The Macro Alibi: Subjective Risk Attribution in Analyst Scenarios

Chen Wang Kangying Zhou*

May 31, 2026

[CLICK HERE FOR THE LATEST VERSION]

Abstract

Sell-side analysts over-attribute stocks' downside scenarios to macroeconomic forces. We document this *Macro Bear Bias* in scenario-based valuation reports: conditional on the same market state, bear-case narratives emphasize aggregate macro risk far more than base- or bull-case narratives, even though realized CAPM R^2 is similar across realized bear, base, and bull outcomes. The bear–base macro-attention gap predicts systematic pessimism in subsequent base-case forecasts, and portfolios formed on a nonlinear bias-adjusted signal earn monthly CAPM alphas of up to 1.9%. Mechanism tests favor a cognitive availability-heuristic explanation—analysts anchor downside narratives on salient macro-crisis templates—over a strategic career-concerns explanation. Narrative templates analysts use to rationalize risk can distort analysts' numerical forecasts and asset prices.

Keywords: macro bear bias, narrative finance, downside risk, analyst beliefs, large language model
JEL: G11, G12, G14, G23

*Chen Wang is at University of Notre Dame, chen.wang@nd.edu. Kangying Zhou is at Texas A&M University, kangying.zhou@tamu.edu. We thank Jinchen Yang for excellent research assistance. We are grateful for helpful comments from Nicholas Barberis, Gerard Hoberg, Theis Jensen, Lawrence Jin, Lorena Keller, Bryan Kelly, Marco Rossi, Davide Tomio, and Wei Wu.

1 Introduction

Risk in financial markets is not only a quantity, but also a story. When sell-side equity analysts publish a price target, they do not simply report a number; they explain it, naming the forces they expect to move cash flows up or down, the channels through which uncertainty operates, and the contingencies that would change their view. These narrative attributions matter. They are read by portfolio managers, cited in investment committees, and embedded—through analysts' own forecasts—into the prices that shape and aggregate market beliefs. Whether the stories that professional information intermediaries tell about risk correctly identify the *sources* of risk is therefore a first-order question for understanding belief formation in financial markets.

One particularly visible expression of this question arises in how equity analysts describe firms' upside and downside states. Bull cases tend to emphasize firm-specific opportunities, such as product cycles, margin expansion, and competitive position. Bear cases, by contrast, routinely emphasize aggregate forces, such as recessions, financial contagion, and monetary tightening. A natural interpretation is that this asymmetry reflects the true structure of risk. Conditional asset pricing argues that systematic risk is elevated in bad states: downside correlations rise and conditional betas widen (Ang and Chen, 2002; Ang, Chen, and Xing, 2006; Lettau, Maggiori, and Weber, 2014). Under this view, an analyst who leans on macro forces in the bear cases may simply be tracking where systematic risk is most concentrated. A skeptic who doubts that analysts track risk so precisely might still expect their macro emphasis to be conditional on another time-varying signal. In canonical behavioral accounts, vivid crises should make macro risks more salient, while their influence should fade as those crises recede (Shiller, 2019; Bordalo, Gennaioli, and Shleifer, 2020; Bordalo et al., 2023). Although these rational and behavioral readings differ, they share a common premise: bear-case macro emphasis should vary with an identifiable state variable.

Three findings, taken together, call this premise into question. First, macro emphasis in base, bear, and bull scenarios is essentially flat across market states. It neither rises in weak markets nor recedes in strong ones, even though the true conditional importance of systematic risk does. This pattern is difficult to reconcile with a rational benchmark: analysts do not appear to track where systematic risk concentrates, because their emphasis is insensitive to the market state on which that concentration depends. Second, in the time series, bear-case macro attention rises only mildly after vivid crises and does not unwind as those crises recede. The direction is consistent with the canonical memory account, but the magnitude is not: salience shifts the emphasis at the margin rather than serving as the first-order force. Once both conditional signals are accounted for, a third fact remains: bear narratives consistently assign a higher macro share than base or bull narratives. This gap is wide, stable, and present across all market states. Belief formation about the downside therefore appears to follow a default script rather than a state-specific calculation. The distortion is behavioral in origin, in the sense that it reflects a story analysts reach for whenever a downside case is required, but it is near-unconditional in form and far stickier than canonical memory models allow.

To study this distortion, we exploit the unique institutional setting of the Morgan Stanley Risk-Reward Framework reports from 2007 to 2025, comprising 63,191 firm-quarter analyst scenario-based reports. The Risk-Reward Framework is unusual in requiring analysts to articulate three explicit scenarios—bear, base, and bull—for the same firm at the same point in time, each accompanied by a structured narrative explaining the conditions under which that scenario obtains. This within-report, across-scenario design holds the analyst, the firm, the calendar date, and the prevailing macro environment fixed, so the primary variation we exploit is the directional state of the world being described. We use large language models to extract, from each scenario narrative, the share of analyst attention allocated to macroeconomic, industry-level, and firm-specific drivers. We then build ex post counterparts for the same reports: we classify which scenario is ultimately realized—bear if the stock subsequently falls below the analyst's bear target, base if it lands between the bear and bull targets, and bull if it exceeds the bull target—and compute realized CAPM R^2 within the corresponding market-state-by-realized-scenario subsample. Under the rational benchmark, the ex ante macro attention an analyst assigns to a given scenario should track the ex post realized macro content of that scenario within the same conditioning set. A persistent disconnect between these two objects is our measure of narrative misattribution.

The data reveal a sharp and systematic asymmetry. Bear-case narratives devote significantly more (6.3 percentage points) attention to macroeconomic forces than base-case narratives—a within-report gap that absorbs all analyst-, firm-, and time-level heterogeneity by construction. The corresponding bull–base gap is essentially zero, an order of magnitude smaller. The bear-narrative tilt is also remarkably stable across market states: it does not collapse in calm years and is not amplified in crises. The ex post benchmark behaves very differently. Following Roll (1988), who established CAPM R^2 as the canonical measure of how much of a firm's return variation is attributable to market-wide forces, we use realized R^2 as the natural quantitative analog to the share of analyst attention devoted to macro forces. Conditional on the same forward market state, the bear–base difference in realized CAPM R^2 is close to zero in weak market states and turns negative in good market states. Realized downside outcomes are therefore not disproportionately explained by market-wide variation; bear narratives describe them as if they were. We refer to this dissociation between ex ante narrative attribution and ex post realized variation as *Macro Bear Bias*: analysts exhibit a tendency to over-attribute downside outcomes to aggregate macroeconomic risks in bear-case narratives, even when the realized data do not support the attribution. In this sense, macro risk can serve as a "macro alibi" for the bear case.

The bias has economic consequences that extend well beyond the bear narrative itself. Scenario narratives serve as the cognitive scaffolding for the analyst's valuation judgment, and a macro-heavy bear case spills over into the base-case forecast on which prices are anchored. The evidence is consistent with a discount-rate channel: a bear narrative dominated by systematic risk leads the analyst to perceive elevated aggregate exposure for the firm, which inflates the implicit hurdle rate applied to cash flows and depresses the base-case price target. The distortion scales multiplicatively with the firm's true systematic exposure, so high-beta firms should display larger forecast-error responses to a given bear–base macro-attention gap than low-beta firms—a cross-sectional prediction the data bear out. Empirically, a one-standard-deviation increase in the bear–base macro-attention gap predicts a positive base-case forecast error on the order of 1–2%: analysts whose bear narratives lean disproportionately on macro forces are systematically too pessimistic about the modal outcome.

The forecast-error pattern is strongly nonlinear, and it translates directly into prices.[1] Sorting on analysts' raw base-case forecasts generates no abnormal performance. Adjusting those forecasts for the predictable bias component using a flexible approximation to the nonlinear conditional bias produces a value-weighted long–short portfolio with an average monthly return of 2.02%, a CAPM alpha of 1.91%, and an annualized Sharpe ratio of 1.10.[2] The narrative asymmetry, in other words, is not a stylistic feature of analyst writing; it carries first-order information about subsequent forecast errors and cross-sectional returns.

We next study the mechanism behind Macro Bear Bias. What cognitive process produces a bear-case template applied so independently of the market state? A purely rational benchmark would predict bear-case macro emphasis that tracks the conditional importance of systematic risk; a canonical memory-based account would predict emphasis that rises and falls sharply with the salience of recent crises. The data are partly consistent with the latter: bear-case macro attention does tick up in the aftermath of vivid crisis episodes such as the Global Financial Crisis (GFC) and the COVID-19 crash. But this salience-driven movement is modest in magnitude and is layered on top of a much stickier baseline that does not unwind as the originating salience fades. The mechanism we identify is therefore best understood as a refinement of standard memory-based accounts rather than an alternative to them.

Its central ingredient is the availability heuristic: analysts anchor downside narratives on a small number of vivid macro-crisis episodes that make macro explanations cognitively accessible whenever a bear narrative is required. But the heuristic operates in three ways that distinguish it from pure individual recall: crisis templates persist long after the news cycle that originally surfaced them, are encoded institutionally rather than only at the analyst level, and diffuse through peer information networks. Together, these features explain why the salience-driven response we do

[1]This is consistent with the broader theoretical observation that belief-driven asset price dynamics are generically nonlinear and state-dependent. Dew-Becker, Giglio, and Molavi (2025) show that under Bayesian filtering—including under behavioral departures from full rationality—the response of subjective beliefs to signals is governed by the higher-order cumulants of agents' posterior distributions, so that the mapping between subjective belief objects and realized outcomes is nonlinear by construction.

[2]The equal-weighted counterparts are 1.80%, 1.74%, and 0.99

detect remains modest relative to the persistent baseline—and why the template is, on net, applied near-unconditionally across market environments.

The data support four implications of this account. First, if templates respond to but outlive salience, increases in Wall Street Journal attention to macro topics should raise bear-case attention in subsequent months and persist after that coverage subsides. They do, and the response is asymmetric: only negative-sentiment macro news activates the templates, and crisis topics linger in bear narratives long after the news cycle has moved on. Second, if templates are deployed where they are most cognitively accessible, high-systematic-risk industries should display larger bear–base macro gaps after the GFC, when a financial-contagion template was most salient. The post-2008 widening matches this prediction. Third, if templates are institutionally encoded rather than only held in individual memory, analysts who worked at Morgan Stanley during the GFC should exhibit persistently larger gaps than otherwise comparable analysts, while industry-wide crisis tenure should be less informative. Both halves of the prediction hold in the data: the Morgan Stanley–specific tenure effect is significant and persistent, whereas broader profession-wide crisis tenure carries little explanatory power. Fourth, if templates diffuse through shared information environments, an analyst's own gap should co-move with the leave-one-out average among peers covering firms in the same sector and quarter. The peer-group co-movement is positive and statistically significant.

A second feature of the mechanism concerns when, in the cross-section of reports, Macro Bear Bias is most actively realized. Analysts constructing a bear case need a coherent explanation for why the stock could underperform, and the explanation they reach for depends on what the rest of the report already supplies. When the Morgan Stanley base-case target is below the IBES consensus, the base case already embeds firm-specific pessimism, the bear narrative has a coherent firm-level downside to draw on, and the macro template is suppressed: the bear–base gap is significantly smaller. When the base case is above consensus and no firm-specific downside is supplied, analysts reach for the macro template to make the bear case coherent. The macro alibi, in this view, fills the explanatory gap where firm-specific downside narratives are unavailable. A strategic career-concerns alternative—in which macro narratives function as reputationally safer deflections from sharp firm-specific calls—finds little support: junior analysts do not exhibit stronger bias, recent Morgan Stanley investment-banking relationships do not amplify it, and standard analyst-level incentive proxies explain little of its variation. Overall, the evidence is more consistent with a cognitive-template account than with a strategic-deflection account.

Related literature. Our paper first contributes to the literature on subjective risk perceptions. A central insight of this literature is that perceived risk need not coincide with statistical risk: investors' beliefs about uncertainty are shaped by experience, memory, attention, and extrapolation, and these subjective perceptions matter for prices and portfolio choices (Malmendier and Nagel, 2011; Guiso, Sapienza, and Zingales, 2018; Greenwood and Shleifer, 2014; Giglio et al., 2021; Lochstoer and Muir, 2022; Jensen, 2022; Bohren et al., 2024). We document a new form of distorted subjective risk perception: analysts systematically misattribute the source of downside risk, over-weighting macroeconomic forces even when realized downside outcomes are not disproportionately explained by market-wide variation. In contemporaneous work, Chen, Li, and Yang (2025) use scenario-based analyst reports to study the higher moments of subjective return distributions. The two papers are complementary: they characterize the shape of the subjective return distribution analysts assign to each firm, while we study how analysts attribute the *source* of risk across scenarios.

A second strand studies downside risk and state-dependent systematic risk. Ang and Chen (2002) and Ang, Chen, and Xing (2006) document and price downside risk in U.S. equities; Lettau, Maggiori, and Weber (2014) show that a conditional CAPM with higher betas in bad states helps explain currency and equity returns; Kacperczyk, Van Nieuwerburgh, and Veldkamp (2014); Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016) show that fund managers reallocate attention between aggregate and firm-specific information over the business cycle, devoting more attention to macro conditions in recessions.[3] We contribute by studying not realized downside comovement itself but analysts' subjective attribution of downside risk. Analysts appear to overgeneralize a

[3] Lewellen and Nagel (2006) and Lewellen, Nagel, and Shanken (2010) caution that time variation in betas is unlikely to fully resolve standard asset-pricing anomalies and that conventional cross-sectional tests have limited power to distinguish among conditional factor models. Our argument does not rest on conditional asset pricing fully accounting for the realized data; we show that analysts' bear-case macro emphasis is largely unresponsive even to the state-dependent variation that conditional asset pricing does identify.

valid empirical regularity: while systematic risk can be more important in bad states, they apply macro-heavy bear narratives near-unconditionally, even when conditional realized R^2 does not justify the emphasis. Macro Bear Bias is therefore not evidence against state-dependent downside risk; it is evidence that analysts fail to condition their risk narratives on the prevailing market state.

Methodologically, our focus on a nonlinear, state-dependent mapping from analyst narratives to forecast errors connects to two complementary literatures. On the theoretical side, Dew-Becker, Giglio, and Molavi (2025) show that under (Bayesian and behavioral) learning belief dynamics are generically nonlinear, so that linear specifications should be expected to understate the predictive content of subjective beliefs. On the empirical side, Rahimi and Recht (2007, 2008) introduce Random Fourier Features (RFF) regression as a tractable approximation to nonlinear kernel methods, and Kelly, Malamud, and Zhou (2024, 2022); Didisheim et al. (2024); Kelly and Malamud (2025) show that highly parameterized RFF specifications can outperform sparse linear models in return prediction.[4] We bring these together by identifying, in a specific institutional setting, a subjective-belief distortion that is invisible to linear specifications but becomes visible under a flexible nonlinear approximation.

A third strand studies the qualitative narratives in analyst reports and the broader role of narratives in shaping beliefs and prices. Joos, Piotroski, and Srinivasan (2016) show that analyst risk scenarios contain information beyond quantitative forecasts; Ke (2024); Bradshaw et al. (2025); Bastianello, Décaire, and Guenzel (2024) apply textual analysis to study how analysts form subjective beliefs, frame risk discussions, and bring mental models to valuation.[5] At a broader level, Shiller (2019) argues that contagious economic narratives shape beliefs, expectations, and economic outcomes. We extend this literature by exploiting within-report, across-scenario variation in narrative content: the Morgan Stanley Risk-Reward Framework lets us compare how the same analyst explains bull, base, and bear outcomes for the same firm at the same point in time, identifying

[4]More broadly, a growing asset-pricing literature shows that nonlinear methods can uncover predictive relations and interaction effects that are difficult to detect with linear models; see, for example, Gu, Kelly, and Xiu (2020); Freyberger, Neuhierl, and Weber (2020); Gu, Kelly, and Xiu (2021); Chen, Pelger, and Zhu (2024); Feng et al. (2024); Bryzgalova, Pelger, and Zhu (2025). For a broad overview, see Kelly and Xiu (2023).

