favorable conditions such as "beat expectations", "surpass expectations", "confidence", "bullish", "optimistic", ...
- "negative": Indicates pessimism or unfavorable conditions such as "uncertainty", "challenges", "caution", "concerns", "bearish", "underperform expectations", ...
- "neutral": Indicates neutrality or conditions in line with expectations such as "in line with", "meet consensus", ...

8. Keyword Identification

- Identify one or two keywords used by the analyst that best characterize the "Sentiment" variable.
- If no keywords can clearly characterize the sentiment, use "".

9. Time Outlook Identification Assign the time outlook based on the report's content. Remember that the report was written in the year {Year}. Possible values:

- "past": Events that have already occurred.

- "present": Events happening now.
- "near-future": Events expected within 1-3 years or short-term guidance.
- "distant-future": Events expected in more than 3 years or long-term guidance.
- "unclear": When the time frame is not clear.

10. Snippet Identification

- Identify a 5 word snippet from the text to capture where the topic is first discussed.

11. Output Format For each topic that is not already included in the previous analysis, provide your answer in the following format:

```
{Topic}//{Category}//{Channel}//{Sentiment}//{Keyword}//{Timeoutlook}//{Snippet}
```

- `{Topic}` is the specific omitted topic.
- `{Category}`: "firm," "industry," or "macroeconomy."
- `{Channel}`: "earnings/cash flow: sales", "earnings/cash flow: costs", "earnings/cash flow: margins", "earnings/cash flow: earnings", "discount rate", "valuation", or "other".
- `{Sentiment}`:"positive", "negative", "neutral", or "unclear."
- `{Keyword}`: the best keyword that captures the sentiment associated with the topic.
- `{Timeoutlook}`:"past", "present", "near-future", "distant-future", or "unclear."
- `{Snippet}`: 5 word snippet.

12. Example Here is an example of a text with the associated output to help you: "Investment summary: Inflation will raise the cost of inputs while also reducing demand. We have a negative outlook. Despite lagging behind its peers in the EV market, we remain optimistic about the firm's ability to increase its market presence over the next 5 years, as reflected in our hold recommendation. Pricing pressure in the company's main segment has compressed the firm's margins." In this example, the output should be:

```
inflation//macroeconomy//earnings/cash flow: costs//negative//negative
↪  outlook//near-future//Inflation will raise the cost
```

```
market share//firm//earnings/cash flow: sales//positive//remain
↪  optimistic//distant-future//Despite lagging behind its peers
pricing pressure//industry//earnings/cash flow:
↪  margins//negative//pressure//past//Pricing pressure in the company
```

This should be the only output you provide.

Step 3: Generate Standardized Labels for the Topics

Prompt input: Equity report, combined previous analysis, standardized topic lists.

Prompt output: standardized labels.

1. Main Task Identify the standardized labels that best match **each** topic from the list below.

- It is crucial to label every topic from the Topic List.

3. Previous Analysis

`<topic_list>`
`{Topic List}`
`</topic_list>`

Previous analysis output structure is:

`{Topic}//{Sentiment}//{Timeoutlook}//{Snippet}`

Where

- `{Topic}` is a topic.
- `{Sentiment}`:"positive", "negative", "neutral", or "unclear."
- `{Timeoutlook}`:"past", "present", "near-future", "distant-future", or "unclear."
- `{Snippet}`: is a 5 word snippet to help you identify where in the report this topic is discussed.

3. Equity Report

`<equityreport>`
`{EquityReport}`
`</equityreport>`

4. Instructions 1

- Use the "Standardized Label List 1" to label each topic if there is a clear and direct match in the following list:

```
<Standardized Label List 1>
{Standardized Label 1}
</Standardized Label List 1>
```

To apply the standardized labels "undervalued," "overvalued," and "fair-valued," the topic must involve a comparison of the current valuation with either the firm's peers or a historical reference point.

- Undervalued: Use this label when the topic suggests the stock is priced lower than its peers or historical benchmarks. Look for expressions such as "inexpensive," "undervalued," "at a discount," "cheap," or "trades below."
- Overvalued: Apply this label when the topic indicates the stock is priced higher than its peers or historical norms. Key expressions include "expensive," "overvalued," "trades at a premium," or "trades above."
- Fair-Valued: Use this label when the topic implies the stock's valuation is in line with its peers or historical standards. Look for phrases like "compares to," "similar to" or "trades on par."

If none of the labels are a clear and direct match for the topic, try this list instead as a last resort:

```
<Standardized Label List 2>
{Standardized Label 2}
</Standardized Label List 2>
```

Important: It is crucial to only use the labels provided in those two lists.

5. Refinement

- Organize the final list logically.

6. Output Format For each topic, provide your answer in the following format:

```
{Topic}//{Standardized Label}
```

- `{Topic}` is the specific omitted topic.
- `{Standardized Label}`: is the label that best matches the topic.

This should be the only thing that you output!

12. Example

```
Increased sales in Asia//sales and revenues
New competitors//new entrants
US GDP growth//gross domestic product (gdp)
```

make sure that you labeled all the topics contained in topic list.

Appendix B List of Categories and Topics

Table B1: Main topics This table presents all topics included in *Standardized Label 1* used in the prompt presented in Step 3 of Appendix A.

Category	Topic	Category	Topic	Category	Topic
Pricing	Pricing power	Corp. investment	Investment opportunities	Legal	Legal environment
Pricing	Pricing pressure	Corp. investment	Automation and robotics	Legal	Lawsuits and settlement
Pricing	Pricing strategy	Prod./Utilization	Operating leverage and scalability	Regulation	Regulatory environment and compliance
Product mkt. comp.	Barriers to entry	Prod./Utilization	Utilization	Regulation	Cybersecurity and data protection
Product mkt. comp.	Collusion	Financing	Credit rating	Regulation	Regulatory approval
Product mkt. comp.	Cost structure and efficiency	Financing	Credit availability and access	Domestic gov.	Political environment
Product mkt. comp.	Customer concentration	Financing	Debt level/leverage	Domestic gov.	Fiscal environment
Product mkt. comp.	Customer loyalty and retention	Financing	Cash and liquidity	Domestic gov.	Government debt
Product mkt. comp.	Economies of scale	Financing	Stock offerings and issuances	Domestic gov.	Taxation
Product mkt. comp.	Market share	Financing	PE/VC/alternative financing	Domestic gov.	Subsidies
Product mkt. comp.	Market size	Financing	Risk management/hedging	Domestic gov.	Government procurement
Product mkt. comp.	Monopoly	Payout	Stock buybacks and repurchases	Domestic gov.	Immigration
Product mkt. comp.	New entrants	Payout	Dividend policy and strategy	Domestic gov.	Infrastructure
Product mkt. comp.	Oligopoly	M&A	Spin-off and divestitures	Housing market	Housing supply
Product mkt. comp.	Threat of substitute	M&A	Synergy and integration	Housing market	Mortgage rates
Product mkt. comp.	Supplier concentration	M&A	Mergers and acquisitions (M&A)	Housing market	Housing demand
Product mkt. comp.	Vertical integration	M&A	Partnership/Joint-venture	Housing market	Home prices
Product mkt. comp.	Market expansion (industry)	M&A	Restructuring	Demographic	Demographic changes
Product mkt. comp.	Market expansion (geography)	M&A	Antitrust	Demographic	Income inequality
Customer	Customer acquisition	Contracting	Licensing and royalty	Crisis	Natural disasters
Customer	Customer demand/spending	Contracting	Contract terms and duration	Crisis	Pandemic and public health issues
Brand	Brand equity and reputation	Contracting	Covenants	Crisis	Wars and terrorism
Brand	Marketing and advertising	Governance	CEO/Senior management	Crisis	Economic sanctions
Product	Product/service mix	Governance	Management guidance	Raw input	Commodity/raw material prices
Product	Product/service quality	Governance	Corporate governance	Raw input	Oil and gas prices
Product	Product/service positioning and different.	Governance	Corporate ownership structure	GDP growth	Gross domestic product (GDP)
Product	Product life cycle and obsolescence	Labor market	Corporate social responsibility (CSR)	GDP growth	Economic recovery
Product	New products/services	Labor market	Employee safety	GDP growth	Economic recession
Product	Product recall	Labor market	Hiring and layoff	Inflation	Inflation and CPI
Supply chain	Production capacity and planning	Labor market	Labor strikes	Foreign exchange	Foreign exchange and currency rates
Supply chain	Transportation and logistics	Labor market	Pensions	International	International trades/import and exports
Supply chain	Inventory management	Labor market	Employee satisfaction	International	Globalization
Supply chain	Order fulfillment/backlog/backorder	Labor market	Corporate culture	International	Tariffs
Supply chain	Supply chain agility and flexibility	Labor market	Outsourcing	International	Geopolitical environment
Supply chain	Supplier relationships and bargaining power	Labor market	Wages	Industry	Industry life cycle
Innovation / R&D	Products pipeline	Labor market	Unions	Discount rate	Discount rates/WACC/Cost of capital
Innovation / R&D	Innovation and R&D	Labor market	Unemployment rates	Systematic risk	Risk premium/Systematic risk
Innovation / R&D	Clinical trials	Labor market	Labor force participation	Interest rate	Interest rate/Treasury yield
Innovation / R&D	Trade secrets and patents	Labor market	Skills shortages	Interest rate	Monetary policy
Innovation / R&D	Intellectual property (IP)	Environment	Environmental policies and ESG	Valuation	Undervalued
Innovation / R&D	AI and cognitive applications	Environment	Pollution and toxic waste	Valuation	Overvalued
Innovation / R&D	Technology lifecycle	Environment	Weather and climate change	Valuation	Fair-valued
Corp. investment	Capital expenditure (Capex)	Environment	Energy transition		

Table B2: Alternative topics This table presents all topics included in *Standardized Label 2* used in the prompt presented in Step 3 of Appendix A.

Category	Topic	Category	Topic	Category	Topic
Undetermined	Profitability Margins	Undetermined	Cash Flows	Undetermined	Discounted Cash Flow (DCF)
Undetermined	Sales And Revenues	Undetermined	Industry Trends	Undetermined	Sum-Of-The-Parts (SOP)
Undetermined	Costs And Expenses	Undetermined	Macroeconomic Trends	Undetermined	Other
Undetermined	Earnings, EBITDA, And EPS	Undetermined	Valuation Multiples		

Appendix C Text Excerpt and Output

Excerpt (Morgan Stanley, Airbus SE, 2005-03-07): "This is the primary reason why a large portion of Airbus' current deliveries reflect orders placed when the dollar was strong and why **future deliveries will be negatively** impacted by the current weakness in the dollar, regardless of where the spot rate actually goes in two to three years' time. Customer Mix: A **higher proportion of deliveries to** LCCs has negative implications for realized pricing. The major low-cost carriers (such as Jetblue, easyJet and Ryanair) have **gained negotiating leverage by placing** large orders for volume discounts and reducing operating costs. Given low-cost airlines' propensity for operating a single-sourced fleet to reduce recurring operating costs, a manufacturer can justify a slim margin on an initial order in the hope of locking in a new long-term customer."

