From Numbers to Words: Breaking Down Institutional Beliefs*

Andrea Andolfatto[†] *Federico Bastianello*[‡]

[Link to the latest version]

February 1, 2026

Abstract

We examine how large asset managers form and justify long-horizon beliefs by analyzing their Capital Market Assumptions (CMAs)—articulated through tables, figures, and narratives. We develop a novel method that transforms CMA text narratives into quantifiable causal networks using large language models, capturing both the complexity of managers' mental models and their allocation of attention across macro-financial topics. Our analysis reveals substantial heterogeneity in asset managers' beliefs, both quantitative and narrative. Using granular numerical data on the building blocks of managers' return expectations, we identify multiple drivers of cross-sectional dispersion in expectations. Text-based measures show that the average coefficient of variation in cognitive complexity exceeds 0.7, while that in topic attention exceeds 1, indicating pronounced dispersion in both how managers reason and how they allocate attention to economic relationships. We further document systematic biases in asset managers' beliefs using both quantitative and textual evidence. Return expectations deviate predictably from objective benchmark forecasts: greater cognitive complexity is associated with larger ex-ante forecast errors, and differences in attention to key building blocks affect the degree of over- or underreaction. Finally, we find evidence of historically anchored expectations in second moments. Overall, while institutional expectations are economically meaningful and linked to objective return predictors, they nonetheless exhibit systematic and predictable deviations from objective benchmarks.

*We thank Svetlana Bryzgalova, John Y. Campbell, Carlo Favero, Nicola Gennaioli, Katrin Gödker, Francisco Gomes, Sam Hanson, Max Miller, Florian Nagler, Anna Pavlova, Marco Sammon, Andrei Shleifer, Adi Sunderam, Jakob Ahm Sørensen, and Claudio Tebaldi for thoughtful comments and discussions. We also thank the participants of the Harvard Behavioral Reading Group and the Bocconi Workshop on Machine Learning and Financial Decision Making (2026).

[†]Bocconi University: andrea.andolfatto@phd.unibocconi.it

[‡]London Business School: fbastianello@london.edu

How do investors form beliefs, and what drives the source of investor heterogeneity? A growing literature studies how investors form, communicate, and act on expectations along three complementary lines. First, surveys of self-reported beliefs document systematic biases—such as extrapolation and both over- and underreaction—that challenge rational expectations and motivate behavioral models of belief formation (Greenwood & Shleifer 2014, Bouchaud et al. 2019, Bordalo et al. 2019, 2024*b*, Bastianello 2025). Second, recent work links self-reported beliefs to portfolio choices and trading, showing that pass-through of beliefs to actions varies across investor types (Giglio et al. 2021, Dahlquist & Ibert 2024, Andonov et al. 2026). Third, text-based approaches use investor-authored narratives and news to extract expectations as well as the underlying economic channels shaping them (Bybee 2023, Décaire & Graham 2024, Bastianello et al. 2024, Sarkar 2025).

We draw on all three tools—numeric expectations, portfolio data, and text—as well as the information flow in news to analyze how large institutional asset managers formulate and rationalize the long-horizon expectations for equity return expectations and volatility reported in their Capital Market Assumptions (CMAs) (Couts et al. 2023, 2024, Dahlquist & Ibert 2024, 2025, Begenau et al. 2025).[1] We "read" their CMAs in two complementary dimensions: in numbers (from tables and figures with forecasts and their decompositions, labeled as "building blocks") and in words (narratives that explain the underlying mechanisms), linking beliefs to portfolio allocations and the arrival of news.

Our research question is straightforward: How does the joint analysis of CMA forecasts, their building blocks and the narratives in CMA documents reveal the mechanisms driving cross-sectional and time-series variation in asset managers' beliefs? Do these beliefs matter investment decisions, and do they exhibit systematic biases?

Previewing our findings, we show that combining granular data on the building blocks of CMA forecasts with their narratives reveals the mechanisms behind asset managers beliefs. We show that cross-sectional variation in return expectations is explained by four main

[1]CMAs are not generic surveys; they are inputs that anchor strategic asset allocation and risk budgeting—indeed, leading institutions describe their CMA framework as "a critical foundation" of multi-asset processes, making them a uniquely informative window into professional belief formation. For example, J.P. Morgan's 2022 Long-Term Capital Market Assumptions (LTCMAs) report explicitly presents its CMA methodology as the cornerstone of its strategic asset-allocation process.

building blocks: expected changes in valuation (40%), expected real growth (22%), dividend yield (17%), and buybacks (10%). Over time, valuation expectations move countercyclically, whereas growth expectations are procyclical. By showing that the growth and valuation building blocks systematically map into portfolio allocations, we highlight the drivers of the pass-through of beliefs-to-action. By contrast, volatility expectations are history-anchored, clustering around long-run averages and adjusting only weakly to realized volatility. We provide a novel approach to extract causal networks from the text of the reports that uncover heterogeneity in asset managers' explanations, highlighting different levels of complexity and attention to topics in their narratives. By quantifying this heterogeneity, we find evidence of systematic biases in return expectations, as ex-ante forecast errors are predictably related to the causal network complexity and attention to topics in the narratives contained in the CMA reports.

We begin by constructing a novel dataset by extracting structured text and return decompositions from Capital Market Assumption (CMA) reports using natural language processing (NLP) and text-parsing techniques. We show that these decompositions can be systematically interpreted through the lens of an approximate accounting identity, revealing that asset managers share mental models grounded in fundamental drivers such as income (dividends and net buybacks), real growth, inflation, and expected valuation changes.

We use granular data on the building blocks provided by asset managers to quantify the components of cross-sectional variation in return expectations and find that four main components dominate, with expected valuation change explaining 40% and expected growth in fundamentals another 22%. This component-level variance decomposition is novel relative to Dahlquist & Ibert (2025). While they identify valuation change as a source of heterogeneity using external proxies for the building blocks, our direct use of CMA disclosures quantifies the relative importance of each component and highlights that growth, in addition to valuation change, contributes substantially to the cross-sectional variation in return expectations.

We then interpret the building blocks through a Campbell–Shiller lens at the asset manager level. Our contribution is to show explicitly the sources of the countercyclicality in

CMAs' total return expectations (Dahlquist & Ibert 2024). Two components primarily drive the cyclical properties of asset managers' return expectations: growth expectations are procyclical, whereas expected valuation changes are countercyclical and, on average, dominate—resulting in overall countercyclical return expectations.

We show the relevance of the building blocks by relating them to the equity allocations of asset managers, which respond more strongly to the growth and valuation-change components than to the income/dividend component. This reveals that the parts of beliefs that differ most across institutions are also the parts that move portfolios. In matched fund-level regressions, the coefficients on nominal growth and valuation change are both significant and economically meaningful, with the coefficient on expected growth nearly twice as large as that on expected valuation change and larger than the univariate coefficient on total return expectations. Taken together, these results confirm that CMA expectations matter for asset allocations and add new evidence on which specific components of return expectations transmit into portfolio allocations.

We complement the numerical evidence by "reading" CMAs in words. Using topic modeling on equity-focused segments of each report, we show that the documents cluster into a small number of stable narrative groups, indicating systematic differences in how institutions frame the macroeconomic environment. We then introduce a novel approach to extract causal networks from asset managers' explanations of their return expectations, using a large-language-model (LLM) pipeline that encodes directed and signed links among macro variables and the building-block drivers of return expectations. We document systematic heterogeneity in both the complexity (path distance, indirect connection ratio, transitivity, or eigenvector centrality) and the attention (the normalized number of causal connections to each topic node) to the building blocks of these causal networks, with average coefficients of variation exceeding 0.7 and 1, respectively—indicating that asset managers differ substantially in the richness, focus, and emphasis of their narratives.

We also study a novel dimension of institutional beliefs: how asset managers incorporate textual news into long-horizon expectations. We treat news as a primitive input to belief updating, rather than as a slow-moving state variable. Unlike valuation ratios or macroeconomic indicators, textual news provides a granular and high-frequency measure of information flow. Using text-based earnings sentiment, we show that long-horizon return expectations increase following positive earnings news, but adjust excessively, revealing systematic overreaction. While prior studies relate beliefs to valuation states or macroeconomic conditions, they do not examine how beliefs respond to incoming textual news. Our findings connect the CMA literature to the news-based literature by using textual news to directly observe how institutional long-horizon expectations update in response to new information.

Having established the properties of return expectations, we then turn to beliefs on volatility. We find that long-horizon volatility expectations are highly conservative and display limited responsiveness to market conditions. Manager identity explains the vast majority of cross-sectional variation in volatility forecasts, indicating that institutions rely on stable, idiosyncratic frameworks for assessing volatility. Reported forecasts cluster tightly around 15–17% per year and adjust only minimally to large swings in realized volatility. Consistent with this anchoring, an equally weighted moving average (EWMA) analysis on the volatility expectations implies that managers place disproportionate weight on very long histories, with a half-life of nearly seven years. Taken together, these findings highlight a sharp contrast between forward-looking return expectations and backward-looking, history-anchored volatility beliefs.

We finally turn to the text data to better identify the sources of heterogeneity and dynamics in institutional investors' return expectations. We show that the structure and content of asset managers' causal narratives (quantified by our extracted causal networks) generate systematic and predictable biases in long-horizon return expectations. Narrative complexity does not shift forecast levels directly, but persistent differences in complexity shape how managers process news: managers with structurally more complex causal networks underreact to positive fundamentals and produce more pessimistic forecasts relative to objective benchmarks. Topic-specific attention is an additional source of bias. Emphasis on valuation-change mechanisms amplifies negative reactions to positive news, generating systematic underreaction, while emphasis on dividend yield, growth, or downturn risks produces the opposite pattern. Temporary deviations in topic attention operate in the same direction as structural differences. Overall, both complexity and topic emphasis in causal narratives are key predictors of systematic forecast bias in asset managers' beliefs.

Literature review

A rapidly growing set of papers has pinned down several first-order facts about institutional CMA beliefs. First, subjective premia co-move closely with simple objective benchmarks and are countercyclical; disagreement across institutions is large and persistent, and different priors about long-run valuation levels plays an important role (Dahlquist & Ibert 2024, 2025). Second, across assets there is a strong positive subjective risk–return trade-off; both the dispersion of beliefs and the strength of the trade-off are greater across asset classes than across institutions, and aggregated subjective returns predict realized returns largely through the risk-premia component (Couts et al. 2023, 2024).

We contribute to this literature by addressing our central question—how managers form and justify long-horizon expectations on both returns and volatility, which constitute key inputs in asset pricing models. Four gaps matter. First, prior work reads CMAs "in numbers" and leaves the mechanism—how inflation, growth, productivity, policy, and valuations are linked in a causal story—implicit. We open that box by extracting manager-authored causal networks from the text of CMAs and relating stated channels to the quantitative building blocks. Second, the literature emphasizes first moments, but has not examined whether, how, and why institutional beliefs depart systematically from objective benchmarks. We document systematic deviations by showing that ex ante forecast errors and changes in return expectations are predictable. In addition, we study second-moment beliefs and document that long-horizon volatility expectations are strikingly sticky relative to large swings in realized volatility. Third, existing studies relate beliefs to state variables but not to directly measured news. We address this link empirically by considering text-based earnings sentiment news and show that long-horizon return expectations overreact to earnings news. Finally, while prior work documents a positive pass-through of institutional beliefs to their actions, we further document that the source of variation in return expectations driving asset managers' equity allocations originates from their change in valuation and growth in fundamentals expectations.

Beyond the CMA literature, our paper speaks to several broader strands of research.

First, it contributes to work documenting that beliefs of both retail and professional investors are heterogeneous and systematically biased (Malmendier & Nagel 2011, Greenwood & Shleifer 2014, Bouchaud et al. 2019, Andre et al. 2022, Nagel & Xu 2023, De la O & Myers 2021, 2024, Bordalo et al. 2024*a*,*b*, Bastianello 2024, 2025) and provides novel evidence on systematic deviations of asset managers' return expectations from objective benchmarks.[2] Second, it relates to the literature at the intersection of machine learning and finance, particularly work using textual information to study beliefs (Asquith et al. 2005, Baker et al. 2021, Bastianello et al. 2024, Bianchi et al. 2024, Bybee et al. 2024, Décaire & Graham 2024, Lopez-Lira 2024, van Binsbergen et al. 2024, Bastianello & Peng 2025, Bianchi et al. 2025, Laarits et al. 2025, Sarkar 2025), and our contribution is to study the textual content of CMA narratives rather than newspapers or analysts' reports. Third, it connects with the narrative economics literature, which highlights the role of stories in shaping expectations in financial markets: narratives propagate, frame causal reasoning, and influence asset prices (Shiller 2019, Goetzmann et al. 2022, Kendall & Charles 2022, Bybee 2023), and we document substantial heterogeneity in asset managers' narratives and mental models. Fourth, it contributes methodologically to the "text-as-data" literature, which has progressed from dictionary-based approaches to modern NLP and large language models capable of extracting forward-looking content and causal structure from financial texts (Gentzkow et al. 2019, Israel et al. 2020, Ke et al. 2021, Li et al. 2023, Chen et al. 2024, Sarkar & Vafa 2024), and we add to this literature by developing a novel approach to extract causal networks from CMA narratives. Finally, our work relates to the growing literature on attention and complexity (Veldkamp 2011, Schwartzstein 2014, Bordalo et al. 2019, Gabaix 2019, Bordalo et al. 2022, Gormsen & Huber 2025, Gabaix & Graeber 2024, Bastianello & Imas 2025, Gonçalves et al. 2025), and we contribute by constructing new text-based measures of attention and complexity using causal networks extracted from CMA narratives, showing that institutional asset managers exhibit systematic heterogeneity in both attention allocation and narrative complexity, which helps explain dispersion in return forecast levels and forecast bias.

At a high level, our contribution is to integrate these strands of literature in a single CMA-centered empirical design: we read beliefs "in numbers", via forecasts in levels and

[2] See Barberis (2018) and Adam & Nagel (2023) for comprehensive reviews.

component decompositions from tables and figures, and “in words”, via an LLM pipeline that recovers the asset managers’ causal networks underlying their mental models. This joint treatment moves us beyond documenting first- and second-moment levels and their co-movements with state variables, to identifying mechanisms within their mental models, and detecting systematic biases in their expectations.

Roadmap of the paper

The remainder of this paper proceeds as follows. In Section 1 we discuss data and methodologies. In Section 2 we study equity return and volatility expectations “in numbers”, starting from accounting identities and valuation models. In Section 3 we study expectations “in words”, leveraging unstructured text data from CMA narratives and extracting their causal content. Section 4 concludes.

1 Data and Methodology

This section outlines our dataset construction. We hand-collect publicly available Capital Market Assumptions (CMAs) from major financial institutions covering the period 2010–2025. These documents contain both qualitative narratives and quantitative return forecasts. To enrich this primary dataset, we incorporate news sentiment data from RavenPack and mutual fund allocation data from Morningstar Direct.

1.1 Capital Market Assumptions

We document the Capital Market Assumptions (CMA) produced by large institutional investors, beginning with an overview of their role and followed by a detailed account of our dataset construction and coverage. A key innovation of our dataset is that it captures both the qualitative inputs—economic models, causal narratives, and analytical frameworks—and the corresponding quantitative outputs—explicit return expectations and second moments—within the same documents for each asset manager. This dual structure approximates a supervised learning setting, allowing us to study how different institutions map a common set of inputs into heterogeneous forecasts.

Capital Market Assumptions are forward-looking projections of return expectations, volatilities, and correlations across asset classes and serve as essential inputs for strategic asset allocation and long-term investment planning. Their institutional significance is well established (Dahlquist & Ibert 2024, 2025, Couts et al. 2023, 2024, Andonov et al. 2026). CMAs are not passive survey responses, but integral business tools produced independently by institutions—typically annually or quarterly—by dedicated teams of economists, quantitative analysts, and senior strategists who synthesize macroeconomic trends, asset-pricing models, and academic research. Forecasting horizons usually range from 3 to 30 years, with 10-year projections the most common.[3]

Many CMA reports reference prominent asset-pricing theories and empirical research, often resembling comprehensive white papers rather than simple forecast tables. This analytical depth suggests that CMAs capture more sophisticated beliefs than those typically elicited in household or professional investor surveys.

We hand-collect a dataset of long-term CMAs from both asset managers and institutional investment consultants.[4] For simplicity, we use the term asset managers to refer to both throughout the paper. Including both reflects their complementary roles in shaping institutional investment decisions. Pension funds and endowments frequently rely on consultants for strategic asset allocation guidance, while asset managers use their published CMAs to inform internal portfolio strategies and provide advisory support to external clients.[5]

[Table 1 here]

Our final sample comprises 78 distinct institutions: 54 asset managers and 24 consultants. Table 1 provides detailed information on the organizations included in the dataset along with their respective coverage periods. The CMA documents span the years 2008 to 2025, coverage varies across institutions due to differences in data availability and publication frequency.

[3]Some institutions focus on shorter horizons (less than 10 years), while others emphasize longer-term projections (20–30 years).

[4]CMAs are manually collected from publicly available sources, including institutional websites and archived versions retrieved via archive.org.

[5]See Andonov et al. (2026) for empirical evidence that consultant recommendations directly affect pension fund investment decisions.

For temporal analysis, each CMA report is assigned a reference date corresponding to the data cutoff used in the institution's forecasting process, as explicitly stated in the report.[6]

[Table 2 here]

Each institution-year CMA provides return forecasts across a range of asset classes. We focus on major asset classes, selected based on three criteria: (i) relevance to institutional portfolios (e.g., U.S. Equity, Developed ex-U.S. Equity, Emerging Market Equity), (ii) consistency of appearance within institutions over time, and (iii) broad coverage across institutions. Because asset class labels vary across managers (e.g., "U.S. Large Cap Equity" vs. "U.S. Equity"), we standardize them using a consistent naming convention.[7]

Our baseline empirical analysis uses expected geometric returns as reported in the CMA documents, without imposing a uniform forecast horizon (the median forecast horizon equal to 10 years across all asset classes). Table 2 summarizes coverage across 10 asset classes, reporting the number of institutions and years with available forecasts.

1.1.1 Return Decomposition

CMA reports typically include sections that describe the forecasting frameworks and component assumptions underlying asset managers' equity return expectations, as illustrated in Figure 1. This qualitative evidence serves as the foundation for our subsequent quantitative analysis, in which we integrate additional datasets to examine what drives these component-level beliefs and how they are formed. The widespread use of the building-block approach—first documented by Dahlquist & Ibert (2025)—motivates our investigation. We hand-collect and construct a new CMA dataset at the asset manager–year level that disaggregates forecasts into seven components: growth, dividends, valuation change, buybacks, share issuance, margin mean-reversion, and currency/inflation adjustments. This structure enables us to measure each component's contribution to belief dispersion and to relate asset manager beliefs to their narratives and subsequent portfolio allocations.

[6] This reference date captures the most recent information incorporated into the institution's subjective expectations and typically precedes the publication date by weeks or months. When a report does not specify a data cutoff date, we use the publication date as a proxy.

[7] If a CMA specifies a benchmark (e.g., S&P500, MSCI EM), we use that information to guide the mapping. When no benchmark is provided, we rely on textual descriptions of regional or style exposures. If a U.S. Cash forecast is not reported, we substitute short-term U.S. Treasury bills as a proxy.

Figure 1 summarizes the equity model formulations employed by asset managers across comparable time periods, revealing substantial heterogeneity that we explore in detail in later sections. Despite institutional differences, most managers rely on a common set of foundational elements - nominal growth, valuation adjustments, and dividend yields. This convergence reflects a shared methodological framework in which return expectations are expressed as linear combinations of fundamental drivers. While these three elements form the analytical core, many institutions incorporate additional components, such as buybacks, share issuance, profit margin normalization, and currency adjustments. This additive decomposition offers a transparent mapping from macro-financial assumptions to return expectations and naturally reflects an approximate return accounting identity. At the same time, the precise structure of the decomposition varies by asset class and across managers, highlighting key sources of disagreement in return expectations.

[Figure 1 here]

1.1.2 Text

We convert each CMA report into a structured digital format with an Optical Character Recognition and machine-learning-based workflow that preserves reading order, recovers document hierarchy, and extracts embedded visuals. The resulting structured data undergoes preprocessing. For topic modeling with Latent Dirichlet Allocation (LDA), we perform additional text preprocessing and construct a document-term matrix, as detailed in Appendix A.1.1. In contrast, the LLM pipeline relies solely on the initial preprocessing stage and requires no further text-specific cleaning (Appendix A.2.1). We also analyze word counts over time and across institutional investors, as described in Appendix A.3, revealing disparities in reporting practices. Finally, Appendix A.4 presents word clouds for key financial categories, highlighting the language asset managers use when discussing different topics.

1.1.3 Summary Statistics

We analyze return forecasts that asset managers publish across multiple asset classes and investment horizons. Figure 2 presents subjective return forecasts for U.S. equity, revealing two key patterns. Panel (a) and (b) show that expectations and volatilities forecasts dispersion remains consistently high throughout our sample period, suggesting that disagreement stems from fundamental differences in methodology rather than transient information asymmetries around specific events. Panel (c) examines a subset of representative investors selected for having the most comprehensive time series in our sample, documenting substantial heterogeneity in both belief persistence and baseline optimism across institutions. In Appendix A.5, Figure A9 displays the time-series variation in asset managers' return expectations across Developed Market Equity, Emerging Market Equity, U.S. Investment Grade, U.S. High Yield, U.S. Government Bond, and U.S. Cash. These patterns are consistent with our findings for U.S. Equity return expectations.

[Figure 2 here]

1.2 Other Data

We complement our primary Capital Market Assumptions dataset by merging several external data sources. First, we incorporate news-sentiment data from RavenPack to capture market information flow and investor sentiment, which we match with market capitalization data from CRSP. Second, we use actively managed U.S. mutual fund data from Morningstar Direct to study how stated market expectations relate to actual portfolio allocation decisions. Third, we obtain monthly return-predictor data from Amit Goyal's webpage to construct objective expected-return forecasts, and we supplement these with earnings, dividends, the CAPE ratio, and stock return data retrieved from Robert Shiller's webpage.

1.2.1 Ravenpack: Earnings Sentiment

Our news sentiment analysis leverages RavenPack, a platform that processes corporate and macroeconomic news from multiple sources in real time, generating quantitative sentiment scores and event classifications.[8] We focus specifically on earnings-related news events in our analysis. Earnings announcements represent the most frequent and systematically scheduled

[8]We utilize RavenPack's structured metadata including sentiment scores, event classifications, and relevance indicators. Details of our processing methodology are provided in Appendix A.6.

firm-specific information releases in RavenPack, providing regular updates on corporate performance. Using this earnings-focused approach, we assess news sentiment impact on the single U.S. Equity asset class rather than individual firms, aggregating firm-level news to the S&P 500 index level, including only companies identified through CRSP monthly market cap data. We define unweighted and weighted sentiment measures:

$$\bar{S}_\tau = \frac{1}{|\mathcal{I}_\tau|} \sum_{i \in \mathcal{I}_\tau} S_{i,t} \tag{1}$$

$$\bar{S}_\tau^{(\text{wtd})} = \frac{\sum_{i \in \mathcal{I}_\tau} w_{i,\tau} \, S_{i,t}}{\sum_{i \in \mathcal{I}_\tau} w_{i,\tau}} \tag{2}$$

where $\mathcal{I}_\tau$ represents news items within interval τ, $S_{i,t}$ is the sentiment score of item i, and $w_{i,\tau}$ represents market-cap weights.

1.2.2 Morningstar: Equity Allocations

We investigate portfolio-allocation decisions by analysing U.S.-domiciled open-end mutual funds drawn from Morningstar Direct and matched to the asset managers in our CMA database, using the "BrandingName" variable. The resulting sample comprises funds that are classified as Equity, Fixed Income, Allocation, or Alternative in Morningstar's Global Broad Category Group. Portfolio holdings are reported monthly or quarterly per SEC requirements.

We assume CMA narratives influence fund managers' decisions within defined temporal windows. Our baseline specification considers a three-month influence window following the CMA embedded date rather than publication date. We also examine alternative windows starting before the CMA embedded date, as managers may process information early. We track net fund allocations in U.S. equity markets.

1.3 Extracting Causal Networks with LLMs

Asset managers approach complex forecasting problems by decomposing them into components and forming beliefs about how these causally interact. Institutional asset managers make these reasoning chains explicit in Capital Market Assumption reports when discussing equity return expectations, articulating how economic variables transmit to returns. We extract these stated causal mechanisms to quantify heterogeneity in managers' mental models, measuring which components receive most attention and the complexity of causal chains.

Our approach shares the intuition of directed acyclic graph (DAG) frameworks that formalize causal reasoning, while adapting to the specific context of extracting stated relationships from text. The networks we extract preserve cycles and feedback loops as managers articulate them, capturing heterogeneity in how different institutions conceptualize the same forecasting problem.

Our extraction pipeline (Figure 3) transforms unstructured narratives into quantifiable networks through systematic text processing. We segment CMA reports to isolate discussions of equity returns and their drivers. We use the large language model Anthropic's Claude Sonnet 4.5 to identify explicit causal claims, extracting three elements: the causing concept, the affected concept, and the directional relationship. We harmonize these concepts across reports using a standardized financial taxonomy derived from RavenPack event categories, assigning each concept to consistent topics and subtopics. Each relationship receives a sign classification—positive when an increase in the cause produces an increase in the effect, negative for inverse relationships. The final dataset contains, for each causal link: source text, standardized cause and effect nodes with their topic classifications, and relationship sign. To isolate the mental model underlying equity forecasts, we retain only nodes connected through causal paths to U.S. equity return expectations, as managers may discuss economic relationships tangentially without incorporating them into their return expectations.[9]

[Figure 3 here]

Table 3 reports summary statistics for causal network complexity and topic attention measures, documenting systematic heterogeneity in how asset managers structure their economic reasoning. We construct these statistics by computing the time-series average of each measure for each asset manager over the sample period, then calculating cross-sectional moments across them. The complexity measures—indirect connection ratio, average eigenvector centrality, average shortest path length, and network transitivity—capture distinct aspects

[9]Additional implementation details and prompt methodologies are documented in Appendix A.2.

of network structure, while the attention measures quantify edge counts to specific economic topics including valuation change, economic growth, and dividend yield (see Appendix A.8.2 for detailed definitions).