[5]A complementary literature studies subjective beliefs in analysts' *numerical* forecasts; see, for example, La Porta (1996), Bordalo et al. (2019), Bordalo et al. (2024), Bouchaud et al. (2019), and De la O and Myers (2021).

the incremental macro attribution analysts add when moving from the expected case to the downside case—the specific miscalibration that single-scenario textual analyses cannot detect.

Finally, our mechanism connects to a long line of work in behavioral economics and finance on the availability heuristic and memory-based belief formation. Tversky and Kahneman (1973) introduce the availability heuristic as a fundamental bias in subjective probability judgments, and Bordalo, Gennaioli, and Shleifer (2020) and Bordalo et al. (2023) formalize how cues from the current environment trigger retrieval of similar past experiences, generating beliefs that systematically over-weight vivid and recently activated memories. We bring this mechanism into a high-stakes professional setting and identify institution-specific and peer-network channels through which crisis templates persist beyond individual recall—extending the canonical individual-memory account to a setting in which narrative conventions are socially and institutionally encoded.[6]

The rest of the paper proceeds as follows. Section 2 documents Macro Bear Bias. We introduce the Morgan Stanley Risk-Reward reports, construct LLM-based narrative attention measures, and compare analysts' ex ante macro attributions with an ex post realized CAPM R^2 benchmark across market states. Section 3 develops a conceptual framework that formalizes Macro Bear Bias and derives its implications for forecasts and prices. We show that the bear–base macro-attention gap predicts base-case forecast errors and that nonlinear bias-adjusted forecasts generate significant return predictability. Section 4 examines mechanisms. We test an availability-heuristic interpretation using macro-news shocks, post-crisis industry exposure, Morgan Stanley-specific crisis experience, and sector-level narrative convergence. We then evaluate a strategic career-concerns explanation, for which we find little support. Section 5 concludes.

[6]Our findings also relate to the literature on career concerns and herding (Holmström, 1999; Scharfstein and Stein, 1990; Hong, Kubik, and Solomon, 2000; Hong and Kubik, 2003; Chevalier and Ellison, 1999). A natural alternative to our cognitive interpretation is that macro narratives function as reputationally safer deflections from sharp firm-specific calls, because aggregate explanations are harder to falsify and less damaging to management relationships. The data provide little support for this strategic view: junior analysts do not display stronger Macro Bear Bias, recent Morgan Stanley investment-banking relationships do not amplify it, and standard analyst-level career-incentive proxies explain little of its variation.

2 Documenting Macro Bear Bias

This section documents the paper's main empirical result: Macro Bear Bias. We begin by introducing the Morgan Stanley Risk-Reward reports and our text-based measures of narrative risk attribution. We then establish that bear scenarios receive disproportionately greater macro emphasis than base or bull scenarios, and compare these ex ante attributions with an ex post realized benchmark to assess whether this emphasis reflects the true driver of downside outcomes. Finally, we characterize the content of this macro emphasis—which topics populate bear narratives and how attention responds to the news cycle.

2.1 Data

Our paper uses data from the Risk-Reward Framework in Morgan Stanley analyst reports downloaded from the LSEG (formerly Refinitiv) Workspace. In total, we collected 63,191 Risk-Reward reports from January 2007 to May 2025. Some examples are shown in Appendix Section A.

Traditional sell-side research often relied on single-point price targets or coarse recommendations (e.g., "buy," "hold," or "sell"), which can convey a false sense of precision and provide limited information about the range of possible future outcomes. To address this limitation, Morgan Stanley introduced the Risk-Reward Framework in 2007, requiring analysts to present and justify 12-month stock-price expectations under three scenarios: bull, base, and bear. This framework was designed to provide a more informative and differentiated research product for institutional investors, who care not only about expected return but also about the distribution of potential outcomes and the sources of upside and downside risk. By explicitly presenting upside and downside scenarios, these reports help investors assess the skewness of expected payoffs, understand the assumptions underlying each case, and incorporate analyst views more directly into portfolio decisions. The reports are particularly valuable for our purposes because they provide both structured scenario-specific price targets and narrative discussion of the reasoning, key drivers, and investment debate behind those targets. [7]

[7] For introductions to the framework, see Weyns et al. (2011) and Srinivasan and Lane (2011).

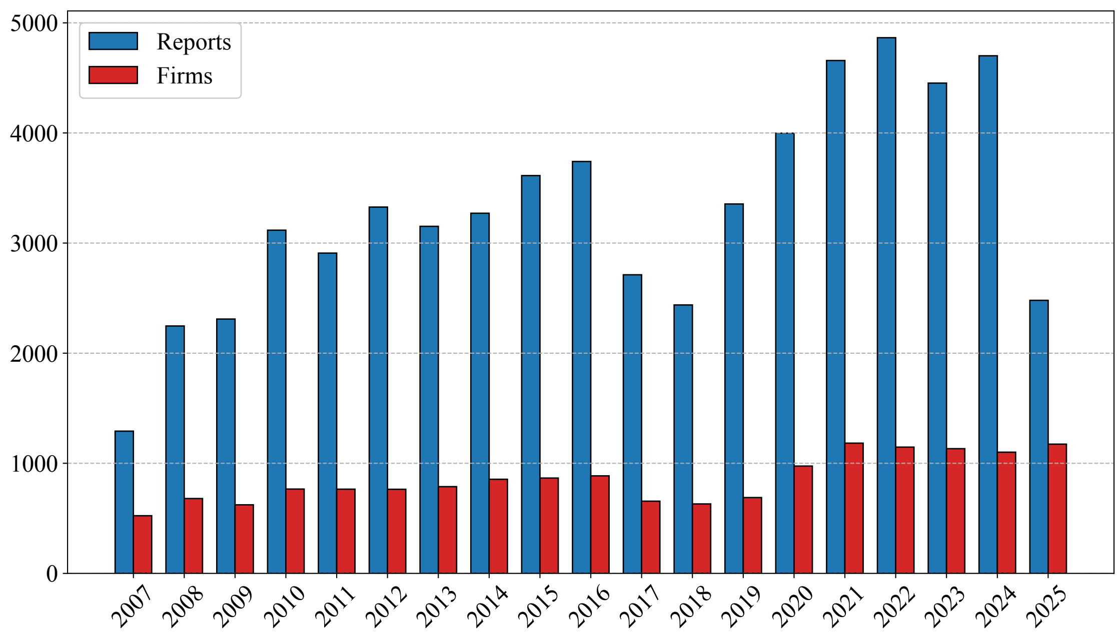

Figure 1 Histogram of Reports and Firms by Year.

Note: This figure shows the yearly count of Morgan Stanley Risk-Reward analyst reports and covered firms from January 2007 to May 2025.

This dataset offers two main advantages. First, in addition to the base-case forecast, it contains analyst expectations under both bull and bear scenarios, which allows us to measure subjective upside and downside risk expectations separately. Second, the accompanying narratives explain the assumptions and reasoning behind each scenario. Using natural language processing (NLP) and large language models (LLMs), we can systematically extract and analyze these texts to identify the drivers of perceived upside and downside risks.

Figure 1 reports the number of available Risk-Reward reports and the number of covered firms by year. Because the Risk-Reward Framework is specific to Morgan Stanley and covers primarily large-capitalization, analyst-covered U.S. equities, the results should be interpreted with this sample scope in mind.

2.2 Textual Identification of Subjective Risks

We focus on the narrative text in each Morgan Stanley Risk-Reward report to understand how analysts justify their bull-, base-, and bear-case expectations. Specifically, we examine whether analysts exhibit a downside–upside asymmetry in the types of risks they emphasize across scenarios, distinguishing among macro-level, industry-level, and firm-level risks. To quantify the content of these narratives, we apply the GPT prompt described in Appendix Section B using the gpt-4o-mini model. Our main finding is that analysts place significantly greater emphasis on macro-level risk in bear scenarios than in base or bull scenarios. As a robustness check, we also examine the distinction between systematic and idiosyncratic risk using the prompt described in Appendix Section B.3.

Figure 2 presents the report-level distributions of risk attention across scenarios. Panel A shows the distribution of attention to macro-level risks in bull, base, and bear cases. Panel B reports the corresponding distribution for firm-level risks, and Panel C does the same for industry-level risks. The figure reveals a clear pattern: relative to bull and base scenarios, bear scenarios assign substantially greater weight to macro-level risks and significantly less weight to firm-level risks. By contrast, the distributions of macro-level attention in bull and base scenarios are similar to each other, as are the distributions of firm-level attention across these two scenarios. Industry-level attention shows broadly similar distributions across all three scenarios.

Figure 3 complements this distribution evidence by showing the time-series evolution of average attention to macro-, firm-, and industry-level risks from 2007 to 2025. Panel A plots macro-level attention, Panel B plots firm-level attention, and Panel C plots industry-level attention. The time-series evidence reinforces the distributional results: throughout the sample period, analysts consistently devote more attention to macro-level risks in bear scenarios than in bull or base scenarios, while placing less attention on firm-level risks in bear scenarios. In contrast, attention patterns in bull and base scenarios remain broadly similar over time for both macro- and firm-level risks. Industry-level attention again appears relatively stable and similar across bull, base, and bear cases. For completeness, Figure A.1 reports the distribution at each point in time.

2.3 Realized CAPM R^2

The evidence so far establishes that analysts assign substantially more macro attention to bear scenarios than to base or bull scenarios. But is this differential attention warranted? To assess whether analysts are correct to place greater weight on macro risks in bear scenarios, we construct

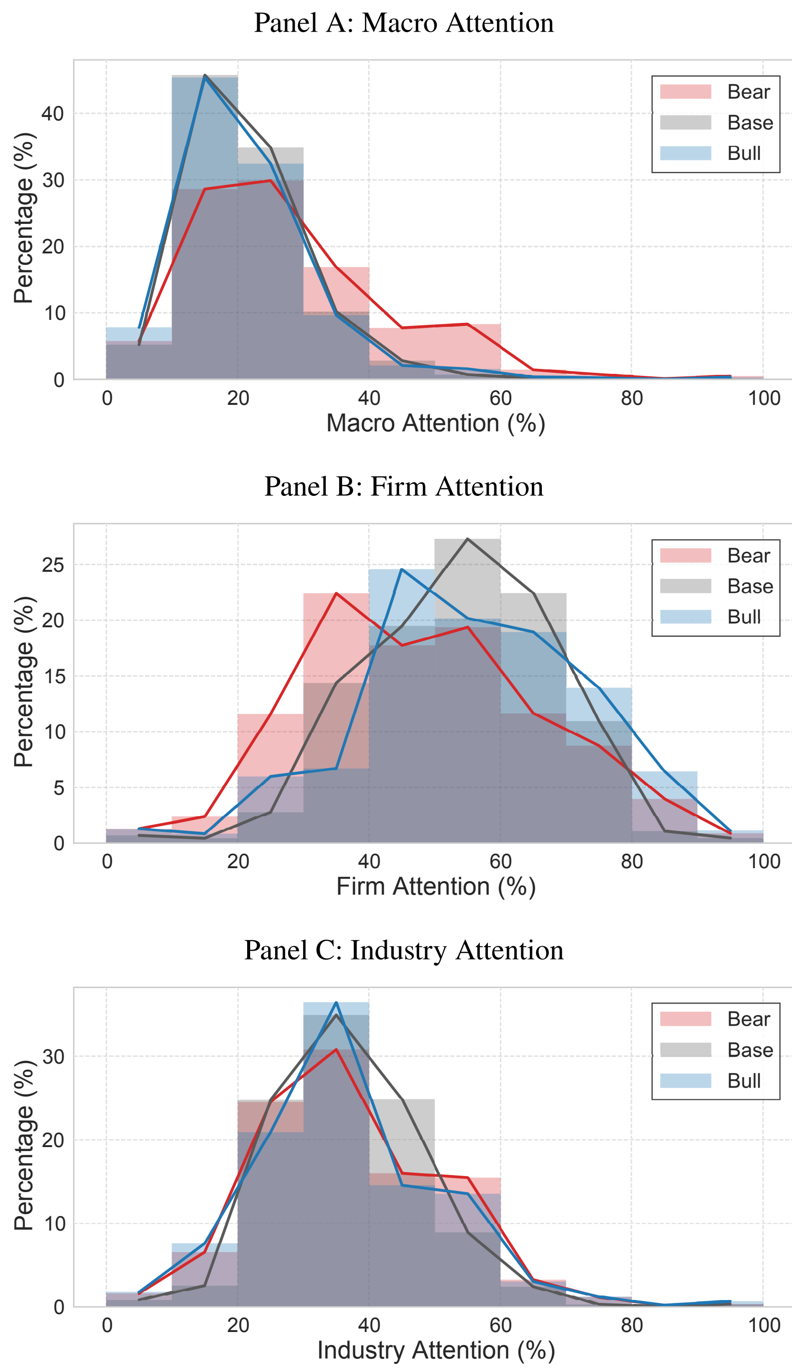

Figure 2 Attention Distribution to Macro, Firm, and Industry News.

Note: Report-level distribution of attention shares to firm-, industry-, and macro-level news across bull, base, and bear cases.

an *ex post* realized benchmark based on the realized CAPM R^2 in each case, following Roll (1988)'s classic argument that R^2 captures the share of a firm's return variation attributable to market-wide

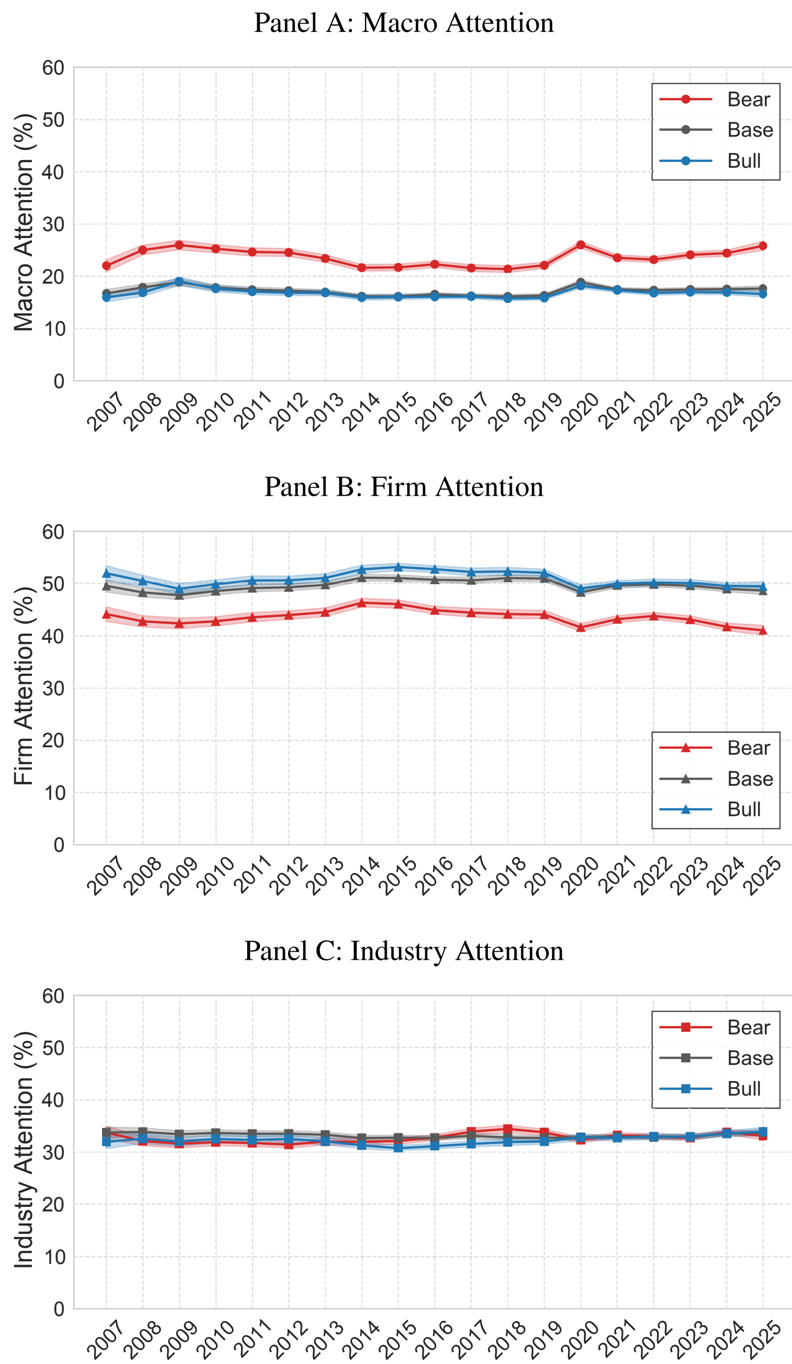

Figure 3 Time Series of Average Attention to Macro, Firm, and Industry News.