Output structure:

- {Currency exchange rate}//{macroeconomy}//{earnings/cash flow: sales}//{negative}//{negative impact}//{near-future}{**future deliveries will be negatively**}
- {Customer mix}//{industry}//{earnings/cash flow: margins}//{negative}//{negative implications}//{present}//{**higher proportion of deliveries to**}
- {Low-cost carrier negotiations}//{industry}//{earnings/cash flow: sales}//{negative}//{volume discounts}//{present}//{**gained negotiating leverage by placing**}

Appendix D Appendix Figures and Tables

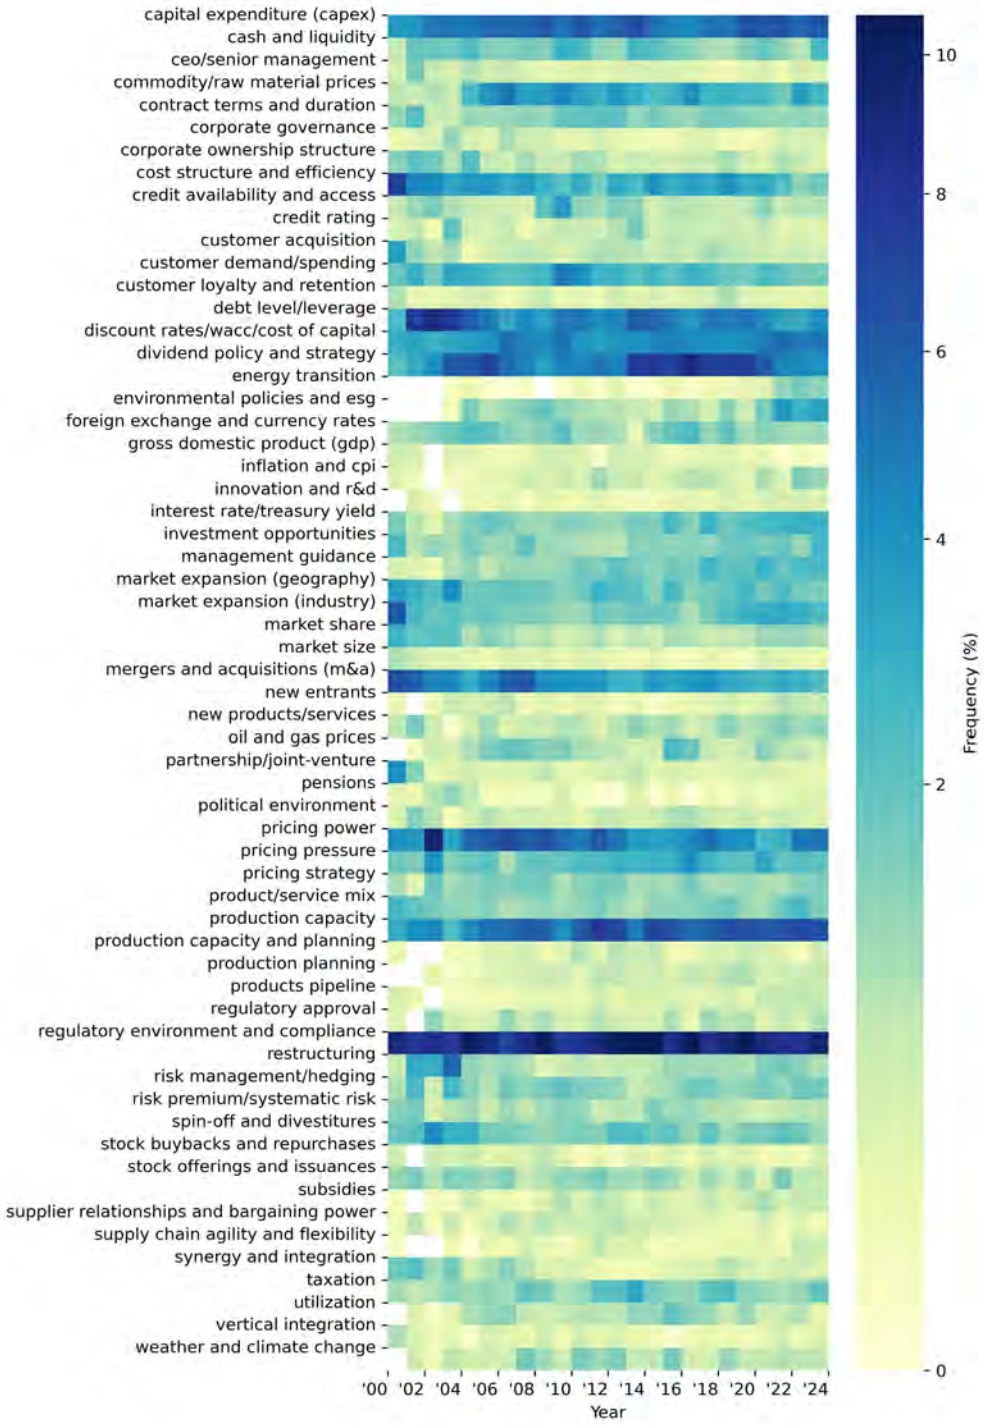


Figure D.1:
Attention Allocation Over Time – Utilities

This figure plots the attention allocation to the 60 most important topics for the utilities firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

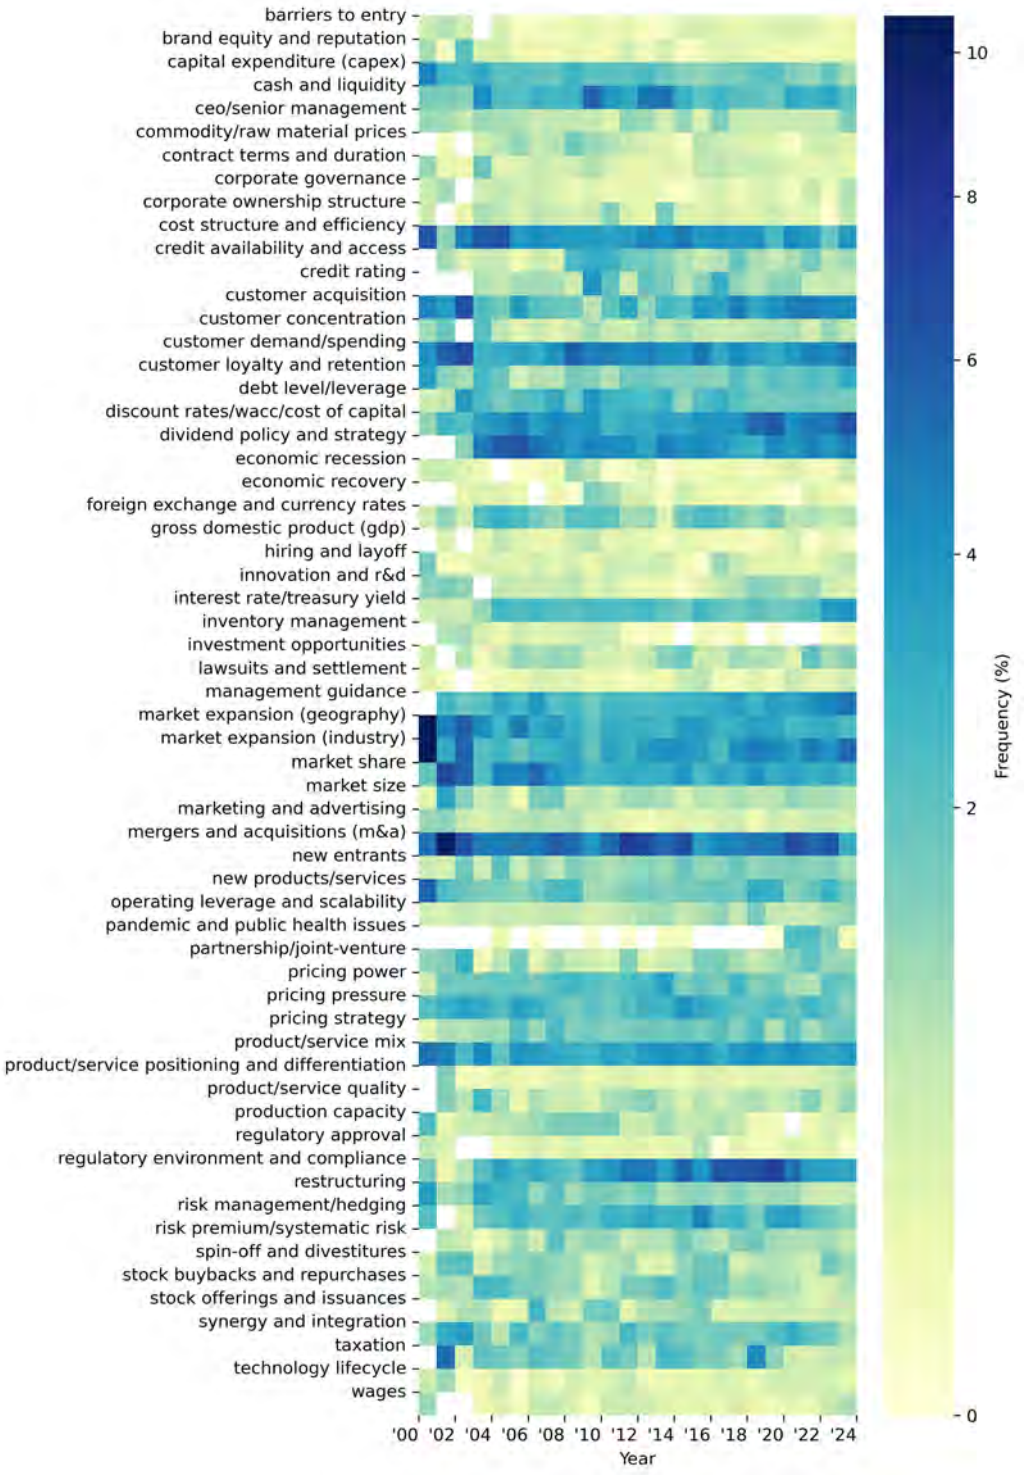


Figure D.2:
Attention Allocation Over Time – Finance and insurance

This figure plots the attention allocation to the 60 most important topics for the finance and insurance firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