The coefficients of variation exceed 0.7 for complexity measures and exceed 1 for attention measures, indicating that asset managers differ substantially in both how they reason and how they allocate attention to economic relationships. This heterogeneity is consistent across metrics, reflecting genuine differences in how managers construct and communicate their economic reasoning frameworks.

[Table 3 here]

The extracted networks enable systematic analysis of how causal reasoning shapes institutional beliefs. Differences in network architecture—which concepts are central, how densely interconnected the relationships, the typical path length from shock to outcome—provide a quantifiable representation of each manager's mental model. This transformation from qualitative narratives to measurable structures opens new avenues for understanding why sophisticated investors persistently disagree despite access to similar information. The approach provides field evidence on how institutions translate complex economic environments into structured forecasting frameworks and where their conceptual models fundamentally diverge.

2 Model Based Analysis

We organize the analysis based on accounting identities and valuation models in two parts: Section 2.1 maps equity return forecasts to a common building-block identity (income, real growth, inflation, valuation change) and studies the building blocks properties, and Section 2.2 examines how long-horizon volatility is formed, how persistent it is, and how informative it is for realized volatility.

2.1 Returns Expectation

The substantial heterogeneity in institutional return expectations raises fundamental questions about how sophisticated investors construct their beliefs. We decompose these expectations through an approximate return identity framework, revealing which economic fundamentals—growth, valuation changes, and shareholder distributions—drive cross-sectional disagreement among asset managers. Despite diverse proprietary models, return expectations map systematically to these common building blocks. We examine these components through a Campbell-Shiller lens to understand how beliefs about fundamentals and valuations vary across institutions and time, then link them to actual portfolio allocations to identify which components most strongly influence investment decisions.

2.1.1 An Approximate Identity Framework

Text analysis of CMA reports reveals frequent references to building-block approaches in the discussion of equity return assumptions. In contrast, mentions of multi-factor risk models are comparatively rare, as most institutions emphasize fundamental drivers such as yields and growth. Consistent with this narrative, the forecast tables published by these institutions typically decompose return expectations into components closely aligned with the Gordon Growth Model—such as income yield, growth, and valuation change. Dahlquist & Ibert (2025) note that CMAs return expectations are decomposed into building blocks, and by using external proxies for the blocks, they infer that change in valuation is the culprit for the observed heterogeneity in subjective equity premia. We take a complementary route: rather than inferring components residually, we extract the components managers explicitly publish in their CMA tables, enabling like-for-like comparisons across asset managers and years and find that while the change in valuation captures the largest variation across the components in explaining the variation in return expectations, it only accounts for 40% of the variation of return expectations, thus highlighting the role of the other components.

This methodology disaggregates total equity return expectations into key components: real growth, inflation, income from shareholder distributions, and valuation adjustments. We examine the structure of these beliefs and their consistency with the standard Campbell–Shiller present-value logic in Section 2.1.2. The novelty is that our component-level data are observed (from the CMA building-block tables) rather than imputed from auxiliary sources, which allows us to quantify component-specific variation and pass-through with minimal modeling assumptions.

We begin by noting that asset managers provide return expectations and decomposition in raw (non-log) units. To align with this convention, we derive a simple approximate return identity that reflects the underlying logic of these decompositions. Starting with a one-period identity

$$1 + R_{t+1} = \frac{D_{t+1} + P_{t+1}}{P_t} = \frac{D_{t+1}}{P_t} + \frac{P_{t+1}}{P_t} \tag{3}$$

where P_{t+1} and D_{t+1} are the price and the dividend paid at time $t+1$. We further decompose the capital and dividend yield terms into nominal growth in earnings, X_{t+1}/X_t, and growth in price-to-earnings ratios, $(P/X)_{t+1}/(P/X)_t$:

$$\begin{aligned} 1 + R_{t+1} &= \frac{D_{t+1}}{P_{t+1}} \frac{(P/X)_{t+1}}{(P/X)_t} \frac{X_{t+1}}{X_t} + \frac{(P/X)_{t+1}}{(P/X)_t} \frac{X_{t+1}}{X_t} && (4) \\ &= (1 + DY_{t+1})(1 + G_{px,t+1})(1 + G_{x,t+1}) && (5) \\ &= (1 + DY_{t+1})(1 + G_{px,t+1})(1 + G_{real\,x,t+1})(1 + \pi_{t+1}) && (6) \end{aligned}$$

where $DY_{t+1} = D_{t+1}/P_{t+1}$, $G_{px,t+1} = (P/X)_{t+1}/(P/X)_t$, $G_{x,t+1} = X_{t+1}/X_t$. The last expression decomposes the nominal growth rate in earnings into its real component, $G_{real\,x,t+1}$, and gross inflation factor, $(1 + \pi_{t+1})$. Some asset managers distinguish between earnings growth due to revenue expansion and that due to changes in profit margins. To capture this distinction, we further decompose real earnings growth as: $(1 + G_{real\,x,t+1}) = (1 + G_{real\,adj.\,x,t+1})(1 + G_{margin,t+1})$, where $G_{real\,adj.\,x,t+1}$ is the earnings growth adjusted for profit margin changes and $G_{margin,t+1}$ captures the contribution of profit margins. Incorporating this refinement leads us to:

$$1 + R_{t+1} = (1 + DY_{t+1})(1 + G_{px,t+1})(1 + G_{real\,adj.\,x,t+1})(1 + G_{margin,t+1})(1 + \pi_{t+1}) \tag{7}$$

We then linearize this expression by dropping higher-order interaction terms, which are typically small in magnitude. This yields an additive approximation:

$$R_{t+1} \approx DY_{t+1} + G_{px,t+1} + G_{real\,adj.\,x,t+1} + G_{margin,t+1} + \pi_{t+1} \tag{8}$$

This identity highlights how the raw return can be expressed as a linear function of raw dividend yields, raw growth rates, and inflation. Some asset managers note that D represents the full cash flow to investors - including dividends and net share repurchases - rather than dividends alone. Following Boudoukh et al. (2007), we therefore refine the payout term as: $DY_{t+1} = DY_{div,t+1} + DY_{\text{buyback},t+1} - DY_{\text{issuance},t+1}$, where $DY_{div,t+1}$ is the dividend yield, $DY_{\text{buyback},t+1}$ the yield from net share repurchases, and $DY_{\text{issuance},t+1}$ the dilution yield from net issuance. Taking subjective expectations and combining all terms, we express the return expectation in a form consistent with the decomposition observed in CMA reports:

$$\begin{aligned} \widetilde{\mathbb{E}}_t^{CMA}[R_{t+1}] &= \widetilde{\mathbb{E}}_t\left[DY_{div,t+1} + DY_{buyback,t+1} - DY_{issuance,t+1}\right] \\ &\quad + \widetilde{\mathbb{E}}_t\left[G_{px,t+1}\right] + \widetilde{\mathbb{E}}_t\left[G_{real\,adj.\,x,t+1}\right] + \widetilde{\mathbb{E}}_t\left[G_{margin,t+1}\right] + \widetilde{\mathbb{E}}_t\left[\pi_{t+1}\right] \end{aligned} \tag{9}$$

where $\widetilde{\mathbb{E}}_t\left[G_{real\,x,t+1}\right] = \widetilde{\mathbb{E}}_t\left[G_{real\,adj.\,x,t+1}\right] + \widetilde{\mathbb{E}}_t\left[G_{margin,t+1}\right]$.

This standardized identity allows us to compare levels and loadings on the same components across asset managers and years. In practice, these components are labeled differently across CMA reports. The term $\widetilde{\mathbb{E}}_t[DY_{div,t+1}]$ is commonly referred to as "income", "net payout yield", "carry", or "dividend yield" (in simpler decompositions). The valuation term $\widetilde{\mathbb{E}}_t[G_{px,t+1}]$ is often labeled as "valuation adjustment" or "change in valuation." For most asset managers, $\widetilde{\mathbb{E}}_t[G_{real\,x,t+1}]$ is captured under "EPS growth" or "nominal growth", though some distinguish between the revenue-driven component $\widetilde{\mathbb{E}}_t[G_{real\,adj.\,x,t+1}]$ and the margin component $\widetilde{\mathbb{E}}_t[G_{margin,t+1}]$. Finally, $\widetilde{\mathbb{E}}_t[\pi_{t+1}]$ represents the "inflation" component.

Harmonizing these labels to a common taxonomy is a data contribution that enables the cross-asset manager variance decompositions and pass-through tests below. This model-based framework for returns is consistent with a broad literature on equity premia, valuation, and returns and cash flows expectations. For example, Fama & French (2002) decompose the historical equity premium into contributions from dividend yields and dividend growth. Campbell & Ammer (1993) show that time variation in stock and bond returns can be attributed to changing expectations of future dividends and discount rates. Cochrane (2011) argues that fluctuations in asset prices are driven primarily by time-varying discount rates rather than by changes in expected cash flows. More recently, a growing behavioral literature has proposed subjective expectations—about cash flows as well as discount rates—as a key source of variation in asset prices (De la O & Myers 2021, Nagel & Xu 2022, Bordalo et al. 2024*a*,*b*, Bretscher et al. 2024, Bastianello 2024).

Nearly all asset managers incorporate nominal growth and dividend yield into their frameworks—Figure A10 in Appendix A.7 summarizes which building blocks each asset manager includes in their CMA decomposition. However, there is substantial variation in the inclusion of other components—particularly margin adjustments and share issuance metrics. Notably, dividend yield, buybacks, and share issuance are conceptually linked, as they jointly determine net shareholder payouts. We quantify this institutional heterogeneity in model structure and show that three blocks—growth, income, and valuation—dominate in practice. This documents the common accounting logic across proprietary frameworks and provides the basis for component-level variance decompositions and pass-through analyses of beliefs to allocations.

Panel (a) in Figure 4 quantifies the diversity of model structures, showing that approximately 60% of CMA reports include three building blocks, while fewer reports incorporate only two or more than four components. Panel (b) in Figure 4 provides additional insight by contrasting the inclusion and exclusion rates for each building block across the full sample. Blue bars indicate the percentage of asset managers that include a given factor, while orange bars represent exclusion rates. Nominal growth and dividend yield emerge as near-universal components, valuation change is included by approximately 85% of managers, and fewer than half incorporate margin adjustments or share issuance.

[Figure 4 here]

Figure 5 presents partial R^2 estimates for equity return expectations under three sets of fixed effects: asset manager, year, and investment horizon. The bar chart indicates that asset manager-specific factors explain the largest share of variation in return forecasts, substantially exceeding the contributions of temporal variation and forecast horizon. This suggests that, while market conditions and forecast periods do play a role, the identity of the forecasting institution is the primary driver of return expectation estimates. This attributes most of the cross-institution dispersion to valuation-change and growth assumptions in the published CMA blocks.

[Figure 5 here]

Figure 6 presents a stacked area chart that decomposes the average subjective equity return expectations into its constituent building blocks over the period 2010-2024. The most prominent and stable contributors are inflation (blue), real growth (dark blue), and dividend yield (brown). In contrast, the valuation change component (red) exhibits notable variation. The average subjective equity return, shown as a dashed-dotted red line, exhibits significant time variation, with a decline around 2016–2018 followed by a recovery in subsequent years.

[Figure 6 here]

Figure 7 reports the yearly standard deviations of individual building blocks from 2015–2024, capturing the cross-sectional dispersion in beliefs across asset managers. The dashed-dotted dark red line represents the overall standard deviation of subjective equity return forecasts during this period. The thickness of each area segment represents its contribution to overall volatility. Valuation changes (red) consistently account for the largest share of volatility, reflecting substantial cross-sectional disagreement among asset managers. Real growth expectations (dark blue) also contribute significantly, followed by dividend yield (brown), while inflation expectations exhibit relatively low volatility.

The overall standard deviation of the equity return equity return expectations (dashed-dotted red line) remains below the cumulative volatility of individual components, illustrating diversification effects - disagreements across components partially offset one another. Notable spikes in volatility occur around 2016 and 2021, potentially reflecting periods of elevated macroeconomic uncertainty or structural market shifts, during which institutional forecasters exhibited more divergent views on future conditions.

[Figure 7 here]

Figure 8 reports partial R^2 values from a regression with time fixed effects that decomposes institutional return expectations into seven explanatory components: real growth, inflation, dividend yield, valuation change, buybacks, share issuance, and margins. By absorbing common time shocks, this specification isolates cross-sectional heterogeneity across asset managers. Valuation change and real growth emerge as the dominant factors, jointly accounting for the largest share of the variation in return forecasts across institutions (40% and 22%, respectively).

Our variance decomposition shows directly from the reported blocks that although valuation change (repricing) explains the largest component of the cross-institution variation in subjective equity return expectations, growth and the other blocks together account for most of the remaining variation (approximately 60%).

[Figure 8 here]

2.1.2 Campbell-Shiller Decomposition

In line with works studying the relationship between return expectations, cash-flow expectations, and valuation ratios (De la O & Myers 2021, Dahlquist & Ibert 2024, Bordalo et al. 2024*b*), we consider a subjective Campbell–Shiller decomposition to study the building blocks of the return expectations of large institutional asset managers. We estimate asset manager-level Campbell–Shiller regressions using each institution's own component series. We find that countercyclicality in return expectations arises from (i) negative loadings of expected valuation change on valuation ratios and (ii) positive, procyclical loadings of growth—thereby locating the source of cyclicality at the component level within each asset manager.

Following Campbell & Shiller (1987, 1988), log-linearizing around the one-period return identity, the log price–dividend ratio pd_t can be written as:

$$pd_t = \kappa + \Delta d_{t+1} - r_{t+1} + \rho\, pd_{t+1} \tag{10}$$

where Δd_{t+1} is the future dividend growth, r_{t+1} is the future return, and pd_{t+1} is the future price-dividend ratio. Taking subjective expectations, then taking the covariance of the left and right-hand sides with respect to pd_t, and dividing by the variance of pd_t we obtain:

$$1 = \underbrace{\frac{Cov(\widetilde{\mathbb{E}}_t[\Delta d_{t+1}], pd_t)}{Var(pd_t)}}_{\beta_{(\widetilde{\Delta d},pd)}} - \underbrace{\frac{Cov(\widetilde{\mathbb{E}}_t[r_{t+1}], pd_t)}{Var(pd_t)}}_{\beta_{(\widetilde{r},pd)}} + \rho \underbrace{\frac{Cov(\widetilde{\mathbb{E}}_t[pd_{t+1}], pd_t)}{Var(pd_t)}}_{\beta_{(\widetilde{pd},pd)}} \tag{11}$$

where $\beta_{(\widetilde{\Delta d},pd)}$ is the regression coefficient of subjective growth in dividends on the price-dividend ratio, $\beta_{(\widetilde{r},pd)}$ is the regression coefficient of subjective return expectations on the price-dividend ratio, and $\beta_{(\widetilde{pd},pd)}$ is the regression coefficient of the subjective expected future price-dividend ratio on the price-dividend ratio. To express (11) in terms of valuation changes, we add and subtract pd_t from the first argument of the last covariance term in (11) and reach:

$$1 - \rho = \beta_{(\widetilde{\Delta d},pd)} - \beta_{(\widetilde{r},pd)} + \rho \underbrace{\frac{Cov(\widetilde{\mathbb{E}}_t[\Delta pd_{t+1}], pd_t)}{Var(pd_t)}}_{\beta_{(\widetilde{\Delta pd},pd)}} \tag{12}$$

where Δpd_{t+1} is the one-period change in the log price-dividend ratio and $\beta_{(\widetilde{\Delta pd},pd)}$ is the regression coefficient of the subjective expected change in future price-dividend ratio on the price-dividend ratio. Equation (12) highlights how the regressions of $\widetilde{\mathbb{E}}_t[\Delta d_{t+1}]$ on pd_t, $\widetilde{\mathbb{E}}_t[r_{t+1}]$ on pd_t, and $\widetilde{\mathbb{E}}_t[\Delta pd_{t+1}]$ on pd_t are connected through the accounting identity.

In order to introduce earnings in the identity, in the spirit of De la O & Myers (2021), we substitute the identity $px_t = pd_t - dx_t$ (where x_t denotes smoothed log earnings and $dx_t = d_t - x_t$ is the smoothed log payout ratio) into (10):

$$px_t = \kappa + \Delta x_{t+1} - r_{t+1} + (1-\rho)dx_{t+1} + \rho\, px_{t+1} \tag{13}$$

Since $1-\rho$ is close to zero, movements in smoothed payout ratio do not play a large role in explaining variation in px_t and we approximate (13) by:

$$px_t \approx \kappa + \Delta x_{t+1} - r_{t+1} + \rho\, px_{t+1} \tag{14}$$

Following the same procedure, taking covariances with px_t and dividing by $Var(px_t)$, we obtain:

$$1-\rho = \underbrace{\frac{Cov(\widetilde{\mathbb{E}}_t[\Delta x_{t+1}], px_t)}{Var(px_t)}}_{\beta_{(\widetilde{\Delta x}, px)}} - \underbrace{\frac{Cov(\widetilde{\mathbb{E}}_t[r_{t+1}], px_t)}{Var(px_t)}}_{\beta_{(\widetilde{r}, px)}} + \rho \underbrace{\frac{Cov(\widetilde{\mathbb{E}}_t[\Delta px_{t+1}], px_t)}{Var(px_t)}}_{\beta_{(\widetilde{\Delta px}, px)}} \tag{15}$$

where $\beta_{(\widetilde{\Delta x}, px)}$ is the regression coefficient of the subjective expected change in long-term earnings on the price-to-long-term-earnings ratio, $\beta_{(\widetilde{r}, px)}$ is the regression coefficient of the return expectation on the price-to-long-term-earnings ratio, and $\beta_{(\widetilde{\Delta px}, px)}$ is the regression coefficient of the subjective expected change in valuation ratio on the price-to-long-term-earnings ratio.

Equations (12) and (15) hold true at the individual or consensus level and motivate our first part of the empirical analysis. We focus our analysis on individual level regressions as this avoids aggregation issues related to different sampling composition and horizons. Specifically, we focus on CMA asset managers for which we collected at least 5 years of data. We construct the log price dividend ratio of the S&P500 from Robert Shiller's website as a proxy for pd_t and the log cyclically adjusted price-earnings ratio from Robert Shiller's website as a proxy for px_t. From the CMA data, we use the return expectations as a proxy for $\widetilde{\mathbb{E}}_t[r_{t+1}]$, the expected growth rate in fundamentals as a proxy for $\widetilde{\mathbb{E}}_t[\Delta d_{t+1}]$ or $\widetilde{\mathbb{E}}_t[\Delta x_{t+1}]$, and the expected change in valuation ratios as a proxy for $\widetilde{\mathbb{E}}_t[\Delta pd_{t+1}]$ or $\widetilde{\mathbb{E}}_t[\Delta px_{t+1}]$. In Figure 9, for each asset manager, we estimate regressions of the form:

$$\widetilde{Y}_{t,t+h} = \alpha + \beta_{pd}\, pd_t + \epsilon_{t,h} \tag{16}$$

$$\widetilde{Y}_{t,t+h} = \alpha + \beta_{px}\, px_t + \epsilon_{t,h} \tag{17}$$

where $\widetilde{Y}_{t,t+h}$ is the subjective nominal return expectation or the subjective expectation of nominal growth in fundamentals or the subjective expectation of growth in valuation ratios. Figure 9 displays the results from our estimations. Consistent with the findings of previous works, the majority of asset managers exhibit a negative relationship between return expectations and CAPE or price-dividend ratio. Notably, the relationship inverts sign when looking at expectations of growth in fundamentals: the majority of asset managers exhibit positive coefficients, implying that on average the growth expectations are pro-cyclical in nature. Finally, when looking at the expected valuation changes, the majority of asset managers exhibit a negative correlation with price-to-fundamentals ratios, which represents the main driver behind the counter-cyclicality of asset managers' return expectations.

[Figure 9 here]

2.1.3 Evaluating CMA Forecasts: An Objective Forecasting Benchmark

To evaluate whether CMA forecasts contain valuable information beyond publicly observable data, we construct an objective forecasting model following Welch & Goyal (2008) and Goyal et al. (2024). The approximate return identity in Section 2.1.1 decomposes expected returns into income, real growth, inflation, and valuation change, the same building blocks that asset managers report in their CMAs. We therefore select predictors that map directly to these components: valuation ratios, per-capita consumption growth, dividend yield, equity issuance, and inflation. A linear specification provides a natural benchmark, as the building-block frameworks in CMAs are themselves linear decompositions.

We estimate h-year-ahead return forecasts using a predictive regression:

$$\hat{R}_{t,t+h} = \hat{\beta}_0 + \sum_{j=1}^{K} \hat{\beta}_j X_{j,t-1} \tag{18}$$

where $\hat{R}_{t,t+h}$ denotes the annualized h-year return forecast formed at time t, and $X_{j,t-1}$ represents predictor variables observed at the end of year $t-1$. The coefficients $\hat{\beta}_j$ are estimated by OLS from historical data, mapping lagged predictors to future returns through their estimated predictive relationships.[10] We use annual data from the Goyal et al. (2024) dataset beginning in 1927.

We evaluate specifications using different combinations of building-block variables, selecting for each horizon $h \in \{1, \ldots, 10\}$ years the model minimizing RMSE over a model-selection

[10]Following Welch & Goyal (2008), we impose a non-negativity constraint $\hat{R}^{CT}_{t,t+h} = \max(0, \hat{R}_{t,t+h})$ to prevent predictions of negative long-run equity returns.

sample (returns ending 1970–2010). Out-of-sample accuracy is measured by:

$$R^2_{OOS} = 1 - \frac{\sum_t (R_{t,t+h} - \hat{R}_{t,t+h})^2}{\sum_t (R_{t,t+h} - \bar{R}_{t-1})^2} \tag{19}$$

where $\bar{R}_{t-1}$ is the expanding historical mean. Best-performing models vary by horizon: consumption growth dominates at short horizons, while valuation measures at long horizons.

We compare the objective model forecasts with CMA forecasts made after 2010, so that model selection precedes evaluation. To formally test for differences in predictive accuracy, we employ the Diebold & Mariano (1995) test:

$$DM = \frac{\bar{d}}{\sqrt{\hat{V}(\bar{d})/n}} \tag{20}$$

where $d_t = e^2_{CMA,t} - e^2_{OBJ,t}$ is the squared-error loss differential. Table 4 reports the results, which favor the objective model at all horizons: RMSE ranges from 3.7% to 10.4% for the objective model versus 6.2% to 12.4% for CMA forecasts, with all DM statistics significant at the 1% level. Simple linear models using publicly available data outperform sophisticated institutional forecasters.

[Table 4 here]

2.1.4 Forecast Dynamics, Errors, and News

Building on our analysis of institutional return expectations and their building-block decomposition, we now examine how forecast changes and ex-ante forecast errors relate to the flow of earnings news. To implement this analysis, we construct an aggregate S&P 500 earnings-news sentiment index following the methodology described in Section 1.2.

We study the relationship between the annual change in average yearly forecasts and earnings-news sentiment with the following regression specification:

$$Y_{i,t} - Y_{i,t-1} = \mu_i + \beta\, S_t + \varepsilon_{i,t}, \tag{21}$$

where $Y_{i,t}$ represents return or building-block forecasts (changes in valuation, growth, and dividend yield) for asset manager i in year t, and S_t denotes the yearly average value-weighted S&P 500 RavenPack sentiment index for earnings-related news. Table 5 reports whether changes in institutional long-horizon forecasts respond systematically to earnings-news sentiment. The results show a clear effect on changes in valuation forecasts and on return forecasts: more positive earnings news leads to upward changes in expected returns and valuation changes. Results are robust to controlling for past market returns, as shown in Table A1.

[Table 5 here]

To determine whether the reaction to positive earnings news is accurate, we next study the relationship between ex-ante CMA forecast errors—defined as the difference between the objective benchmark, G_t (described in Section 2.1.3), and the asset managers' return forecasts, $F_{i,t}$—and earnings-news sentiment:

$$G_t - F_{i,t} = \mu_i + \beta\, S_t + \varepsilon_{i,t}. \tag{22}$$

The results in Table 6 show that earnings-news sentiment strongly and negatively predicts ex-ante forecast errors, implying that asset managers' biases in return expectations are predictable. We interpret this as evidence that asset managers react positively to earnings news, but do so excessively: in other words, they overreact to earnings news. Results are robust to controlling for past market returns, as shown in Table A2. Taken together, these results provide evidence of systematic deviations of asset managers' forecasts from objective return forecast benchmarks, which we examine further in Section 3.2.

[Table 6 here]

2.1.5 Fund Allocation

Having established how institutional return expectations decompose into fundamental building blocks, we now examine whether these stated beliefs translate into actual investment decisions. To investigate this relationship, we link Capital Market Assumption (CMA) forecasts to fund allocation data from Morningstar Direct (described in Section 1.2), testing whether managers' published expectations align with their portfolio positioning.

We estimate the following panel specification:

$$w_{i,t} = \mu_i + \tau_t + \beta X_{i,t} + \varepsilon_{i,t}, \tag{23}$$

where $w_{i,t}$ denotes fund i's U.S. equity net allocations measured between the previous and next quarter of the CMA reporting date, μ_i are fund fixed effects, and τ_t are year fixed effects.

Table 7 presents three specifications. The first examines CMA-implied subjective excess returns across equity asset classes, the second focuses on subjective U.S. equity returns, and the third decomposes these returns into their constituent building blocks: nominal growth, valuation changes, and dividends. Following Dahlquist & Ibert (2024), column 1 shows that institutional investors exhibit systematic relationships between their stated return expectations and portfolio allocations. Our results replicate their key finding: higher subjective U.S. equity return expectations are associated with increased domestic equity weights, while more optimistic projections for foreign markets correspond to reduced U.S. allocations, as shown in Column 3. This confirms that CMA expectations are economically relevant for portfolios.