Note: Time series of report-level annual average attention shares to macro-, firm-, industry-level news. The shaded areas are 99% confidence interval of the average.

forces. If macro risks truly matter more in downside states, a firm's realized return variation should be more strongly explained by market-wide movements when the report ultimately realizes a bear

Table 1 Scenario Dummies and Attention Shares

Note: This table reports panel regressions of attention shares on scenario indicators. The dependent variables are macro attention (cols 1–2), firm attention (cols 3–4), and industry attention (cols 5–6). Base scenario is the omitted category; the constant reports the mean attention share for the Base scenario. Bear and Bull coefficients measure the difference relative to Base. All specifications include firm fixed effects; even columns add year-month fixed effects. Standard errors are clustered by firm and month. t-statistics in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1) Macro	(2) Macro	(3) Firm	(4) Firm	(5) Industry	(6) Industry
Bear	6.4687***	6.4687***	-6.0672***	-6.0672***	-0.4027**	-0.4027**
	(34.09)	(34.07)	(-31.54)	(-31.53)	(-2.55)	(-2.55)
Bull	-0.3982***	-0.3982***	1.1609***	1.1609***	-0.7654***	-0.7654***
	(-5.54)	(-5.54)	(10.20)	(10.19)	(-7.21)	(-7.21)
Constant	17.1843***	17.1843***	49.6454***	49.6454***	33.1654***	33.1654***
	(443.27)	(443.27)	(858.57)	(858.57)	(722.96)	(722.96)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE		Yes		Yes		Yes
Observations	187,917	187,917	187,917	187,917	187,917	187,917
Within R^2	0.0781	0.0863	0.0431	0.0521	0.0006	0.0053

outcome than when it realizes a base or bull outcome. We test whether the realized CAPM R^2 is significantly higher for realized-bear reports than for realized-base or realized-bull reports. A higher realized R^2 in bear outcomes would support analysts' macro emphasis in downside scenarios; otherwise, it would suggest overweighting.

Our empirical procedure has three steps. First, for each report, we compute the realized ex post CAPM R^2 by regressing the firm's daily excess returns on daily market excess returns over the 252 trading days following the report release. This R^2 measures the extent to which the firm's subsequent stock return variation is explained by systematic market movements.

Second, we classify each report as realized-bear, realized-base, or realized-bull according to the relationship between the analyst's target prices and the stock's realized price one year after the report release. A report is classified as realized-bear if the realized price is less than or equal to the bear-case target price. It is classified as realized-bull if the realized price is greater than or equal to the bull-case target price. If the realized price lies strictly between the bear and bull targets, the report is classified as realized-base.

Third, we aggregate these report-level realized R^2 measures to the firm level. For each firm, the realized-bear R^2 is the average realized CAPM R^2 across all of its realized-bear reports. Likewise, the realized-base and realized-bull R^2 are defined as the average realized CAPM R^2 across the firm's realized-base and realized-bull reports, respectively. Constructing the measures at the firm level allows us to compare whether the realized R^2 is higher in bear outcomes than in base or bull outcomes for the same firm, thereby reducing concerns that cross-firm differences drive the results.

To make these comparisons within a common market environment, we further construct firm-level realized R^2 measures separately by market condition. Specifically, under a given market state, we restrict the sample to firms that experience all three realized report types—bear, base, and bull—within that same market condition. For example, when analyzing periods in which the market excess return over the 252 trading days after the report release lies between −5% and 5%, we include only firms that have realized-bear, realized-base, and realized-bull reports within that market-return range. This restriction ensures that comparisons across realized cases are not mechanically driven by differences in overall market environments.

We obtain stock return data from the CRSP database. After matching Morgan Stanley reports to 252-trading-day forward returns following the report release date, the sample contains 45,348 reports in total. Figure 4 presents the annual number of Morgan Stanley Risk-Reward analyst reports and the number of covered firms in the matched sample.

To mitigate distortions from stock splits and reverse splits, we construct an alternative measure of the one-year-ahead realized price by compounding daily returns excluding dividends, denoted by $r_{t,x}$, over the 252 trading days after the report release:

$$\text{constructed realized price} = \text{current price} \times \prod_{t=1}^{252} (1 + r_{t,x}).$$

If the realized price in one year differs from the constructed realized price by more than 10% in either direction, we replace the observed realized price with the constructed realized price. This adjustment helps correct cases in which stock splits or reverse splits mechanically distort the reported price level.

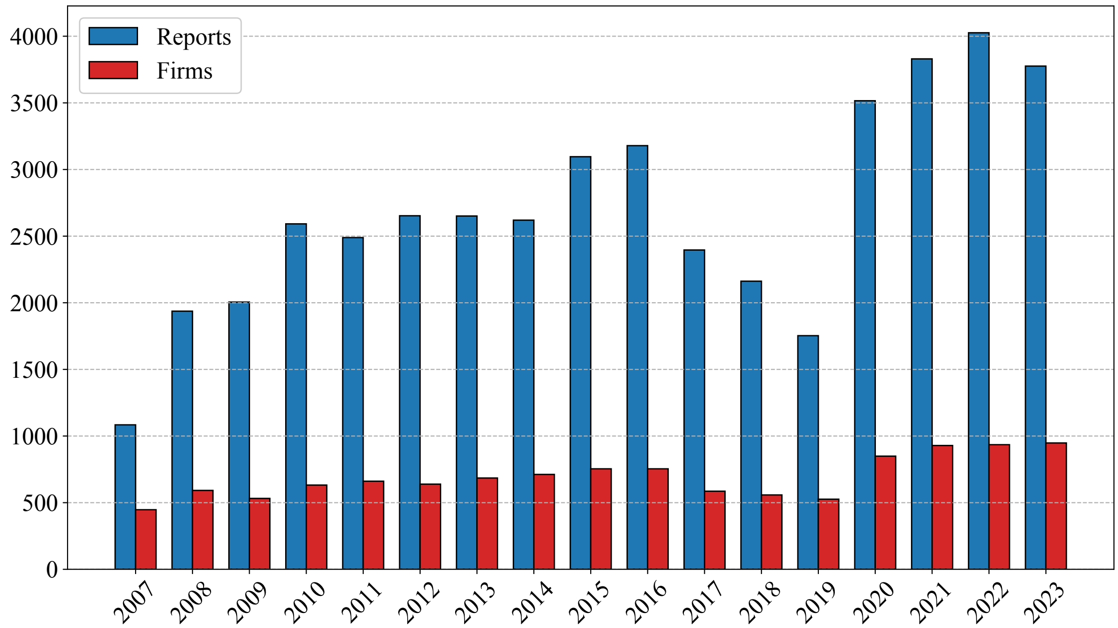

Figure 4 Yearly Distribution of Reports and Firms after Matching with Returns.

Note: This figure shows the annual number of Morgan Stanley Risk-Reward analyst reports and the number of covered firms after matching the reports with one-year-ahead firm returns.

Table 2 Summary Statistics

Note: This table reports summary statistics for the Morgan Stanley Risk-Reward report sample from 2007 to 2023. Panel A reports report-level attention shares (percentages) by scenario. Panel B reports firm characteristics at the report date. All attention shares are extracted by GPT using the prompt in Appendix B.

Variable	N	Mean	Median	Std	Min	P25	P75	Max
Panel A: Narrative Attention								
Bull: Macro	45765	16.70	10.00	11.11	0.00	10.00	20.00	100.00
Bull: Industry	45765	32.29	30.00	14.33	0.00	20.00	40.00	100.00
Bull: Firm	45765	51.01	50.00	17.32	0.00	40.00	60.00	100.00
Base: Macro	45765	17.02	15.00	9.46	0.00	10.00	20.00	100.00
Base: Industry	45765	33.14	30.00	11.49	0.00	25.00	40.00	100.00
Base: Firm	45765	49.84	50.00	14.43	0.00	40.00	60.00	100.00
Bear: Macro	45765	23.33	20.00	15.10	0.00	10.00	30.00	100.00
Bear: Industry	45765	32.64	30.00	14.06	0.00	20.00	40.00	100.00
Bear: Firm	45765	44.03	40.00	18.32	0.00	30.00	60.00	100.00
Panel B: Returns and Forecast Variables								
Return	45765	0.12	0.07	0.74	-3.29	-0.17	0.31	74.14
Subjective Return	45765	0.20	0.10	1.79	-4.67	-0.01	0.22	115.33
Forecast Error	45765	-0.08	-0.02	1.91	-115.26	-0.30	0.26	74.03
$\hat{\beta}_{\text{CAPM}}$	43008	1.13	1.08	0.44	-1.04	0.83	1.37	5.29
R^2_{CAPM}	43008	0.32	0.30	0.18	0.00	0.18	0.44	0.86

Table 2 reports summary statistics of core narrative and firm-level data. Panel A highlights the narrative asymmetry across scenarios. Bear-case macro attention averages 23.4%, compared with 17.0% and 16.7% in the base and bull cases, respectively—a bear−base gap of approximately 6.3 percentage points. The mirror image appears in firm-level attention: bear-case firm attention averages 44.0%, roughly 6 percentage points below the base-case mean of 49.9%, while industry attention is nearly identical across all three scenarios. Panel B shows that analyst base-case targets are slightly optimistic on average (mean forecast error of −0.08), and that the coverage universe skews toward high-beta firms (mean $\hat{\beta}_{\text{CAPM}} = 1.13$), consistent with Morgan Stanley's focus on large-cap stocks.

Figure 5 illustrates how realized bear, base, and bull reports are distributed across different market states. To characterize market conditions, we group reports according to the forward 252-day market excess return following the analyst report release. For each point on the horizontal axis, we consider reports whose forward market excess returns fall within a centered 10% window around that point.

The two panels present the same data from complementary perspectives. Panel A focuses on the composition of realized outcomes within each market state. For a given market-return state, it reports the shares of realized bear, base, and bull reports among all reports that fall into that state. This perspective shows how the relative likelihood of each realized outcome changes as aggregate market conditions become weaker or stronger. The figure shows a clear pattern: as future market excess returns rise, the share of realized bear reports declines, while the share of realized bull reports increases. In contrast, the share of realized base reports remains comparatively stable across market states. This pattern is intuitive. When the overall market performs poorly, more stocks are likely to end up near analysts' downside scenarios, making realized bear outcomes more common. When the market performs strongly, more stocks are likely to reach their upside scenarios, increasing the frequency of realized bull outcomes.

Panel B instead focuses on how each realized outcome is distributed across market states. For realized bear, base, and bull reports separately, it plots the fraction of that report type that falls into each market-return state. This perspective highlights the market environments in which each

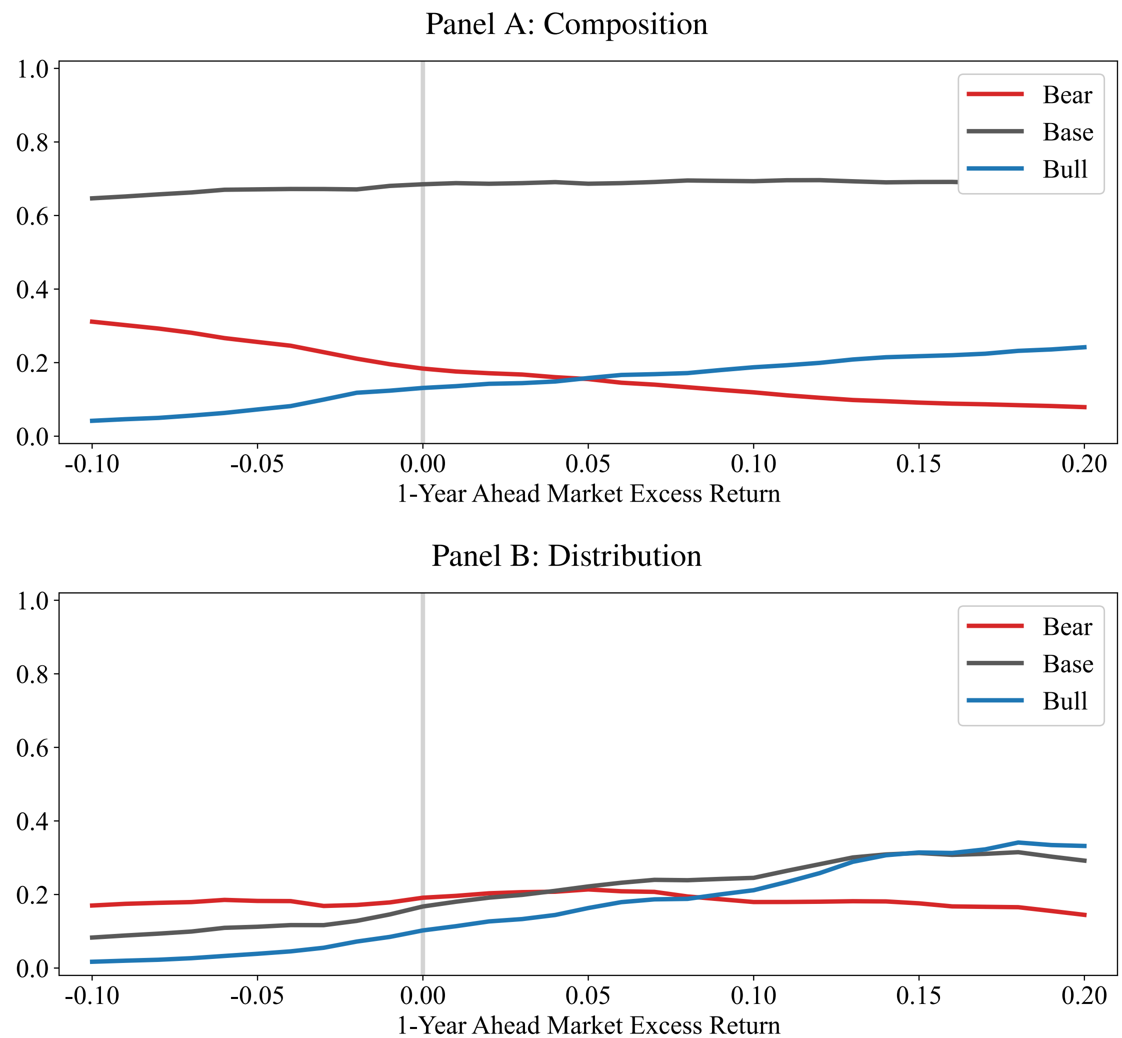

Figure 5 Composition and Distribution of Realized Bear, Base, and Bull Reports Across Market States.

Note: This figure plots the composition and distribution of realized bear, base, and bull reports across market states. For each point x on the horizontal axis, we consider the reports within 10% bins ($[x-5\%, x+5\%]$) of forward 252-day market excess returns following the analyst report release. Panel A shows the composition of bear, base, and bull reports within each return state:

$$\frac{\text{Number of realized bear/base/bull reports in a return state}}{\text{Total number of reports in that market excess return}}.$$

Panel B shows the distribution for realized bear, base, and bull reports:

$$\frac{\text{Number of realized bear/base/bull reports in a market excess return state}}{\text{Total number of realized bear/base/bull reports}}.$$

realized outcome is most likely to occur. Realized bear reports are spread more broadly across market states, resulting in a relatively flat profile. By contrast, realized base reports, and especially realized bull reports, become more concentrated in stronger market states. This pattern suggests that the distribution of forward market excess returns is tilted toward positive values, so realized base and bull outcomes are more likely to occur when subsequent market conditions are favorable.

Read jointly, the two panels provide a more complete picture of the relation between realized report outcomes and subsequent market conditions. The figure shows that realized outcomes are systematically linked to the broader market environment: realized bear outcomes are more prevalent when subsequent market conditions are weak, while realized bull outcomes become both more common and more concentrated when subsequent market conditions are strong. This implies that, when comparing realized bear, base, and bull benchmarks, it is important to condition on market state; otherwise, observed differences may conflate firm-level realization with variation in aggregate market performance.