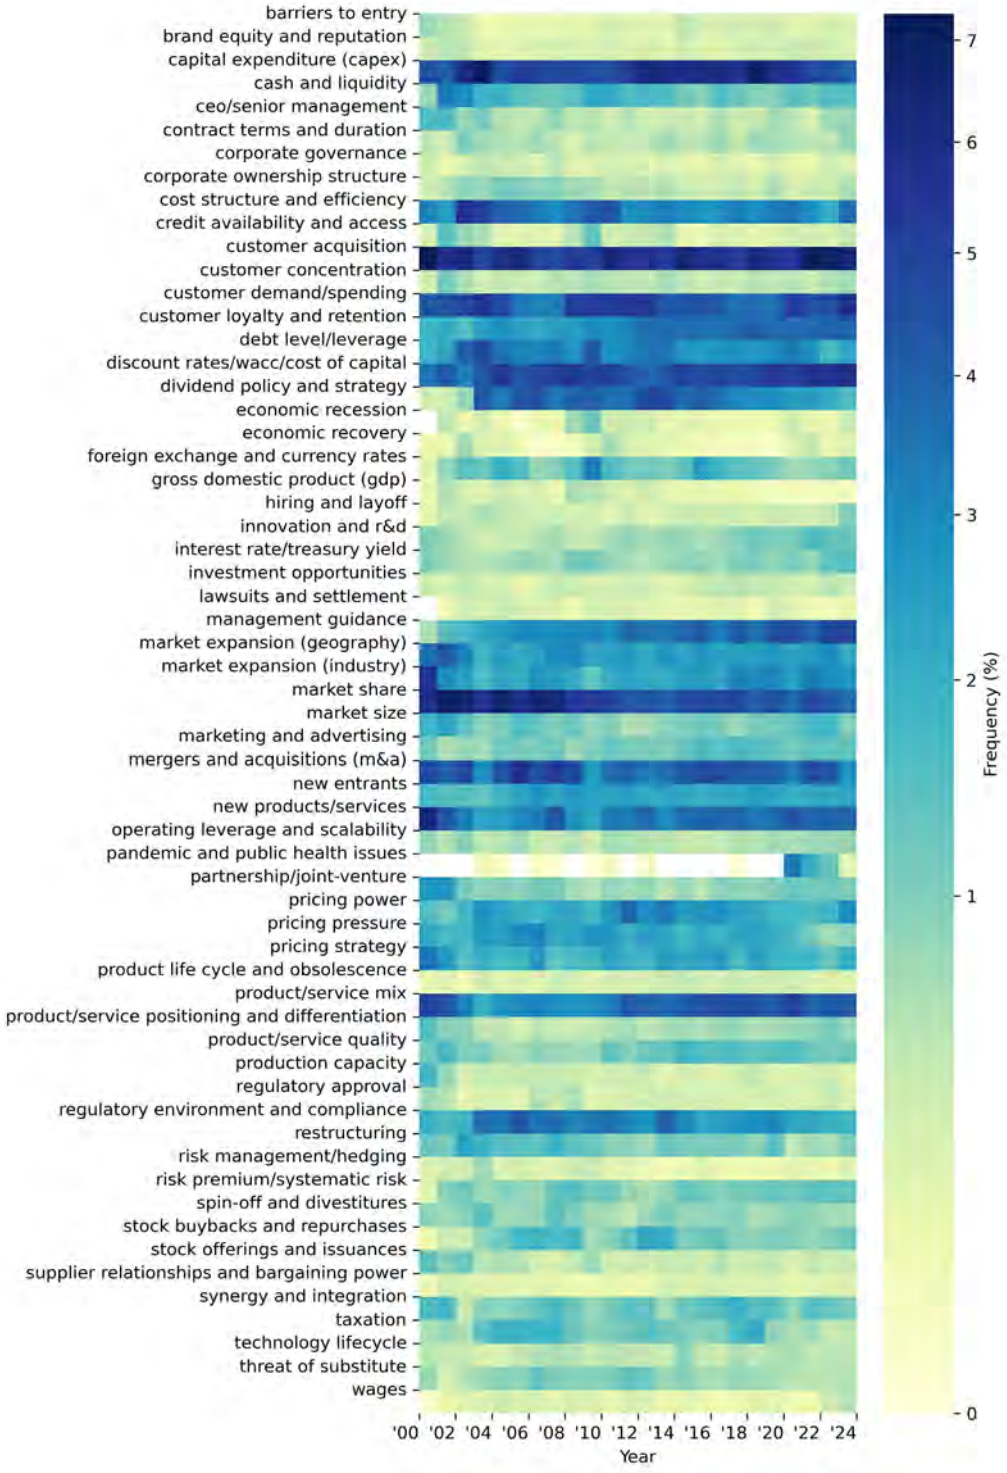


Figure D.3:
Attention Allocation Over Time – Information

This figure plots the attention allocation to the 60 most important topics for the information firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

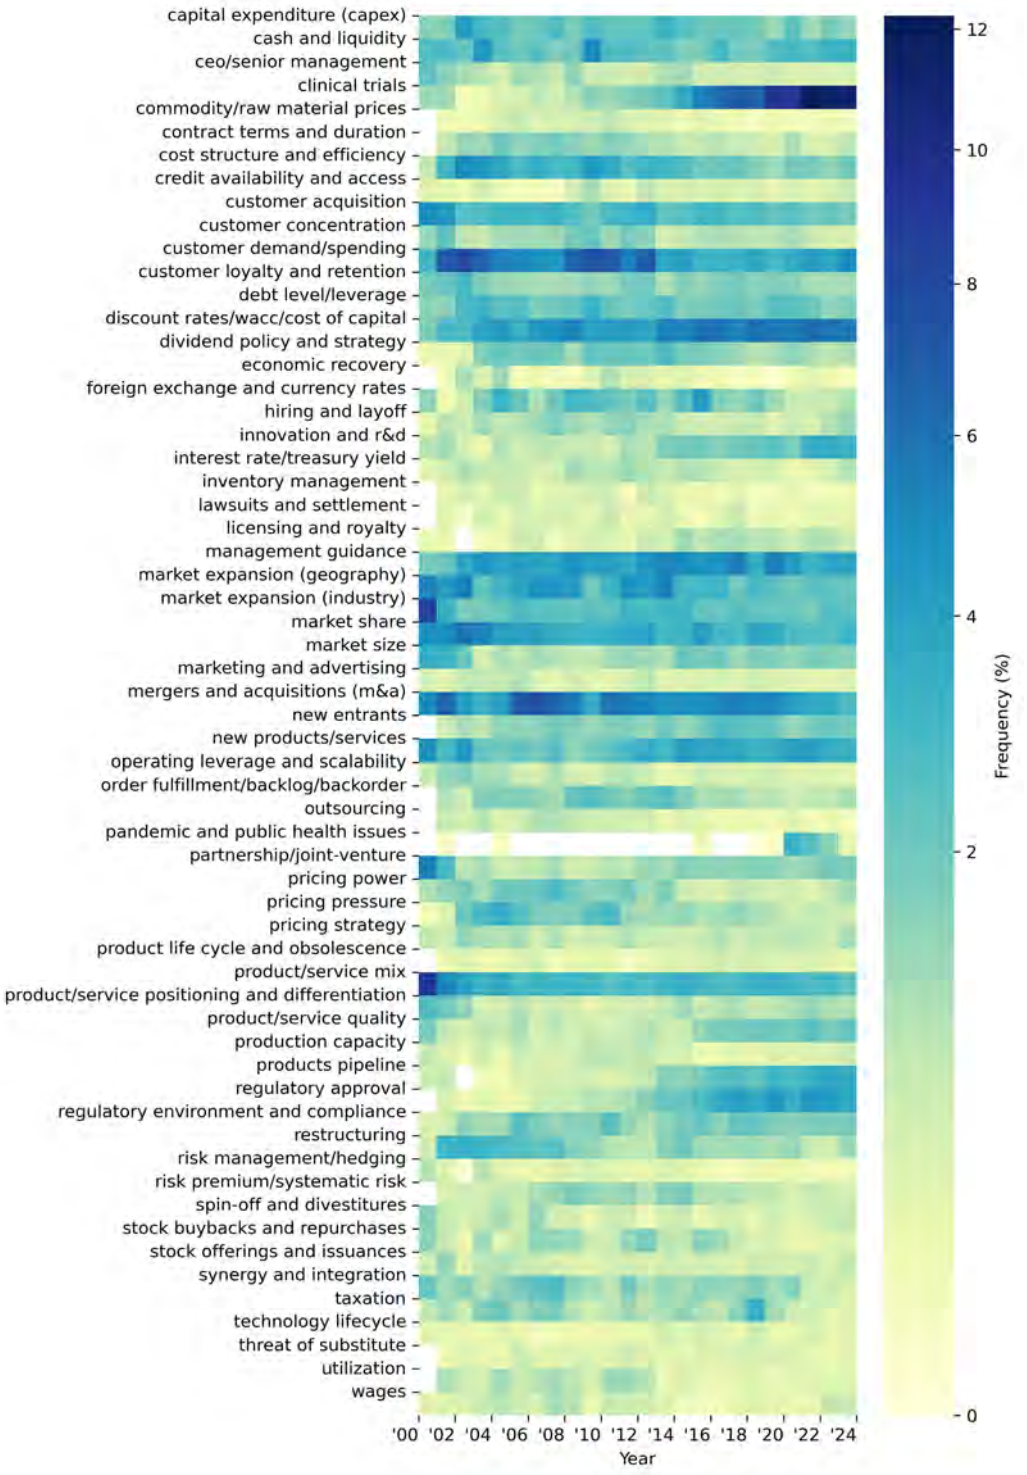


Figure D.4:
Attention Allocation Over Time – Professional, Scientific, and Technical Services

This figure plots the attention allocation to the 60 most important topics for the professional services firms in our sample for the period 2000–2023.The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

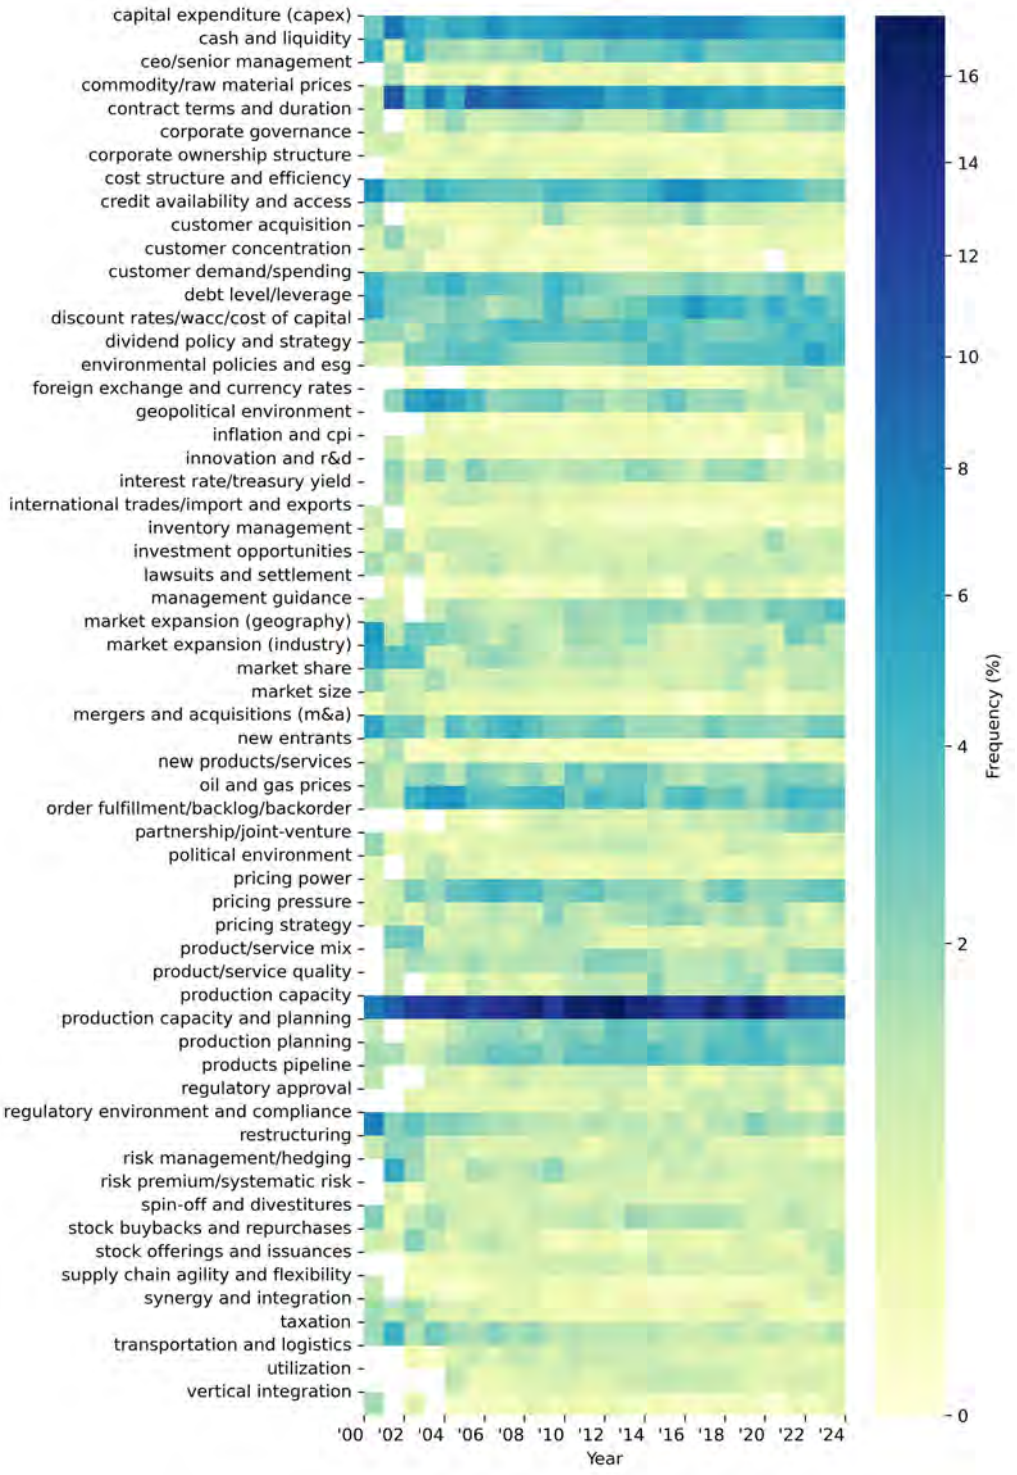


Figure D.5:
Attention Allocation Over Time – Mining / Oil & Gas

This figure plots the attention allocation to the 60 most important topics for the mining and oil & gas firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