Using our building-block decomposition data, Column 3 examines whether valuation expectations, growth assumptions, or income components differentially influence portfolio weights. The results show that while the dividend yield component has limited explanatory power, both nominal growth and valuation change exhibit strong and positive coefficients - 0.46 and 0.23, respectively. This component-level pass-through is new: prior CMA papers document overall pass-through or identity-based mechanisms but do not identify which published components drive allocations. We show that allocations load on growth and valuation—not on dividends—thereby highlighting the specific channels through which beliefs translate into weights. This suggests that both the pro-cyclical growth component and the counter-cyclical valuation component play significant roles in asset allocation decisions.

These findings underscore the importance of understanding how the individual components of CMA return expectations translate into portfolio choices. If one interprets the valuation change as capturing perceived mispricing in equity markets, then the strong positive coefficient on this term aligns with Bastianello & Peng (2024), who document that global fund managers' asset allocation intentions are positively correlated with their perceived mispricings.

[Table 7 here]

2.2 Volatility Forecasts

We next examine volatility forecasts. We decompose the variation in subjective volatility forecasts to identify whether heterogeneity stems from persistent manager characteristics or time-varying market conditions. We then test whether disagreement in volatility perceptions drives the cross-sectional dispersion in return expectations. Finally, we investigate how managers form volatility forecasts—specifically, whether they anchor to long-term historical averages or incorporate recent market information—and assess whether their subjective forecasts predict future realized volatility relative to statistical benchmarks.

We decompose the variation in volatility forecasts $F_{i,t,h}$ to identify whether heterogeneity stems from persistent manager characteristics or time-varying market conditions. We estimate a saturated fixed-effects model:

$$F_{i,t,h} = \alpha + \mu_i + \tau_t + \eta_h + \varepsilon_{i,t,h}, \tag{24}$$

where μ_i captures manager-specific effects, τ_t represents common time variation, and η_h accounts for forecast-horizon effects.

To quantify each component's contribution, we apply a Shapley-value decomposition of the R^2, which provides order-invariant variance contributions by averaging marginal effects across all possible variable orderings. Figure 10 reveals that the model explains 88.1% of total variation in volatility forecasts. Strikingly, manager fixed effects account for 72.1% of explained variance, while time effects contribute 25.3% and horizon effects only 2.6%.

[Figure 10 here]

This decomposition shows that volatility forecasts are predominantly shaped by persistent manager-specific characteristics rather than shared responses to market conditions. The dominance of manager effects suggests that institutional investors maintain stable, heterogeneous frameworks for assessing market risk that persist through time.

To understand how asset managers form volatility forecasts, we examine their relationship with historical realized volatility. We employ exponentially-weighted moving average (EWMA) models with endogenously estimated decay parameters to quantify the extent to which volatility expectations are backward-looking. Our analysis reveals that managers place disproportionate weight on long-term historical averages while exhibiting limited responsiveness to recent volatility innovations.

Figure 11 reveals that managers' volatility forecasts cluster tightly within a narrow band of 15–17% annually, displaying minimal variation despite substantial fluctuations in realized volatility over the sample period. The subjective forecasts (blue box plots) remain remarkably stable across time, even as three EWMA realized volatility measures with different decay parameters (0.75, 0.98, 1) exhibit considerable variability. This compression of forecasts suggests that managers anchor their volatility forecasts to long-term historical norms rather than incorporating current market conditions.

[Figure 11 here]

To investigate how asset managers weight past volatility when forming forecasts, we estimate an optimal exponentially weighted moving average (EWMA) model following the approach in Greenwood & Shleifer (2014). We fit the nonlinear regression:

$$F_t = \alpha + \beta \cdot \mathrm{EWMA}_t(\gamma) + \varepsilon_t, \tag{25}$$

where F_t is the subjective volatility forecast and $\mathrm{EWMA}_t(\gamma) = \sum_{j=1}^{L} \omega_j \cdot \mathrm{RV}_{t-j}$ represents an exponentially weighted average of past realized volatilities with weights $\omega_j = \gamma^{j-1} / \sum_{k=1}^{L} \gamma^{k-1}$. The decay parameter $\gamma \in [0, 1]$ governs the persistence of memory in the averaging process: values near unity indicate that distant observations retain substantial influence, while lower values imply rapid decay in the relevance of past information. We estimate this model using quarterly realized volatility computed from daily S&P500 returns over a 10-year lookback window.

Table 8 presents the optimal EWMA estimates using quarterly data. The estimated decay parameter is $\hat{\gamma} = 0.975$ (s.e. $= 0.019$), corresponding to a half-life of 27.7 quarters—approximately 7 years. This high persistence parameter indicates that managers place substantial weight on long-term historical volatility when forming forecasts, consistent with our earlier finding that subjective volatility is compressed relative to realized measures.

[Table 8 here]

3 Model Free Analysis

Building on the framework-based decomposition of return expectations, we explore the underlying mechanisms through which institutional investors process information and form beliefs. We employ text analysis to uncover thematic narratives that shape CMA reports, using topic modeling techniques to identify recurring patterns in how managers discuss economic and financial themes. Through large language models, we then extract and analyze the causal networks embedded in these reports, revealing how managers explicitly connect economic phenomena through causal reasoning. We then examine how institutional return expectations relate to past earnings news sentiment, capturing how managers process information flows from earnings signals.

3.1 Topic Text Analysis

To understand the narrative structures underlying institutional return expectations, we analyze the thematic content of CMA reports using topic modeling techniques. Following the methodology detailed in Appendix A.1.1, we segment CMA reports to isolate discussions of equity returns and related macroeconomic drivers, then apply Latent Dirichlet Allocation (LDA) to extract latent topics from these equity-focused narratives. This approach builds on the narrative asset pricing framework of Bybee et al. (2023), who show that textual narratives contain systematic risk factors beyond traditional quantitative measures. The choice of 60 topics, determined through Bayes factor optimization, balances model complexity with interpretability while capturing the primary narrative dimensions in our corpus.

Figure 12 visualizes the resulting narrative landscape across institutional investors. Each CMA report is represented as a 60-dimensional vector of topic proportions, which we project into two dimensions using t-SNE to preserve local similarities in the high-dimensional topic space. The application of K-Means clustering to these projections reveals 14 distinct narrative groups, suggesting that asset managers organize their economic reasoning around coherent thematic frameworks. The clear separation between clusters indicates systematic differences in how institutions frame market dynamics and structure their investment narratives. This heterogeneity in narrative focus provides a foundation for understanding the cross-sectional variation in return expectations documented in Section 2.1.

[Figure 12 here]

3.2 Causal Network

Asset managers explicitly articulate causal mechanisms in their CMA reports, documenting how economic variables transmit effects to equity returns. Using the LLM methodology from Section 1.3, we extract these stated causal relationships to construct directed networks that reveal each institution's economic reasoning. Unlike correlation analysis, these networks capture managers' beliefs about transmission channels and causal ordering—distinguishing between primary drivers, intermediate mechanisms, and final outcomes in their mental models of market dynamics.

Figure 13 illustrates this approach through DWS's 2015–2025 causal network, where node size reflects a topic's interconnectedness and edge colors encode directional effects. The network reveals that growth and valuation change emerge as the most central topics, with multiple causal pathways connecting policy and regulations, interest rates, and dividend yield. Figure 14 isolates first-order drivers of U.S. equity return expectations—concepts directly linked through a single causal step—highlighting how DWS emphasizes valuation adjustments, growth, dividend yield, and inflation. This structural mapping enables systematic comparison across institutions to identify where sophisticated investors fundamentally disagree about market mechanics. In both networks, growth, inflation, dividends, and valuation adjustments—the four building blocks that our decomposition analysis (Section 2.1) identified as primary sources of forecast disagreement—appear with distinct coloring to facilitate the identification of their causal pathways. Appendix A.8.1 presents additional network visualizations for four asset managers and reports the number of topics and causal links per report for all institutions in our sample.

[Figures 13 and 14 here]

Figure 15 presents two complementary perspectives on the causal networks. Panel (a) displays how frequently economic topics appear as causal agents across all relationships, with growth, valuation changes, and policy/regulatory factors dominating managers' reasoning frameworks. Panel (b) isolates topics directly linked to U.S. equity return expectations, revealing a concentrated set of drivers-valuation changes and economic growth appear in 59% and 44% of asset manager-years, respectively.

[Figure 15 here]

Figure 16 reveals substantial heterogeneity in asset managers' causal frameworks. The stacked bars show average causal links per year by topic, with marked differences between asset managers. JP Morgan and Robeco construct dense networks exceeding 200 links per year, while others maintain smaller frameworks. Topics with most links also vary, some managers emphasize growth and valuation, others prioritize policy and inflation.

[Figure 16 here]

Attention, Bias, and Complexity

Having extracted causal networks from CMA narratives, we now examine whether the structure of these networks explains cross-sectional heterogeneity in institutional beliefs. Our analysis tests whether managers who assign different levels of attention to specific economic mechanisms—as revealed by their causal reasoning—produce systematically different return expectation forecast levels, changes, and ex-ante errors. We examine how institutional return expectations relate to the structure and content of managers' causal reasoning networks using panel regressions at the asset manager-year level.

To investigate whether managers with more complex causal reasoning structures produce different return forecasts and whether this relationship varies with earnings sentiment, we employ a within-between decomposition of causal network measures (Mundlak 1978, Chamberlain 1982). This approach separates temporary deviations from a manager's typical complexity level (within-manager variation) from their persistent complexity characteristics (between-manager variation). We estimate the following specification:

$$F_{i,t} = \mu_i + \tau_t + \beta_w(C_{i,t} - \overline{C}_i) + \gamma_w(C_{i,t} - \overline{C}_i) \times S_t + \gamma_b\overline{C}_i \times S_t + \varepsilon_{i,t} \tag{26}$$

where $F_{i,t}$ denotes the return forecast for manager i in year t, μ_i and τ_t represent asset manager and year fixed effects, $\overline{C}_i$ represents the between-manager component (manager i's average complexity), $(C_{i,t} - \overline{C}_i)$ captures the within-manager component (deviation from this average), S_t is the aggregate sentiment measure, constructed as a rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index over a 1-year window, centered around the time-series mean, and $\varepsilon_{i,t}$ is the error term. The coefficient β_w identifies the direct effect of temporary increases in causal complexity on forecasts at average sentiment levels, γ_w reveals whether managers adjust their news sensitivity when their networks become temporarily more complex, and γ_b captures whether structurally complex managers react differently to news. Note that the between-manager component $\overline{C}_i$ is absorbed by the asset manager fixed effects μ_i, hence it only appears in interaction with sentiment. This specification allows us to identify the effect of temporary complexity changes within managers while controlling for time-invariant manager characteristics and common time trends. We examine three complexity measures: average shortest path length, indirect connection ratio, network transitivity (see Appendix A.8.2 for definitions of all network measures).

Table 9 shows that causal network complexity is not directly associated with the level of return expectations: the coefficients on the within component $(C_{i,t} - \overline{C}_i)$ are small and statistically insignificant across all three complexity measures. Instead, complexity matters by shaping how managers react to news. The interaction between temporary complexity and sentiment, $(C_{i,t} - \overline{C}_i) \times S_t$, is weakly significant when complexity is measured by network transitivity. Conversely, the interaction between structural (between-manager) complexity $\overline{C}_i \times S_t$ is negative and strongly significant, suggesting that managers with persistently complex causal networks react more negatively to news. This heterogeneity supports the view that differences in narrative structures help explain cross-manager variation in long-horizon return expectations.

[Table 9 here]

We also examine whether managers' attention to specific topics influences return expectations. Topic attention is measured by the edge count—the (normalized) number of causal connections to each topic node—thus capturing its role as a critical transmission mechanism in the manager's economic reasoning. For this analysis, we estimate:

$$F_{i,t} = \mu_i + \tau_t + \beta_w(A_{i,t} - \overline{A}_i) + \gamma_w(A_{i,t} - \overline{A}_i) \times S_t + \gamma_b \overline{A}_i \times S_t + \varepsilon_{i,t} \tag{27}$$

where $\overline{A}_i$ is manager i's time-series mean topic-specific edge count (between-manager component), $(A_{i,t} - \overline{A}_i)$ is the deviation from this average (within-manager component). We examine five topics: valuation change, dividend yield, economic growth, inflation, and economic downturn. The coefficient β_w captures the direct relationship between topic attention and forecasts, γ_w reveals how within-manager variation in topic attention moderates sensitivity to aggregate sentiment, and γ_b captures the between-manager interaction with sentiment. Note that the between-manager component $\overline{A}_i$ is absorbed by the asset manager fixed effects μ_i, hence it only appears in interaction with sentiment.

Table 10 shows that topic-specific attention in managers' causal networks affects return forecasts primarily through its interaction with news sentiment rather than through direct level effects. The within component $(A_{i,t} - \overline{A}_i)$ is small and statistically insignificant across all five topics, indicating that temporary increases in attention to a given mechanism do not mechanically shift long-horizon expectations. In contrast, the interaction terms $(A_{i,t} - \overline{A}_i) \times S_t$ reveal economically meaningful heterogeneity for valuation change expectations: managers become substantially more sensitive to sentiment when they temporarily devote more attention to valuation adjustments (column 1) and also mildly to economic downturn narratives (column 5). Between-manager components $\overline{A}_i \times S_t$ display similarly heterogeneous signs, indicating that managers whose narratives structurally emphasize certain topics (e.g., downturns or growth) react persistently more to news than others. Overall, Table 10 shows that the content of a manager's causal narrative—not only its complexity—shapes how sentiment is incorporated into long-horizon return forecasts, with different economic mechanisms amplifying or dampening news sensitivity in systematic ways.

[Table 10 here]

We next examine whether the structural complexity of managers' causal networks influences how beliefs change over time. Complex mental models may exaggerate reaction to news. We test this by analyzing how network structure predicts year-over-year forecast changes.

We employ three complementary complexity measures: average shortest path length, indirect connection ratio, and network transitivity. We estimate the following specification:

$$F_{i,t} - F_{i,t-1} = \mu_i + \tau_t + \beta_w(C_{i,t} - \overline{C}_i) + \gamma_w(C_{i,t} - \overline{C}_i) \times S_t + \gamma_b \overline{C}_i \times S_t + \varepsilon_{i,t} \tag{28}$$

where $F_{i,t} - F_{i,t-1}$ is the year-over-year forecast change for manager i, $\overline{C}_i$ is manager i's time-series mean complexity, $(C_{i,t} - \overline{C}_i)$ is the temporary deviation from this mean. The coefficient β_w identifies how temporary complexity changes affect forecast change magnitude, while γ_w and γ_b capture differential news sensitivity.

Table 11 examines whether managers' causal-network complexity predicts the magnitude of their year-over-year forecast changes. Across all three complexity measures, the within component $(C_{i,t} - \overline{C}_i)$ is small and statistically insignificant, indicating that temporary increases in network complexity do not mechanically generate larger forecast adjustments. The between–sentiment interaction $\overline{C}_i \times S_t$ is negative, economically large, and statistically significant across all three measures. Managers with persistently more complex causal networks revise their long-horizon forecasts more negatively in response to positive earnings-news innovations. Overall, the results show that structural complexity, more than temporary changes in narrative complexity, governs how managers update beliefs in response to news.

[Table 11 here]

We separately examine whether attention to specific economic topics influences forecast changes:

$$F_{i,t} - F_{i,t-1} = \mu_i + \tau_t + \beta_w(A_{i,t} - \overline{A}_i) + \gamma_w(A_{i,t} - \overline{A}_i) \times S_t + \gamma_b \overline{A}_i \times S_t + \varepsilon_{i,t} \tag{29}$$

where $\overline{A}_i$ is manager i's time-series mean topic-specific edge count (between-manager component), $(A_{i,t} - \overline{A}_i)$ represents within-manager variation in topic attention. We examine five topics: valuation change, dividend yield, economic growth, inflation, and economic downturn. Table 12 studies whether managers' attention to specific economic mechanisms—as captured by topic-specific causal network edge counts—helps explain year-over-year changes in long-horizon return expectations. Across all topics, the within component $(A_{i,t} - \overline{A}_i)$ is small and statistically insignificant, indicating that temporary shifts in attention do not mechanically induce forecast changes. Forecast changes instead respond to news conditional on topic attention. The interaction terms $(A_{i,t} - \overline{A}_i) \times S_t$ is significant for valuation changes, dividend yield, and growth, whereas the between-manager interactions $\overline{A}_i \times S_t$ is positively and significant for dividend yield and economic downturn. Managers who are more attentive to valuation changes exhibit a negative revision in response to positive news, whereas those who place greater weight than usual on dividend yields and growth expectations show a positive revision to the same news. In addition, managers who are more attentive than others to dividend yields and economic downturns react more positively to earnings-sentiment news. The results suggest that forecast revisions are driven not only by how a manager's attention to different topics varies over time, but also by how their attention compares with that of other managers.

Overall, the table shows that the content of causal narratives is a key determinant of how managers update long-horizon expectations. Topic emphasis—both temporary and structural—modulates the responsiveness of forecast changes to fundamentals news, creating systematic and persistent heterogeneity in belief updating across institutions.

[Table 12 here]

To investigate whether network complexity predicts systematic deviations from objective forecasts, we employ a within-between decomposition and estimate the following specification:

$$G_t - F_{i,t} = \mu_i + \tau_t + \beta_w(C_{i,t} - \overline{C}_i) + \gamma_w(C_{i,t} - \overline{C}_i) \times S_t + \gamma_b \overline{C}_i \times S_t + \varepsilon_{i,t} \qquad (30)$$

where $G_t - F_{i,t}$ is the ex-ante forecast error for manager i in year t relative to the benchmark forecast G_t (described in Section 2.1.3), $\overline{C}_i$ is manager i's time-series mean complexity (between-manager component), and $(C_{i,t} - \overline{C}_i)$ is the temporary deviation from this mean (within-manager component). The coefficient β_w identifies whether temporary changes in complexity lead to more optimistic or pessimistic forecasts relative to the benchmark. The interaction coefficients γ_w and γ_b reveal whether complexity moderates how news sentiment affects forecast bias, with γ_b specifically capturing whether managers with structurally complex networks exhibit different behavioral responses to news. We examine three complexity measures: average shortest path length, indirect connection ratio, network transitivity.

Table 13 examines whether the structural complexity of managers' causal networks predicts ex-ante forecast errors relative to the objective benchmark. The within-manager component $(C_{i,t} - \overline{C}_i)$ is statistically insignificant across all three complexity measures, indicating that temporary changes in narrative complexity do not drive forecast errors. Instead, persistent differences in complexity across managers shape how they process information.

The between-manager interaction with sentiment $\overline{C}_i \times S_t$ is positive and highly significant across all measures. The ex-ante forecast errors of managers' forecasts with structurally more complex causal networks exhibit greater sensitivity to earnings news, reacting less strongly to information flows while generating systematically more pessimistic forecasts following positive news. During periods of favorable sentiment, these managers produce larger forecast errors relative to the objective benchmark, underreacting to positive information more than managers with simpler causal structures.

[Table 13 here]

We separately examine whether attention to specific economic topics predicts systematic deviations from the objective benchmark. For this analysis, we estimate:

$$G_t - F_{i,t} = \mu_i + \tau_t + \beta_w(A_{i,t} - \overline{A}_i) + \gamma_w(A_{i,t} - \overline{A}_i) \times S_t + \gamma_b \overline{A}_i \times S_t + \varepsilon_{i,t} \tag{31}$$

where $\overline{A}_i$ captures the manager i's average topic attention level (measured by edge count), and $(A_{i,t} - \overline{A}_i)$ is the temporary deviation from this mean (within-manager component). The coefficient β_w identifies whether temporal increases in topic-specific attention lead to systematic optimism or pessimism relative to the rational benchmark, while γ_w and γ_b reveal how topic attention moderates the influence of news sentiment on forecast bias through within and between variation, respectively. We examine five economic topics: valuation change adjustment, dividend yield, economic growth, inflation, and economic downturn.

Table 14 examines whether topic-specific attention in managers' causal networks predicts ex-ante forecast errors relative to the objective benchmark. The results show that the centrality of key topics systematically predicts forecast errors, as topic attention shapes how asset managers react to information flows. Asset managers with greater attention to valuation change in their causal networks react more negatively to positive fundamentals news, generating systematically more pessimistic forecasts relative to the objective benchmark. Managers who on average emphasize valuation change (the between-manager component $\overline{A}_i \times S_t$) exhibit this pattern persistently and significantly. When individual managers temporarily increase their attention to valuation change (the within-manager component $(A_{i,t} - \overline{A}_i) \times S_t$), the effect operates in the same direction. Higher valuation attention amplifies sensitivity to news in a way that generates underreaction relative to the objective benchmark.

Asset managers with greater attention to dividend yield and economic downturn produce the opposite pattern. The between-manager interaction with sentiment are negative and highly significant. The results show that topic-specific attention modulates news sensitivity: managers emphasizing valuation adjustments underreact to positive news, while those emphasizing dividend yield and downturn risks overreact to news, introducing systematic and predictable deviations from the objective benchmark.

[Table 14 here]

4 Conclusion

We study how large asset managers form and justify long-horizon beliefs by combining the quantitative and narrative content of their Capital Market Assumptions (CMAs). Using hand-collected data, we map return expectations to a common set of building blocks—growth, valuation change, income, and inflation—and show that although expected valuation changes capture the largest share of cross-sectional variation in return expectations across asset managers, the remaining components capture the majority of the cross-sectional variation. Return expectations are countercyclical mainly because valuation changes are negatively related to valuation ratios, while growth expectations are procyclical. These components also explain differences in equity allocations across managers, confirming that the parts of beliefs that differ most across institutions are those that drive portfolio positions.

We complement this numerical analysis with text-based evidence. Causal networks extracted from CMA narratives reveal substantial heterogeneity in both the complexity and topic attention of managers' reasoning. Differences in these narrative structures explain variation in return forecasts, belief changes, and systematic forecast errors relative to objective benchmarks. Volatility expectations show a contrasting pattern: they are highly conservative, dominated by manager-specific frameworks, and adjust only slowly to realized volatility. Overall, institutional expectations are economically meaningful but display systematic and predictable deviations from objective benchmarks, reflecting heterogeneity in both the quantitative building blocks and the causal reasoning underlying asset managers' beliefs.

References

Adam, K. & Nagel, S. (2023), Expectations data in asset pricing, *in* 'Handbook of Economic Expectations', Elsevier, pp. 477–506.

Andonov, A., Couts, S. J., S Gonçalves, A., Loudis, J. & Rossi, A. (2026), 'Subjective beliefs and the portfolio allocations of institutional investors', *Working Paper* .

Andre, P., Pizzinelli, C., Roth, C. & Wohlfart, J. (2022), 'Subjective models of the macroeconomy: Evidence from experts and representative samples', *The Review of Economic Studies* **89**(6), 2958–2991.

Asquith, P., Mikhail, M. B. & Au, A. S. (2005), 'Information content of equity analyst reports', *Journal of financial economics* **75**(2), 245–282.

Baker, S., Bloom, N., Davis, S. J. & Sammon, M. C. (2021), What triggers stock market jumps?, Technical report, National Bureau of Economic Research Cambridge, MA.

Barberis, N. (2018), Psychology-based models of asset prices and trading volume, *in* 'Handbook of behavioral economics: applications and foundations 1', Vol. 1, Elsevier, pp. 79–175.

Bastianello, F. (2024), 'Time-series and cross-section of risk premia expectations: Evidence from financial analysts', *Available at SSRN 4204968* .

Bastianello, F. (2025), 'Bias and predictability in analysts' beliefs', *Available at SSRN 5733302* .

Bastianello, F., H Décaire, P. & Guenzel, M. (2024), 'Mental models and financial forecasts', *Marius, Mental Models and Financial Forecasts (October 30, 2024)* .

Bastianello, F. & Imas, A. (2025), 'Biases in belief updating within and across domains', *Available at SSRN 5413920* .

Bastianello, F. & Peng, C. (2024), 'Global fund managers' beliefs, perceived mispricing, and asset allocation', *Perceived Mispricing, and Asset Allocation (December 15, 2024)* .

Bastianello, F. & Peng, C. (2025), 'The term structure of return expectations', *Available at SSRN 5390270* .

Begenau, J., Liang, P. & Siriwardane, E. (2025), 'The rise of alternatives', *Available at SSRN 4940886* .

Bianchi, F., Gómez-Cram, R. & Kung, H. (2024), 'Using social media to identify the effects of congressional viewpoints on asset prices', *The Review of Financial Studies* **37**(7), 2244–2272.

Bianchi, F., Lee, D. Q., Ludvigson, S. C. & Ma, S. (2025), The prestakes of stock market investing, Technical report, National Bureau of Economic Research.

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

Bordalo, P., Gennaioli, N., La Porta, R. & Shleifer, A. (2024*a*), 'Belief overreaction and stock market puzzles', *Journal of Political Economy* **132**(5), 1450–1484.

Bordalo, P., Gennaioli, N., La Porta, R. & Shleifer, A. (2024*b*), Finance without exotic risk, Technical report, National Bureau of Economic Research.

Bordalo, P., Gennaioli, N., Porta, R. L. & Shleifer, A. (2019), 'Diagnostic expectations and stock returns', *The Journal of Finance* **74**(6), 2839–2874.

Bordalo, P., Gennaioli, N. & Shleifer, A. (2022), 'Salience', *Annual Review of Economics* **14**(1), 521–544.

Bouchaud, J.-P., Krueger, P., Landier, A. & Thesmar, D. (2019), 'Sticky expectations and the profitability anomaly', *The Journal of Finance* **74**(2), 639–674.

Boudoukh, J., Michaely, R., Richardson, M. & Roberts, M. (2007), 'On the importance of measuring payout yield: Implications for empirical asset pricing', *Journal of Finance* **62**(2), 877–915.

Bretscher, L., Malkhozov, A., Tamoni, A. & Yang, H. (2024), 'Distorted beliefs and asset prices', *Available at SSRN 4947547* .

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A. et al. (2020), 'Language models are few-shot learners', *Advances in Neural Information Processing Systems (NeurIPS)* **33**, 1877–1901.

Bybee, J. L. (2023), 'The ghost in the machine: Generating beliefs with large language models', *arXiv preprint arXiv:2305.02823* .