2.4 Empirical Evidence on Macro Bear Bias

With the realized CAPM R^2 benchmark and market-state conditioning framework in hand, we now examine whether the elevated macro emphasis in bear-case narratives reflects the actual macro content of bear-case outcomes. We show that it does not. Instead, analysts systematically overweight macro risk in bear scenarios. We refer to this distortion as Macro Bear Bias: the tendency to over-attribute downside outcomes to aggregate macro risks in bear-case narratives.

To ensure an apples-to-apples comparison with realized R^2, we construct firm-level ex ante macro attention for each case (bear, base, and bull) as the average macro attention across all reports of the firm that fall within the same market state. Thus, the comparison with the ex post realized R^2 is based on the exact same subset of reports. Figure 6 plots realized R^2 and macro attention conditional on market state. For each point x on the horizontal axis, we consider reports whose forward 252-day market excess returns fall within a centered 10% window, $[x-5\%, x+5\%]$, around that point. Panel A shows the average CAPM R^2 over the 252 trading days after report release

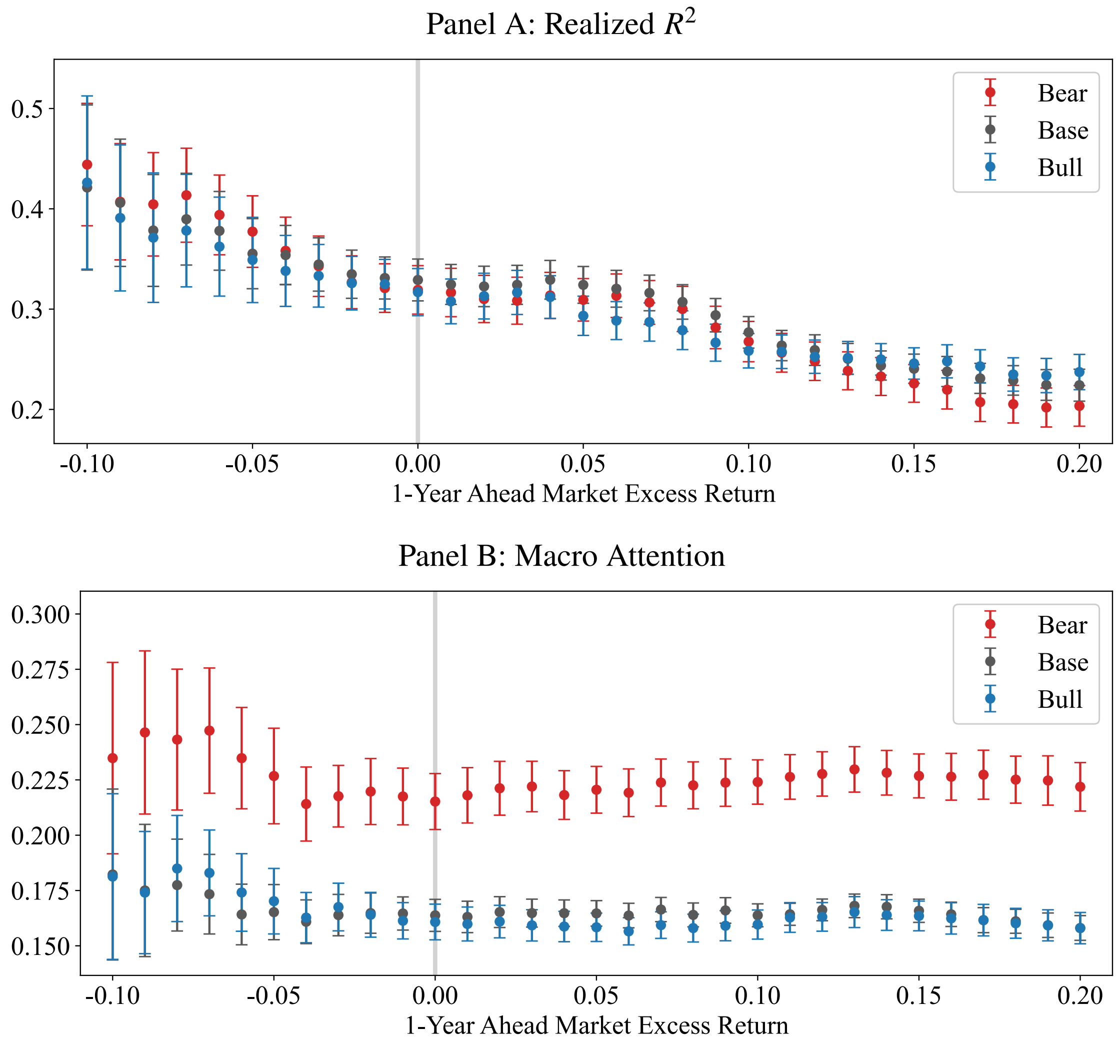

Figure 6 Realized R^2 and Macro Attention Conditional on Market States.

Note: This figure shows the realized R^2 and macro attention conditional on different market states. For each point x on the horizontal axis, we consider the reports within 10% bins ($[x - 5\%, x + 5\%]$) of forward 252-day market excess returns following the analyst report release. We apply the same-firms rule within each bin, requiring firms to have realized-bear, realized-base, and realized-bull reports. Macro attention is measured as the firm-level average macro attention within each realized scenario (bear, base, and bull). Realized CAPM R^2 is computed as the firm-level average R^2 over the 252 trading days following the report release for reports ex post classified as realized bear/base/bull cases. Error bars show 95% confidence intervals from constant-only regressions.

for realized bear, realized base, and realized bull cases. Panel B shows the average ex ante macro attention that analysts assign to the bear, base, and bull cases. [8]

Two findings stand out. First, macro attention in the bear case is consistently and substantially higher than in the base or bull case, even though the realized R^2 is similar across the realized

[8] As a robustness check, we find similar results when classifying risks into systematic and idiosyncratic components; see Figure A.3.

bear, base, and bull cases. This bear-base attention gap is the empirical signature of Macro Bear Bias: even when the realized macro component of risk is similar across scenarios conditional on market state, analysts continue to describe downside outcomes as disproportionately macro-driven. Second, the three attention series are relatively flat across market states, in sharp contrast to the realized R^2 benchmarks, which decline monotonically with market states. This suggests that analysts do not sufficiently adjust their attribution of macro risk to the future one-year realized market environment.

Figure 7 makes this pattern especially clear by plotting the bear–base gap in macro attention together with the corresponding bear–base gap in realized R^2 across market states. The realized R^2 gap is close to zero when the subsequent annual market excess return is below 12%, and turns negative when the market excess return is above 12%. In contrast, the bear–base gap in macro attention remains significantly positive throughout. In other words, analysts consistently portray the bear case as much more macro-driven than the base case, even though the realized difference in macro risk is economically small or even negative, depending on the market state. This divergence provides direct evidence of Macro Bear Bias.[9]

These patterns point to two possible explanations for Macro Bear Bias. The first is that analysts implicitly treat realized macro risk as roughly constant across market states, while at the same time mechanically assigning extra weight to macro risk in bear cases. Under this explanation, analysts do condition on downside scenarios, but they fail to recognize that the importance of macro risk varies with the market environment. The second possibility is that analysts understand that realized R^2 declines with market returns, but still attach a persistently high probability to a severely negative future market outcome regardless of the prevailing market state. Under this interpretation, the elevated macro attention in bear cases would reflect a belief that downside scenarios are tightly linked to large aggregate market declines.

The evidence is more consistent with the first explanation. If the second explanation were correct, two implications should follow. First, base-case macro attention should track the slope of realized base-case R^2 across market states, in whichever direction Figure 6 establishes that slope.

[9]Figure A.4 shows that the results are similar when risks are classified into systematic and idiosyncratic.

Second, bull-case macro attention should be materially lower than base-case macro attention, since the second explanation implies analysts assign less macro risk to more favorable scenarios. The data support neither implication. Base-case macro attention is largely flat across market states (Panel B of Figure 6), and bull-case macro attention is not materially below base-case macro attention—the bull–base gap in Table 1 is −0.40 percentage points, an order of magnitude smaller than the bear–base gap of 6.47. Viewed in concert, these findings suggest that analysts are not merely mapping different scenarios into different expected market states; rather, Macro Bear Bias reflects a systematic tendency to over-attribute downside outcomes to macro risks in bear-case narratives.

The first explanation—that analysts apply a near-unconditional macro narrative to the bear case—points to a memory-based mechanism (Bordalo, Gennaioli, and Shleifer, 2020; Bordalo et al., 2023). In constructing a bear case, an analyst may reach for resemblance to a canonical macro-crisis prototype—a severe, overwhelmingly systematic downturn—rather than form a state-contingent assessment of systematic risk. Because prototype matching is unconditional, it naturally predicts the regime-independent pattern we document, and it explains the contrast with idiosyncratic framing: firm-specific disasters do not share a common prototype and are therefore less likely to anchor bear-case narratives. We develop this availability-heuristic interpretation, and test it directly, in Section 4.1. The second mechanism is anchoring with insufficient adjustment (Tversky and Kahneman, 1974): even when firm-specific information calls for a different framing, analysts may adjust too little away from the initial macro template.

A natural objection is that the pattern could be mechanical: bear cases describe worse outcomes, which tend to co-occur with macro downturns, so bear narratives should naturally emphasize macro factors more. However, the regime-conditional analysis in Figures 6 and 7 directly addresses this. Conditional on the same forward market state, realized CAPM R^2 is similar across bear, base, and bull outcomes, yet the macro-attention gap remains large and positive. If elevated bear-case macro attention were a rational response to genuinely higher macro content in bear outcomes, the gap should shrink toward zero once we condition on market state—it does not. In unreported tests, we find no economically or statistically significant relationship between downside beta and

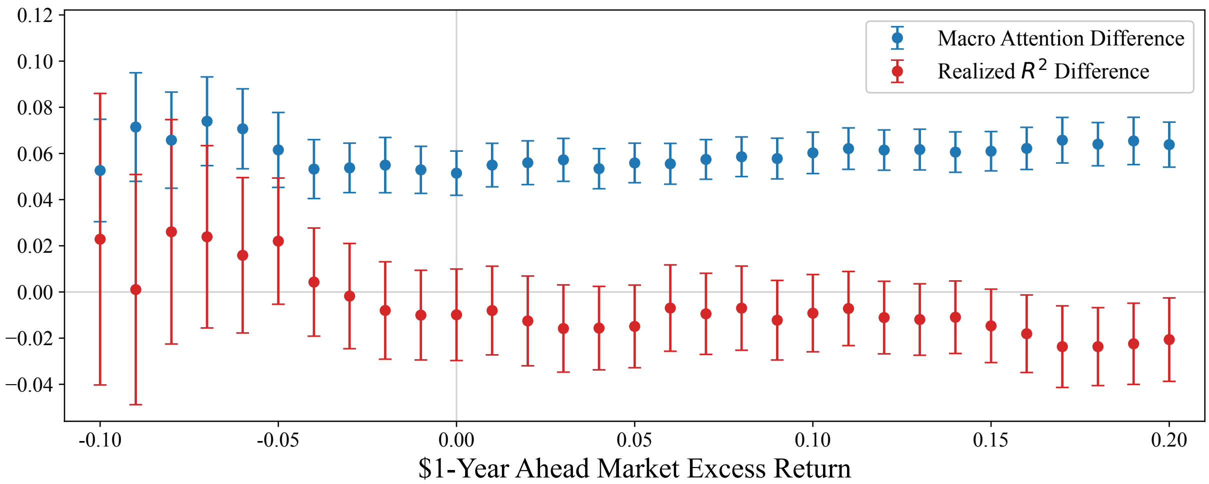

Figure 7 Bear – Base Gaps in Macro Attention and Realized CAPM R^2 Conditional on Market States.

Note: This figure plots firm-level bear–base differences in macro attention and realized CAPM R^2 on different market states. For each point x on the horizontal axis, we consider the reports within 10% bins ($[x - 5\%, x + 5\%]$) of forward 252-day market excess returns following the analyst report release. We apply the same-firms rule within each bin, requiring firms to have realized-bear, realized-base, and realized-bull reports. Macro attention is measured as the firm-level average macro attention in the bear and base cases. Realized CAPM R^2 is computed as the firm-level average R^2 over the 252 trading days following the report release for reports ex post classified as realized bear or realized base cases. The plotted values represent the bear-minus-base differences. Error bars show 95% confidence intervals from constant-only regressions.

the bear–base macro-attention gap conditional on market state, further confirming that the pattern reflects misattribution rather than rational conditioning.

2.5 Bear-Case Macro Topics

The preceding evidence establishes that bear-case narratives systematically over-emphasize macro risk relative to the realized ex post benchmark. We now characterize the *content* of this emphasis: which macroeconomic topics populate bear narratives, and do analysts update their topic-level attention in response to news or anchor to persistent templates?

To identify topic-level narrative content, we apply a second LLM extraction step that classifies each macro-risk mention extracted from the bear-case narrative into one of 52 macro topic categories drawn from Zhou (2025), excluding 7 corporate finance topics from their original 59-topic taxonomy. A single bear-case narrative may therefore contribute multiple topic labels, one per mention. Appendix B.2 describes the prompt and Table A.1 reports each topic's description

Table 3 Bear Analyst Attention and WSJ Attention Across Key Topics

Note: This table reports macro topics in analysts' bear-case narratives and corresponding attention to the same topics in Wall Street Journal articles. Bear analyst attention for a given topic is measured as the share of topic mentions in the bear-case narrative associated with this topic across analyst reports. WSJ attention for a given topic is measured as the average monthly share of Wall Street Journal articles associated with that topic over the sample period. Panel A shows the top five topics ranked by average monthly bear analyst attention. Panel B shows the top five topics with the highest correlation between monthly bear analyst attention and monthly WSJ attention. Bear Rank and WSJ Rank report the rankings of each topic by average attention within analyst bear narratives and the Wall Street Journal, respectively.

Panel A: Top Five Topics by Average Monthly Bear Analyst Attention

Topic	Bear Avg.	Bear Rank	WSJ Avg.	WSJ Rank
Economic Growth and Recovery Outlook	10.66%	1	2.36%	12
COVID-19 Pandemic and Vaccine Developments	5.36%	2	5.34%	2
Market Activity and Financial Performance	4.75%	3	5.44%	1
Consumer Price Index and Inflation Trends	2.40%	4	0.86%	27
International Trade Relations and Economic Policies	2.31%	5	1.60%	16

Panel B: Top Five Topics by Correlation Between Bear Analyst Attention and WSJ Attention

Topic	Corr.	Bear Avg.	Bear Rank	WSJ Avg.	WSJ Rank
COVID-19 Pandemic and Vaccine Developments	0.76	5.36%	2	5.34%	2
Economic Growth and Recovery Outlook	0.49	10.66%	1	2.36%	12
Consumer Behavior Amid Economic Strains	0.38	1.91%	7	0.24%	49
Geopolitical Tensions and Economic Impacts	0.38	0.46%	20	4.40%	3
Economic Stimulus and Government Interventions	0.37	0.40%	25	0.46%	40

and summary statistics. Following Zhou (2025), bear-case analyst attention to a topic in a given month is the share of mentions associated with that topic across all bear-case narratives released in that month. We benchmark this against Wall Street Journal (WSJ) coverage. WSJ attention to a topic is the monthly share of WSJ articles tagged with that topic. The two series therefore use different denominators—mention shares on the analyst side and article shares on the WSJ side—so the comparison in Figure 8 and Table 3 is best read as comovement in standardized time-series variation rather than a level-for-level match.

Table 3 reports average topic-level attention and time-series co-movement with the WSJ. Panel A ranks topics by average monthly bear-case analyst attention. *Economic Growth and Recovery Outlook* is the dominant bear-narrative topic, accounting for 10.66% of bear-case attention on average—more than four times its WSJ share of 2.36% (WSJ rank: 12). This disproportion suggests that macroeconomic growth narratives function as a canonical bear-case template: analysts invoke them persistently, at far higher rates than contemporaneous news coverage would predict. *COVID-19 Pandemic and Vaccine Developments* ranks second at 5.36%. *Market Activity and Financial Performance* ranks third (4.75%), despite being the single most-covered topic in the WSJ (rank 1, 5.44%), suggesting analysts allocate less narrative emphasis to general market conditions than the news cycle does. Inflation (*Consumer Price Index and Inflation Trends*, rank 4) and trade policy (*International Trade Relations and Economic Policies*, rank 5) round out the top five, each receiving materially more analyst attention than their WSJ ranks imply.