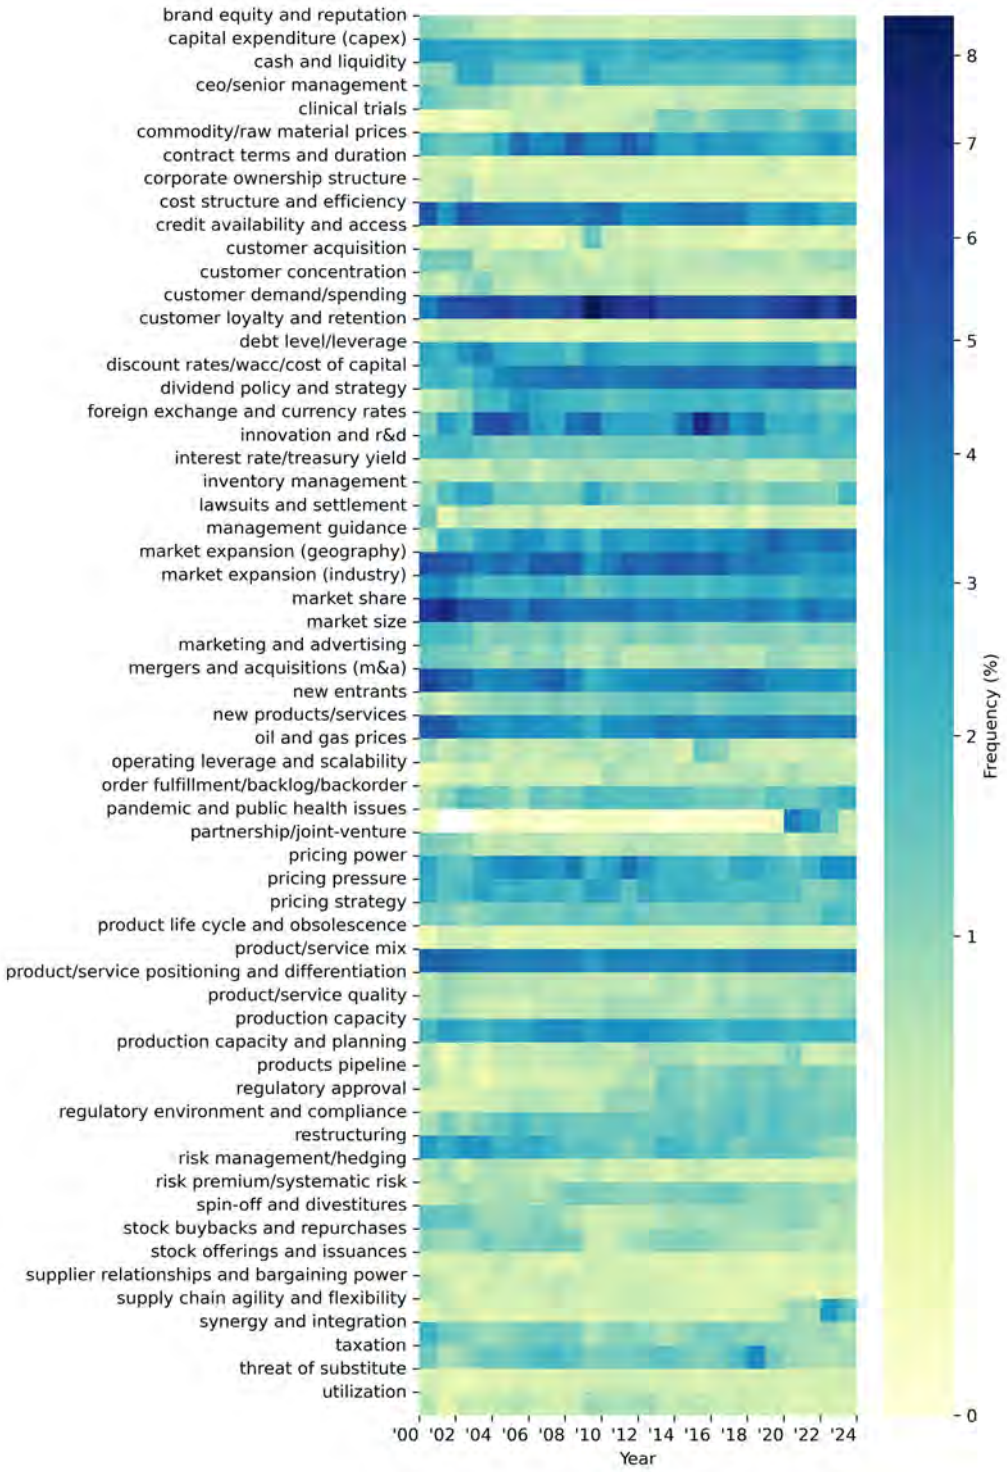


Figure D.6:
Attention Allocation Over Time – Manufacturing

This figure plots the attention allocation to the 60 most important topics for the manufacturing firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

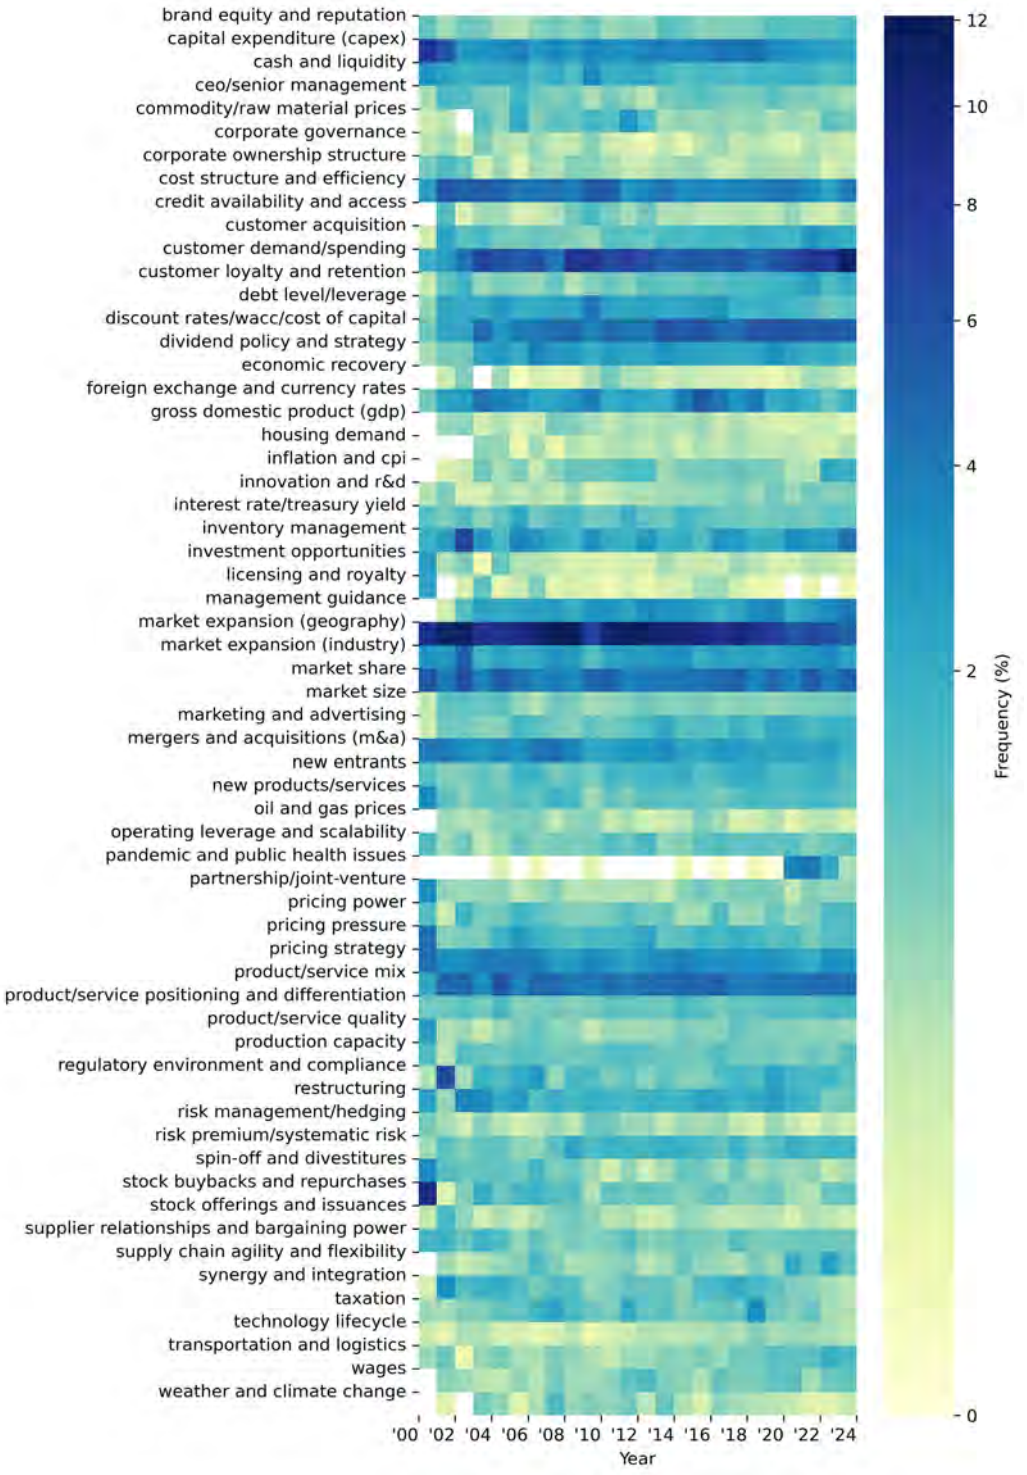


Figure D.7:
Attention Allocation Over Time – Retail Trade

This figure plots the attention allocation to the 60 most important topics for the retail trade firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