Bybee, L., Kelly, B., Manela, A. & Xiu, D. (2024), 'Business news and business cycles', *The Journal of Finance* **79**(5), 3105–3123.

Bybee, L., Kelly, B. & Su, Y. (2023), 'Narrative asset pricing: Interpretable systematic risk factors from news text', *The Review of Financial Studies* **36**(12), 4759–4787.

Campbell, J. Y. & Ammer, J. (1993), 'What moves the stock and bond markets? a variance decomposition for long-term asset returns', *The Journal of Finance* **48**(1), 3–37.

Campbell, J. Y. & Shiller, R. J. (1987), 'Cointegration and tests of present value models', *Journal of political economy* **95**(5), 1062–1088.

Campbell, J. Y. & Shiller, R. J. (1988), 'The dividend-price ratio and expectations of future dividends and discount factors', *Review of Financial Studies* **1**(3), 195–228.

Chamberlain, G. (1982), 'Multivariate regression models for panel data', *Journal of Econometrics* **18**(1), 5–46.

Chen, Y., Kelly, B. & Xiu, D. (2024), 'Expected returns and large language models', *Working Paper* . University of Chicago and Yale University.

Cochrane, J. H. (2011), 'Presidential address: Discount rates', *The Journal of Finance* **66**(4), 1047–1108.

Couts, S. J., Gonçalves, A. S. & Loudis, J. (2023), 'The subjective risk and return expectations of institutional investors', *Fisher College of Business Working Paper* **14**.

Couts, S. J., S Gonçalves, A., Liu, Y. & Loudis, J. (2024), 'Institutional investors' subjective risk premia: Time variation and disagreement'.

Dahlquist, M. & Ibert, M. (2024), 'Equity return expectations and portfolios: Evidence from large asset managers', *The Review of Financial Studies* **37**(6), 1887–1928.

Dahlquist, M. & Ibert, M. (2025), 'Institutions' return expectations across assets and time'.

Dang, T. L., Moshirian, F. & Zhang, B. (2015), 'Commonality in news around the world', *Journal of Financial Economics* **116**(1), 82–110.

De la O, R. & Myers, S. (2021), 'Subjective cash flow and discount rate expectations', *The Journal of Finance* **76**(3), 1339–1387.

De la O, R. & Myers, S. (2024), 'Which subjective expectations explain asset prices?', *The Review of Financial Studies* **37**(6), 1929–1978.

Diebold, F. X. & Mariano, R. S. (1995), 'Comparing predictive accuracy', *Journal of Business & Economic Statistics* **13**(3), 253–263.

Décaire, P. H. & Graham, J. R. (2024), 'Valuation fundamentals', *Working Paper* . Available at SSRN.

Fama, E. F. & French, K. R. (2002), 'The equity premium', *The Journal of Finance* **57**(2), 637–659.

Fu, Y., Peng, H., Sabharwal, A., Clark, P. & Khot, T. (2023), 'Complexity-based prompting for multi-step reasoning', *arXiv preprint arXiv:2210.00720* .

Gabaix, X. (2019), Behavioral inattention, *in* 'Handbook of behavioral economics: Applications and foundations 1', Vol. 2, Elsevier, pp. 261–343.

Gabaix, X. & Graeber, T. (2024), The complexity of economic decisions, Technical report, National Bureau of Economic Research.

Gao, A. K. (2023), 'Prompt engineering for large language models', *SSRN Electronic Journal* .

Gao, P., Parsons, C. A. & Shen, J. (2017), 'Global relation between financial distress and equity returns', *The Review of Financial Studies* **31**(1), 239–277.

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

Giglio, S., Maggiori, M., Stroebel, J. & Utkus, S. (2021), 'Five facts about beliefs and portfolios', *American Economic Review* **111**(5), 1481–1522.

Goetzmann, W. N., Kim, D. & Shiller, R. J. (2022), Crash narratives, Technical report, National Bureau of Economic Research.

Gonçalves, D., Libgober, J. & Willis, J. (2025), 'Retractions: Updating from complex information', *Review of Economic Studies* p. rdaf032.

Gormsen, N. J. & Huber, K. (2025), 'Corporate discount rates', *American Economic Review* **115**(6), 2001–2049.

Goyal, A., Welch, I. & Zafirov, A. (2024), 'A comprehensive 2022 look at the empirical performance of equity premium prediction', *The Review of Financial Studies* .

Greenwood, R. & Shleifer, A. (2014), 'Expectations of returns and expected returns', *The Review of Financial Studies* **27**(3), 714–746.

Hafez, P. (2009), 'Construction of market sentiment indices using news sentiment', *RavenPackInternationalS.L.* .

Israel, R., Kelly, B. & Moskowitz, T. (2020), 'Can machines "learn" finance?', *SSRN Working Paper* .

Ke, Z. T., Kelly, B. & Xiu, D. (2021), 'Predicting returns with text data', *Working Paper* . University of Chicago, Yale University, and Harvard University.

Kendall, C. W. & Charles, C. (2022), Causal narratives, Technical report, National Bureau of Economic Research.

Laarits, T., Matthies, B., Vasudevan, K. & Yang, W. (2025), 'Policy by committee'.

Li, X., Chan, S., Zhu, X., Pei, Y., Ma, Z., Liu, X. & Shah, S. (2023), 'Are chatgpt and gpt-4 general-purpose solvers for financial text analytics? a study on several typical tasks', *arXiv preprint arXiv:2305.05862* .

Lopez-Lira, A. (2024), 'Can chatgpt forecast stock price movements?'.

Malmendier, U. & Nagel, S. (2011), 'Depression babies: Do macroeconomic experiences affect risk taking?', *The quarterly journal of economics* **126**(1), 373–416.

Mundlak, Y. (1978), 'On the pooling of time series and cross section data', *Econometrica* **46**(1), 69–85.

Nagel, S. & Xu, Z. (2022), 'Asset pricing with fading memory', *The Review of Financial Studies* **35**(5), 2190–2245.

Nagel, S. & Xu, Z. (2023), 'Dynamics of subjective risk premia', *Journal of Financial Economics* **150**(2), 103713.

RavenPack (2020), *RavenPack Analytics—User Guide and Service Overview*, RavenPack.

Sarkar, S. K. (2025), 'Economic representations'.

Sarkar, S. K. & Vafa, K. (2024), 'Lookahead bias in pretrained language models', *Available at SSRN* .

Schwartzstein, J. (2014), 'Selective attention and learning', *Journal of the European Economic Association* **12**(6), 1423–1452.

Shiller, R. J. (2019), *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*, Princeton University Press.

Steyvers, M. & Griffiths, T. (2007), Probabilistic topic models, *in* T. K. Landauer, D. S. McNamara, S. Dennis & W. Kintsch, eds, 'Handbook of Latent Semantic Analysis', Lawrence Erlbaum Associates, pp. 427–440.

van Binsbergen, J. H., Bryzgalova, S., Mukhopadhyay, M. & Sharma, V. (2024), (almost) 200 years of news-based economic sentiment, Technical report, National Bureau of Economic Research.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. & Polosukhin, I. (2017), 'Attention is all you need', *Advances in Neural Information Processing Systems* **30**.

Veldkamp, L. L. (2011), *Information choice in macroeconomics and finance*, Princeton University Press.

Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., Narang, S., Chowdhery, A., Zhou, D. et al. (2022), 'Self-consistency improves chain of thought reasoning in language models', *arXiv preprint arXiv:2203.11171* .

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E. H., Le, Q. V. & Zhou, D. (2022), 'Chain-of-thought prompting elicits reasoning in large language models', *Advances in Neural Information Processing Systems (NeurIPS)* **35**.

Welch, I. & Goyal, A. (2008), 'A comprehensive look at the empirical performance of equity premium prediction', *Review of Financial Studies* **21**(4), 1455–1508.

Zhao, T. Z., Wallace, E., Feng, S., Klein, D. & Singh, S. (2021), 'Calibrate before use: Improving few-shot performance of language models', *arXiv preprint arXiv:2102.09690* .

Figure 1: Building-Block Decompositions of Equity Return Expectations. This figure presents three asset managers' approaches to decomposing equity return forecasts into fundamental components for 2023. Panel (a): BNY Mellon separates returns into inflation, real earnings growth, income, valuation, and currency effects. Panel (b): PGIM uses income, real earnings growth, inflation, valuation adjustment, and currency. Panel (c): T. Rowe Price employs valuation change, expected inflation, dividend/buyback yield, and real EPS growth.

(a) BNY Mellon (30 September 2023)

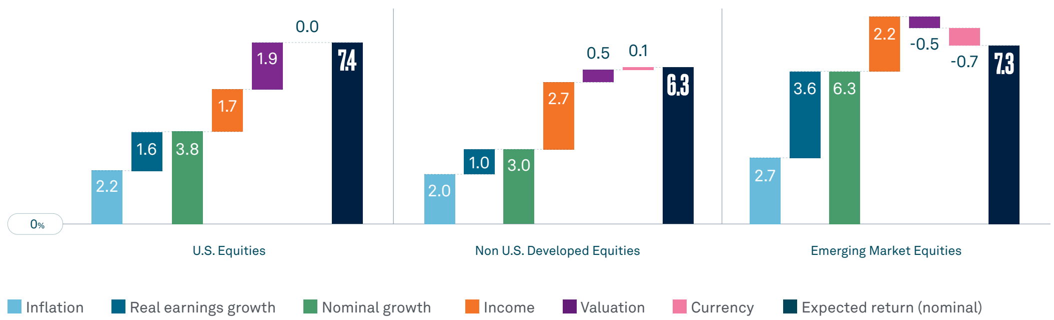

(b) PGIM (30 September 2023)

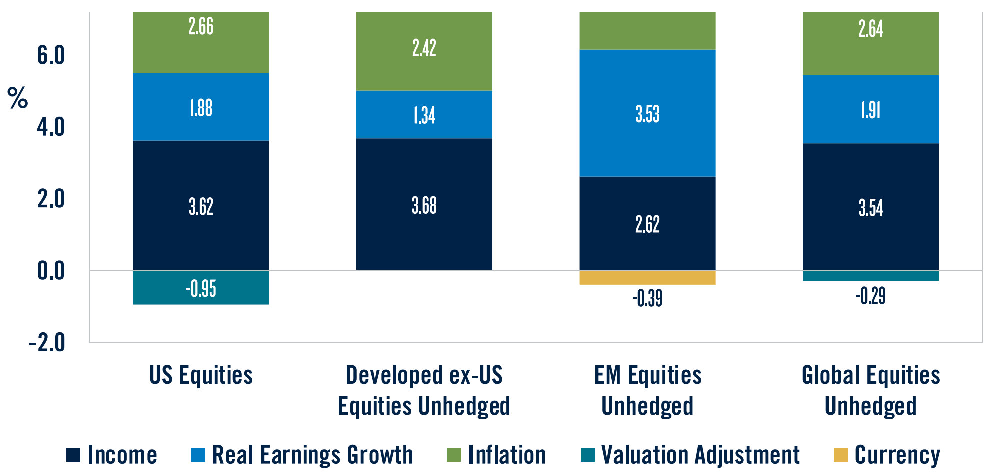

(c) T. Rowe Price (31 January 2023)

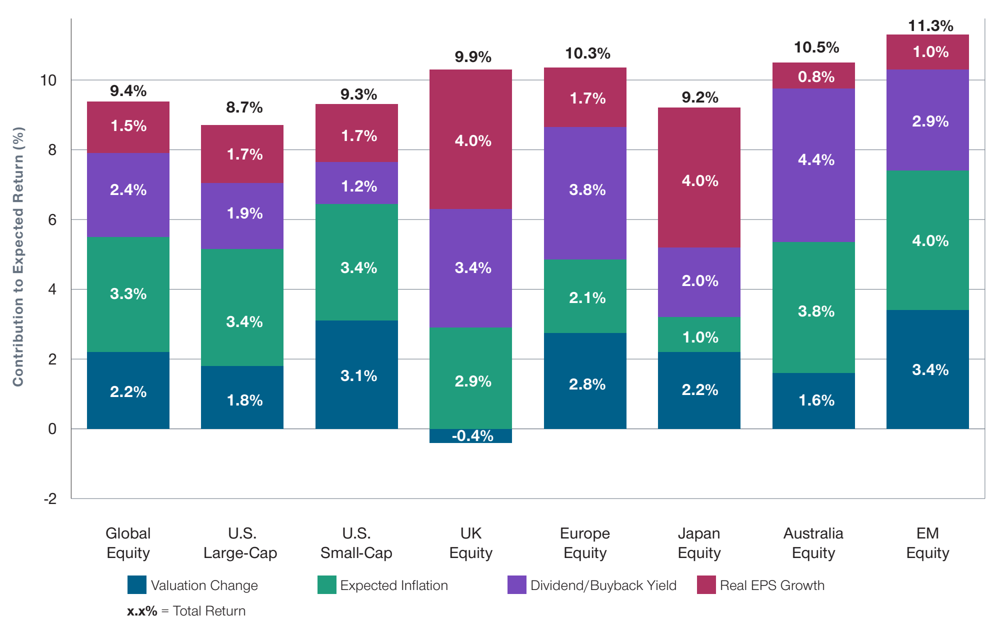

Figure 2: Subjective Return Beliefs for U.S. Equity. This figure shows forecasts for U.S. equity from CMA reports. Panel (a) shows subjective return expectations. Panel (b) shows subjective volatility forecasts. Panel (c) displays time-series of return expectations from five representative asset managers (JP Morgan, Sellwood, AQR, Northern Trust, Callan). Sample: CMA reports, 2015–2025.

(a) Subjective Return Expectation

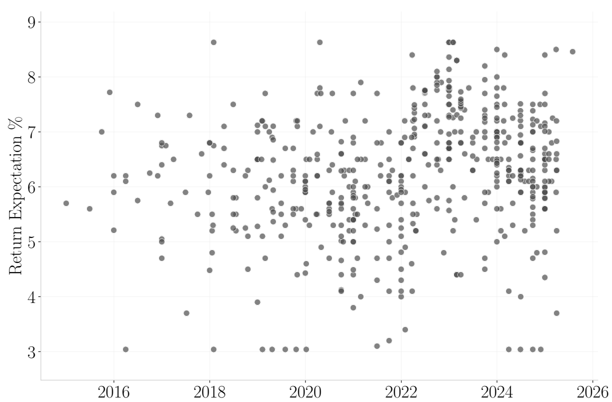

(b) Subjective Return Volatility

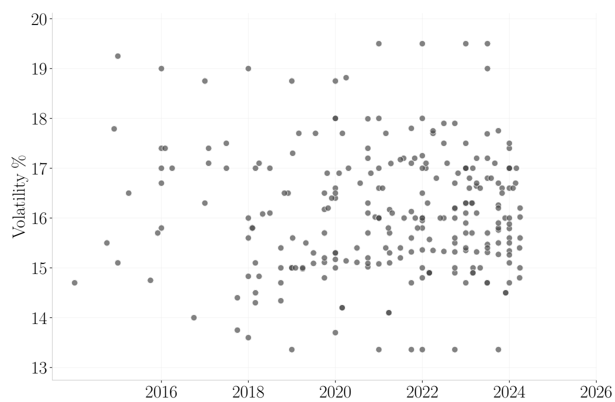

(c) Selected Managers: Subjective Return Expectation

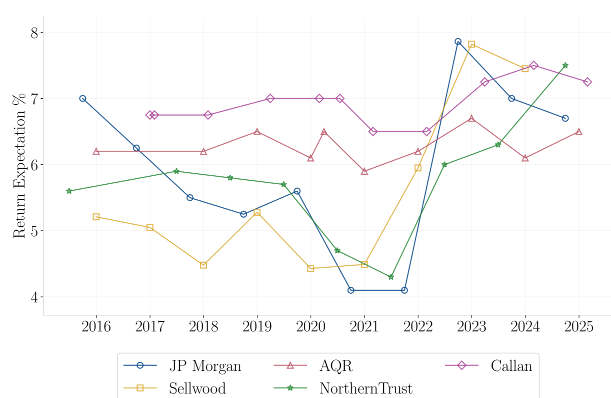

Figure 3: LLM Pipeline for Extracting Causal Networks. This figure illustrates the three-stage pipeline for extracting causal relationships from Capital Market Assumptions (CMAs). Relevant text sections are first extracted with surrounding context, then processed through a large language model: (i) identification of causal links between economic variables, (ii) classification into topics and subtopics, and (iii) determination of directional polarity (positive/negative) and intensity. Each causal link identifies a causal agent, an affected agent, their associated topics, and the directional relationship between them.

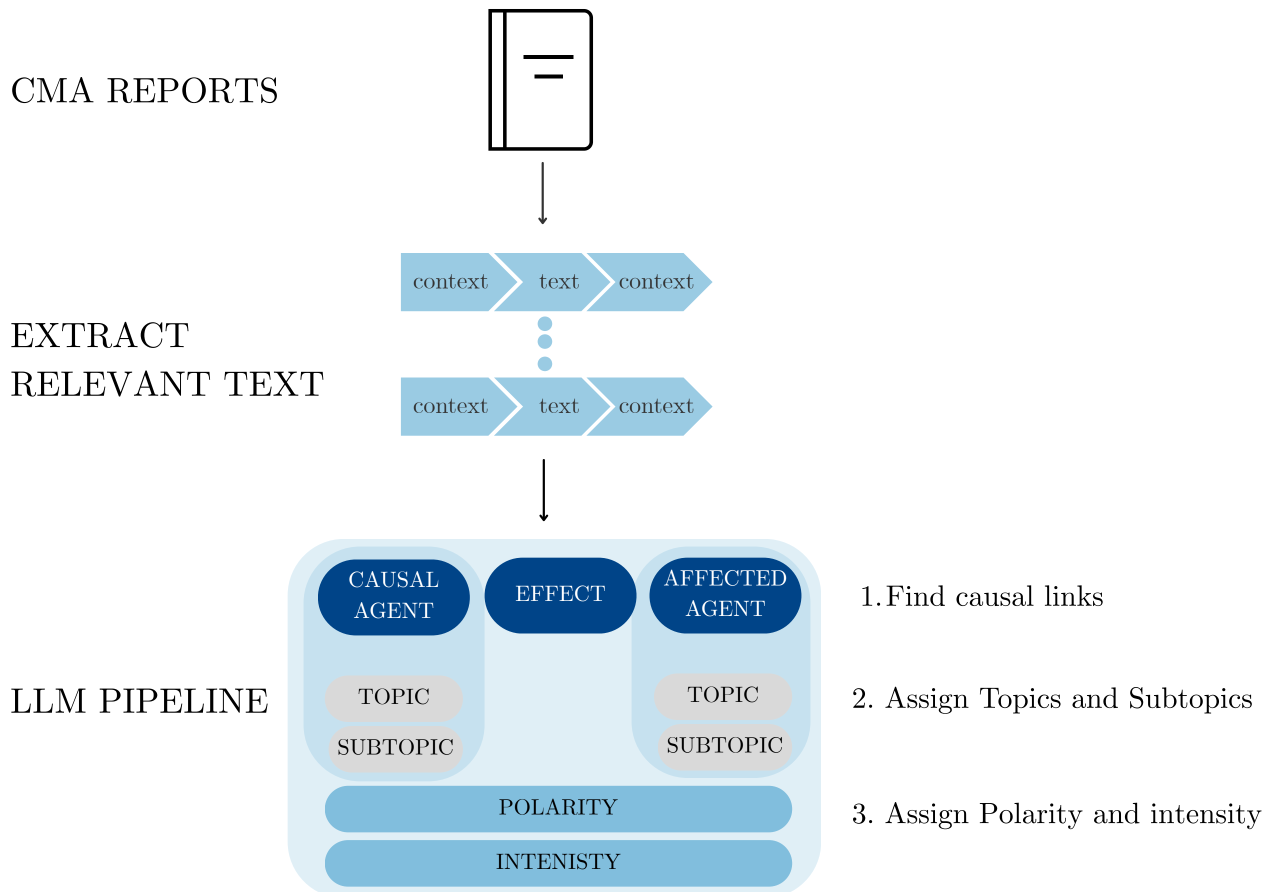

Figure 4: Building Block Structure in Institutional Equity Return Models. This figure shows the structure of return decompositions across asset managers using building blocks (described in Section 2.1.1). Panel (a) displays the distribution of reports by number of components used in their return decomposition. Panel (b) presents the adoption rate of each component across institutions: nominal growth, dividend yield, valuation change, buybacks, margin adjustments, and share issuance. Sample: CMA reports, 2010–2025.

(a) Distribution of building blocks per model

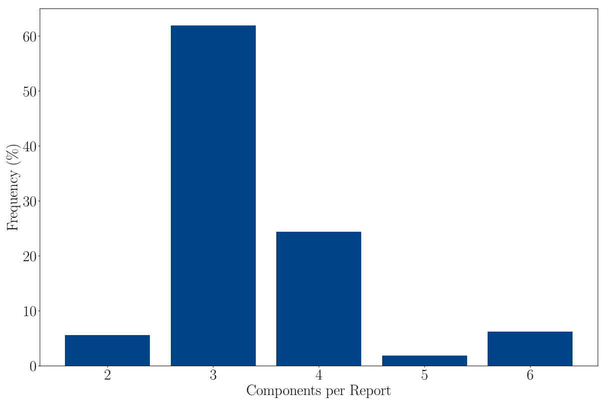

(b) Institutional adoption by component type

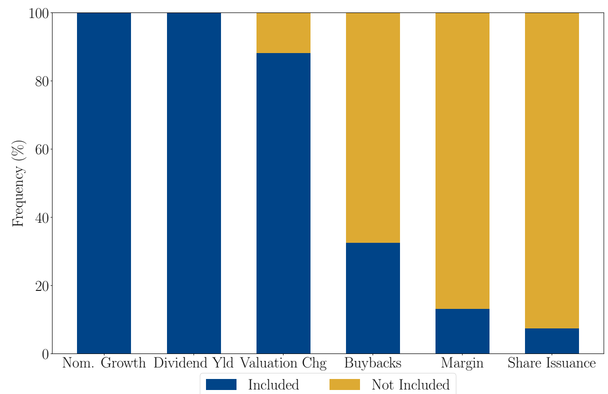

Figure 5: Sources of Variation in Subjective Equity Return Forecasts. This figure shows a variance decomposition using partial R^2 from a saturated fixed-effects model of U.S. equity return expectations. Each bar represents the percentage of total variance explained by asset manager, year, and forecast horizon fixed effects, respectively. Sample: CMA reports, 2010–2025.

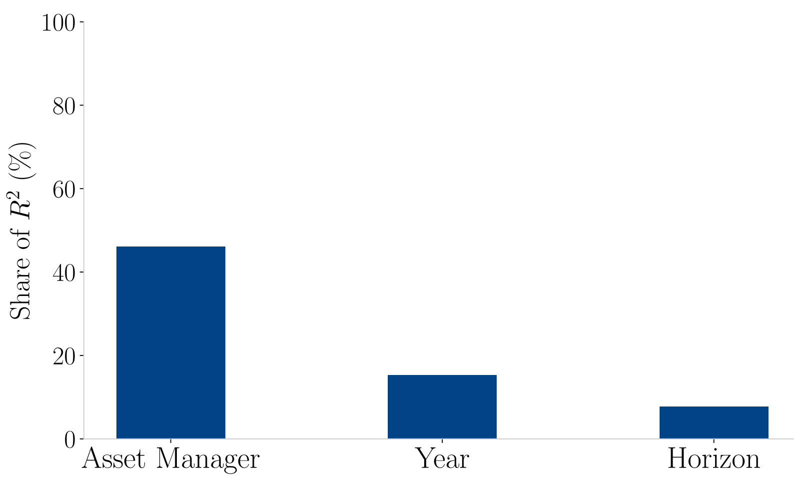

Figure 6: Time-Series of Return Expectation Components. This figure shows average subjective U.S. equity return expectations decomposed into constituent building blocks (described in Section 2.1.1). Each colored area represents one component's contribution: dividend yield, inflation, real growth, share issuance, buybacks, margins, and valuation change. The red line indicates the average subjective equity return across all asset managers. Sample: CMA reports, 2012–2024.

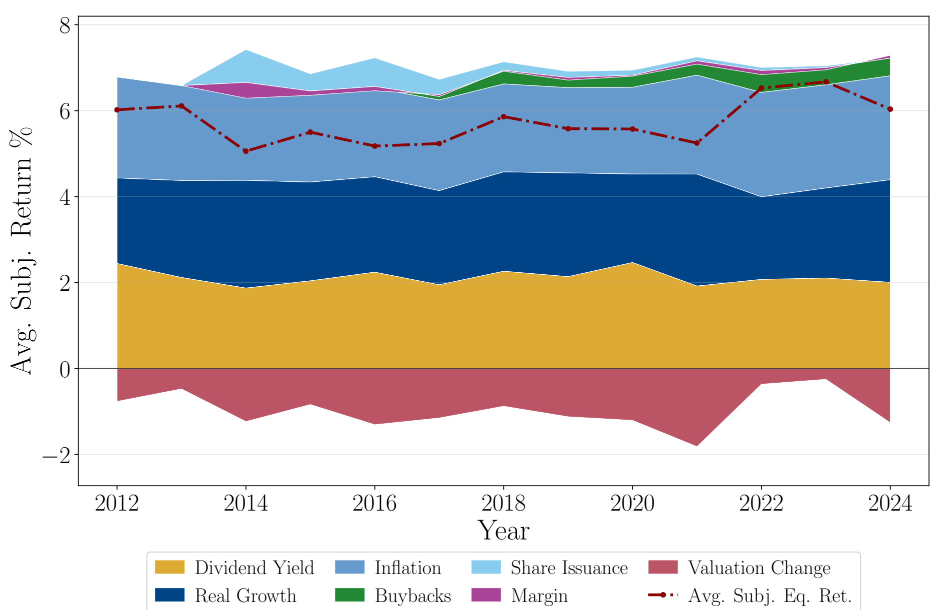

Figure 7: Time-Series of Dispersion in Return Expectation Components. This figure shows annual cross-sectional standard deviations of each building block component (described in Section 2.1.1), measuring disagreement among asset managers. Each colored area represents the dispersion of one component: valuation change, dividend yield, real growth, and inflation. The red dashed line shows the standard deviation of total subjective equity returns across asset managers. Sample: CMA reports, 2012–2024.