The reverse asymmetry is equally informative. The previous paragraph shows that analysts devote unusually high bear-case attention to recession-and-growth topics even when those topics are not among the most prominent in the WSJ. The opposite pattern appears for geopolitics: *Geopolitical Tensions and Economic Impacts* is the third most-covered WSJ topic, with a 4.40% share, but it ranks only 20th in bear-case analyst attention, with a 0.46% share. Thus, analysts are not simply importing the most visible macro topics from the news into their bear narratives. If availability operated through generic media exposure, geopolitical risk should be much more prominent in bear cases. Instead, analysts selectively emphasize a narrower set of recession-and-growth topics while largely excluding geopolitical risk despite its media salience. This pattern suggests that the relevant form of availability is not "whatever is currently in the news," but a familiar macro-crisis template for explaining downside risk. Put differently, analysts appear to reach for the kind of macro story that resembles past market-wide crises. Section 4.1 tests this interpretation directly by showing that Macro Bear Bias is stronger in post-GFC high-systematic-risk sectors and among analysts with Morgan Stanley-specific GFC experience.

Panel B ranks topics by the time-series correlation between monthly bear-case analyst attention and WSJ attention. COVID-19 leads (0.76), followed by economic growth (0.49). High correlations confirm that analysts respond to the news environment in their narrative topic choices. The critical question, however, is whether that response is symmetric: do analysts reduce topic emphasis as

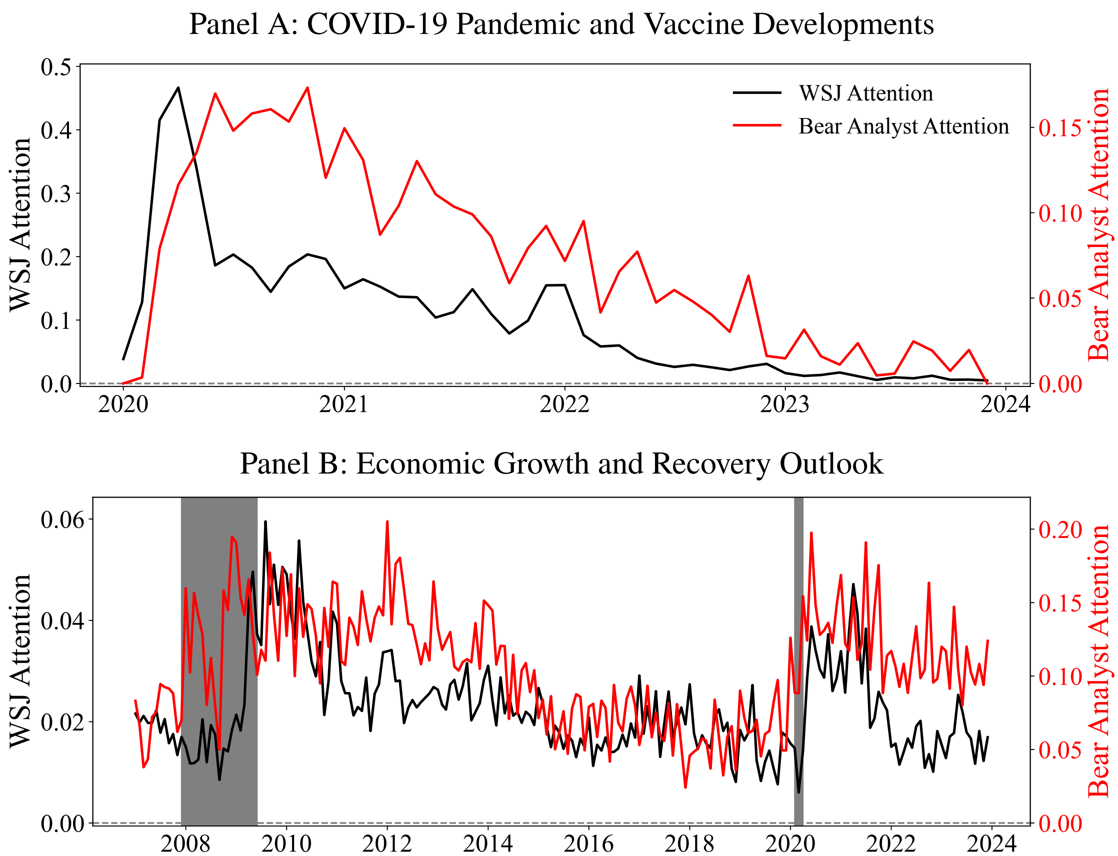

Figure 8 Bear Analyst Attention and WSJ Attention Over Time for Selected Topics.

Note: This figure plots monthly bear analyst attention (red, right axis) alongside monthly Wall Street Journal (WSJ) attention (black, left axis) for two representative macro topics. Bear analyst attention for a given month is the share of bear-case narrative mentions across all reports in that month associated with the topic. WSJ attention is the monthly share of WSJ articles associated with the topic. Panel A covers *COVID-19 Pandemic and Vaccine Developments*. Panel B covers *Economic Growth and Recovery Outlook*; gray bands indicate the Global Financial Crisis recession period (2008-2009) and the COVID-19 recession period (2020).

quickly as they raise it, or do bear-case templates persist long after the originating events have receded?

Figure 8 examines this asymmetry directly. Each panel overlays monthly bear-case analyst attention (red, right axis) against WSJ attention (black, left axis) for the two highest-correlation topics. Panel A, covering COVID-19, shows a sharp joint spike in early 2020. WSJ attention subsequently declines steeply, reaching near zero by 2023; bear-case analyst attention remains materially elevated for roughly two additional years. Panel B, covering economic growth, replicates this pattern across both the Global Financial Crisis and the COVID-19 period: analyst attention rises alongside WSJ coverage during each crisis but reverts far more slowly, remaining elevated long after media attention has subsided. The asymmetry is the same in both panels: bear-case attention rises with WSJ coverage as each crisis breaks but reverts far more slowly once coverage fades. This asymmetric response is consistent with template-based updating: salient crises make particular macro explanations easy to recall, and analysts keep deploying them as default bear-case narratives long after the originating shock has passed. Bear narratives therefore do not simply mirror contemporaneous news salience; salient macro episodes leave durable templates that continue to shape how analysts describe downside risk.

3 Conceptual Framework and Forecast Error Evidence

We now formalize the bias documented in Section 2. Analysts routinely describe bear, base, and bull cases to explain the range of possible future outcomes. In principle, these narratives should reflect how different sources of risk contribute to firm value conditional on the prevailing market environment. In practice, however, analysts appear to assign disproportionately high macro emphasis to downside scenarios, even when the conditional macro component of risk is not especially large.

A central implication of this framework is that the bias does not remain confined to the bear narrative itself. Because scenario narratives shape analysts' overall valuation judgments, excessive macro emphasis in the bear case can spill over into the analyst's base-case forecast, making it too pessimistic relative to subsequently realized outcomes. This link connects Macro Bear Bias to the prediction of analyst forecast errors.

3.1 Benchmark: Conditional Attribution of Risk

Consider firm i's excess return:

$$R_{i,t}^{\text{realized}} = \alpha_i + \beta_i M_t + \varepsilon_{i,t}, \tag{1}$$

where M_t denotes the future market excess return, β_i measures the firm's exposure to aggregate macro conditions, and $\varepsilon_{i,t}$ captures firm-specific shocks. The key feature of the benchmark is that the importance of macro risk is conditional: it should be evaluated relative to the market state prevailing when the analyst forms the forecast.

Let s_t denote the information available at time t about the market environment. When an analyst constructs scenario $k \in \{\text{bear, base, bull}\}$, the rational benchmark requires that the macro emphasis assigned to that scenario reflect the analyst's expectation of how important macro forces would be if scenario k were to materialize. A bear narrative should emphasize macro risk to the extent that bear-case outcomes are genuinely macro-driven; a bull narrative should do so to the extent that bull-case outcomes are. Empirically, we evaluate this benchmark by classifying reports ex post—k = bear if the firm's stock price ultimately falls below the analyst's bear-case target, k = base if it falls between the bear and bull targets, and k = bull if it exceeds the bull target—and comparing the ex ante attention to the ex post macro content of each realized outcome.

Formally, let

$$\text{Macro Importance}_{i,t} \equiv \frac{\beta_i^2 \, \text{Var}(M_t \mid s_t, k)}{\text{Var}(R_{i,t} \mid s_t, k)}$$

denote the fraction of firm-level return variance attributable to market movements, conditional on the prevailing market state s_t and the realized scenario k. Here β_i is defined as the population coefficient from projecting $R_{i,t}$ on M_t within the conditioning set defined by (s_t, k). Equivalently, the residual $\varepsilon_{i,t}$ is orthogonal to M_t within that conditioning set, so the ratio has a natural interpretation as a variance share.[10] The empirical CAPM R^2 computed in Section 2.4 is the sample analogue of this conditional object within the corresponding market-state–realized-scenario subsample.

The time-t rational expectation of macro importance at $t + 1$ is:

$$A_{i,k,t} = g\left(\mathbb{E}_t\left[\text{Macro Importance}_{i,t+1} \mid s_t, k\right]\right), \tag{2}$$

[10]Without this conditional orthogonality, the variance decomposition would also include a $2\beta_i \text{Cov}(M, \varepsilon \mid s_t, k)$ term. The conditioning therefore makes the object a clean variance share rather than a covariance-laden index.

where $g(\cdot)$ is an increasing mapping from the underlying importance of macro risk into narrative emphasis. Under this benchmark, two properties should hold. First, macro attention should vary with the prevailing market environment, since the contribution of macro forces to firm outcomes is state-dependent. Second, differences in macro attention across scenarios should reflect genuine differences in the conditional relevance of macro risk, not a mechanical template that treats bad outcomes as macro-driven. The bear case may legitimately receive more macro emphasis, but only when conditional macro exposure truly justifies it.

This benchmark provides the reference for interpreting the data: if analysts are well calibrated, ex ante macro attention should track the ex post macro-related component of realized variation conditional on market state. Divergence between the two signals misattribution.

3.2 Macro Bear Bias: Over-attributing Macro Risks in Bear Narratives

Macro Bear Bias arises when analysts depart from this benchmark by assigning excessive macro weight specifically to the bear narrative. Formally, let observed macro attention satisfy

$$\tilde{A}_{i,k,t} = A_{i,k,t} + b_{i,t} \cdot \mathbf{1}\{k = \text{bear}\}, \qquad b_{i,t} \geq 0, \tag{3}$$

where attention shares lie in $[0, 100]$, so the inequality $0 \leq b_{i,t} \leq 100 - A^*_{i,\text{bear},t}$ ensures feasibility and the bear-case excess is offset by a corresponding reduction in non-macro attention.[11] The parameter $b_{i,t}$ captures the analyst's excess tendency to explain downside scenarios using aggregate macro forces, beyond what is justified by the conditional importance of macro risk. We allow $b_{i,t}$ to vary across firms and over time, and the empirical hypothesis is that $\mathbb{E}[b_{i,t}] > 0$ in the relevant population.

Under this formulation, the distortion is not that analysts discuss macro risk per se. Rather, it is that they over-explain downside outcomes through macro forces. The bear narrative becomes disproportionately macro-heavy even when, conditional on the same market state, the realized

[11] For interior values of $A^*_{i,\text{bear},t}$ the upper bound is slack; the constraint binds only when the unbiased bear-case macro share is already near the ceiling.

macro component of outcomes is similar to that in the base or bull case. This is why Macro Bear Bias is best understood as a misattribution bias.

The empirical counterpart of this bias is the bear–base macro-attention gap:

$$\text{MacroBearGap}_{i,t} = \text{MacroAttention}^{\text{bear}}_{i,t} - \text{MacroAttention}^{\text{base}}_{i,t}. \tag{4}$$

A larger $\text{MacroBearGap}_{i,t}$ indicates that the analyst places substantially greater macro emphasis on the downside scenario than on the base case. Under the benchmark, this gap should reflect genuine conditional differences in macro relevance. Under Macro Bear Bias, however, the gap is too large because the bear narrative receives an additional, bias-driven macro overlay.

3.3 Economic Intuition

The framework distinguishes between two objects: the true conditional importance of macro risk for firm outcomes, and the analyst's narrative attribution of macro risk across scenarios. If analysts are well calibrated, these should move together. Macro Bear Bias breaks this alignment. Analysts apply a coarser rule: downside scenarios are treated as macro-driven almost by construction, regardless of the prevailing market state. The distortion is not that analysts mismeasure macro risk on average but that they do so *asymmetrically*, loading macro explanations disproportionately into bear narratives.

This asymmetry matters because scenario narratives are not merely descriptive—they organize analysts' thinking about valuation. A macro-heavy bear case can contaminate the base-case forecast itself, making it too pessimistic relative to subsequently realized outcomes. We develop the microfoundation of this spillover through a discount-rate channel and derive a cross-sectional prediction linking this channel to the direction of analyst conviction.

The primary channel operates through discount rates. A bear case described as predominantly macro-driven leads the analyst to perceive elevated systematic risk, implicitly adopting a higher equity risk premium ($\mu_t \equiv \mathbb{E}_t(M_{t+1})$) when discounting base-case cash flows. Let the base-case price target be $P^{\text{base}} = \sum_\tau CF_\tau/(1 + r_\tau)^\tau$, where $r_\tau = r_f + \tilde{\beta}_i \cdot \mu_t$. Macro Bear Bias inflates the analyst's perceived beta *multiplicatively*, so that the distortion scales with the firm's true systematic exposure:

$$\tilde{\beta}_{i,t} = \beta_i \left(1 + \phi \cdot \text{MacroBearGap}_{i,t}\right), \qquad \phi > 0. \tag{5}$$

Under this specification, the incremental discount-rate wedge is $\beta_i \phi \cdot \text{MacroBearGap}_{i,t} \cdot \mu_t$: a high-beta firm is hit by a larger absolute wedge than a low-beta firm at the same MacroBearGap. The wedge raises the discount rate and depresses P^{base} even if point estimates of base-case cash flows are unchanged, and the scaling by β_i generates the $\beta_i \times \text{MacroBearGap}_{i,t}$ interaction tested in the empirical specifications below. This microfoundation maps directly onto the reduced-form spillover parameter θ in Eq. (9): θ captures the product of ϕ and the sensitivity of the price target to the discount rate.[12] [13]

A complementary narrative-substitution logic predicts that the direction of an analyst's base-case conviction relative to the consensus should predict the structure of the bear-case narrative. An optimistic analyst—whose base-case price target exceeds the consensus—needs a bear case sufficiently distant from the elevated base case. Firm-specific risks (product delays, margin compression) may not generate enough downside distance, because they are bounded by the firm's own cash-flow distribution. A macro bear narrative, by contrast, can invoke a market-wide downturn that depresses the stock regardless of firm fundamentals, producing the requisite spread. This narrative-substitution reasoning runs alongside the discount-rate channel: the optimistic analyst

[12] Some DCF implementations weight scenarios explicitly: $P^{\text{reported}} = w_{\text{bear}} P^{\text{bear}} + w_{\text{base}} P^{\text{base}} + w_{\text{bull}} P^{\text{bull}}$. If availability inflates w_{bear} or macro narratives compress P^{bear}, the composite target shifts down. This is a reduced-form version of the discount-rate channel.