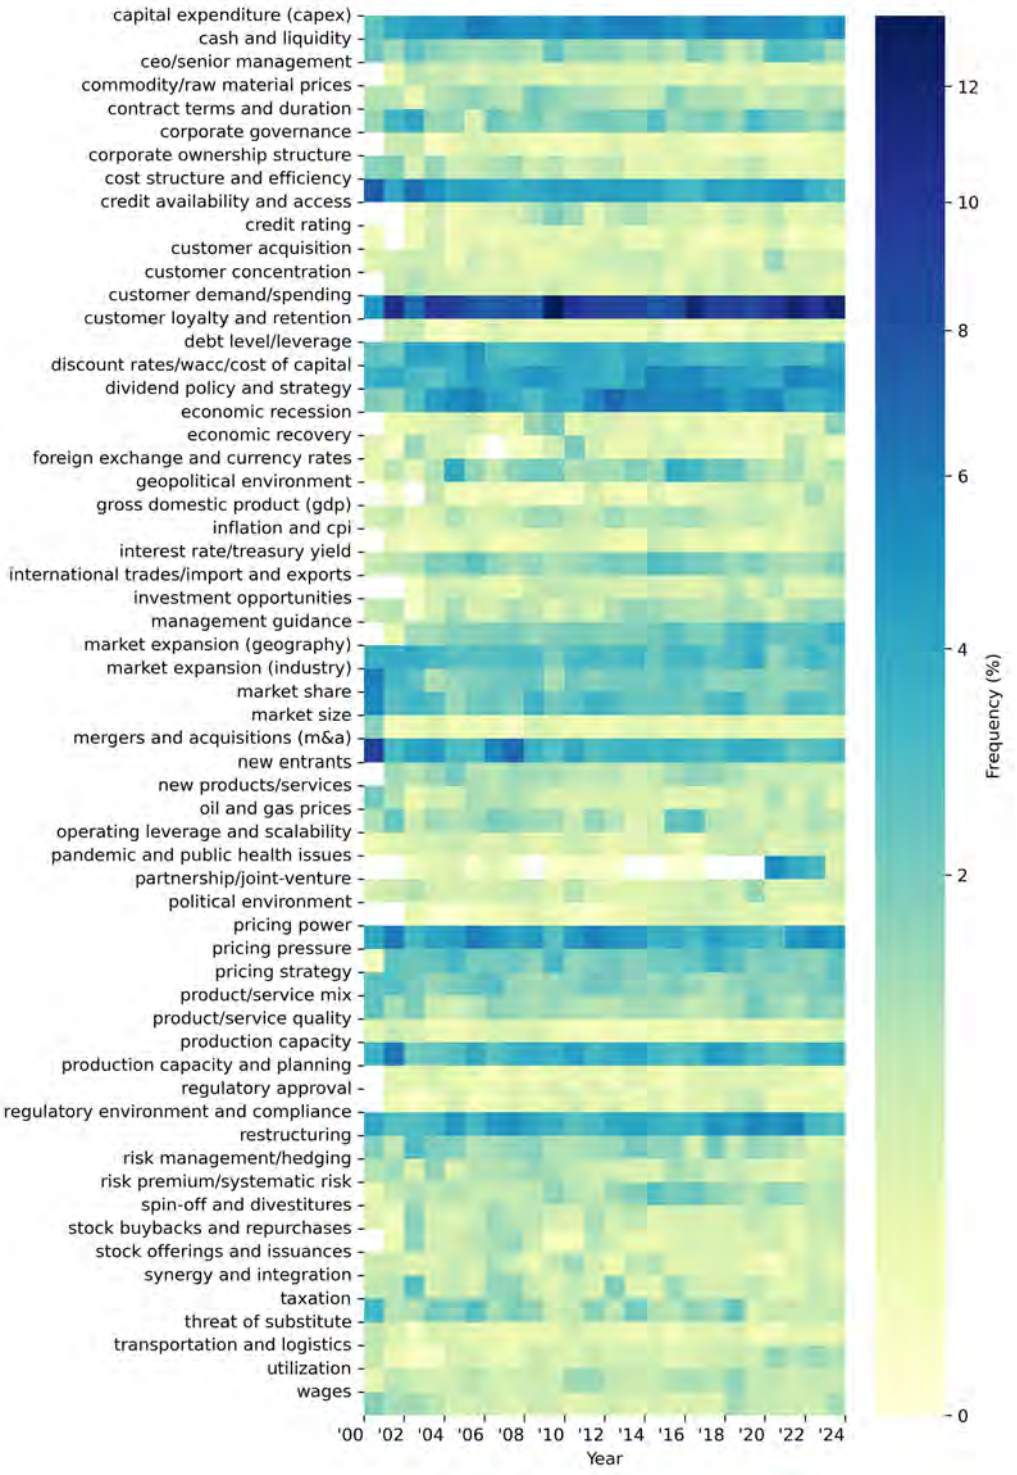


Figure D.8:
Attention Allocation Over Time – Transportation

This figure plots the attention allocation to the 60 most important topics for the transportation firms in our sample for the period 2000–2023. The x-axis is the year the report is produced, and the y-axis denotes the top 60 topics (out of 139 possible topics pre-selected by the research team). The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

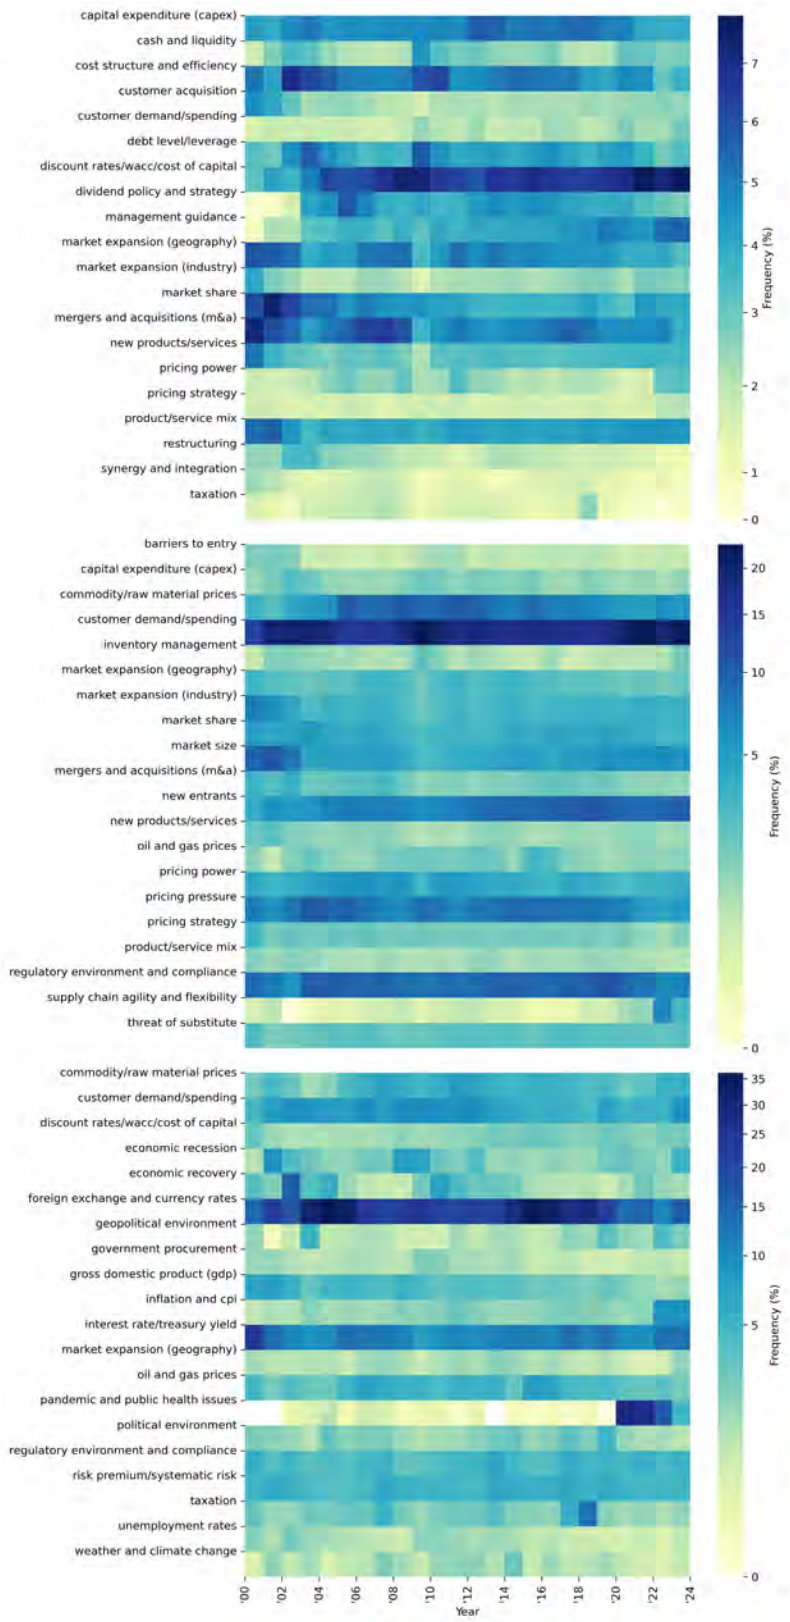


Figure D.9:
Attention Allocation Over Time – Entities

This figure plots the attention allocation to the 20 most important topics for the three types of entities we focus on in the sample–firm, industry, and macroeconomy–for the period 2000–2023. The first block denote the Firm topics. The second block denotes the industry topics. The third block denote the macroeconomy topics. The color-coded heat map is shown on the left, with yellow shades indicating lower values, and dark blue shades the largest magnitude measure for the sample. The heat map uses a non-linear gamma scale to better illustrate the gradients of intensity over the entire range of possible values. The Frequency (%) corresponds to the share of the average report that discusses a given topic in a specific year. For example, a value of 5% for topic A in 2005 indicates that the average report written in 2005 would have allocated 5% of its focus to topic A.

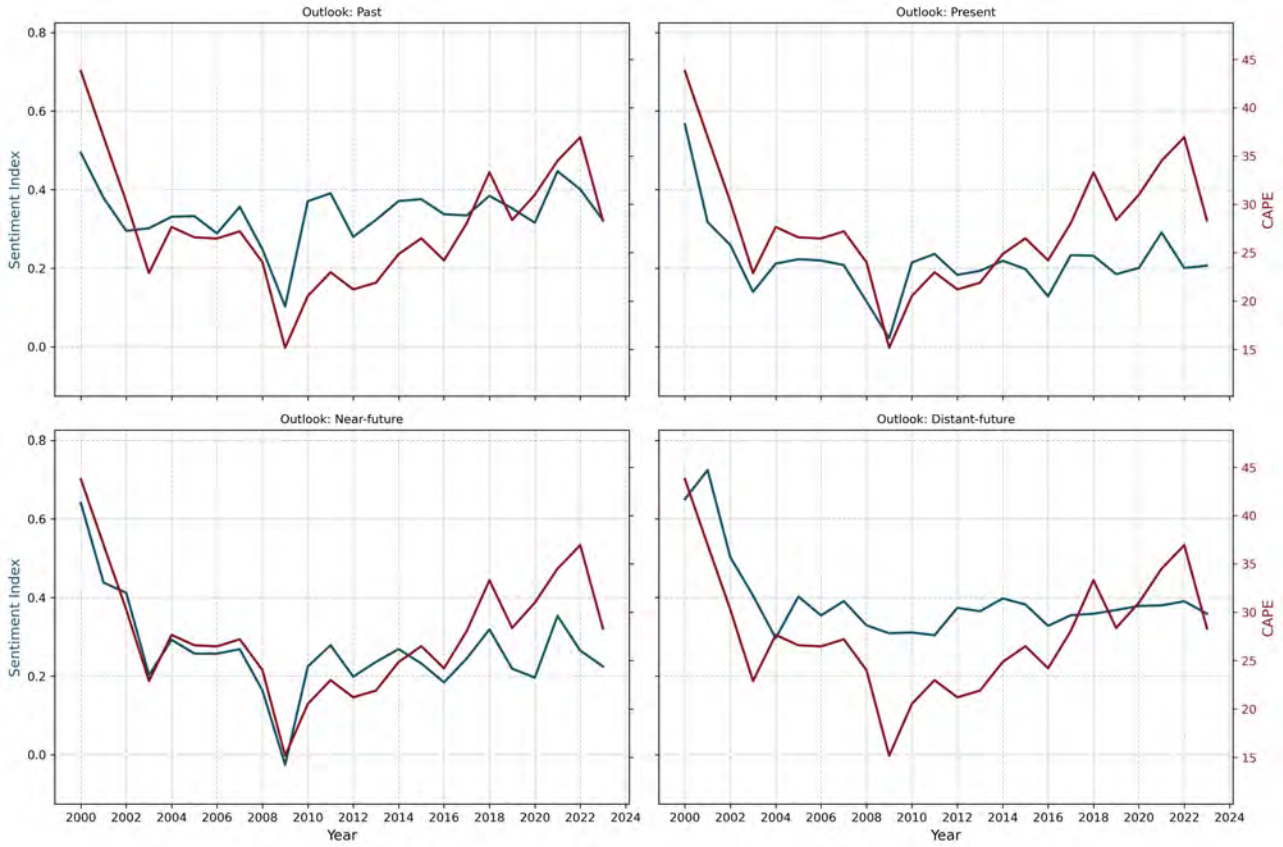


Figure D.10:
Average Sentiment by Time Outlook - US Firms

This figure plots Shiller's CAPE and the average sentiment over time for each time outlook in our sample: past, present, near-future (1-3 years), distant-future (> 3 years) for the US firms in our sample for the period 2000–2023. The x-axis is expressed in years. The y-axis denotes the sentiment associated with the average report written in a given year (the average is measure using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports). The correlation between the past series and CAPE is 0.72. The correlation between the present series and CAPE is 0.80. The correlation between the near-future series and CAPE is 0.83. The correlation between the distant-future series and CAPE is 0.68.