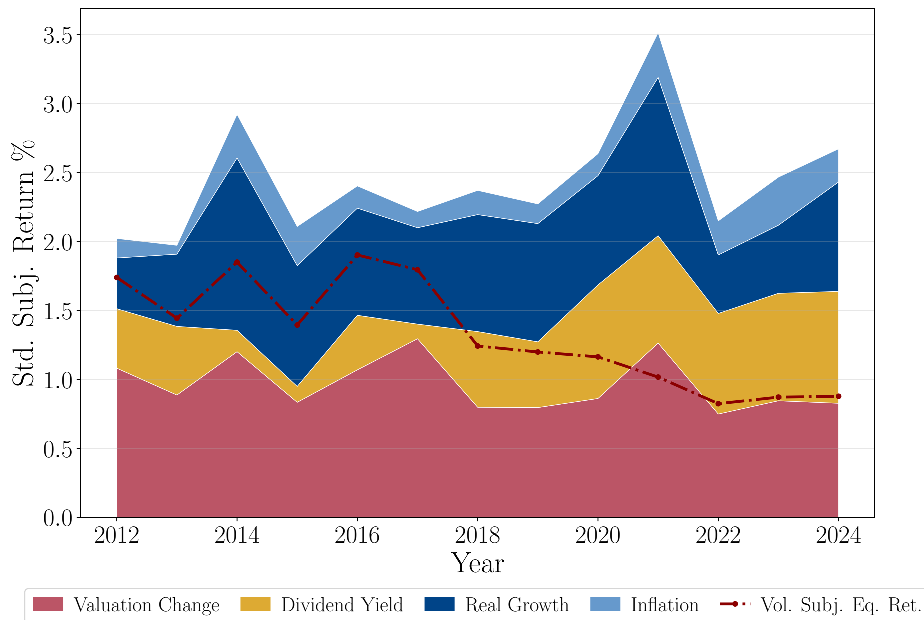

Figure 8: Variance Decomposition of Equity Return Forecasts by Building Block Components. This figure shows partial R^2 estimates from a regression of subjective U.S. equity return expectations on seven building block components (described in Section 2.1.1) with time fixed effects. Each bar represents the percentage of cross-sectional variance in return forecasts explained by valuation change, real growth, dividend yield, buybacks, inflation, share issuance, and margin adjustments. Sample: CMA reports with explicit building block decompositions, 2010–2025.

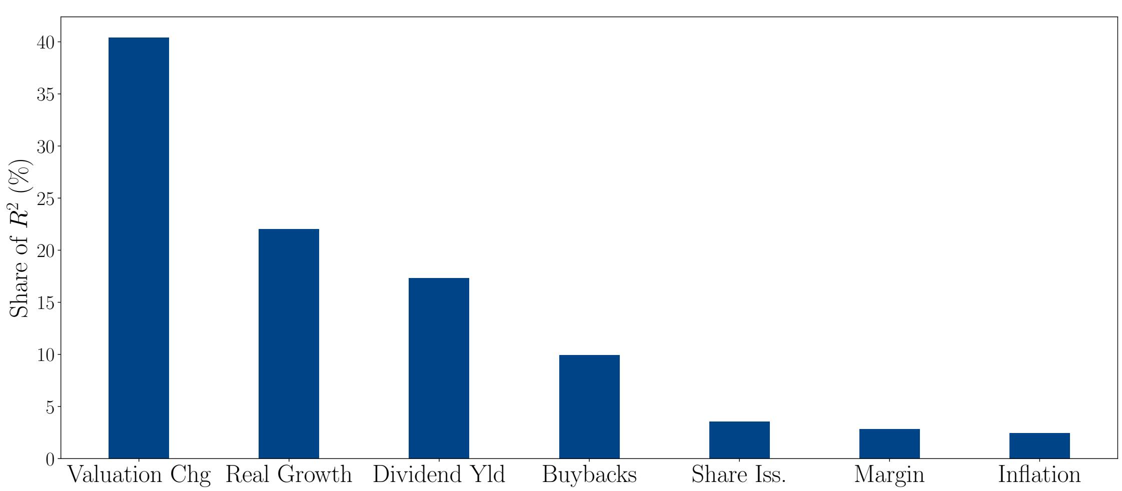

Figure 9: Campbell-Shiller Decomposition: Asset Manger-Level Loadings on Valuation Ratios. This figure shows regression coefficients from asset manager level specifications where ten year U.S. equity forecasts are regressed on valuation ratios. In each panel, $Y_{i,t}$ represents return expectations (left), nominal growth component (middle), or valuation change component (right). Panel (a) regresses forecasts on the log price-dividend ratio ($\ln(\text{PD}_t)$). Panel (b) regresses forecasts on the log cyclically-adjusted price-earnings ratio ($\ln(\text{CAPE}_t)$). Dots denote point estimates, horizontal bars are 95% confidence intervals. Sample: CMA reports by asset managers with at least five years of forecasts.

(a) $\quad Y_{i,t} = \alpha_i + \beta_i \ln(\text{PD}_t) + \varepsilon_{i,t}$ $\qquad$ (b) $\quad Y_{i,t} = \alpha_i + \beta_i \ln(\text{CAPE}_t) + \varepsilon_{i,t}$

(a) CMA forecasts on ln(PD).

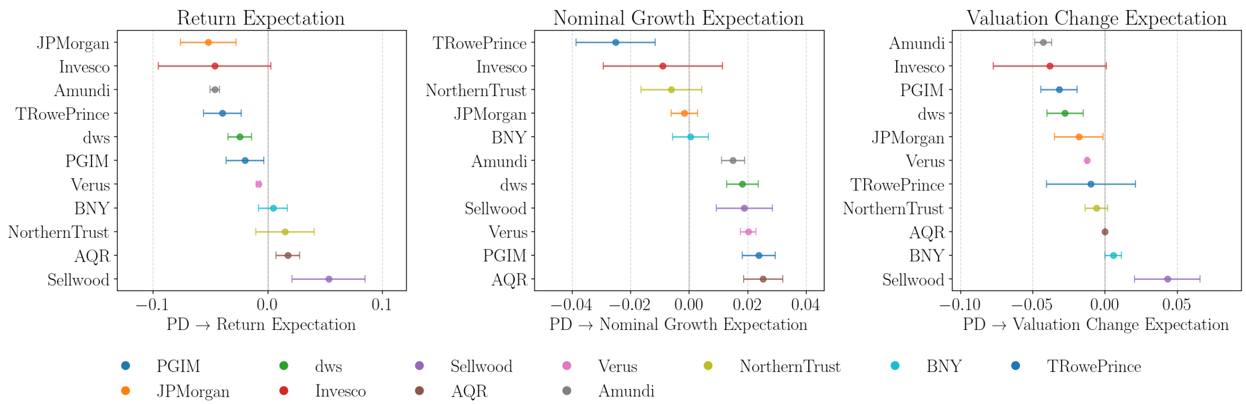

(b) CMA forecasts on ln(CAPE).

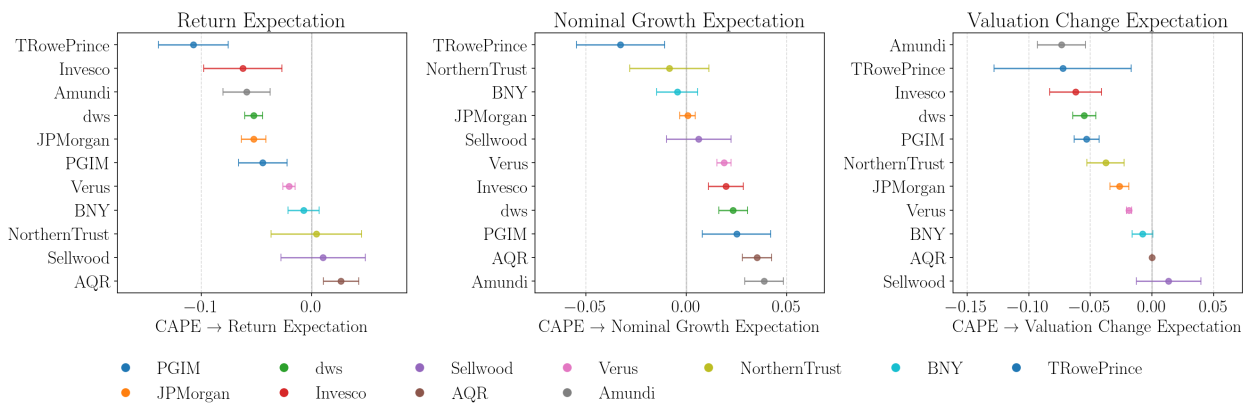

Figure 10: Sources of Variation in Volatility Forecasts. This figure shows Shapley values decomposing the R^2 from a saturated fixed-effects regression of U.S. equity volatility forecasts. Each bar represents the percentage of total variance explained by asset manager, year, and horizon fixed effects, respectively. Sample: CMA volatility forecasts, 2010–2025.

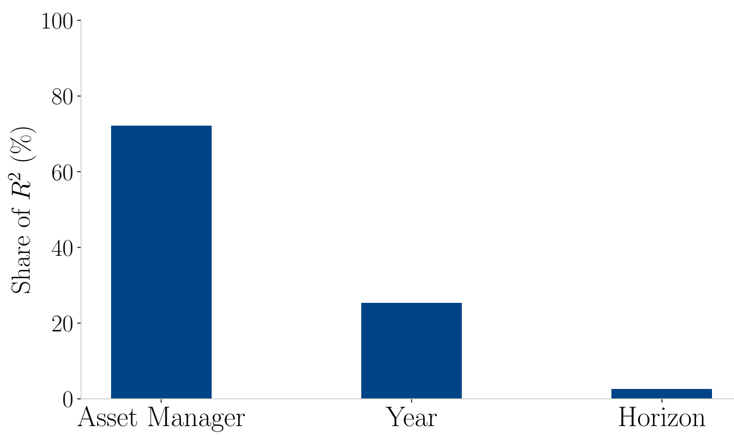

Figure 11: Subjective Volatility Forecasts versus Realized Volatility Measures. This figure compares asset managers' subjective volatility forecasts (blue box plots) with exponentially-weighted moving averages (EWMA) of trailing realized volatility for U.S. equity, quarterly frequency. Box plots show cross-sectional distribution of forecasts across managers at each point in time. Lines represent EWMA of 10-year trailing realized volatility with decay parameters: $\lambda = 0.98$, $\lambda = 0.75$, and $\lambda = 1.00$. Forecasts are winsorized at 1st and 99th percentiles. Sample: CMA reports and CRSP S&P 500 index, 2012–2025.

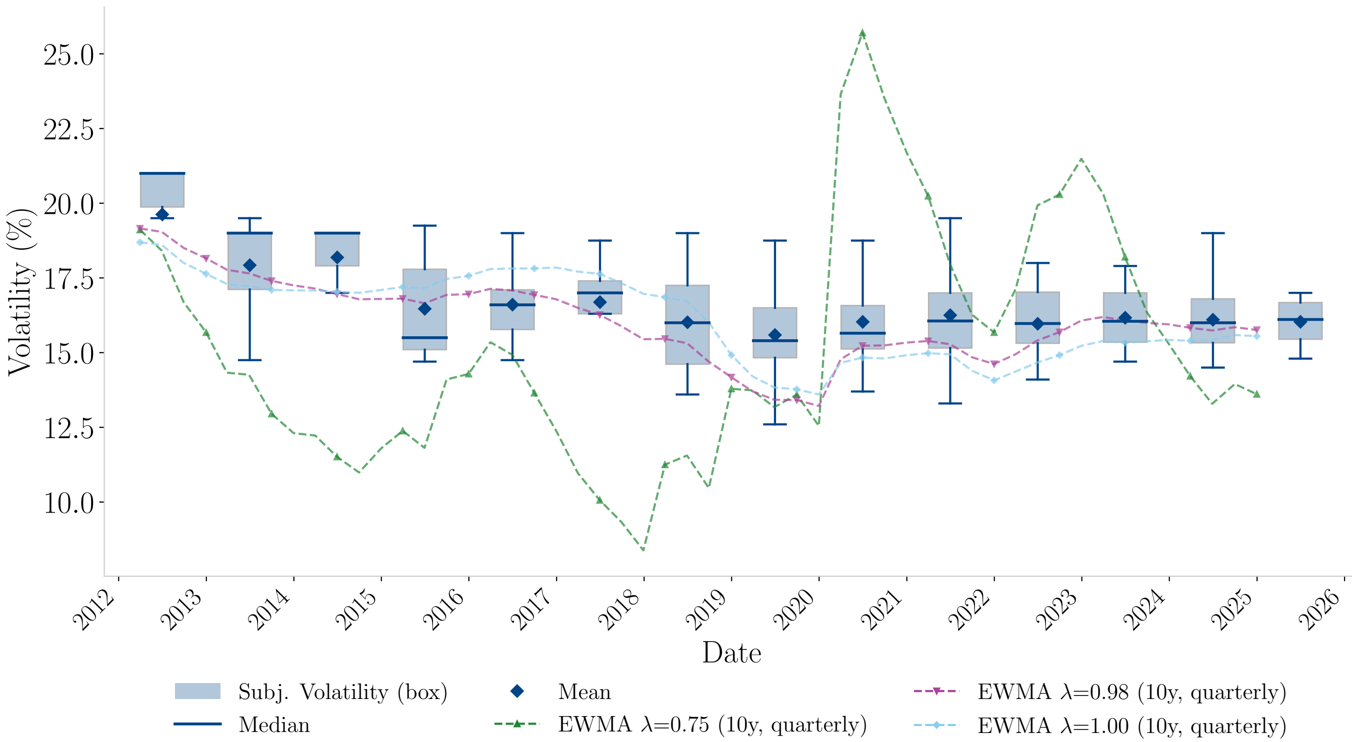

Figure 12: Topic-Based Clustering of CMA Reports. This figure shows a two-dimensional t-SNE projection of document-topic distributions from CMA reports. Each point represents one institution-report characterized by its LDA topic distribution. We first estimate a 60-topic LDA model on the reports, then apply K-Means to group them into 14 distinct clusters (shown by colors). For visualization, t-SNE projects the high-dimensional data onto two dimensions. Sample: CMA reports, 2010–2025.

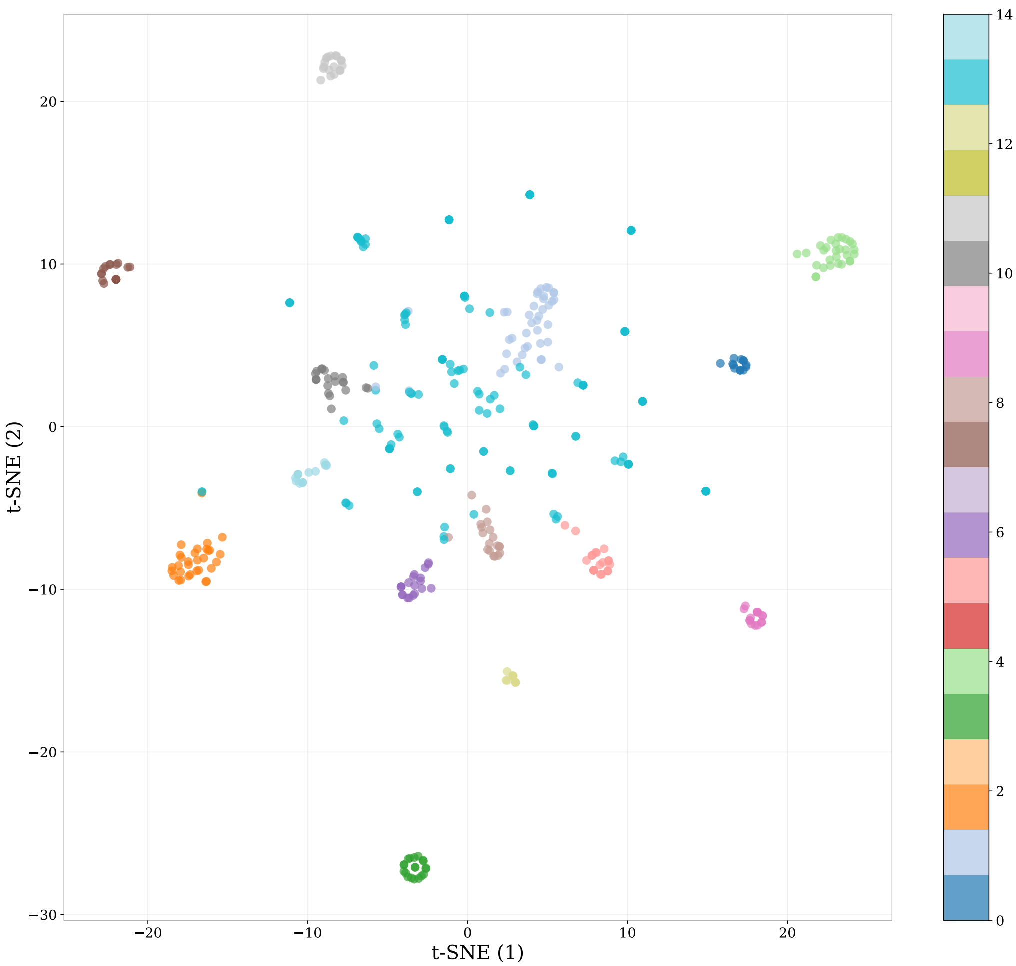

Figure 13: Asset Manager Causal Network of U.S. Equity Return Expectations. This figure shows the directed network of causal relationships (described in Section 1.3) extracted from DWS Capital Market Assumptions. Nodes represent topics identified in the text, with size reflecting centrality (number of connections). Edges indicate the direction of causality, with colors showing positive (green), negative (red), or neutral (gray) causal effects as stated in the reports. U.S. Equity Return and the four building blocks from the return decomposition (described in Section 2.1.1)—growth, valuation change, dividend yield, and inflation—are highlighted with distinct colors. Sample: DWS causal network, 2015–2025.

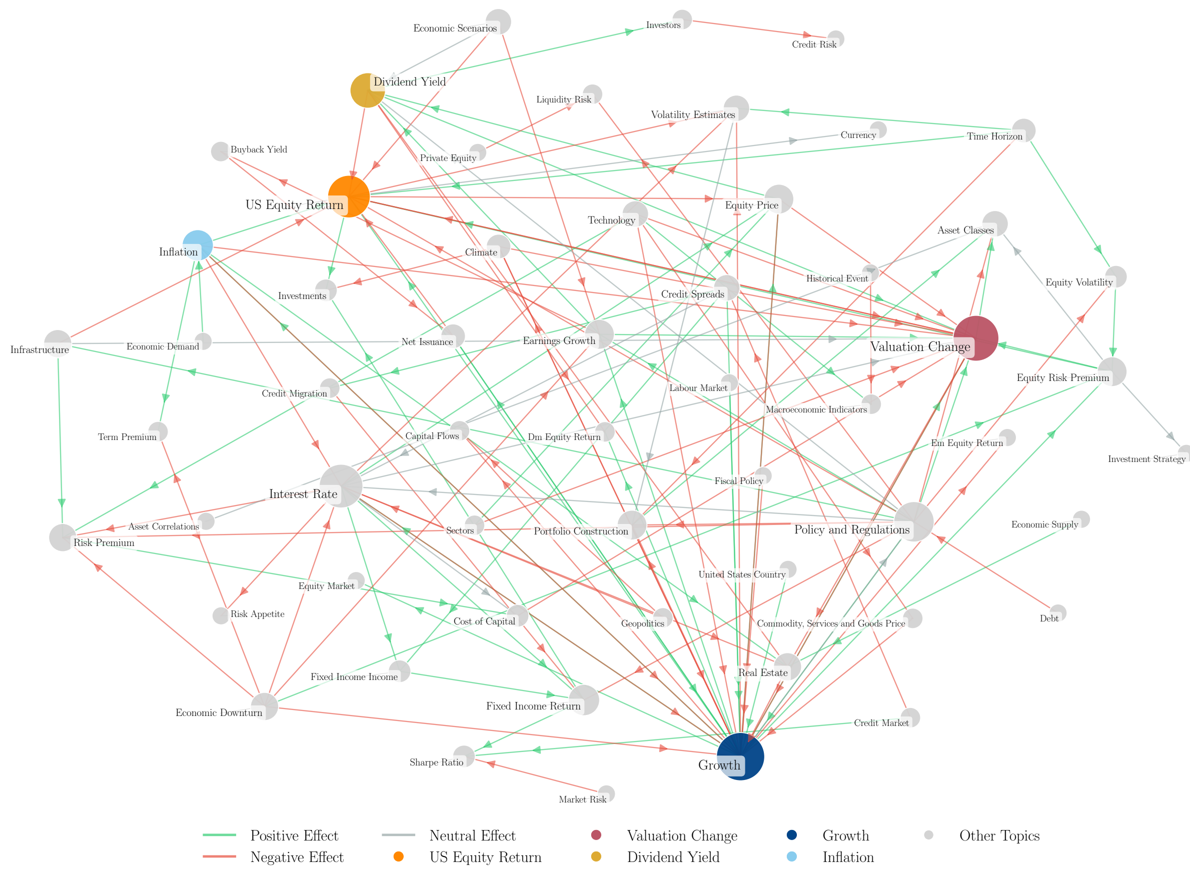

Figure 14: Asset Manager Direct Causal Drivers of U.S. Equity Return Expectations. This figure shows all topics directly linked to U.S. Equity Return in DWS Capital Market Assumptions. Node size reflects the number of causal connections (described in Section 1.3). Edges indicate the direction of causality, with colors showing positive (green), negative (red), or neutral (gray) effects as stated in the text. U.S. Equity Return and the four building blocks from the return decomposition (described in Section 2.1.1)—growth, valuation change, dividend yield, and inflation—are highlighted with distinct colors. Sample: DWS causal network, 2015–2025.

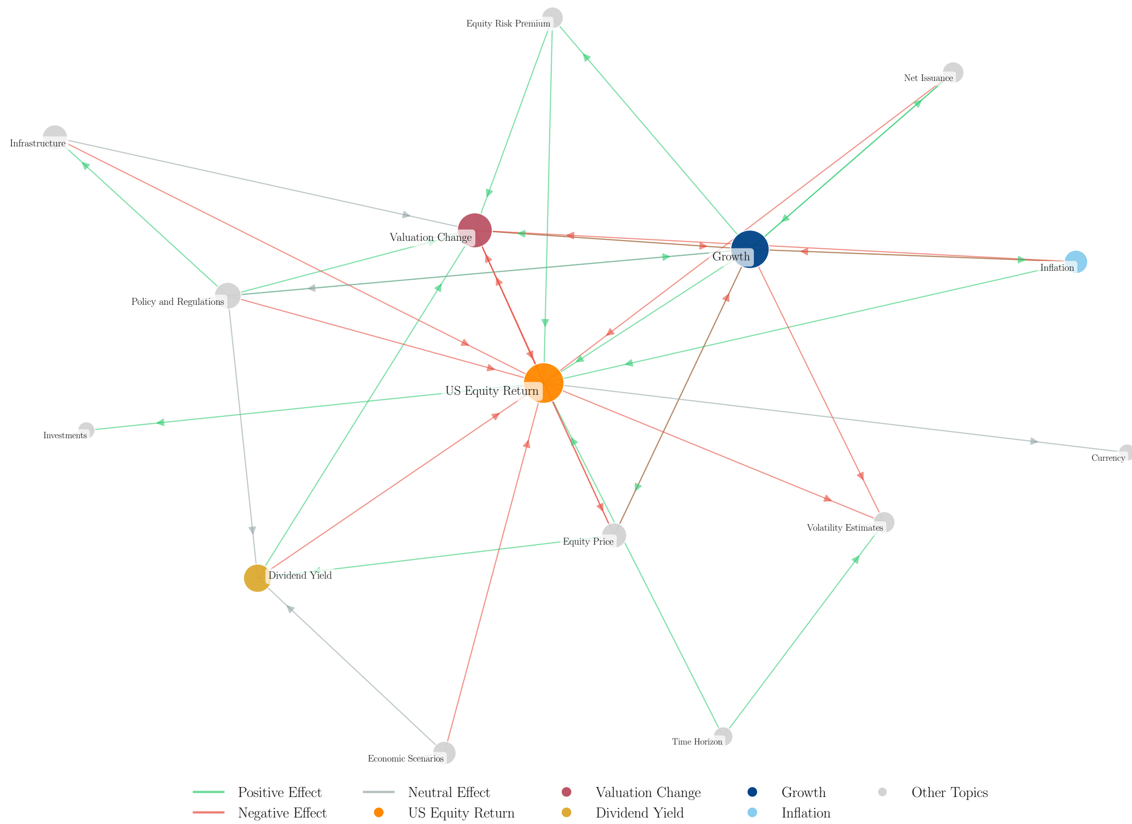

Figure 15: Causal Network Topic Composition. This figure shows the prevalence of topics in asset managers' causal networks (described in Section 1.3). Panel (a) shows the percentage of asset manager-year observations where each topic appears in any causal relationship (as cause or effect). Panel (b) shows the percentage of asset manager-year observations where each topic directly links to U.S. equity return expectations (as driver or consequence). Sample: Asset manager-year causal networks, 2015–2025.