[13] A numerical example anchors the economic magnitude. At the sample-mean values $\hat{\beta}_{\text{CAPM}} = 1.13$ (Table 2), a six-percent equity risk premium, and a bias coefficient $\phi = 0.005$, a one-standard-deviation MacroBearGap shift of about 6.3 percentage points raises the perceived beta from 1.13 to $1.13 \times (1 + 0.005 \times 6.3) = 1.166$. The incremental discount-rate wedge is therefore $\hat{\beta}_{\text{CAPM}} \phi \, \text{MacroBearGap} \cdot \mu = 1.13 \times 0.005 \times 6.3 \times 6\% \approx 21$ bp. At an effective valuation duration of ten years, modified duration is $10/(1 + r) \approx 9.4$, so the implied compression of the base-case target price is -9.4×21 bp $\approx -2.0\%$. This is comparable to the median 1–2% standardized base-case forecast error documented in Table 4, suggesting the discount-rate channel can plausibly carry the magnitude of the observed bias. The same calculation under $\beta = 2.0$ (high-beta firm) yields a wedge of $2.0 \times 0.005 \times 6.3 \times 6\% \approx 38$ bp and a target-price compression of about -3.5%; under $\beta = 0.5$ (low-beta firm) the wedge is about 9 bp and the compression is about -0.9%. The wedge therefore scales roughly four-to-one between extreme high- and low-beta firms in the sample, which delivers the cross-sectional content of the $\beta_i \times \text{MacroBearGap}_{i,t}$ interaction in Eq. (10): high-beta firms should display the largest forecast-error response to MacroBearGap.

signals confidence in firm-specific fundamentals, so the residual explanation for downside must be systematic. Conversely, an analyst whose base case is below consensus has already embedded pessimism into the base case; firm-specific downside narratives suffice, and the macro alibi is less needed. This reasoning predicts that base-case position relative to consensus should covary with the macro content of the bear case: above-consensus base targets associated with larger MacroBearGap, below-consensus base targets with smaller MacroBearGap. We test this prediction in Section 4.1.2; the data support it but in an asymmetric form—the cross-section is driven primarily by the below-consensus margin.

These microfoundations also clarify why comparing ex ante macro attention with ex post realized R^2 is informative, and why a linear specification understates the predictive content of Macro Bear Bias. As documented in Section 2.4, bear-case macro attention is essentially flat across market environments. But the consequences of this flat gap are state-dependent: in strong markets, the same macro-heavy bear narrative diverges more sharply from realized conditions, generating larger forecast errors; in weak markets, the template is closer to correct, and the distortion is smaller. When macro attention in bear narratives remains elevated even though conditional realized R^2 is similar across scenarios, that wedge is direct evidence of Macro Bear Bias.

3.4 Forecast Error Prediction

The key implication of the framework is that Macro Bear Bias should predict subsequent errors in analysts' base-case forecasts. Let the realized return and the analyst's base-case forecasted return be

$$R^{\text{realized}}_{i,t+1} \equiv \frac{P^{\text{realized}}_{i,t+1}}{P_t}, \qquad R^{\text{base}}_{i,t+1} \equiv \frac{P^{\text{base}}_{i,t+1}}{P_t}, \tag{6}$$

and define the base-case forecast error as

$$\text{Forecast Error}_{i,t+1} = R^{\text{realized}}_{i,t+1} - R^{\text{base}}_{i,t+1}. \tag{7}$$

Under this definition, a positive forecast error indicates that the analyst's base-case forecast is too pessimistic relative to the subsequently realized return.

Absent Macro Bear Bias, the analyst's base-case forecast would reflect expected market conditions:

$$R^{\text{base}}_{i,t+1} = \alpha_i + \beta_i \tilde{\mu}_t, \tag{8}$$

where $\tilde{\mu}_t$ is the analyst's expectation of the future market excess return. Macro Bear Bias enters when excessive macro emphasis in the bear narrative spills over into the base-case forecast. We capture this spillover through a reduced-form specification: the base-case forecast is shifted downward in proportion to both the firm's systematic exposure β_i—because the bias operates through the macro channel—and the size of the bear–base attention gap. The parameter $\theta > 0$ is a reduced-form coefficient. The sign prediction in Eq. (11) holds for any spillover function that is monotonically increasing in MacroBearGap; we adopt the linear form below for analytical transparency, but the empirical content is the sign of the derivative, not the specific functional form.

$$R^{\text{base}}_{i,t+1} = \alpha_i + \beta_i \tilde{\mu}_t - \theta \beta_i \,\text{MacroBearGap}_{i,t}, \qquad \theta > 0. \tag{9}$$

This specification formalizes the idea that stronger Macro Bear Bias makes the analyst's central forecast too low. The distortion operates through the macro channel: when analysts over-attribute downside outcomes to aggregate macro forces, they embed too much downside macro risk into the base case.

Substituting the realized-return process in Eq. (1) and the biased base-case forecast in Eq. (9) into the forecast-error definition in Eq. (7) yields

$$\text{Forecast Error}_{i,t+1} = \beta_i (M_{t+1} - \tilde{\mu}_t) + \theta \beta_i \,\text{MacroBearGap}_{i,t} + \varepsilon_{i,t+1}. \tag{10}$$

The intercept α_i enters both Eq. (1) and Eq. (9) identically and therefore cancels in the forecast error. This equation generates a sharp empirical prediction. Holding fixed the realized market shock and other determinants of firm outcomes, a larger bear–base macro-attention gap should predict a larger, more positive forecast error. Equation (10) implies

$$\frac{\partial\,\text{Forecast Error}_{i,t+1}}{\partial\,\left(\beta_i\,\text{MacroBearGap}_{i,t}\right)} \;=\; \theta > 0. \tag{11}$$

The intuition is straightforward. A larger β_i MacroBearGap$_{i,t}$ indicates that the analyst places unusually strong macro emphasis on the downside scenario relative to the base case. If this reflects Macro Bear Bias rather than a rational assessment of conditional macro risk, then the analyst sets the base-case forecast too pessimistically. When outcomes are later realized, the forecast tends to be too low, producing a positive forecast error.

We now test this prediction empirically. Equation (10) motivates a panel regression of forecast errors on MacroBearGap with firm fixed effects and a market-state control. Because α_i has already cancelled in the derivation of Eq. (10), the role of firm fixed effects is to absorb any residual persistent firm-level forecast-error component (e.g., time-invariant analyst optimism on a given firm) that is not modeled in the framework but could correlate with MacroBearGap. Because β_i enters the slope multiplicatively, firm fixed effects do not remove cross-sectional variation in β_i; the pooled OLS slope on MacroBearGap instead recovers a variance-weighted average of the firm-specific products $\theta\beta_i$. The forward-looking market surprise $\beta_i(M_{t+1} - \tilde{\mu}_t)$ enters the residual, so identification requires that MacroBearGap does not predict future market returns. We examine this prediction in two steps. Section 3.5 uses a linear panel regression; Section 3.6 uses a flexible nonlinear specification that allows the effect to vary across market states.

3.5 Empirical Evidence: Forecast Error Prediction and Portfolio Performance

Guided by the framework in Section 3.4, our central empirical prediction is that stronger Macro Bear Bias should be associated with more positive subsequent base-case forecast errors. Intuitively, when analysts place excessive macro emphasis on downside scenarios, they embed too much downside macro risk into the central forecast. As a result, the base-case forecast becomes too pessimistic relative to subsequently realized returns.

To test this prediction, we estimate the following specification:

$$\text{Forecast Error}_{i,t+1} = a + b_1\, \hat{\beta}_{i,t} \times \text{MacroBearGap}_{i,t} + b_2\, M_t + \mathbf{b}' X_{i,t} + \gamma_i + u_{i,t+1}, \tag{12}$$

where $\text{MacroBearGap}_{i,t}$ is the bear–base gap in macro attention, $\hat{\beta}_{i,t}$ is the past 252-day CAPM beta, M_t is the past 252-day market excess return prior to report release, $X_{i,t}$ is a vector of time-varying firm characteristics including the stock's own past 252-day return $R_{i,t}$ and CAPM beta $\hat{\beta}_{i,t}$ before report release, and γ_i denotes firm fixed effects. We include M_t to control for the market state prevailing when the analyst issues the report. The framework predicts $b_1 > 0$: firms with stronger Macro Bear Bias should exhibit larger positive subsequent forecast errors.

For the empirical analysis, we go beyond the bear–base macro-attention gap and also examine the bull–base macro-attention gap, as well as the corresponding gaps in firm-specific attention as placebo tests:

$$\text{Forecast Error}_{i,t+1} = a + b_1\, \hat{\beta}_{i,t} \times \text{AttentionGap}_{i,t} + b_2\, M_t + b_3 \hat{\beta}_{i,t} + b_4 R_{i,t} + \gamma_i + u_{i,t+1}, \tag{13}$$

where $\text{AttentionGap}_{i,t}$ denotes the narrative attention-gap variable. This design helps distinguish Macro Bear Bias from more general differences in narrative intensity across scenarios. If forecast errors are specifically driven by Macro Bear Bias, then the relevant predictor should be the bear–base gap in macro attention, rather than scenario differences in upside narratives or in firm-level risk attention.

The regression results are reported in Table 4. Across columns (1), (5), and (6), the coefficient on the bear–base macro-attention gap is positive and statistically significant, with t-statistics of 2.5, 2.3, and 3.2, respectively. These results indicate that firms for which analysts place relatively greater macro emphasis in the bear case than in the base case subsequently exhibit larger base-case forecast errors. Under our definition of forecast error, this means that analysts set base-case forecasts that are too low for these firms. The evidence is therefore consistent with Macro Bear Bias distorting analysts' base-case expectations in a systematically pessimistic direction.[14] Additionally, because

[14] We also conduct a robustness check in which we regress forecast errors directly on attention gaps, rather than on

the predictor enters multiplicatively as $\hat{\beta}_{i,t} \times \text{MacroBearGap}_{i,t}$, the positive and significant first coefficient confirms the cross-sectional content of Eq. (10): the forecast-error response to a given macro-attention gap is larger for high-beta firms.

Table 4 Forecast Error Regressions

Note: This table reports empirical results from the following forecast-error prediction regression (13):

$$\text{Forecast Error}_{i,t+1} = a + b_1\,\hat{\beta}_{i,t} \times \text{AttentionGap}_{i,t} + b_2\,M_t + b_3\hat{\beta}_{i,t} + b_4 R_{i,t} + \gamma_i + u_{i,t+1},$$

where each narrative attention-gap variable is interacted with stock i's CAPM beta, $\hat{\beta}_{i,t}$, estimated using daily returns over the 252 trading days prior to the report release. M_t and $R_{i,t}$ denote the market excess return and stock i's excess return over the same 252-day pre-release window. Columns (1)–(4) include each beta-weighted attention gap individually. Column (5) includes the two macro beta-weighted gaps jointly. Column (6) includes all four beta-weighted attention-gap variables simultaneously. Firm fixed effects are included in all specifications. Standard errors are clustered by firm and month. t-statistics are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
Attention Gaps						
$\hat{\beta}_{i,t} \times$ (Macro Bear – Base)	0.0011**				0.0011**	0.0022***
	(2.49)				(2.29)	(3.21)
$\hat{\beta}_{i,t} \times$ (Macro Bull – Base)		0.0005			0.0001	-0.0002
		(0.75)			(0.16)	(-0.25)
$\hat{\beta}_{i,t} \times$ (Firm Bear – Base)			0.0006			0.0018**
			(1.06)			(2.16)
$\hat{\beta}_{i,t} \times$ (Firm Bull – Base)				-0.0001		-0.0005
				(-0.16)		(-0.60)
Control Variables						
M_t	-0.1449	-0.1457	-0.1461	-0.1456	-0.1449	-0.1452
	(-1.51)	(-1.52)	(-1.52)	(-1.51)	(-1.51)	(-1.52)
$\hat{\beta}_{i,t}$	-0.0392	-0.0316	-0.0280	-0.0315	-0.0390	-0.0351
	(-0.68)	(-0.55)	(-0.49)	(-0.55)	(-0.68)	(-0.61)
$R_{i,t}$	-2.5075	-2.4404	-2.2543	-2.3862	-2.5196	-2.2828
	(-0.15)	(-0.15)	(-0.13)	(-0.14)	(-0.15)	(-0.14)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43,306	43,306	43,306	43,306	43,306	43,306
Within R^2	0.0005	0.0004	0.0004	0.0003	0.0005	0.0007

By contrast, the bull–base macro-attention gap is not statistically significant, which is consistent with the view that the bias is specific to downside macro narratives rather than to scenario differentiation more broadly. The placebo tests based on firm-specific attention provide weaker

the interaction between $\hat{\beta}_{i,t}$ and attention gaps. The results, reported in Table A.3, are very similar to those in Table 4. This is expected because the vast majority of estimates satisfy $\hat{\beta}_{i,t} > 0$, as shown in Table 2.

and less consistent evidence. Taken together, the results suggest that the predictive content comes primarily from the downside macro component of analyst narratives, in line with the mechanism of Macro Bear Bias.

The univariate-to-multivariate comparison reveals a suppression effect. The Macro Bear−Base coefficient rises from 0.0011 ($t = 2.5$) in column (1) to 0.0022 ($t = 3.2$) in column (6). Because the macro and firm bear–base gaps are mechanically negatively correlated—the compositional constraint $A^{\text{macro}} + A^{\text{firm}} + A^{\text{industry}} \leq 100$ means more macro bear-case attention implies less firm-specific attention—the univariate estimate is attenuated by the omitted Firm Bear−Base gap. Column (6) recovers the partial effect and is our preferred linear specification.

We next examine whether correcting for Macro Bear Bias improves the out-of-sample cross-sectional ranking of stock returns. The key idea is that the analyst's base-case forecast may contain useful firm-level information but may also be distorted by the predictable pessimism induced by macro-heavy bear narratives. Rather than replacing the analyst forecast with a purely statistical return forecast, we therefore estimate the component of the analyst's forecast error that is predictable from Macro Bear Bias and use it to adjust the original base-case forecast. This exercise asks whether the narrative-based bias documented above contains return-relevant information beyond the analyst's raw valuation signal.

To implement this test, we use the specification in Column (6) of Table 4, but exclude the time-varying firm characteristics $R_{i,t}$ and $\hat{\beta}_{i,t}$. This restriction ensures that out-of-sample predictability comes from the interaction between firms' systematic exposure and the narrative attention gaps, rather than from standalone sorts on past returns or beta. We estimate the following regression separately for each firm:

$$
\begin{aligned}
\text{Forecast Error}_{i,t+1} = a_i &+ b_{i,1}\,\hat{\beta}_{i,t} \times \text{MacroBearGap}_{i,t} + b_{i,2}\,\hat{\beta}_{i,t} \times \text{MacroBullGap}_{i,t} \\
&+ b_{i,3}\,\hat{\beta}_{i,t} \times \text{FirmBearGap}_{i,t} + b_{i,4}\,\hat{\beta}_{i,t} \times \text{FirmBullGap}_{i,t} \\
&+ b_{i,5}\,M_t + u_{i,t+1}, \qquad (14)
\end{aligned}
$$

where $\text{MacroBearGap}_{i,t}$ is the bear–base macro-attention gap, $\text{MacroBullGap}i, t$ is the bull–base macro-attention gap, FirmBearGapi,t is the bear–base firm-attention gap, and FirmBullGapi,t is the bull–base firm-attention gap. Because (14) is estimated for each firm individually rather than as a panel regression, all slope coefficients carry a firm index i in contrast to (13).

We construct firm-level out-of-sample forecasts using a rolling training window of 60 observations. At each forecast date, the training sample consists of the most recent 60 reports for the same firm whose outcomes are observable, meaning that the reports were released at least one year before the test date.[15] Each firm is therefore forecasted using only its own historical relation between beta-weighted attention gaps, the aggregate market state, and subsequent forecast errors.

We then evaluate portfolio performance using two sorting signals. The first signal is the Bias-adjusted Forecast, defined as the analyst's base-case forecast return plus the out-of-sample forecast error predicted from (14). This adjusted measure corrects the analyst's base-case forecast for the predictable bias component associated with Macro Bear Bias. The second signal is the analyst's raw base-case forecast return, labeled Base Forecast, which serves as the benchmark for evaluating portfolio performance before correcting for Macro Bear Bias.

For each signal, we form long-short portfolios using an annual rebalancing strategy implemented separately at the end of each calendar month. Specifically, for each calendar month, we construct a portfolio that is rebalanced annually in that month, yielding 12 distinct annually rebalanced strategies. To focus on the part of the cross-section where the signal is strongest, we go long stocks in the top decile of the sorting signal and short stocks in the bottom decile. We construct both value-weighted (VW) and equal-weighted (EW) portfolios. We then average performance across the 12 calendar-month implementations and report the annualized Sharpe ratio, average monthly return, and CAPM alpha.