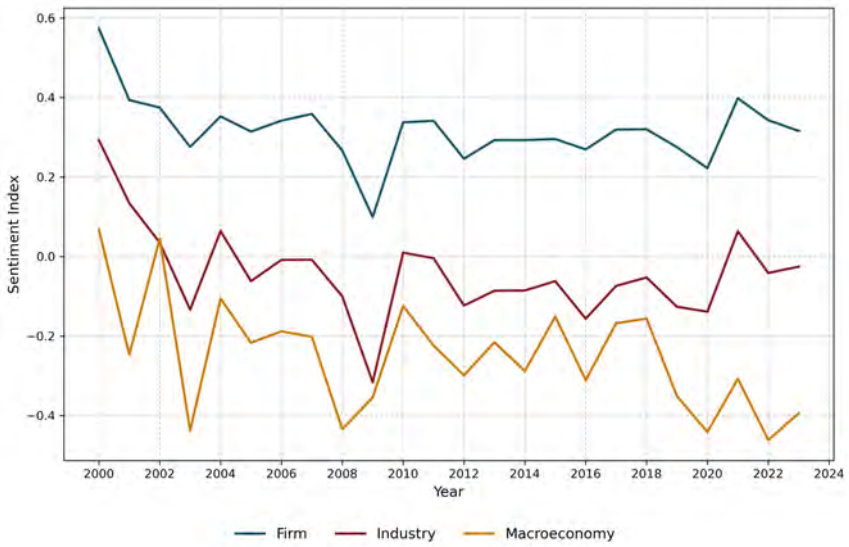


Figure D.11:
Average Sentiment by Entities

This figure plots the average sentiment for the three type of entity evaluated in our date collection (i.e., firm, industry, and macroeconomy). The x-axis is expressed in years. The y-axis denotes the sentiment associated with the average report written in a given year (the average is measure using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports).

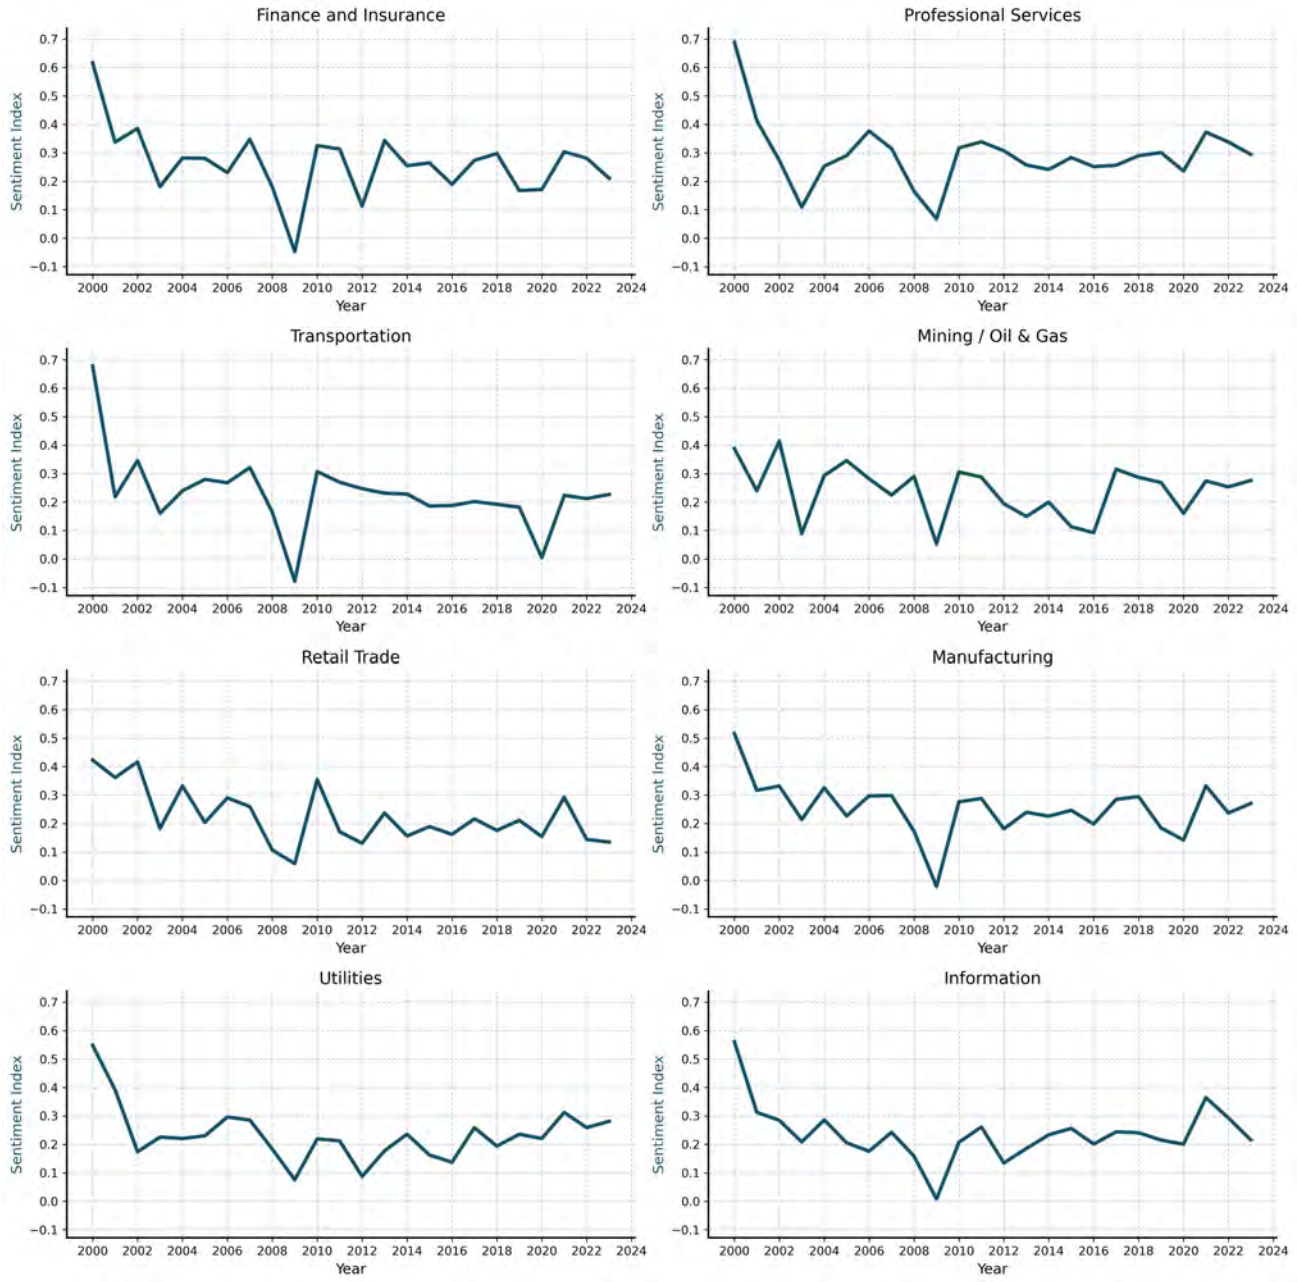


Figure D.12:
Average Sentiment by Industry

This figure plots the average sentiment over time for the 8 major industries in our sample. The x-axis is expressed in years. The y-axis denotes the sentiment associated with the average report written in a given year (the average is measure using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports).

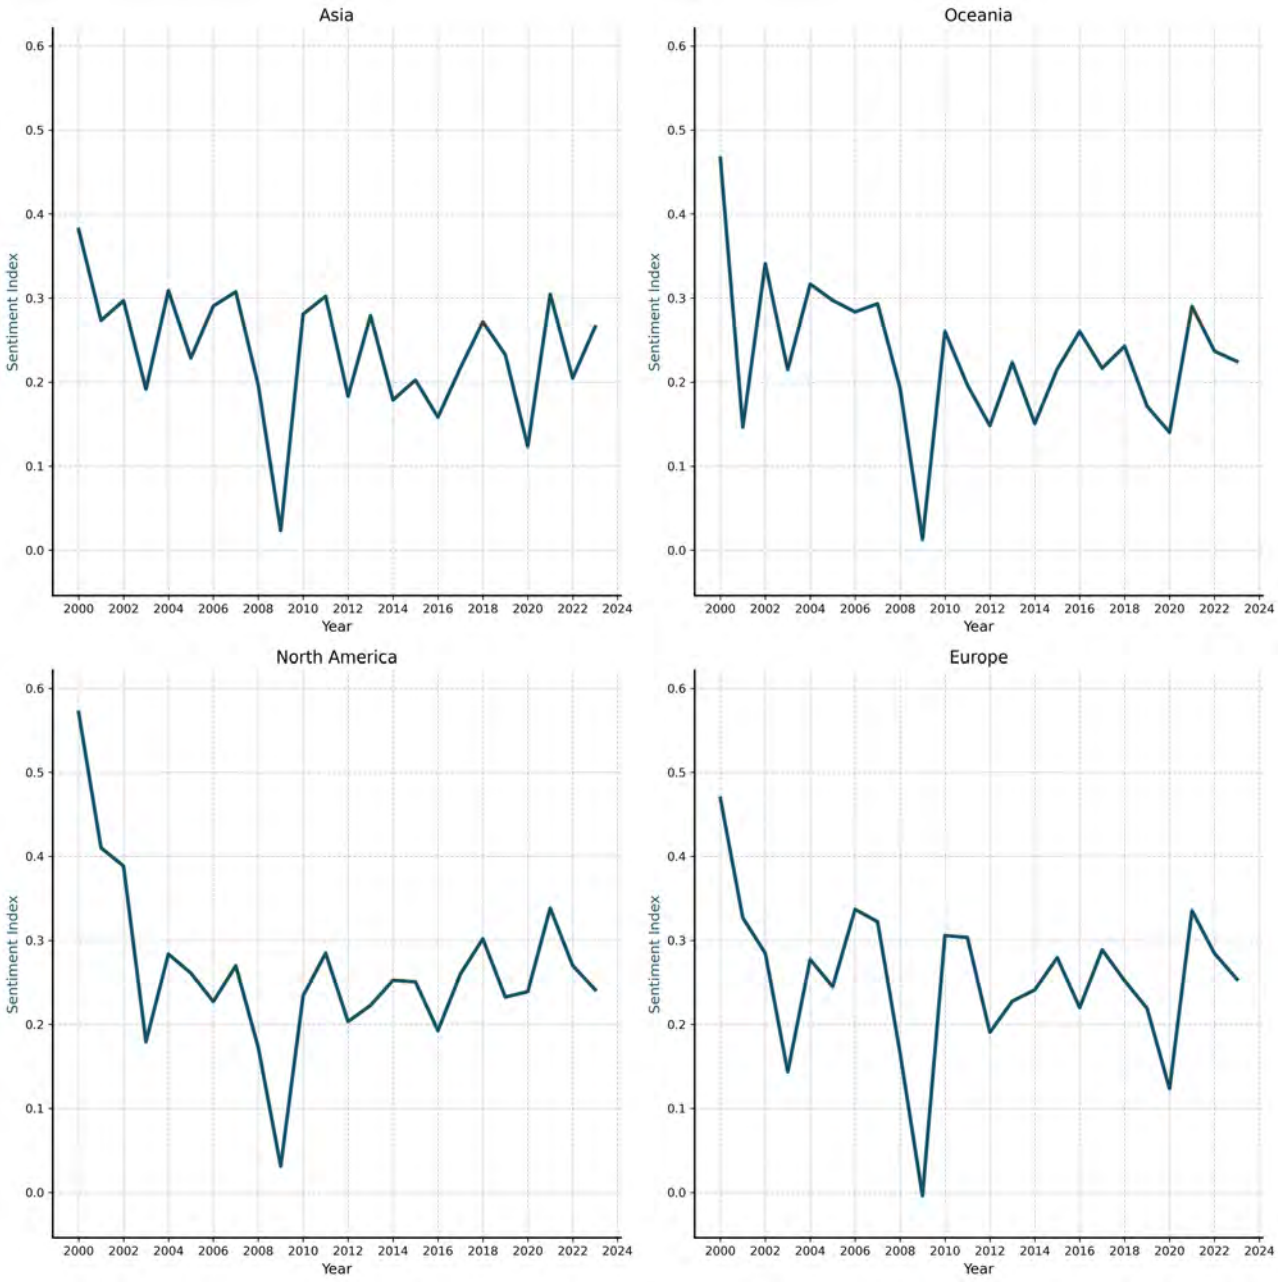


Figure D.13:
Average Sentiment by Continent

This figure plots the average sentiment over time for the 4 major continents in our sample. The x-axis is expressed in years. The y-axis denotes the sentiment associated with the average report written in a given year (the average is measure using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports).