(a) Topic frequency

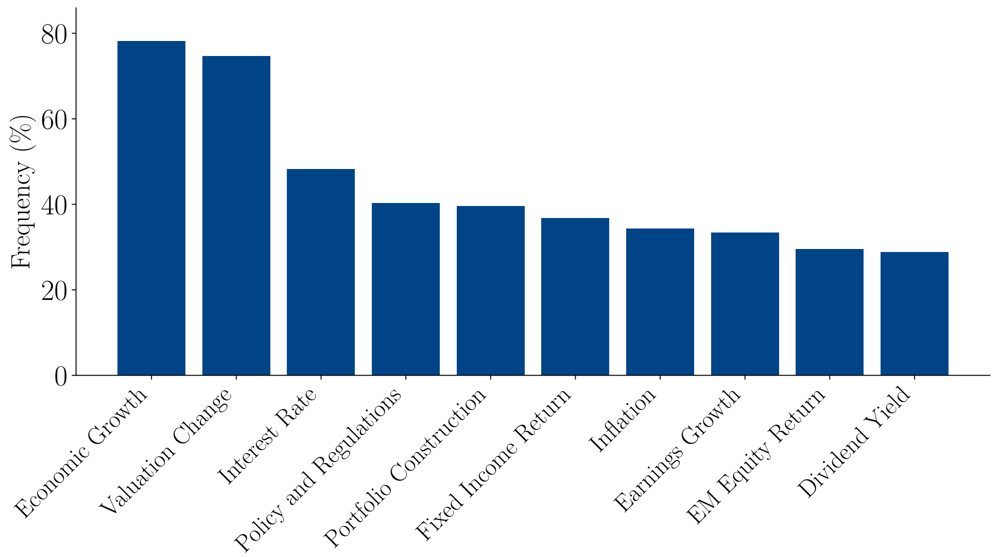

(b) Topics directly linked to U.S. equity return expectations

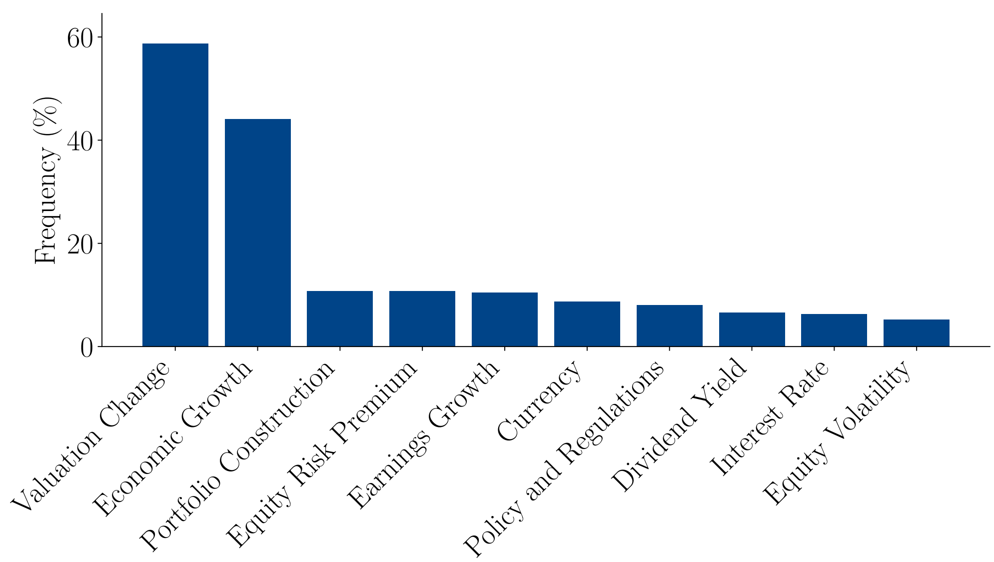

Figure 16: Heterogeneity in Causal Network Composition Across Asset Managers. This figure shows the average number of causal links (described in Section 1.3) per asset manager-year, decomposed by topic. Each horizontal bar represents one asset manager, with colored segments showing the contribution of the top 10 most frequent topics and gray segments showing all remaining topics. Sample: CMA causal networks, 2015–2025.

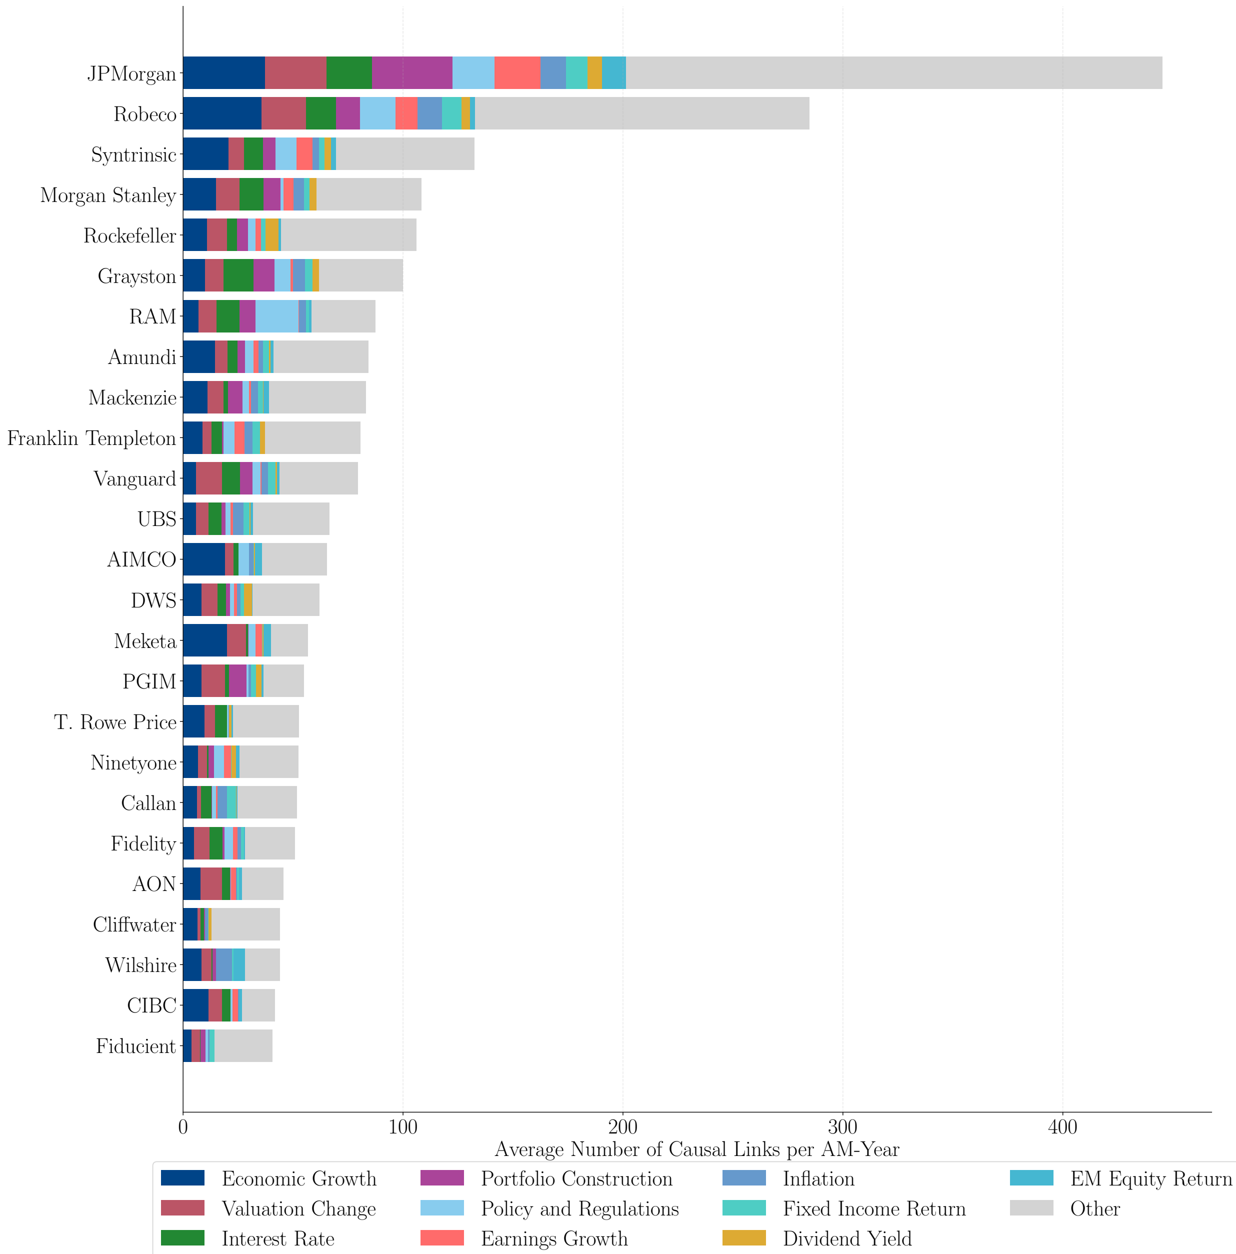

Table 1: CMA Institutional Investors. This table lists the 78 institutions in our Capital Market Assumptions dataset (54 asset managers and 24 consultants). The period column indicates the years of CMA coverage for each institution, depending on data availability and publication schedules. Sample: CMA reports 2008–2025.

Institution	Type	Years	Institution	Type	Years
ACG	Consultant	2023–2023	Invesco	Manager	2018–2024
Advent	Manager	2024–2024	Janney	Manager	2022–2025
AIMCo	Manager	2021–2023	Janus Henderson	Manager	2020–2022
Allianz	Manager	2021–2024	JP Morgan	Manager	2009–2024
American Century	Manager	2024–2024	Mackenzie	Manager	2019–2023
Amundi	Manager	2018–2025	Manulife	Manager	2020–2020
Angeles	Consultant	2018–2022	Meketa	Consultant	2020–2024
AON	Consultant	2012–2024	Merrill Lynch	Consultant	2018–2018
AQR	Manager	2013–2024	MFS	Manager	2024–2024
Baillie Gifford	Manager	2018–2018	Morgan Stanley	Consultant	2019–2025
Barclays	Manager	2021–2023	NEPC	Consultant	2016–2024
Benjamin F.E.	Manager	2024–2024	Neuberger	Manager	2022–2023
BlackRock	Manager	2019–2024	Ninety One	Manager	2023–2025
BNP	Manager	2020–2024	Northern Trust	Manager	2013–2024
BNY Mellon	Manager	2013–2024	Nuveen	Manager	2022–2025
Callan	Consultant	2016–2025	OneAscent	Consultant	2024–2024
CalPERS	Manager	2021–2021	PGIM	Manager	2018–2024
Candriam	Manager	2024–2024	PIMCO	Manager	2022–2024
Capital Group	Manager	2022–2024	RAM	Manager	2020–2024
CIBC	Manager	2018–2025	RBC	Manager	2022–2025
Cliffwater	Consultant	2019–2025	Robeco	Manager	2019–2024
Cohen Steers	Manager	2022–2024	Rockefeller	Consultant	2019–2020
Columbia Thread.	Manager	2017–2025	RVK	Consultant	2021–2024
Credit Suisse	Manager	2020–2020	SEI	Consultant	2025–2025
Crescent Capital	Manager	2020–2024	Schroders	Manager	2017–2025
CWO	Manager	2019–2020	Sellwood	Consultant	2012–2023
Deutsche Bank	Manager	2023–2025	SJS	Consultant	2019–2020
DSA	Consultant	2011–2011	State Street	Manager	2016–2025
DWS	Manager	2017–2024	Syntrinsic	Consultant	2016–2024
Edward Jones	Manager	2020–2025	TD	Manager	2022–2024
EGF	Manager	2024–2024	T. Rowe Price	Manager	2019–2025
Ellwood	Consultant	2017–2017	UBS	Manager	2019–2023
Envestnet	Consultant	2013–2024	Vanguard	Manager	2014–2025
FI3	Manager	2022–2022	Verus	Consultant	2008–2024
Fidelity	Manager	2020–2024	Voya	Manager	2016–2024
Fiducient	Consultant	2021–2024	Wealthspire	Consultant	2023–2023
Franklin Temp.	Manager	2016–2024	Wellington	Manager	2021–2024
Graystone	Consultant	2021–2023	Wells Fargo	Manager	2023–2023
HSBC	Manager	2024–2024	Wilshire	Consultant	2021–2024

Table 2: Asset Class Coverage and Return Expectations. This table shows coverage statistics for the 10 core asset classes in the CMA dataset. For each asset class: number of institutions providing forecasts, years of coverage, median forecast horizon, and average subjective return expectation and volatility. Sample: CMA reports, 2008–2025.

Asset Class	No. of Inst.	Years	Horizon (median)	Avg. Subj. Exp.	Avg. Subj. Vol.
U.S. Equity	76	2008 - 2025	10 yrs	6.28%	16.25%
Developed Equity	70	2008 - 2025	10 yrs	6.94%	17.52%
Emerging Equity	74	2008 - 2025	10 yrs	8.19%	22.70%
U.S. Govt Long Bond	25	2009 - 2025	10 yrs	2.57%	13.29%
U.S. Govt Bond	45	2008 - 2025	10 yrs	3.00%	5.82%
U.S. IG Bond	41	2009 - 2025	10 yrs	3.83%	7.39%
U.S. HY Bond	68	2008 - 2025	10 yrs	5.14%	9.97%
U.S. TIPS	47	2008 - 2025	10 yrs	3.25%	6.09%
U.S. Cash	64	2008 - 2025	10 yrs	2.41%	0.83%
U.S. Inflation	45	2008 - 2025	10 yrs	2.32%	2.15%

Table 3: Causal Network Complexity and Topic Attention Statistics. This table reports summary statistics for causal network measures (described in Section 1.3) aggregated at the asset manager level. Panel A presents statistics for network complexity measures including indirect connection ratio, average eigenvector centrality scores, average shortest path length and network transitivity. Panel B presents statistics for topic-specific attention measures based on edge counts for valuation change/adjustment, economic growth, and dividend yield/income topics. For each measure, we report the mean, standard deviation, and coefficient of variation. Data are winsorized at the 1% and 99% levels before calculating statistics. Sample: CMA causal networks, 2015–2025.

Panel A: Network Complexity Measures

	Mean	Std. Dev.	Std/Mean
Indirect Ratio	0.097	0.090	0.92
Avg Eigenvector Centrality	0.217	0.151	0.70
Avg Shortest Path Length	2.284	0.744	0.33
Network Transitivity	0.033	0.039	1.17

Panel B: Topic Attention Measures

	Mean	Std. Dev.	Std/Mean
Valuation Change Edge Count	4.387	5.147	1.17
Economic Growth Edge Count	6.164	7.649	1.24
Dividend Yield Edge Count	0.732	1.218	1.66

Table 4: Diebold-Mariano Test: CMA vs. Objective Model Forecasts. This table reports Diebold-Mariano test statistics comparing the predictive accuracy of CMA return forecasts with the objective model (described in Section 2.1.3). For each horizon (3, 5, and 10 years) and pooled across all horizons, we report the number of paired observations (N), root mean squared error (RMSE) for CMA and objective model forecasts, the DM test statistic, and p-values. A positive DM statistic indicates that the objective model has lower forecast error than CMA forecasts. Sample period: 2010 to 2025.

Horizon	N	$RMSE_{CMA}$	$RMSE_{Obj}$	DM Stat	p-value
3 years	13	12.40%	10.38%	3.06***	0.002
5 years	22	8.82%	7.14%	4.98***	<0.001
10 years	22	6.17%	3.75%	9.39***	<0.001
Pooled	65	8.83%	7.01%	6.14***	<0.001

Table 5: Forecast Changes and News Sentiment. This table reports regression results examining the relationship between year-over-year forecast changes and the aggregate earnings-news sentiment metric. The dependent variable in each column is the annual forecast change, defined as $Y_{i,t} - Y_{i,t-1}$, where $Y_{i,t}$ denotes either the annual average return forecast or a return-forecast building block from asset manager i in year t. Columns (1) and (2) present results for overall U.S. equity return forecasts (ER) for the full sample and for the subsample restricted to years in which all building blocks—nominal growth (g), dividend yield (DY), and valuation change (VC)—are available. Columns (3)–(5) report results from decomposing equity return forecasts into these building blocks. The key independent variable is an aggregate sentiment measure, S_t, constructed as a one-year rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1). All regressions include asset manager fixed effects (μ_i). Robust standard errors, clustered by asset manager, are reported in parentheses. Sample period: 2010 to 2025.

$$Y_{i,t} - Y_{i,t-1} = \mu_i + \beta\, S_t + \varepsilon_{i,t}$$

$\Delta Y_{i,t} = Y_{i,t} - Y_{i,t-1}$	(1) ΔER	(2) ΔER	(3) Δg	(4) ΔDY	(5) ΔVC
Avg. Earnings Sentiment (1y)	5.8918***	7.4205***	-2.3056	0.4619	10.0454**
	(1.7166)	(1.8159)	(2.9165)	(0.5675)	(4.3422)
Asset Manager FE	Y	Y	Y	Y	Y
R^2	0.221	0.286	0.204	0.248	0.342
Observations	181	81	81	81	67

Table 6: Ex-ante Forecast Errors and News Sentiment. This table presents regression results examining the relationship between CMA return expectations ex-ante forecast errors and earnings-news sentiment metrics. The dependent variable is the ex-ante forecast error, defined as $G_t - F_{i,t}$, where G_t is the benchmark forecast (described in Section 2.1.3) and $F_{i,t}$ is the annual average return forecast from asset manager i in year t. We use an aggregate sentiment measure, S_t, constructed as a one-year rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1). The regression includes asset manager fixed effects. Robust standard errors, clustered by asset manager, are reported in parentheses. Sample period: 2010 to 2025.

$$G_t - F_{i,t} = \mu_i + \beta\, S_t + \varepsilon_{i,t}$$

	$G_t - F_{i,t}$
Avg. Earnings Sentiment (1y)	-4.2896***
	(1.4230)
Asset Manager FE	Y
R^2	0.432
Observations	401

Table 7: Portfolio Allocations and Return Expectations. This table reports regressions of U.S. equity allocations on CMA subjective returns and their building-block components (described in Section 2.1.1), where $w_{i,t}$ is fund i's net U.S. equity allocation at time t, and $F_{i,t}$ is its CMA subjective return, decomposed as dividend yield ($DY_{i,t}$), nominal growth ($g_{i,t}$), and valuation change ($VC_{i,t}$). Column (1): subjective excess returns across asset classes. Column (2): U.S. equity return expectations only. Column (3): decomposed components. The regressions include asset manager fixed effects (μ_i) and year fixed effects (τ_t). Standard errors (in parentheses) clustered by fund and year. Sample period: 2010 to 2025.

$$w_{i,t} = \mu_i + \tau_t + \alpha + \beta X_{i,t} + \varepsilon_{i,t}$$

	Asset Alloc U.S. Equity Net ($w_{i,t}$)		
$X_{i,t}$	(1)	(2)	(3)
Subj. Exc. Developed Eq.	-0.4624**		
	(0.2334)		
Subj. Exc. Emerging Eq.	-1.3595***		
	(0.3356)		
Subj. Exc. U.S. Eq.	0.5983***		
	(0.2252)		
Subj. U.S. Eq.		0.2365***	
		(0.0695)	
Dividend Yield			-0.3622
			(0.2863)
Nominal Growth			0.4582***
			(0.1445)
Valuation Change			0.2336**
			(0.1111)
Fund FE	Y	Y	Y
Year FE	Y	Y	Y
R^2	0.063	0.036	0.034
Observations	3251	25765	25765

Table 8: EWMA Estimation of Volatility Forecasts. This table reports nonlinear least squares estimates of the exponentially-weighted moving average model for subjective volatility forecasts. The decay parameter γ governs the persistence of memory in the averaging process, with the half-life indicating how quickly past observations lose influence. Dependent variable is forecasted annualized volatility; regressor is EWMA of quarterly realized volatilities with endogenously estimated γ.

$$F_t = \alpha + \beta \cdot \text{EWMA}_t(\gamma) + \varepsilon_t$$

$$\text{EWMA}_t(\gamma) = \sum_{j=1}^{L} \omega_j \cdot RV_{t-j} \quad \text{where} \quad \omega_j = \frac{\gamma^{j-1}}{\sum_{k=1}^{L} \gamma^{k-1}}$$

	F_t
a (Intercept)	8.3672***
	(1.5118)
b (Slope)	0.4837***
	(0.0915)
γ (Decay)	0.9814***
	(0.0103)
Half-life	36.82 quarters
R^2	0.0902
Observations	276

Table 9: Equity Return Expectations and Network Complexity. This table presents regression results of CMA U.S. equity return expectations on the causal network complexity measures (described in Section 1.3) using a within–between decomposition (described in Section 3.2). Each column corresponds to a different network complexity measure: (1) average path length (APL), (2) indirect connection ratio (ICR), and (3) network transitivity (NT). $F_{i,t}$ is the annual average CMA U.S. equity return forecast (in percentage points) for asset manager i in year t. $C_{i,t}$ denotes the causal network complexity measure for asset manager i in year t. $\overline{C}_i$ is manager i's time-series mean (between component), and $(C_{i,t} - \overline{C}_i)$ captures the temporary deviation (within component). We use an aggregate sentiment measure, S_t, constructed as a one-year rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1), centered around its time-series mean. The sample includes U.S. equity forecasts at the asset manager–year level. All regressions include asset manager fixed effects (μ_i) and year fixed effects (τ_t). Robust standard errors clustered at the asset-manager level are reported in parentheses. Sample period: 2015 to 2025.

$$F_{i,t} = \mu_i + \tau_t + \beta_w(C_{i,t} - \overline{C}_i) + \gamma_w(C_{i,t} - \overline{C}_i) \times S_t + \gamma_b\overline{C}_i \times S_t + \varepsilon_{i,t}$$

	(1) APL	(2) ICR	(3) NT
$C_{i,t} - \overline{C}_i$	0.2544	-0.2893	-0.4051
	(0.4773)	(0.3657)	(0.2500)
$(C_{i,t} - \overline{C}_i) \times S_t$	17.8837	-20.4201	30.4977*
	(20.1366)	(16.4193)	(17.0434)
$\overline{C}_i \times S_t$	-17.3872***	-19.6916**	-23.4760***
	(6.4102)	(7.7341)	(8.4797)
Asset Manager FE	Y	Y	Y
Year FE	Y	Y	Y
Within R^2	0.220	0.210	0.139
R^2	0.761	0.771	0.762
Observations	173	173	173

Table 10: Equity Return Expectations and Topic Attention. This table presents regression results of CMA U.S. equity return expectations on topic-specific causal network attention (described in Section 1.3) using a within–between decomposition (described in Section 3.2). Each column corresponds to the asset managers' causal network edge count for a specific topic: (1) valuation change (VC), (2) dividend yield (DY), (3) economic growth (g), (4) inflation (π), and (5) economic downturn (ED). $F_{i,t}$ denotes the annual average CMA U.S. equity return forecast (in percentage points) for asset manager i in year t. $A_{i,t}$ denotes the topic-specific edge count for manager i in year t. $\overline{A}_i$ is manager i's time-series mean (between component), and $(A_{i,t} - \overline{A}_i)$ captures the temporary deviation (within component). We use an aggregate sentiment measure, S_t, constructed as a one-year rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1), centered around its time-series mean. The sample includes U.S. equity forecasts at the asset manager–year level. All regressions include asset manager fixed effects (μ_i) and year fixed effects (τ_t). Robust standard errors clustered at the asset-manager level are reported in parentheses. Sample period: 2015 to 2025.

$$F_{i,t} = \mu_i + \tau_t + \beta_w(A_{i,t} - \overline{A}_i) + \gamma_w(A_{i,t} - \overline{A}_i) \times S_t + \gamma_b\overline{A}_i \times S_t + \varepsilon_{i,t}$$

	(1) VC	(2) DY	(3) g	(4) π	(5) ED
$A_{i,t} - \overline{A}_i$	0.0876	-0.1174	-0.0542	-0.1765	-0.8261*
	(0.1260)	(0.3838)	(0.1298)	(0.3064)	(0.4389)
$(A_{i,t} - \overline{A}_i) \times S_t$	-15.2588***	20.5545	0.5411	7.8141	39.0372*
	(5.6974)	(22.0601)	(5.2792)	(13.5185)	(20.3779)
$\overline{A}_i \times S_t$	-12.4515***	11.6953**	-4.1668	16.1492	14.5876***
	(4.1448)	(5.2850)	(3.4444)	(10.8206)	(5.4906)
Asset Manager FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Within R^2	0.227	0.178	0.081	0.096	0.279
R^2	0.788	0.754	0.738	0.746	0.777
Observations	173	173	173	173	173

Table 11: Forecast Changes and Network Complexity. This table presents regression results of CMA U.S. equity return expectation changes on the causal network complexity measures using a within–between decomposition (described in Section 3.2) and news sentiment innovations. The dependent variable in each column is the forecast change, defined as $F_{i,t} - F_{i,t-1}$, where $F_{i,t}$ is annual average CMA return forecast of asset manager i in year t. Each column corresponds to a different network complexity measure: (1) average shortest path length (APL, normalized, weak component), (2) indirect ratio across all topics (IRC, normalized), and (3) network transitivity (NT, normalized). $C_{i,t}$ denotes the complexity measure for manager i in year t. $\overline{C}_i$ is manager i's time-series mean (between component, capturing structural complexity), and $(C_{i,t} - \overline{C}_i)$ is the temporary deviation from this mean (within component, capturing time-varying complexity changes). We use an aggregate sentiment measure, S_t, constructed as a one-year rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1), centered around its time-series mean. ΔS_t denotes the demeaned news sentiment innovation. The sample includes U.S. equity forecasts at the asset manager–year level. All regressions include asset manager fixed effects (μ_i) and year fixed effects (τ_t). Robust standard errors clustered at the asset-manager level are reported in parentheses. Sample period: 2015 to 2025.

$$F_{i,t} - F_{i,t-1} = \mu_i + \tau_t + \beta_w(C_{i,t} - \overline{C}_i) + \gamma_w(C_{i,t} - \overline{C}_i) \times S_t + \gamma_b \overline{C}_i \times S_t + \varepsilon_{i,t}$$

	(1) IRC	(2) APL	(3) NT
$C_{i,t} - \overline{C}_i$	0.1675	0.1072	-0.0606
	(0.3115)	(0.8584)	(0.2013)
$(C_{i,t} - \overline{C}_i) \times S_t$	-34.2382*	0.6198	18.7242**
	(19.4139)	(39.8761)	(8.9155)
$\overline{C}_i \times S_t$	-22.2396***	-24.6361**	-27.3063***
	(3.5139)	(10.1928)	(5.7615)
Asset Manager FE	Y	Y	Y
Year FE	Y	Y	Y
Within R^2	0.158	0.133	0.101
R^2	0.551	0.554	0.525
Observations	122	119	122