Table 5 reports portfolio performance for the two sorting signals. The Base Forecast portfolio, which sorts stocks solely on analysts' raw base-case forecast returns, does not deliver significant performance. The VW long-short portfolio has a negative average return and a negative CAPM

[15] Because the forecast error is defined as the difference between the realized one-year-ahead return and the analyst's base-case forecast return, we use only training observations released at least one year before the test report. This timing convention avoids using realized returns that would not have been observable at the forecast date. For example, to predict the forecast error for a report released on June 1, 2021, the training sample consists of reports released before June 1, 2020.

alpha, and the EW portfolio is similarly weak. Thus, analysts' raw base-case forecasts do not provide meaningful out-of-sample cross-sectional return predictability.

By contrast, the Bias-adjusted Forecast delivers economically large and statistically significant abnormal returns in the VW portfolio. The VW long-short portfolio earns an average monthly return of 1.52% (t = 2.31) and a CAPM alpha of 1.36% (t = 2.00), with an annualized Sharpe ratio of 0.83. The EW portfolio also earns positive returns and positive CAPM alpha in contrast to the Base Forecast. These results indicate that the predictable forecast errors associated with Macro Bear Bias contain substantial return-relevant information. The stronger VW results further suggest that this return-relevant bias is concentrated among larger, more investable firms, where analyst reports are more likely to be read by investors and incorporated into prices.

These findings provide a direct economic interpretation of Macro Bear Bias. The bias is not merely a textual regularity in analyst narratives. It predicts systematic errors in analysts' numerical base-case forecasts, and those forecast errors contain information about subsequent stock returns. Stocks whose base-case forecasts are too pessimistic after accounting for macro-heavy bear narratives subsequently outperform stocks whose base-case forecasts are less favorably adjusted. Thus, the bias-adjusted signal recovers return-relevant information obscured in raw analyst forecasts, suggesting that narrative misattribution can spill over into valuation and affect the cross-section of expected returns.

3.6 Nonlinear Forecast Error Prediction and Portfolio Performance

The linear correction in Section 3.5 generates portfolio returns in the predicted direction, but the economic magnitudes are modest. This pattern suggests that a single linear slope captures only part of the return-relevant information contained in Macro Bear Bias. The conceptual framework explains why: Macro Bear Bias reflects a failure to condition appropriately on aggregate market states. The same bear–base macro-attention gap may therefore imply different forecast distortions depending on whether market conditions are weak or strong. In weak market states, a given amount of downside macro emphasis may translate into one degree of pessimism in the base-case forecast; in stronger market states, the same narrative gap may imply a different distortion. More generally,

Table 5 Portfolio Performance of Linear Forecast Error Prediction

Note: This table reports monthly long-short portfolio performance for Bias-adjusted Forecast and Base Forecast. For each signal, the portfolio goes long stocks in the top decile and short stocks in the bottom decile. Bias-adjusted Forecast sorts on the analyst's base-case forecast return plus the out-of-sample forecast error predicted from (14). Base Forecast sorts on the analyst's raw base-case forecast return and serves as the benchmark before bias adjustment. Portfolios are rebalanced annually across 12 calendar-month implementations. Annual SR is the annualized Sharpe ratio. Avg Return is the average monthly return of the long-short portfolio. CAPM α is the intercept from a CAPM time-series regression. VW and EW denote value-weighted and equal-weighted portfolios, respectively. t-statistics are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Return Forecast	VW			EW		
	Annual SR	Avg Return	CAPM α	Annual SR	Avg Return	CAPM α
Bias-adjusted	0.8348	0.0152**	0.0136**	0.4864	0.0080	0.0064
		(2.31)	(2.00)		(1.35)	(1.05)
Base	-0.1330	-0.0019	-0.0043	-0.0424	-0.0005	-0.0034
		(-0.37)	(-0.86)		(-0.12)	(-0.79)

the relationship between subjective belief objects and realized outcomes should be expected to be nonlinear by construction. Dew-Becker, Giglio, and Molavi (2025) establish this point formally, showing that under Bayesian filtering—including under behavioral departures such as misspecified priors—belief dynamics are generically nonlinear and state-dependent. The effect of Macro Bear Bias on forecast errors is therefore likely governed by nonlinear interactions between narrative-based attention measures and market conditions, rather than by their separate linear effects.

To capture these nonlinearities, we follow Kelly, Malamud, and Zhou (2024, 2022) and use the Random Fourier Features (RFF) approach of Rahimi and Recht (2007, 2008) to predict out-of-sample forecast error. The raw predictors are the macro bear–base gap, macro bull–base gap, firm bear–base gap, firm bull–base gap, and the cumulative market excess return over the 252 trading days prior to the report release. Let $G_{i,t}$ denote the 5×1 predictor vector for a given firm-date observation:

$$G_{i,t} = \left[\text{MacroBearGap}_{i,t}, \text{MacroBullGap}_{i,t}, \text{FirmBearGap}_{i,t}, \text{FirmBullGap}_{i,t}, M_t\right]. \tag{15}$$

Compared to the linear specification in (14), we exclude the CAPM beta $\hat{\beta}_{i,t}$ from $G_{i,t}$. This choice makes the nonlinear specification deliberately conservative: firm-specific information enters only through the analyst's own narrative attention measures.

The RFF procedure maps $G_{i,t}$ into nonlinear features,

$$S_{i,t,p} = \left[\sin(\gamma\omega_p' G_{i,t}), \cos(\gamma\omega_p' G_{i,t})\right]', \qquad \omega_p \sim i.i.d.\, N(0, I), \tag{16}$$

where each ω_p defines a random projection of the original predictors, which is then transformed using sine and cosine functions. This transformation generates a rich set of nonlinear combinations of attention gaps and market states. As emphasized by Kelly, Malamud, and Zhou (2024), RFF can be interpreted as a wide two-layer neural network with random first-layer weights and estimated second-layer coefficients. Its advantage in our setting is that it introduces flexible nonlinearities while remaining computationally tractable and straightforward to evaluate out of sample. We set the bandwidth parameter to $\gamma = 2$ and draw up to $P_{\max} = 180{,}000$ random projections for each firm. To study the role of model complexity, we estimate forecasts over a grid of model sizes P ranging from 6 to 180,000, using nested subsets of these projections and holding the ridge penalty fixed at $z = 1{,}000$.

The corresponding linear benchmark of the RFF regression is

$$\begin{aligned}
\text{Forecast Error}_{i,t+1} = {} & a_i + b_{i,1}\text{MacroBearGap}_{i,t} + b_{i,2}\text{MacroBullGap}_{i,t} \\
& + b_{i,3}\text{FirmBearGap}_{i,t} + b_{i,4}\text{FirmBullGap}_{i,t} \\
& + b_{i,5}\, M_t + u_{i,t+1},
\end{aligned} \tag{17}$$

which feeds the same five predictors $G_{i,t}$ from Eq. (15) into a linear specification, and therefore omits the $\hat{\beta}_{i,t}$ interaction used in the linear out-of-sample model of Eq. (14). Holding the predictor set fixed at $G_{i,t}$, the comparison between this linear benchmark and the RFF forecast isolates the incremental contribution of nonlinearity.

Following the out-of-sample design in Section 3.5, we estimate firm-level RFF forecasts and the corresponding linear benchmark in (17) using a 60-observation rolling window. To avoid look-ahead bias, each training observation must have been released at least one year before the test date. Forecasts are estimated separately by firm, using the firm's attention-gap variables and the aggregate market state.

We then examine whether the nonlinear forecast-error correction improves the cross-sectional portfolio performance more than the linear forecast-error correction. In addition to the analyst's raw base-case forecast return (denoted as the Base Forecast), we construct a Nonlinearity-adjusted Forecast, defined as the base-case forecast return plus the out-of-sample forecast error predicted by the RFF model. To compare the performance improvement contributed by nonlinearity, we also construct the portfolio by sorting on the corresponding Linearly-adjusted Forecast, which is the base-case forecast return plus the out-of-sample forecast error predicted by (17).

Portfolio construction follows the same procedure as in Section 3.5. For each signal, we form top-minus-bottom decile long-short portfolios, construct both equal-weighted (EW) and value-weighted (VW) versions, implement the annual rebalancing strategy separately across the 12 calendar months, and report average performance across these implementations using the annualized Sharpe ratio, average monthly return, and CAPM alpha.

Table 6 shows a sharp contrast across the three signals. The Base Forecast signal, which sorts on the analyst's base-case forecast return alone, generates no abnormal return. [16] By contrast, the Nonlinearity-adjusted Forecast delivers economically large and statistically significant abnormal returns. The VW long-short portfolio earns an average monthly return of 2.02% (t = 3.04) and a CAPM alpha of 1.91% (t = 2.84), with an annualized Sharpe ratio of 1.10. The EW portfolio earns 1.80% per month (t = 2.74) with a CAPM alpha of 1.74% (t = 2.55) and an annualized Sharpe ratio of 0.99. In comparison, the Linearly-adjusted Forecast produces substantially weaker performance: the average monthly return is 1.25% (t = 1.80) in the VW portfolio and 0.94% (t = 1.43) in the EW portfolio, while the corresponding CAPM alpha is statistically insignificant in both cases.

The nonlinear adjustment also delivers economically meaningful gains relative to the linear

[16] The numbers are slightly different from those for the Base Forecast portfolio performance in Table 5 because Table 5 removes observations with no $\hat{\beta}_{i,t}$.

Table 6 Portfolio Performance of Nonlinear Forecast Error Prediction

Note: This table reports monthly long-short portfolio performance for Nonlinearity-adjusted Forecast, Linearly-adjusted Forecast, and Base Forecast. Each portfolio goes long the top decile and short the bottom decile. Nonlinearity-adjusted Forecast (Nonlinear) sorts on the analyst's base-case forecast return plus the RFF-predicted forecast-error correction. Linearly-adjusted Forecast (Linear) sorts on the analyst's base-case forecast return plus the out-of-sample forecast error predicted from (17). Base Forecast sorts on the analyst's raw base-case forecast return and serves as the benchmark before bias adjustment. RFF specification uses bandwidth $\gamma = 2$, shrinkage $z = 1{,}000$, and complexity $c = 3{,}000$. Portfolios are rebalanced annually across 12 calendar-month implementations. Annual SR is the annualized Sharpe ratio. Avg Return is the average monthly return of the long-short portfolio. CAPM α is the intercept from a CAPM time-series regression. VW and EW denote value-weighted and equal-weighted portfolios, respectively. t-statistics are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Return Forecast	VW				EW			
	Annual SR	Avg Return	CAPM α	vs. Linear	Annual SR	Avg Return	CAPM α	vs. Linear
Nonlinear	1.0985	0.0202***	0.0191***	0.0111**	0.9901	0.0180***	0.0174**	0.0110**
		(3.04)	(2.84)	(2.51)		(2.74)	(2.55)	(2.46)
Linear	0.6518	0.0125*	0.0107		0.5168	0.0094	0.0075	
		(1.80)	(1.53)			(1.43)	(1.13)	
Base	-0.0663	-0.0010	-0.0030		-0.0198	-0.0003	-0.0028	
		(-0.18)	(-0.61)			(-0.05)	(-0.63)	

adjustment. The nonlinearity-adjusted portfolio outperforms the linearly-adjusted portfolio significantly, with a monthly alpha of 1.11% and t-statistics of 2.51 in the VW portfolio and a monthly alpha of 1.10% and t-statistics of 2.46 in the EW portfolio. [17] These results indicate that the return-relevant component of Macro Bear Bias is not fully captured by a linear correction. Allowing forecast-error predictions to vary nonlinearly with narrative gaps and market conditions substantially improves out-of-sample portfolio performance.

Figure 9 shows that portfolio performance improves as model complexity increases. Model complexity c is defined as the number of RFF features divided by the 60-observation training window, so higher values of c correspond to a richer and more flexible nonlinear feature space. Consistent with Kelly, Malamud, and Zhou (2024, 2022), the RFF Predicted Forecast Error signal, shown in the right column, becomes more powerful as c rises: the annualized Sharpe ratio increases steadily, and both the CAPM alpha and the corresponding t-statistics display a similar

[17] The nonlinearity-adjusted portfolio also outperforms the bias-adjusted portfolio in Table 5. The VW monthly alpha is 0.9% with a t-statistic of 1.92. The EW monthly alpha is 1.15% with a t-statistic of 2.53.

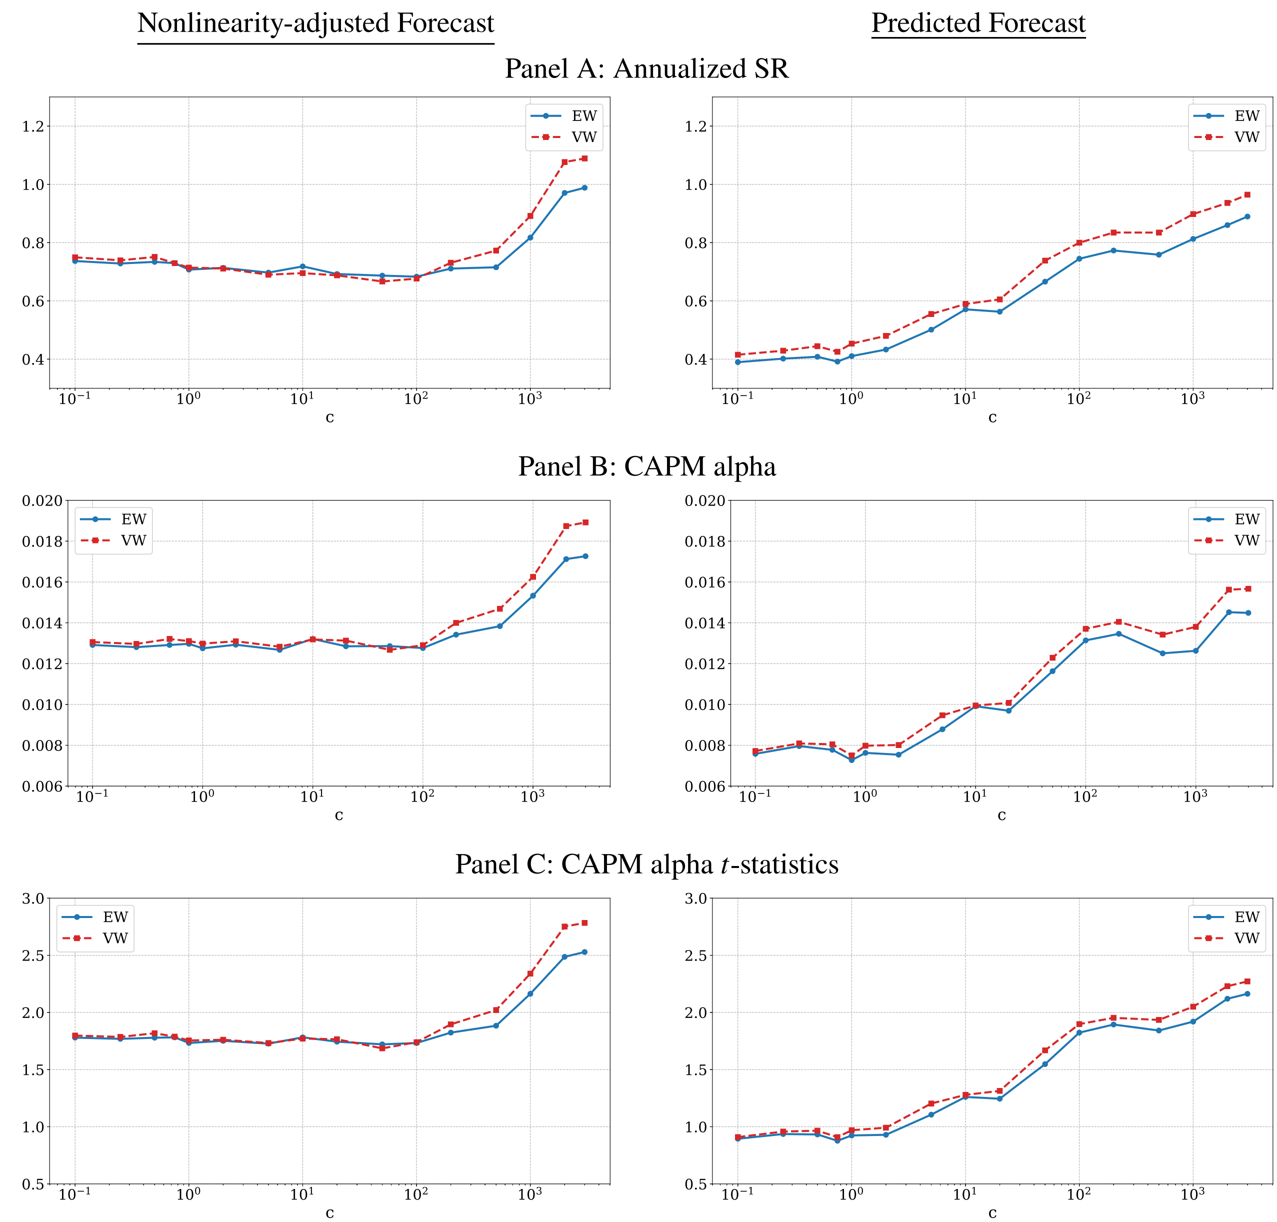

Figure 9 Out-of-sample Long-short Portfolio Performance.