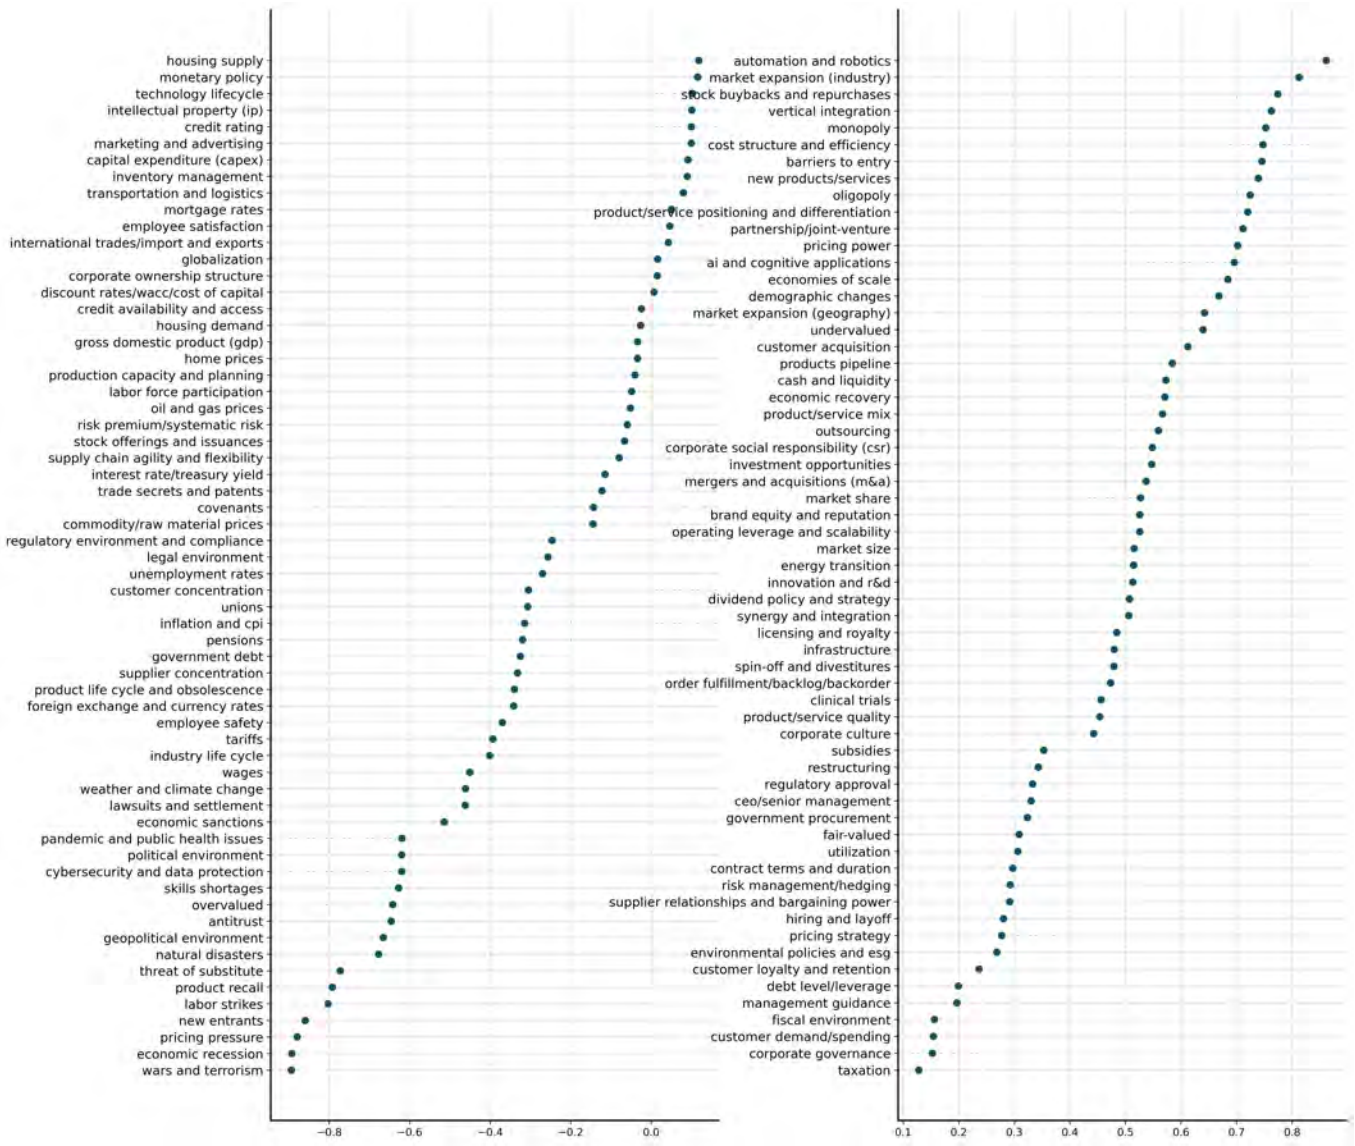


Figure D.14: Topic Sentiment

This figure plots the average sentiment for the most common topics in our sample for the period 2000–2023. The x-axis represents the average sentiment index for those topics across the entire sample. The y-axis indicates the topic label. The sentiment shown in this figure is the average sentiment calculated using values of -1, 0, and 1 for negative, neutral, and positive sentiment, respectively.

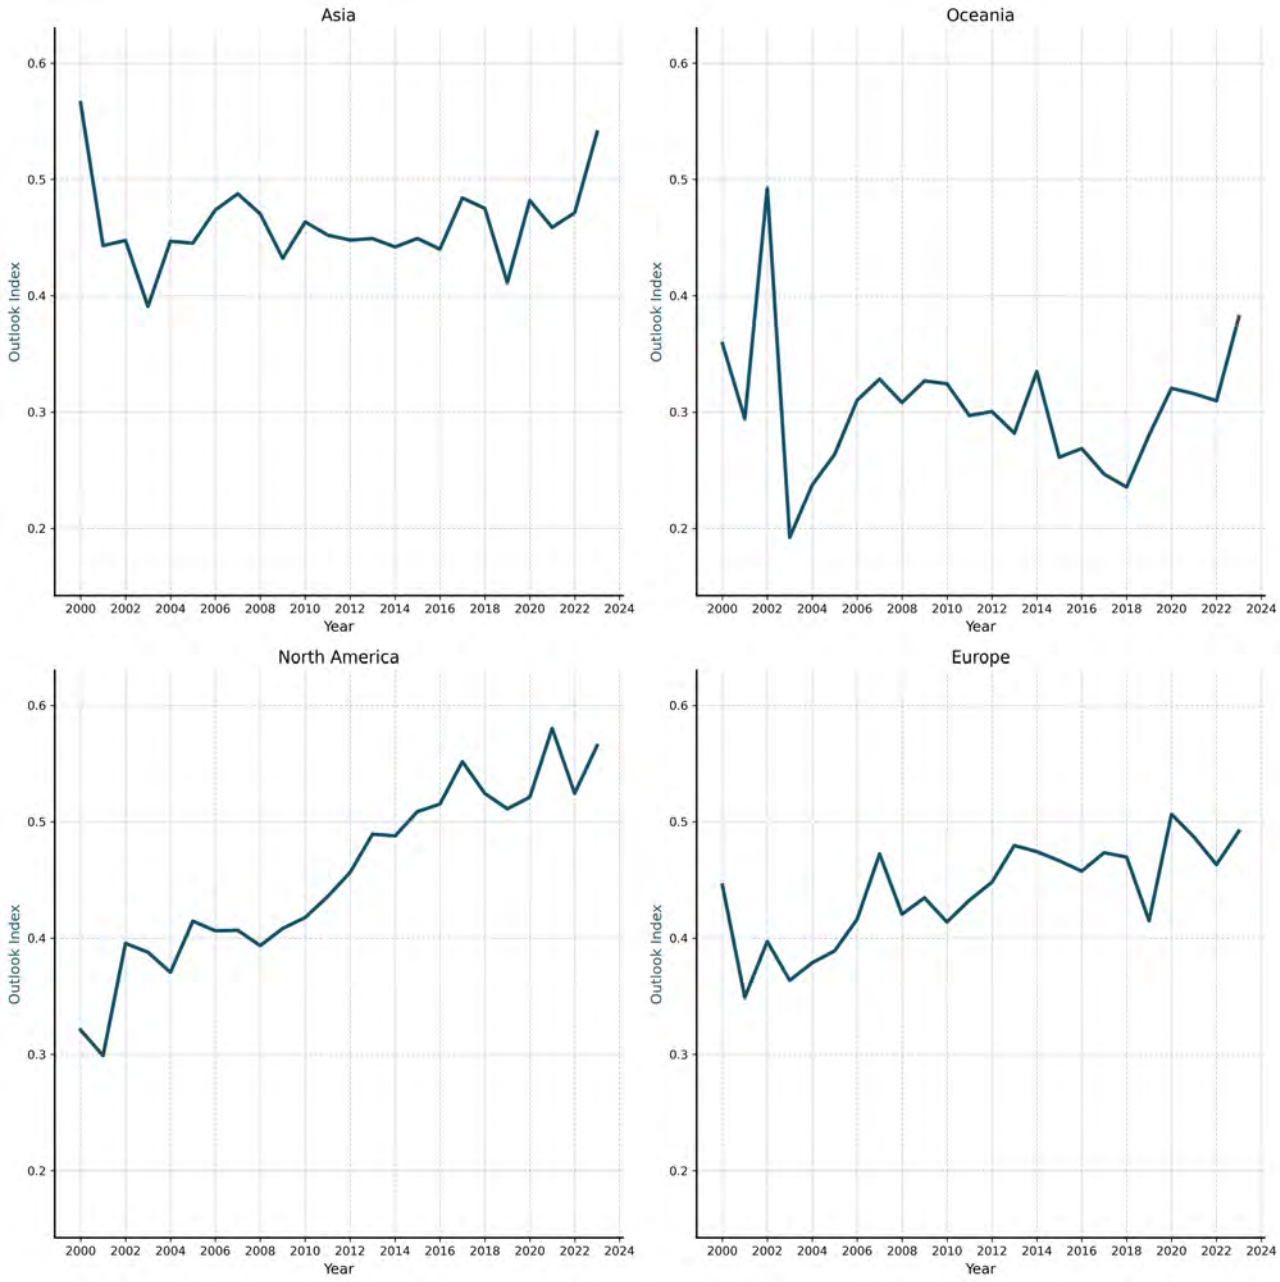


Figure D.15:
Average Outlook by Continent

This figure plots the average sentiment over time for the 4 major continents in our sample. The x-axis is expressed in years. The y-axis denotes the sentiment associated with the average report written in a given year (the average is measure using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports).