Table 12: Forecast Changes and Topic Attention. This table presents regression results of CMA U.S. equity return expectations changes on topic-specific causal network attention (described in Section 1.3) using within-between decomposition (described in Section 3.2) and news sentiment innovations. The dependent variable in each column is the forecast change, defined as $F_{i,t} - F_{i,t-1}$, where $F_{i,t}$ is annual average CMA return forecast of asset manager i in year t. Each column represents standardized edge counts for a specific topic: (1) valuation change (VC), (2) dividend yield (DY), (3) economic growth (g), (4) inflation (π), (5) economic downturn (ED). $A_{i,t}$ denotes the topic-specific edge count (standardized). $\overline{A}_i$ is manager i's time-series mean (between component, representing structural attention), and $(A_{i,t} - \overline{A}_i)$ is the temporary deviation from mean (within component, representing temporary attention shifts). We use aggregate sentiment measure, constructed as a rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1) over 1 year window, S_t, centered around time-series mean. Sample includes U.S. equity forecasts at asset manager-year level, asset manager FE (μ_i) and year FE (τ_t) included. Robust standard errors clustered at the asset-manager level are reported in parentheses. Sample period: 2015 to 2025.

$$F_{i,t} - F_{i,t-1} = \mu_i + \tau_t + \beta_w(A_{i,t} - \overline{A}_i) + \gamma_w(A_{i,t} - \overline{A}_i) \times S_t + \gamma_b \overline{A}_i \times S_t + \varepsilon_{i,t}$$

	(1) VC	(2) DY	(3) g	(4) π	(5) ED
$A_{i,t} - \overline{A}_i$	0.1885	-0.6019	0.0494	0.0734	-0.5892
	(0.1899)	(0.8080)	(0.1864)	(0.1289)	(0.8949)
$(A_{i,t} - \overline{A}_i) \times S_t$	-27.4131***	55.4955*	12.0692***	-17.5182	59.4755
	(8.2100)	(28.8609)	(1.2149)	(16.5137)	(47.0127)
$\overline{A}_i \times S_t$	-10.8218	15.1918**	-9.4599	24.9757	18.1438***
	(7.4986)	(6.5528)	(6.9585)	(22.6393)	(6.2544)
Asset Manager FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Within R^2	0.174	0.092	0.114	0.098	0.138
R^2	0.590	0.581	0.526	0.525	0.589
Observations	122	122	122	122	122

Table 13: Ex-ante Forecast Errors and Network Complexity. This table presents regression results examining the relationship between CMA return expectations ex-ante forecast errors and causal network complexity measures (described in Section 1.3) using within-between decomposition (described in Section 3.2). The dependent variable is the ex-ante forecast error, defined as $G_t - F_{i,t}$, where G_t is the benchmark forecast (described in Section 2.1.3) and $F_{i,t}$ is annual average CMA return forecast from asset manager i in year t. Each column represents a different network complexity measure: (1) average path length (APL), (2) indirect connection ratio (ICR), (3) network transitivity (NT). $C_{i,t}$ denotes the causal network complexity measure. $\overline{C}_i$ is manager i's time-series mean (between component), and $(C_{i,t} - \overline{C}_i)$ is the temporary deviation (within component). We use an aggregate sentiment measure, constructed as a rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index over 1 year window, S_t, (described in Section 1.2.1), centered around time-series mean. Sample includes U.S. equity forecast deviations at asset manager-year level, asset manager FE (μ_i) and year FE (τ_t) included. Robust standard errors clustered at the asset-manager level are reported in parentheses. Sample period: 2015 to 2025.

$$G_t - F_{i,t} = \mu_i + \tau_t + \beta_w(C_{i,t} - \overline{C}_i) + \gamma_w(C_{i,t} - \overline{C}_i) \times S_t + \gamma_b\overline{C}_i \times S_t + \varepsilon_{i,t}$$

	(1) APL	(2) IRC	(3) NT
$C_{i,t} - \overline{C}_i$	-0.4280	0.4602	0.2520
	(0.4373)	(0.3148)	(0.2433)
$(C_{i,t} - \overline{C}_i) \times S_t$	-14.1797	18.9543	-34.8117*
	(20.1620)	(15.9554)	(17.6208)
$\overline{C}_i \times S_t$	19.1554***	21.6563***	27.5102***
	(6.5659)	(8.0022)	(9.3612)
Asset Manager FE	Y	Y	Y
Year FE	Y	Y	Y
Within R^2	0.179	0.161	0.122
R^2	0.803	0.810	0.806
Observations	178	178	178

Table 14: Ex-ante Forecast Errors and Topic Attention. This table presents regression results examining the relationship between asset manager i's CMA return expectations ($F_{i,t}$) ex-ante forecast errors and topic-specific causal network attention (described in Section 1.3) using within-between decomposition (described in Section 3.2). The dependent variable is the ex-ante forecast error, defined as $G_t - F_{i,t}$, where G_t is the benchmark forecast (described in Section 2.1.3) and $F_{i,t}$ is the annual average CMA return forecast from asset manager i in year t. Each column represents edge counts for a specific topic: (1) valuation change (VC), (2) dividend yield (DY), (3) economic growth (g), (4) inflation (π), (5) economic downturn (ED). $A_{i,t}$ denotes the topic-specific edge count. $\overline{A}_i$ is manager i's time-series mean (between component), and $(A_{i,t} - \overline{A}_i)$ is the temporary deviation (within component). We use an aggregate sentiment measure, constructed as a rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index over 1 year window, S_t, (described in Section 1.2.1), centered around time-series mean. Sample includes U.S. equity forecast deviations at asset manager-year level, asset manager FE (μ_i) and year FE (τ_t) included. Robust standard errors clustered at the asset-manager level are reported in parentheses. Sample period: 2015 to 2025.

$$G_t - F_{i,t} = \mu_i + \tau_t + \beta_w(A_{i,t} - \overline{A}_i) + \gamma_w(A_{i,t} - \overline{A}_i) \times S_t + \gamma_b \overline{A}_i \times S_t + \varepsilon_{i,t}$$

	(1) VC	(2) DY	(3) g	(4) π	(5) ED
$A_{i,t} - \overline{A}_i$	0.0228	0.0322	-0.0059	0.2762	0.5630
	(0.1230)	(0.3477)	(0.1300)	(0.3068)	(0.3850)
$(A_{i,t} - \overline{A}_i) \times S_t$	14.3986**	-19.1431	0.5988	-12.9757	-35.6984*
	(5.7784)	(22.3625)	(5.4543)	(13.7056)	(19.8175)
$\overline{A}_i \times S_t$	12.9546***	-13.3078**	4.4112	-17.5701	-16.2326***
	(4.5272)	(5.5291)	(3.4933)	(10.8060)	(5.8299)
Asset Manager FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Within R^2	0.125	0.141	0.058	0.060	0.200
R^2	0.820	0.796	0.781	0.792	0.813
Observations	178	178	178	178	178

Appendix A

A.1 LDA Textual Data Cleaning and Document-Term Matrix Construction

Latent Dirichlet Allocation (LDA) is a widely used technique for discovering latent themes in textual data by grouping words into topics based on their co-occurrence patterns (Blei et al. 2003). This approach enables a numerical representation of the text and facilitates the extraction of narratives, aligning with the methodology of Bybee et al. (2023).

A.1.1 Text preprocessing

Our preprocessing pipeline includes two stages: general text extraction followed by LDA-specific preprocessing.

The first stage includes the following steps: (i) filter to text element types (paragraphs, headings, etc.); (ii) replace tabs with spaces and remove malformed Unicode characters; (iii) require minimum word count with alphabetic characters, at least one alphabetic character per element, and maximum numeric token ratio; (iv) keep elements containing keywords with adaptive context windows around matches and tag matches as direct or context; and (v) output structured JSON of filtered elements.

The second stage includes the following steps: (i) remove non-alphabetical characters, convert text to lowercase, and tokenize words while excluding single-character tokens; (ii) remove stopwords including domain-specific terms like CMA and estimate; (iii) apply light lemmatization by replacing trailing sses with ss, ies with y, and removing s, ly, ed, and ing endings while excluding words with fewer than three letters; (iv) generate bigrams and trigrams to capture multi-word expressions; (v) exclude terms appearing in fewer than 0.1% of reports; and (vi) construct a document-term matrix with term counts.

A.1.2 Latent Dirichlet Allocation for Narratives

The LDA model assumes that a corpus of text can be represented by L topics, each characterized by a probability distribution over terms (ϕ_l). This approach effectively reduces the dimensionality of complex textual datasets while preserving essential thematic content. We estimate LDA topics using Gibbs sampling (Steyvers & Griffiths 2007), selecting $L = 30$ topics based on the Bayes factor criterion. The model outputs:

- ϕ_l: a V-dimensional probability vector representing the term composition of topic l.
- θ_m: an L-dimensional vector capturing topic attention for each report m.

To measure shifts in topic attention, we compute narrative attention shocks as deviations from a moving average:

$$z_\tau = \theta_\tau - \frac{1}{k}\sum_{j=1}^{k}\theta_{\tau-j}, \tag{A1}$$

where k is the averaging window.

[Figure A1 here]

A limitation of LDA is its inability to provide a natural ordering of topics. Because common terms appear in many topics, the most frequent words dominate certain aspects of LDA outputs. To address this, we rescale the topic-term weights by the inverse of term frequency:

$$\tilde{\phi}_{k,v} = \frac{\hat{\phi}_{k,v}}{f_v}, \tag{A2}$$

where f_v is the frequency of term v across all topics. This transformation emphasizes terms that have unusually large weight in a given topic while suppressing generally frequent terms.

A.2 LLM and Text Analysis

We developed a structured interface to use Anthropic's API for processing CMA reports. Our implementation divides the analysis into distinct subtasks, with each handled through separate API calls.

A.2.1 Text preprocessing

Our text preprocessing pipeline for LLM analysis follows the same general text extraction stage described in Appendix A.1.1 and does not require additional preprocessing steps beyond the initial filtering and cleaning.

A.2.2 LLM Pipeline Architecture

We construct causal networks from financial reports through a three-stage natural language processing pipeline. Our approach extracts explicit causal relationships, classifies agents, and determines effect sentiment to create directed networks representing causal reasoning in asset manager communications.

Causal relationship extraction. We identify explicit causal relationships within report texts using Claude Sonnet 4.5 via batch API processing. The model analyzes each text segment with surrounding context for pronoun resolution and extracts causal triplets: (i) causal agent, (ii) effect, and (iii) affected agent. Each relationship receives a confidence score based on marker clarity and relationship explicitness. The batch processing system adaptively groups text segments into API requests to efficiently handle documents of varying lengths, with prompt caching for repeated instructions.

Topic classification. We assign financial topics and subtopics to each causal and affected agent using a predefined taxonomy covering major financial domains. The classification leverages both sentence-level context (where the agent appears) and broader context to disambiguate agent meanings. Context grouping optimization ensures that agents sharing identical sentence and full text contexts are processed together, avoiding redundant transmission of duplicate contexts. Each agent occurrence is classified independently, receiving a topic assignment with a match score indicating classification confidence. Agents matching no predefined category are classified as "other".

Effect sentiment analysis. We determine the directional impact of each causal relationship by analyzing the effect description as the primary signal. Effects containing terms such as "increase," "improve," "boost," or "enhance" receive positive sentiment; those with "decrease," "reduce," "weaken," or "damage" receive negative sentiment; effects with unclear directionality (e.g., "affect," "influence") or offsetting impacts receive neutral classification. Each sentiment assignment includes a strength score based on language intensity in the effect description, relationship directness, and certainty indicators.

A.2.3 Prompt Design

We use pretrained models without fine-tuning with carefully designed prompts to improve task performance (Gao 2023). Our prompts employ a step-by-step structure informed by chain-of-thought methodology (Wei et al. 2022, Fu et al. 2023). This approach decomposes complex analytical tasks into discrete sequential steps within the prompt, enhancing logical coherence and reducing ambiguity, though overly rigid task decomposition can limit adaptability. Each prompt includes a self-consistency verification component, instructing the model to review its outputs before finalizing responses. This technique improves accuracy with minimal computational overhead (Wang et al. 2022).

We incorporate few-shot examples within each prompt to demonstrate expected output formats and reasoning depth (Brown et al. 2020). This approach allows the model to match fine-tuned performance without parameter updates, while the natural-language format of examples enables rapid adaptation to new tasks (Zhao et al. 2021). Our prompts explicitly direct the model to rely only on the provided text, avoiding external knowledge. This approach reduces hallucination risks and ensures responses are based exclusively on the provided text content, though it may limit the model's ability to resolve ambiguous references or fill contextual gaps. We configure API parameters to encourage deterministic and precise outputs. A temperature setting of 0.2 prioritizes accuracy over response variability. These configurations align with established practices for high-precision analytical tasks (Vaswani et al. 2017).

A.3 CMA Reports Text

We examine the evolution of Capital Market Assumption (CMA) reports by analyzing their textual characteristics, using word count as a proxy for report length and depth. Figure A2 summarizes these patterns. Panel (a) shows substantial variability in CMA report lengths over time, with no clear trend in the median until a sharp increase in 2025. However, the spread of word counts widens in later years, indicating increasing heterogeneity in report length Panel (b) highlights considerable variation in typical report length across asset management firms, pointing to differences in communication strategies.

[Figure A2 here]

Figure A3 visualizes the most frequent titles, reference authors and journals, and years using word clouds extracted from the report text.

[Figure A3 here]

A.4 Sentiment Word Clouds

A.4.1 Aggregated Sentiment Keywords

We applied the LLM framework from Section 1.3 to measure sentiment toward valuation ratios, interest rates, GDP growth, and earnings. We then collect the keywords that drove these ratings, figures A4 through A7 show word clouds for each factor. Larger words appear more frequently across our dataset, highlighting the most frequently used descriptive words.

These visualizations highlight the specific terminology and phrasing used to convey optimistic or pessimistic market signals. Figure A4 shows valuation language with negative terms like ”overvalued” and ”expensive” versus positive descriptors like ”cheap” and ”undervalued.” Figure A5 captures monetary policy signals through negative keywords such as ”tightening” and ”hikes” compared to positive terms like ”accommodative” and ”supportive.” Figure A6 reflects economic conditions with negative words including ”recession” and ”contraction” against positive terms like ”robust” and ”expanding.” Finally, Figure A7 highlights corporate performance language, contrasting negative descriptors such as ”declining” and ”disappointing” with positive terms like ”growing” and ”strong.”

In future iterations, we will build on these insights to better understand the heterogeneity in asset managers' communication strategies when conveying their beliefs to financial practitioners.

[Figures A4, A5, A6, and A7 here]

A.4.2 Event Sentiments

We present three word clouds that visualize the sentiment derived from the extracted events. Figure A8 groups sentiment keywords according to three overall sentiment topics: Positive/Very Positive, Neutral/NaN, and Negative/Very Negative.

[Figure A8 here]

A.5 Empirical Facts Subjective Returns

Figure A9 presents return expectations across Developed Market Equity, Emerging Market Equity, U.S. Investment Grade, U.S. High Yield, U.S. Government Bond, and U.S. Cash.

[Figure A9 here]

A.6 RavenPack Data Processing

RavenPack provides comprehensive structured metadata that includes sentiment scores, event classifications, and relevance indicators for news stories across multiple financial and economic domains. We employ two primary sentiment indicators in our analysis. The Event Sentiment Score (ESS) rates each firm-level news event on a -1 to +1 scale based on experts with extensive experience and backgrounds in linguistics, finance, and economics. The Composite Sentiment Score (CSS) indicates how the market responds to news articles. The CSS variable is estimated based on stock price reactions, which are empirically modeled using intraday data from a portfolio of approximately one hundred large-cap stocks (RavenPack 2020, Dang et al. 2015).

We categorize news into various topics using alternative classifications based on the variables *GROUP* and *TYPE*, allowing for different levels of granularity in topic analysis. To mitigate the risk of overcounting due to repeated coverage of the same events by multiple news outlets, we introduce a mandatory one-day gap to classify stories as distinct new events (i.e., setting EVENT SIMILARITY DAYS ≥ 1). To ensure the selected news has substantial implications for the companies under study, we apply two relevance thresholds simultaneously: entity relevance must exceed 75, and event relevance must surpass 90. Additionally, we exclude social media platforms and low-quality news sources to preserve data integrity and emphasize professionally curated content. These filtering criteria follow established practices documented in the literature and RavenPack's user guidelines (RavenPack 2020, Dang et al. 2015, Gao et al. 2017, Hafez 2009).

In constructing news horizons aligned with the CMA report date, we aggregate news based on their publication dates at the daily frequency, following equations (1) or (2). Subsequently, using the resulting daily time series, we calculate average news measures over specified time intervals. We focus on news about S&P 500 companies included in CRSP monthly market cap data. We first aggregate news at the parent company level, then perform fuzzy matching between company names in RavenPack and CRSP.

A.7 Additional Details on the Building-Block Decomposition

Asset managers' return forecasts employ a "building-block" decomposition that breaks expected equity returns into fundamental drivers. This approach parallels the logic of the dividend discount model and has been widely adopted by both practitioners and academics. In its simplest form, the expected long-term equity return can be approximated as the sum of expected real growth, expected inflation, and the initial income yield. Formally, one can write:

$$E(R) \approx R_g + \pi + Y_{\text{div}} \ ,$$

where R_g denotes real earnings (or economic) growth, π is expected inflation, and Y_{div} is the dividend yield. To capture additional drivers of returns, more elaborate decompositions include terms for share repurchases, net issuance, profit-margin mean reversion, and valuation changes. Accordingly, we define an aggregate net distribution yield as

$$Y_{\text{net}} = Y_{\text{div}} + Y_{\text{buyback}} - \text{Issuance} \ ,$$

which augments the traditional dividend yield with buyback payouts and subtracts the dilution from net share issuance. Similarly, if current profit margins are elevated relative to historical norms, one can include a term ΔM to represent expected margin reversion (a downward adjustment to earnings growth to reflect the assumption that unusually high margins will partially revert toward their mean). A more comprehensive additive model for expected returns is:

$$E(R) = \underbrace{(R_g + \pi)}_{\text{Growth and Inflation}} + \underbrace{\Delta M}_{\text{Margin Reversion}} + \underbrace{Y_{\text{net}}}_{\text{Net Yield}} + \underbrace{V}_{\text{Valuation Change}} . \tag{A3}$$

Each component can be adjusted or grouped depending on the forecaster's preferences. For example, some forecasters subsume ΔM into an adjusted growth term (treating $R_g + \pi + \Delta M$ as a single "nominal growth" projection), while others combine dividends and buybacks into one total yield. In all cases, the building-block decomposition provides a structured framework that ensures all key drivers of long-run equity returns are considered. It is worth noting that one could also express the decomposition in a multiplicative form to account for compounding interactions among factors.

Institution-specific building blocks. Figure A10 summarizes the components included by each asset manager in their CMA equity-return decompositions. Each colored bar denotes one of six building blocks—nominal growth, dividend yield, valuation change, buybacks, margin adjustments, and share issuance.

[Figure A10 here]

A.8 Causal Networks

Each Capital Market Assumption report generates a directed graph $G = (V, E)$ where nodes $v \in V$ represent economic topics and directed edges $(u, v) \in E$ represent causal relationships from topic u to topic v extracted via the LLM pipeline described in Section 1.3.

A.8.1 Network Visualizations

Figure A11 presents complete causal networks for four representative asset managers, constructed by aggregating all available CMA reports for each institution across our sample period.

[Figure A11 here]

Table A3 reports the average number of edges and nodes per report for each asset manager in our dataset.

[Table A3 here]

A.8.2 Network Measures

Average Path Length: Mean shortest path length between node pairs in the largest connected component:

$$\text{Avg. Path Length} = \frac{1}{|P|} \sum_{(i,j) \in P} d(i, j) \tag{A4}$$

where $d(i, j)$ is the shortest path length from node i to node j, and P is the set of node pairs in the largest connected component. The strong component version considers directed paths, while the weak component version considers undirected paths.

Indirect Connection Ratio: Proportion of edges involving focal topics that participate in multi-step causal chains to U.S. equity return:

$$\text{Indirect Ratio} = \frac{|\mathcal{E}indirect|}{|\mathcal{E}total|} \tag{A5}$$

where $\mathcal{E}indirect$ is the set of edges involving the focal topics that are part of indirect causal paths leading to (or from) U.S. equity return, and $\mathcal{E}total$ is the total number of edges in the network. This measure captures the extent to which these topics affect U.S. equity through multi-step causal chains rather than direct connections.

Average Eigenvector Centrality: Mean eigenvector centrality across all nodes:

$$\text{Avg. Eigenvector Centrality} = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{A6}$$

where x_i is the eigenvector centrality of node i, which assigns higher scores to nodes connected to other high-scoring nodes.

Network Transitivity: Global clustering coefficient measuring the tendency for triangles to form:

$$\text{Transitivity} = \frac{3 \times \text{number of triangles}}{\text{number of connected triples}} \tag{A7}$$

where a triangle consists of three nodes with edges $(A \to B, B \to C, A \to C)$, and a connected triple has at least two edges $(A \to B, B \to C)$. Values range from 0 (no transitive closures) to 1 (complete transitivity).

Figure A1: Evaluation Metrics for LDA Topic Selection: Bayes Factor, Perplexity, and Coherence. The figure plots three metrics used to assess the quality and interpretability of LDA models estimated on CMA reports. The *Bayes factor* (left panel) identifies the specification that best explains the data relative to a baseline, *perplexity* (middle panel) measures out-of-sample predictive performance (lower values indicate better fit), and *coherence* (right panel) reflects how interpretable or semantically consistent the extracted topics are. We evaluate LDA models with the number of topics ranging from 10 to 250 in increments of 10. Both Bayes-factor and cross-validation criteria suggest that $L \approx 30$–40 yields the best overall balance among these metrics. Sample: CMA reports, 2008–2025.

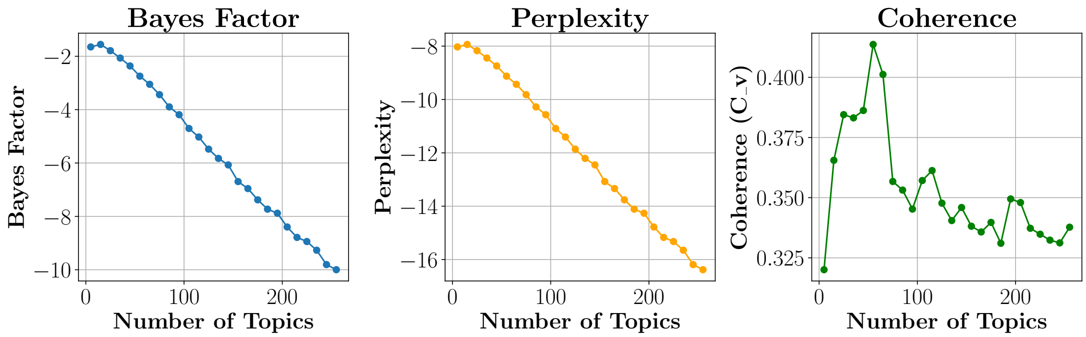

Figure A2: Evolution of CMA Report Textual Characteristics. This figure shows the evolution of CMA report lengths measured by word count. Panel (a) displays the temporal evolution of CMA report lengths measured by word count from 2015 to 2025. Panel (b) shows the histogram of word count distribution across all CMA reports in the sample. Sample: CMA reports, 2015–2025.

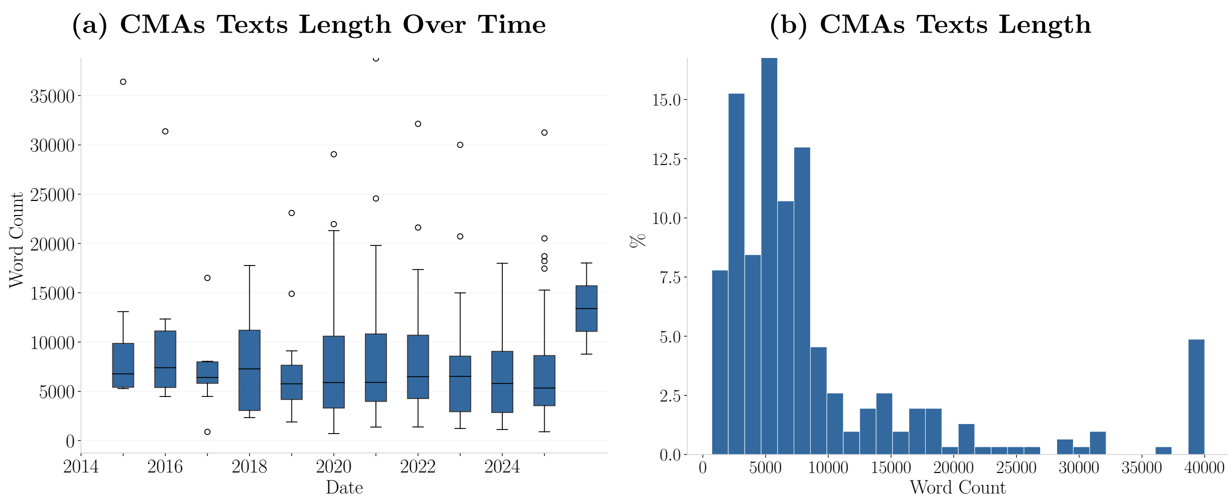

Figure A3: References in CMA Reports. This figure shows word clouds showing frequently occurring *titles* (a), *authors* (b), *journals* (c), and *years* (d) among the extracted references. Larger words indicate higher frequency. Sample: CMA reports, 2008–2025.

(a) Titles

(b) Authors

(c) Journals

(d) Years

Figure A4: Valuation Ratios Sentiment Keywords. This figure shows word clouds depicting the most frequently extracted keywords for negative (a) and positive (b) sentiment in the valuation ratios category. Larger words indicate higher observed frequency across multiple simulation runs and analyst reports. Sample: CMA reports, 2008–2025.

(a) Negative Sentiment **(b)** Positive Sentiment

Figure A5: Interest Rates Sentiment Keywords. This figure shows word clouds depicting the most frequently extracted keywords for negative (a) and positive (b) sentiment in the interest rates topic. Larger words indicate higher observed frequency across multiple simulation runs and analyst reports. Sample: CMA reports, 2008–2025.

(a) Negative Sentiment **(b)** Positive Sentiment

Figure A6: GDP Growth Sentiment Keywords. This figure shows word clouds depicting the most frequently extracted keywords for negative (a) and positive (b) sentiment in the GDP growth topic. Larger words indicate higher observed frequency across multiple simulation runs and analyst reports. Sample: CMA reports, 2008–2025.

(a) Negative Sentiment **(b)** Positive Sentiment

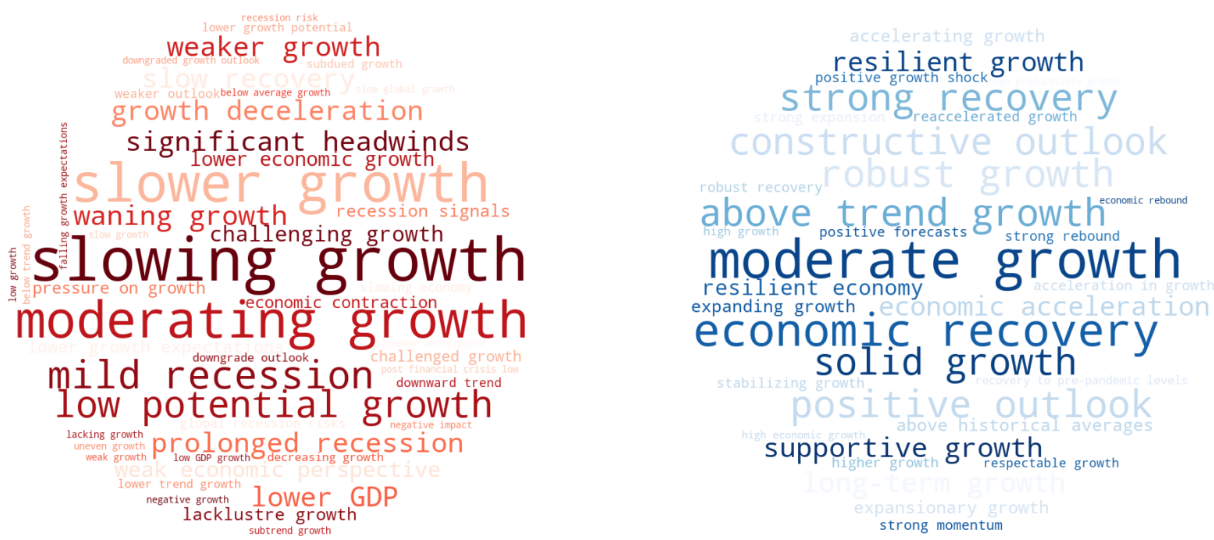

Figure A7: Earnings Sentiment Keywords. This figure shows word clouds depicting the most frequently extracted keywords for negative (a) and positive (b) sentiment in the earnings topic. Larger words indicate higher observed frequency across multiple simulation runs and analyst reports. Sample: CMA reports, 2008–2025.

(a) Negative Sentiment **(b)** Positive Sentiment

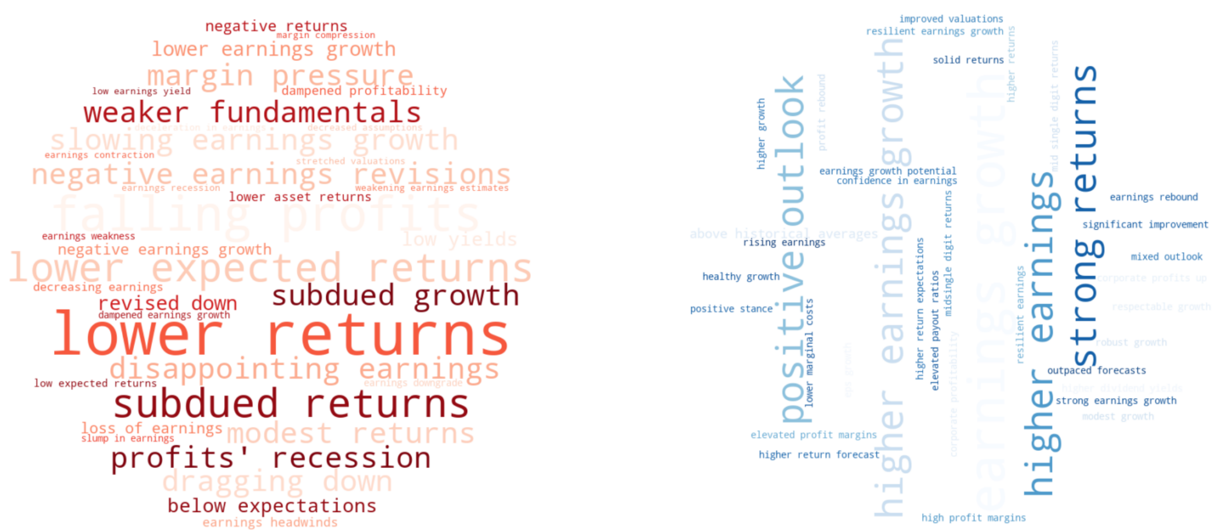

Figure A8: Word clouds of Event sentiment keywords. This figure shows word clouds of the most frequently extracted keywords for each sentiment category: (a) Positive, (b) Neutral, and (c) Negative. Larger words indicate more frequent occurrences across CMA reports. Sample: CMA reports, 2008–2025.

(a) Positive & Very Positive.

(b) Neutral or NaN.

(c) Negative & Very Negative.

Figure A9: Subjective Return Expectations Across Asset Classes. This figure shows subjective long-term return forecasts across report release dates for six asset classes: (a) Developed Market Equity, (b) Emerging Market Equity, (c) U.S. Investment Grade, (d) U.S. High Yield, (e) U.S. Government Bond, and (f) U.S. Cash. Each observation represents a single asset manager-report pair, with forecast horizons ranging from 3 to 30 years. Sample: CMA reports, 2015–2025.

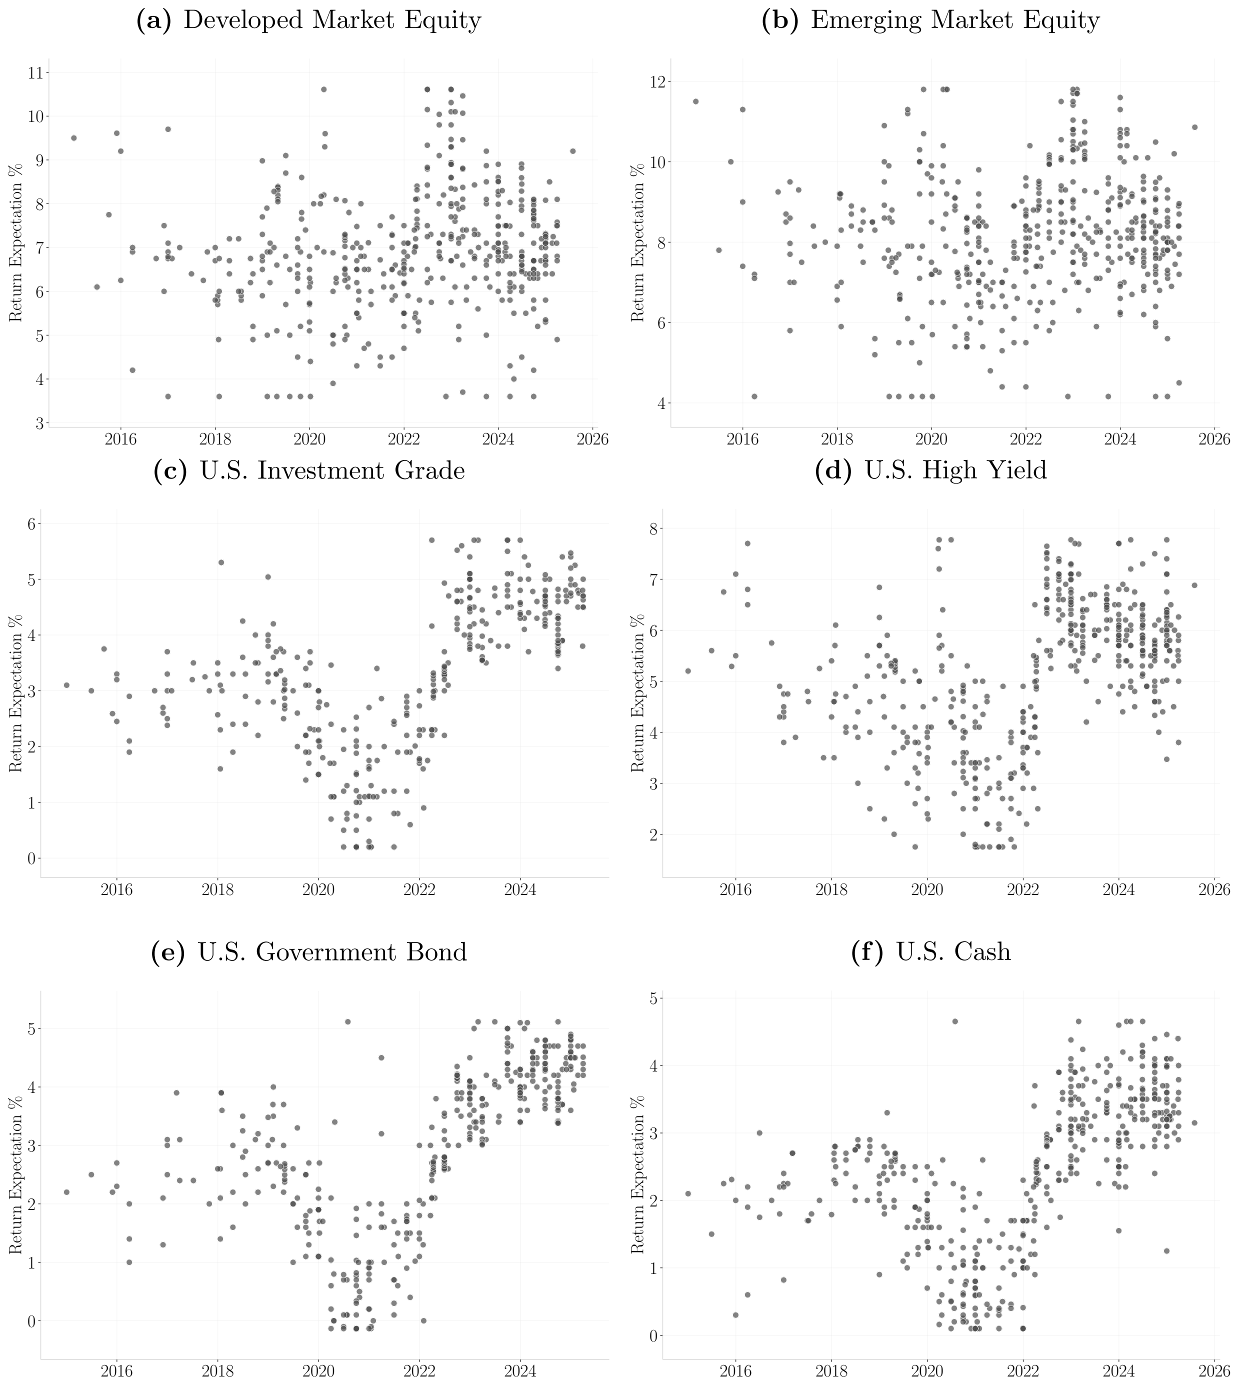

Figure A10: Cross-Sectional Variation in Return Decompositions. This figure shows building blocks (described in Section 2.1.1) employed by 28 asset managers in their U.S. equity return decompositions. Each horizontal bar represents one asset manager's framework, with colored segments indicating the inclusion of six components: nominal growth, dividend yield, valuation change, buybacks, margin adjustments, and share issuance. Sample: CMA reports, 2010–2025.

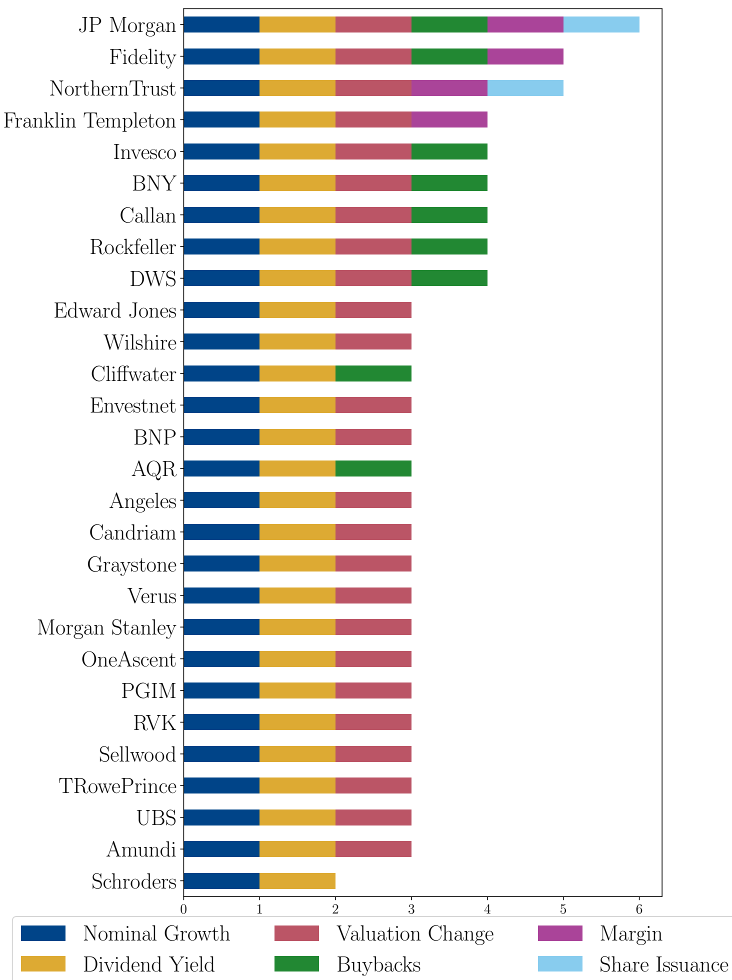

Figure A11: Asset Manager Causal Networks of U.S. Equity Return Expectations. This figure shows directed networks of causal relationships (described in Section 1.3) extracted from CMA reports. Each panel displays the complete causal network for a different asset manager: (a) Robeco, (b) Morgan Stanley, (c) T. Rowe Price, (d) Vanguard. Nodes represent topics identified in the text, with size reflecting centrality (number of connections). Edges indicate the direction of causality, with colors showing positive (green), negative (red), or neutral (gray) causal effects as stated in the reports. U.S. Equity Return and the four building blocks from the return decomposition (described in Section 2.1.1)—growth, valuation change, dividend yield, and inflation—are highlighted with distinct colors. Sample: asset manager causal networks, 2015–2025.

(a) Robeco

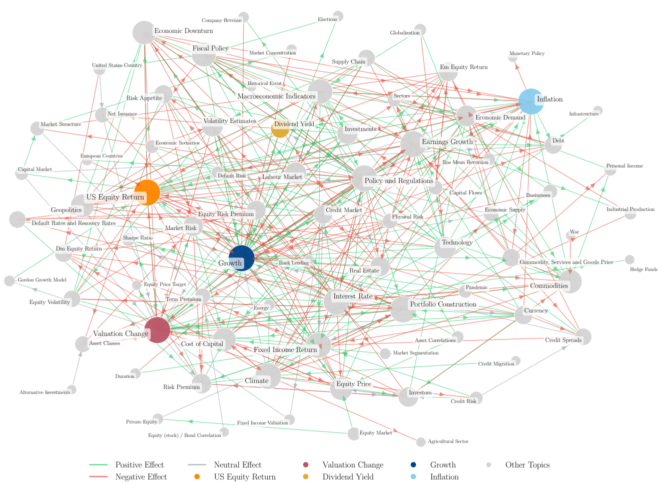

(b) Morgan Stanley

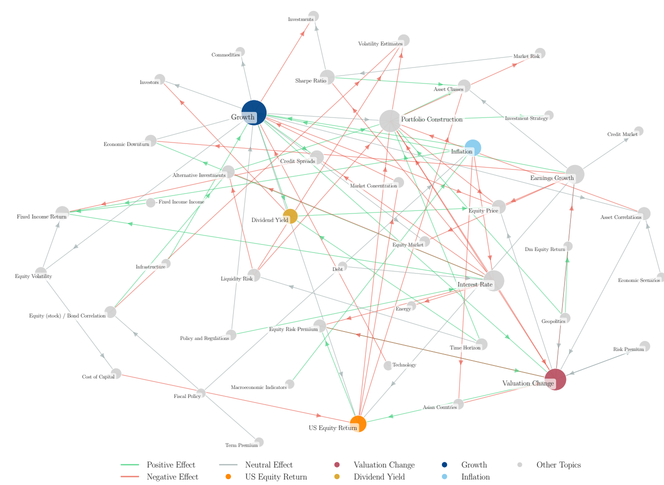

(c) T. Rowe Price

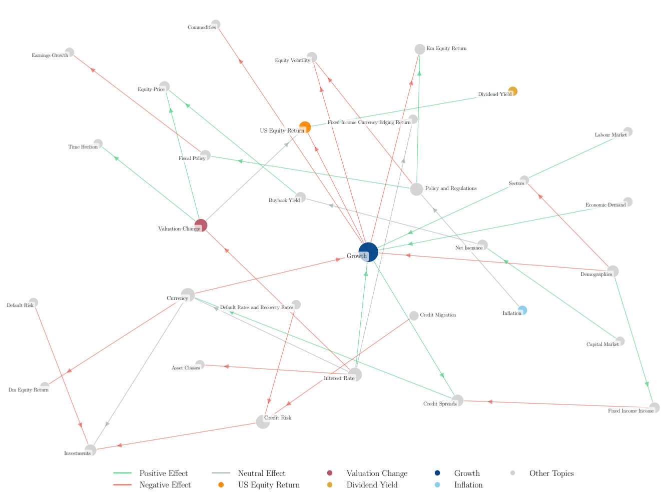

(d) Vanguard

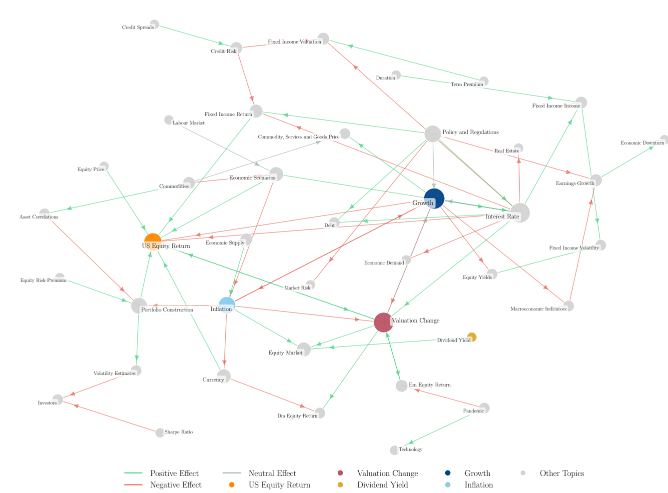

Table A1: Forecast Changes and News Sentiment. This table reports regression results examining the relationship between year-over-year forecast changes and the aggregate earnings-news sentiment metric, controlling for trailing market returns. The dependent variable in each column is the annual forecast change, defined as $Y_{i,t} - Y_{i,t-1}$, where $Y_{i,t}$ denotes either the annual average return forecast or a return-forecast building block from asset manager i in year t. Columns (1) and (2) present results for overall U.S. equity return forecasts (ER) for the full sample and for the subsample restricted to years in which all building blocks—nominal growth (g), dividend yield (DY), and valuation change (VC)—are available. Columns (3)–(5) report results from decomposing equity return forecasts into these building blocks. The key independent variable is an aggregate sentiment measure, S_t, constructed as a one-year rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1). We control for trailing 12-month S&P 500 returns, $R_{t-1,t}$, computed as the geometric compounded return over the prior year. All regressions include asset manager fixed effects (μ_i). Robust standard errors, clustered by asset manager, are reported in parentheses. Sample period: 2010 to 2025.

$$Y_{i,t} - Y_{i,t-1} = \mu_i + \beta\, S_t + \gamma\, R_{t-1,t} + \varepsilon_{i,t}$$

$\Delta Y_{i,t} = Y_{i,t} - Y_{i,t-1}$	(1) ΔER	(2) ΔER	(3) Δg	(4) ΔDY	(5) ΔVC
Avg. Earnings Sentiment (1y)	4.1524**	6.6295***	-2.1300	0.2409	9.5292**
	(1.8722)	(1.9969)	(2.9072)	(0.4783)	(4.4240)
$R_{t-1,t}$	-3.9239***	-4.0368***	0.8959***	-1.1279***	-4.3419***
	(0.6126)	(1.1327)	(0.3356)	(0.2450)	(0.8118)
Asset Manager FE	Y	Y	Y	Y	Y
R^2	0.468	0.503	0.242	0.457	0.568
Observations	175	81	81	81	67

Table A2: Ex-ante Forecast Errors and News Sentiment. This table presents regression results examining the relationship between CMA return expectations ex-ante forecast errors and earnings-news sentiment metrics, controlling for trailing market returns. The dependent variable is the ex-ante forecast error, defined as $G_t - F_{i,t}$, where G_t is the benchmark forecast (described in Section 2.1.3) and $F_{i,t}$ is the annual average return forecast from asset manager i in year t. We use an aggregate sentiment measure, S_t, constructed as a one-year rolling average of the value-weighted S&P 500 RavenPack earnings-news sentiment index (described in Section 1.2.1). We control for trailing 12-month S&P 500 returns, $R_{t-12,t}$, computed as the geometric compounded return over the prior year. The regression includes asset manager fixed effects. Robust standard errors, clustered by asset manager, are reported in parentheses. Sample period: 2010 to 2025.

$$G_t - F_{i,t} = \mu_i + \beta\, S_t + \gamma\, R_{t-12,t} + \varepsilon_{i,t}$$

	$G_t - F_{i,t}$
Avg. Earnings Sentiment (1y)	-4.9604***
	(1.3712)
$R_{t-12,t}$	0.2403
	(0.4251)
Asset Manager FE	Y
R^2	0.445
Observations	394

Table A3: Causal Network Metrics by Asset Manager. This table presents average number of nodes and edges per report for each asset manager in our Capital Market Assumptions causal networks dataset (described in Section 1.3). Edges represent the total number of causal relationships identified, and Nodes represent the number of distinct topics appearing in the causal network. Sample: CMA causal networks, 2008–2025.

Asset Manager	Edges	Nodes	Asset Manager	Edges	Nodes
ACG	3.0	5.0	Janney	5.5	6.0
AIMCo	32.0	23.3	Janus Henderson	8.5	8.5
AON	21.5	15.5	MFS	15.7	17.3
AQR	7.5	7.7	Mackenzie	23.7	17.7
Advent	37.0	18.0	Manulife	13.0	12.0
Allianz	1.8	3.1	Meketa	9.0	6.4
American Century	6.0	8.0	Merrill Lynch	11.0	10.0
Amundi	21.1	15.3	Morgan Stanley	51.0	29.8
Angeles	12.8	14.2	NEPC	3.7	3.8
BNP	9.2	9.4	Neuberger	6.5	9.0
BNY Mellon	16.3	15.4	Ninety One	22.4	18.8
Baillie Gifford	8.0	5.0	Northern Trust	20.1	16.6
Barclays	7.0	8.0	Nuveen	1.0	2.0
CIBC	16.3	15.7	OneAscent	14.0	13.0
CWO	4.5	6.5	PGIM	16.2	11.2
Callan	17.7	14.8	PIMCO	18.5	20.5
Candriam	9.0	12.0	Pinebridge	10.2	14.0
Capital Group	20.7	17.7	RAM	46.2	28.5
Cliffwater	33.5	24.5	RBC	9.3	6.3
Cohen Steers	20.7	19.3	RVK	4.7	4.0
Columbia Thread.	4.2	4.5	Robeco	98.7	42.7
Credit Suisse	47.0	28.0	Rockefeller	57.5	40.5
Crescent Capital	14.5	14.8	SEI	4.0	5.0
DSA	8.0	12.0	SJS	9.0	9.5
DWS	17.8	14.8	Schroders	7.1	6.4
Deutsche Bank	8.0	11.0	Sellwood	4.6	5.2
EGF	15.0	11.0	State Street	7.9	9.4
Edward Jones	4.2	6.2	Syntrinsic	32.6	22.5
Ellwood	2.0	2.0	T. Rowe Price	12.7	13.7
Envestnet	2.7	3.7	TD	3.0	4.3
FI3	17.0	17.0	UBS	20.4	19.0
Fidelity	21.0	16.0	US Bank	14.5	14.0
Fiducient	15.6	15.6	Vanguard	34.0	23.0
Franklin Temp.	22.7	21.4	Verus	13.3	11.3
Graystone	55.0	30.5	Voya	9.8	10.2
HSBC	13.0	16.0	Wealthspire	13.0	11.0
Invesco	11.9	12.1	Wellington	9.3	8.0
JP Morgan	90.5	39.5	Wilshire	28.0	19.5

Table A4: Cross-Sectional Return-Volatility Relationship. This table presents annual cross-sectional regressions of managers' U.S. equity return expectations on their volatility forecasts. Each row reports the yearly estimates of the intercept α_t and slope coefficient β_t with standard errors in parentheses. Sample: CMA reports, 2018-2025.

$$R_{i,t} = \alpha_t + \beta_t \sigma_{i,t} + \varepsilon_{i,t}$$

Year	α_t	β_t	R^2	Observations
2018	6.965*	-0.036	0.003	11
	(3.310)	(0.206)		
2019	6.866***	-0.056	0.008	21
	(2.292)	(0.147)		
2020	5.428**	0.021	0.001	30
	(2.476)	(0.154)		
2021	0.946	0.287**	0.155	30
	(2.066)	(0.127)		
2022	5.687***	0.066	0.016	40
	(1.310)	(0.082)		
2023	3.796*	0.185	0.047	42
	(2.133)	(0.132)		
2024	3.772	0.144	0.016	59
	(2.401)	(0.149)		
2025	11.480	-0.291	0.066	10
	(6.233)	(0.388)		