Figure D.16:
Average Outlook by Industry

This figure plots the average sentiment over time for the 8 major industries in our sample. The x-axis is expressed in years. The y-axis denotes the sentiment associated with the average report written in a given year (the average is measure using values of -1, 0, 1 for negative, neutral, and positive sentiment in the reports).

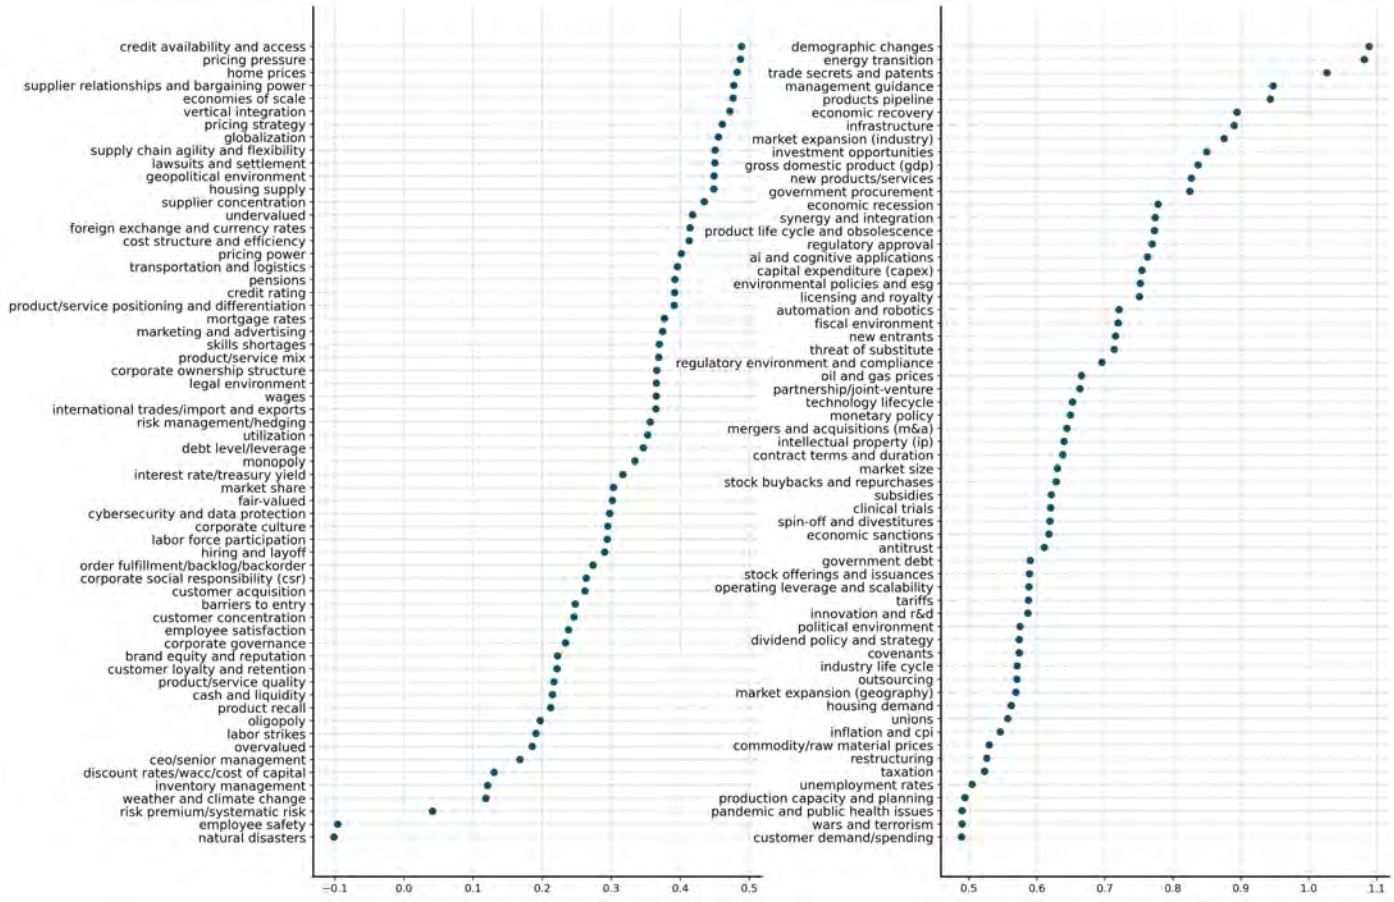
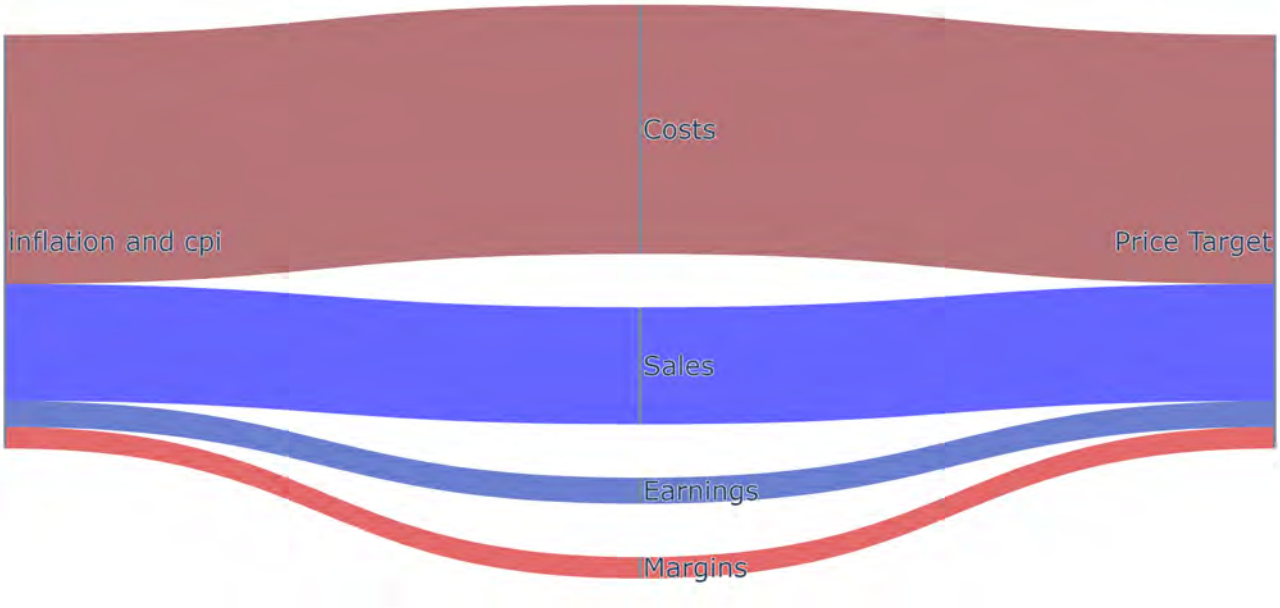


Figure D.17: Topic Outlook

This figure plots the average outlook for the most common topics our sample for the period 2000–2023. The x-axis is the average outlook index measure for that topics in the entire sample. The y-axis indicate the topic label. The outlook shown in this figure is the average outlook measured in the entire sample (the average is measure using values of -1, 0, 1, 2 for past, present, near-future, distant-future outlook).

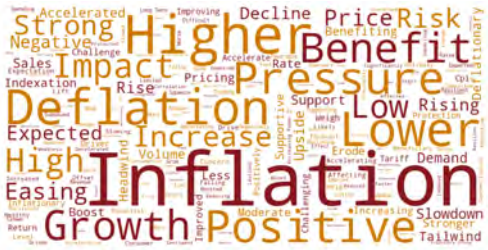
(A) Causal Effect of Inflation on Valuation: Decomposing Cash Flow Channels.



(B) Word Cloud - Costs.



(C) Word Cloud - Sales.



(D) Word Cloud - Margins.



(E) Word Cloud - Earnings.



Figure D.18:
Causal Effect of Inflation

This figure illustrates the narrative that analysts use when discussing how inflation is expected to impact cash flows. In Panel A, the leftmost node indicates inflation, the middle nodes denote the four valuation channels explored in our analysis, and the rightmost node is the valuation outcome. The width of each branch represents the relative importance of the four channels when analysts discuss inflation. Wider branches indicate that this channel receives greater focus. The color of each branch denotes the average sentiment that inflation has on the valuation outcome for that topic. Red means that, on average, analysts believe that the effect of inflation through the specific channel will be negative for the valuation outcome. Conversely, blue indicates that the effect through the specific channel will be beneficial for the valuation outcome. Panels B, C, D, and E present word clouds of the keywords associated with the sentiment variables when analysts discuss inflation for each valuation channel, respectively.

Table D.1:
Sentiment, Location, and Informativeness

This table presents results related to the role of sentiment on the 12-month price target appreciation of forecasters. Sentiment is broken down into quintiles based on the location of the measure in reports. The dependent variable, *ln(Price target-to-earnings ratio)*, corresponds to the natural logarithm of the analyst 12-month price target divided by the earnings per share at the time of the report. $Sentiment^{\text{Quintile}}i, j, t$ denotes the sentiment associated with the part of the reports located in a given quintile. For example, $Sentiment^{\text{0-20th pct}}i, j, t$ corresponds to the sentiment measure in the first quintile of a report. The regressions are estimated using ordinary least squares. Standard errors (in parentheses) are estimated using heteroskedasticity-consistent estimators clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Dependent variable:	ln(Price target-to-earnings ratio)$_{i,j,t}$								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment$^{0-20thpct}_{i,j,t}$	0.13***					0.10***	0.09***	0.08***	0.08***
	(0.01)					(0.01)	(0.01)	(0.01)	(0.01)
Sentiment$^{21-40thpct}_{i,j,t}$		0.10***				0.06***	0.05***	0.05***	0.05***
		(0.00)				(0.00)	(0.00)	(0.01)	(0.01)
Sentiment$^{41-60thpct}_{i,j,t}$			0.09***				0.04***	0.04***	0.03***
			(0.00)				(0.00)	(0.01)	(0.01)
Sentiment$^{61-80thpct}_{i,j,t}$				0.08***				0.03***	0.03***
				(0.00)				(0.00)	(0.01)
Sentiment$^{81-100thpct}_{i,j,t}$					0.07***				0.03***
					(0.00)				(0.01)
Firm*Contributor FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	39,674	39,298	38,536	37,735	37,411	38,975	37,619	35,817	34,020
F Statistics	638.543	515.678	369.769	319.236	184.909	374.409	257.724	189.043	153.414
R^2	0.866	0.864	0.864	0.865	0.865	0.867	0.868	0.870	0.873
Within R^2	0.031	0.021	0.016	0.013	0.008	0.037	0.040	0.041	0.042

Table D.2:

Explanatory Power of Topics Excluded from Reports But Discussed by Peer Analysts

This table presents results about the incremental explanatory power of topics discussed by other analysts evaluating the same firm in the same year, but that are excluded from the analyst's report. The topics discussed by the analysts are obtained from the LLMs strategy. The topics excluded from the report but discussed by peers correspond to all the topics included in the reports of peers evaluating the same firm in the same year, but that are not included in the report of interest.

	Step 1	Step 2
Dependent Variable:	Price-to-Earnings Ratio$_{i,j,t}$	Residual from Step 1
Covariate:	Topic Included in Report	Topic Discussed by Peers, but Excluded from Report
Avg. Adjusted R^2	0.243	-0.177