Note: This figure shows the out-of-sample long-short portfolio performance when the model complexity c increases. The training window is 60 observations for each firm, and RFF count P ranges from 6 to 180,000 with bandwidth $\gamma = 2$ and shrinkage $z = 1000$.

upward pattern. The Nonlinearity-adjusted Forecast, shown in the left column, exhibits the same pattern. Its annualized Sharpe ratio and CAPM alpha are weak at low levels of c, indicating that a sparse RFF feature space is unable to capture the relevant interactions among narrative gaps and market conditions. As c increases, however, the RFF approximation becomes richer and recovers more of the state-dependent predictive content embedded in analyst narratives. Near $c = 3{,}000$, the VW portfolio reaches an annualized Sharpe ratio of approximately 1.1, and the CAPM alpha has a t-statistic above 2.7.

Taken together, these results show that Macro Bear Bias has economically meaningful predictive power for both forecast errors and cross-sectional returns, but that this predictive content is recovered more effectively through a nonlinear specification. A linear correction compresses state-dependent forecast distortions into an average slope and therefore captures only limited portfolio performance. By contrast, the RFF specification allows the forecast-error correction to vary flexibly with the interactions among macro-heavy bear narratives, firm-level narrative content, and market conditions. The resulting improvement in out-of-sample performance indicates that the pricing content of Macro Bear Bias depends not only on whether analysts over-attribute downside risk to macro forces, but also on the states in which this misattribution occurs.

4 Mechanisms

We test two candidate mechanisms for Macro Bear Bias. The first is cognitive: the availability heuristic predicts that analysts anchor bear-case narratives to salient macro-crisis templates because such templates are vivid, widely shared, and easy to retrieve. The second is strategic: career concerns predict that analysts emphasize macro risk in bear scenarios because doing so is professionally safer than making falsifiable firm-specific claims. The evidence supports availability; career concerns receive little support.

4.1 Availability Heuristic

The availability heuristic (Tversky and Kahneman, 1973, 1974) holds that people assess the probability of an event by the ease with which instances come to mind: events that are vivid, widely shared, or recently experienced are recalled more readily, inflating their judged likelihood. Building on this canonical heuristic, modern memory-based theories of belief formation (Bordalo, Gennaioli, and Shleifer, 2020; Bordalo et al., 2023) formalize how cues from the current environment trigger retrieval of similar past experiences, generating beliefs that over-weight vivid and recently activated memories. Our setting brings these mechanisms into a high-stakes professional context and identifies the institutional and peer-network channels through which crisis narrative persists beyond individual recall.

4.1.1 Identification

If Macro Bear Bias arises through this channel, then macro emphasis in bear narratives should respond systematically to factors that heighten the cognitive accessibility of macro-crisis templates. We test this via four predictions: one at the topic level using news coverage shocks (Table 7), and three cross-sectional predictions at the analyst and industry level (Table 8).

Topic salience. The availability heuristic predicts that a sudden burst of news coverage of a macro topic should make that topic cognitively accessible, causing analysts to draw on it more readily when constructing bear-case narratives.

We test this prediction at the topic–month level by regressing bear-case analyst attention to topic k at time t on the WSJ topic shocks, its interaction with WSJ sentiment, and control WSJ sentiment and lagged WSJ attentions:

$$\begin{aligned}\text{BearAttention}_{t,k} &= \beta_1 \hat{\varepsilon}^{\text{WSJAttention}}_{t-\tau,k} + \beta_2 \hat{\varepsilon}^{\text{WSJAttention}}_{t-\tau,k} \times \text{WSJSentiment}_{t-\tau,k} \\ &\quad + \beta_3 \text{WSJSentiment}_{t-\tau,k} + \sum_{j=1}^{3} \phi_j \text{WSJAttention}_{t-j,k} + \alpha_t + u_{t,k}, \qquad (18)\end{aligned}$$

where $\hat{\varepsilon}^{\text{WSJAttention}}_{t-\tau,k}$ denotes the autoregressive (AR) residual of WSJ attention to topic k at lag τ, capturing the unexpected component of topic attention.[18] $\text{WSJSentiment}_{t-\tau,k}$ is the average WSJ news sentiment for topic k at lag τ, and α_t is the time fixed effect. The three lags of $\text{WSJAttention}_{t-j,k}$ control for persistence in coverage so that $\hat{\varepsilon}^{\text{WSJAttention}}_{t-\tau,k}$ isolates only the transient, unexpected increase in attention. The specification nests two predictions: a level effect ($\beta_1 > 0$),

[18]We use an AR(3) model for each topic, since the Akaike information criterion (AIC) selects AR(3) for 40 of the 52 topics.

whereby coverage surprises raise analyst bear-case attention, and an asymmetry effect ($\beta_2 < 0$), whereby negative-sentiment coverage amplifies this response.

Table 7 reports results for $\tau \in \{1, 2, 3, 4\}$. The level effect is positive and statistically significant at short lags: a one-unit attention shock raises bear-case macro attention by 0.064 percentage points at $\tau = 1$ ($t = 1.85$) and 0.052 points at $\tau = 2$ ($t = 1.72$), declining to an insignificant 0.042 and 0.021 at $\tau = 3$ and $\tau = 4$. The interaction coefficient is negative and statistically significant through $\tau = 3$: $\hat{\beta}_2 = -0.188$, -0.237, and -0.186 at $\tau = 1$, 2, and 3 (t-statistics -1.81, -2.43, -2.10), attenuating to an insignificant -0.038 at $\tau = 4$. Baseline sentiment alone is economically negligible and statistically insignificant across all columns, indicating that the direction of coverage matters only in amplifying surprises, not independently.

These results are consistent with the availability heuristic. Unexpected increases in macro news coverage shift bear-case analyst attention toward the newly salient topic, and this shift is amplified when coverage is negative—precisely the threatening, crisis-type signal identified by the availability literature as most cognitively accessible (Tversky and Kahneman, 1973). The attenuation of both effects at $\tau \geq 3$ mirrors the transient nature of news salience: availability fades as coverage recedes.

Industry-crisis salience. A direct cross-sectional prediction of the availability heuristic is that crisis templates should be deployed differentially across industries, depending on the perceived relevance of a given macro crisis to each sector. The 2008 Global Financial Crisis was widely framed as a systematic-risk event centered on financial contagion, so the availability of macro-downside narratives should have increased most for firms in high-systematic-exposure sectors (e.g., financials, energy, and materials) relative to defensive sectors such as utilities, consumer staples, and healthcare. We test this by interacting a post-GFC indicator with a high-systematic-industry dummy.

Column (1) of Table 8 reports the results. The interaction is economically and statistically significant: after the GFC, MacroBearGap increases by 4.9 percentage points more for firms in high-systematic sectors than for other firms ($t = 2.61$). Combining the main effect and the interaction, high-systematic sectors had on average -8.9 percentage-point *lower* MacroBearGap

Table 7 Bear-Case Macro Attention and Topic Salience

Note: This table reports panel regressions of bear-case analyst attention to topic k at time t on the autoregressive (AR) residual of WSJ news coverage (attention shock), its interaction with WSJ news sentiment, the sentiment level, and three lags of WSJ topic attention; see Equation (18). Columns correspond to lag $\tau \in \{1, 2, 3, 4\}$. All specifications include time fixed effects. t-statistics in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$
WSJ Attention Shock$_{t-\tau,k}$	0.0642*	0.0524*	0.0417	0.0207
	(1.85)	(1.72)	(1.54)	(1.50)
WSJ Attention Shock$_{t-\tau,k}$ × WSJ Sentiment$_{t-\tau,k}$	-0.1876*	-0.2374**	-0.1855**	-0.0377
	(-1.81)	(-2.43)	(-2.10)	(-0.55)
WSJ Sentiment$_{t-\tau,k}$	0.0001	0.0000	0.0001	0.0001
	(0.72)	(0.24)	(0.56)	(0.64)
WSJ Attention$_{t-1,k}$	0.3085***	0.2998***	0.3004***	0.3011***
	(11.02)	(10.87)	(10.48)	(10.14)
WSJ Attention$_{t-2,k}$	0.2835***	0.2868***	0.2862***	0.2867***
	(10.33)	(10.44)	(10.20)	(10.16)
WSJ Attention$_{t-3,k}$	0.3826***	0.3876***	0.3873***	0.3865***
	(11.94)	(12.58)	(12.12)	(11.68)
Time FE	Yes	Yes	Yes	Yes
Observations	11,440	11,388	11,336	11,284
R^2	88.90%	88.96%	88.95%	88.93%

than other sectors before the GFC; post-GFC, the differential narrows by 4.9 points to about −4.0 percentage points before any common post-GFC effect is added. The crisis therefore did not flip high-systematic sectors to higher MacroBearGap; it narrowed a pre-existing negative gap, which is the within-firm pattern the availability heuristic predicts when the crisis differentially raises the cognitive accessibility of macro-downside templates for these sectors.[19]

Analyst crisis experience. The crisis experience tests require an analyst-level identifier and therefore restrict the sample to the *analyst-matched subsample*. We match each Morgan Stanley report to IBES via the analyst's name (cleaned of titles and suffixes) and the covered firm's name within a 90-day window around the report date; a report enters the matched subsample if at least one IBES analyst record satisfies the join, and we retain the longest-tenured matching analyst when

[19] We also test whether the post-COVID increase in MacroBearGap is concentrated in COVID-sensitive industries (airlines, hotels, retail). It is not: the COVID interaction with a COVID-sensitive industry indicator is essentially zero (−0.06, $t = -0.06$). This contrast suggests that COVID activated a universal "shutdown" narrative across all sectors, whereas the GFC activated a sector-specific "financial contagion" template.

multiple candidates appear. The matched subsample contains 17,257 report observations, a 45% reduction from the full panel of 31,445 observations used in column (1).[20]

If availability operates through personal exposure to crisis episodes, then analysts active during the GFC should carry stronger macro-crisis anchors. We construct two GFC-experience indicators: one for analysts with a record at Morgan Stanley before March 2009, and one for analysts with a record at *any* brokerage before that date, each proxied by the analyst's earliest report date in the respective database. This pair of indicators distinguishes two channels: the MS-specific measure captures exposure to the firm's crisis-era narrative conventions, while the profession-wide measure captures generic crisis memory.

Columns (2) and (3) of Table 8 report the results. Analysts who experienced the GFC at Morgan Stanley exhibit a 1.2-point higher MacroBearGap ($t = 1.65$); analysts who experienced the GFC elsewhere show no discernible effect. The divergence is consistent with availability operating through firm-specific narrative templates: the pattern is more consistent with exposure to a particular institution's crisis-era framing conventions than with a generic crisis-memory proxy, though it does not fully separate institutional template exposure from selection into Morgan Stanley or differences in the matched analyst populations.[21]

Sector-level narrative herding. A final prediction follows from the shared information environment of sector analysts. If analysts covering firms in the same sector draw on the same pool of salient crisis templates—through common media exposure, industry conferences, and peer discussion—their MacroBearGap values should covary within sector-quarters, even after absorbing firm fixed effects. We construct a leave-one-out (LOO) sector–quarter average of MacroBearGap and test whether an analyst's own MacroBearGap covaries with this peer average.

[20]The reduction is driven primarily by IBES coverage gaps (about 32%) and ambiguous name matches that the join rule discards conservatively (about 13%). Matched and unmatched reports differ on observable dimensions of analyst tenure and firm coverage: matched reports skew toward larger-cap, longer-tenured analyst coverage, which is exactly the segment in which the GFC-experience and IB-relationship tests have the most variation.

[21]We also examine whether the MS-specific GFC effect reflects generic MS tenure rather than crisis-specific experience. In a horse-race regression including both the GFC-experience indicator and years of MS tenure, the GFC coefficient remains positive (+1.27) while the tenure coefficient is essentially zero (-0.02, $t = -0.14$). The point estimate is directionally consistent but imprecisely estimated in the horse-race specification; we therefore read the standalone result as suggestive rather than definitive.

Table 8 Availability Heuristic: Determinants of Macro Bear Bias

Note: This table reports panel regressions of MacroBearGap (bear minus base macro attention, 0–100 scale) on proxies for the availability heuristic. Column (1) interacts a post-GFC indicator with a high-systematic-industry dummy (financials, energy, materials). Columns (2)–(3) test whether analysts who experienced the GFC at Morgan Stanley (col 2) or at any brokerage (col 3) exhibit larger MacroBearGap. Column (4) tests sector-level narrative herding using a leave-one-out sector–quarter average of MacroBearGap. Columns (2)–(3) are restricted to the analyst-matched subsample. All specifications include firm and year fixed effects. Standard errors are clustered by firm. t-statistics in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	MacroGap			
	(1)	(2)	(3)	(4)
Post-GFC × High Systematic	4.9229***			
	(2.61)			
High Systematic	−8.9346***			
	(−4.30)			
GFC Experienced (MS)		1.1544*		
		(1.65)		
GFC Experienced (Prof.)			0.1594	
			(0.25)	
LOO Sector MacroGap				0.0828*
				(1.75)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	31,445	17,257	17,257	30,979
R^2	0.0008	0.0007	0.0000	0.0010

Column (4) of Table 8 is consistent with this prediction. The leave-one-out sector MacroBearGap positively predicts the analyst's own MacroBearGap (0.083, $t = 1.75$), suggesting that analysts covering firms in the same sector–quarter converge on similar levels of macro emphasis beyond what firm fixed effects explain.

Taken together, the four tests in Tables 7 and 8 are consistent with an availability-heuristic channel behind Macro Bear Bias. Bear-case macro attention rises in response to negative news coverage surprises, is larger in high-systematic sectors following the GFC, is elevated for analysts who experienced the crisis at Morgan Stanley, and co-moves within sector–quarter peer groups—a pattern consistent with shared narrative templates rather than independent judgment.

4.1.2 When Macro Bear Bias Is Realized

The availability heuristic explains why analysts default to macro-heavy bear narratives. It does not identify when, in the cross-section of reports, the bias is most actively invoked. A simple prediction follows from how the bear case is constructed: analysts need a coherent explanation for why the stock could underperform, and the explanation they reach for should depend on what the base case has already established. The bear–base macro-attention gap should therefore be smaller when the base case already supplies firm-specific pessimism, and larger when it does not.

Simply put: an analyst whose base-case price target exceeds the consensus has committed to a bullish firm-specific thesis. When the Morgan Stanley Risk-Reward framework requires this analyst to construct a bear case, the most cognitively available source of downside is external—macroeconomic deterioration, recession, or broad multiple compression—because the firm-specific view is favorable. An analyst whose base-case target falls below the consensus has already embedded firm-specific problems into the base-case forecast (execution risk, competitive threats, margin pressure); idiosyncratic downside is already supplied, and macro narratives are less salient.

We operationalize the state of the base-case forecast with two indicators: $\mathbf{1}\{P_{MS}^{\text{base}} > P^{\text{cons}}\}$ flags reports in which the Morgan Stanley base-case target exceeds the IBES consensus, and $\mathbf{1}\{P_{MS}^{\text{base}} < P^{\text{cons}}\}$ flags reports in which it falls below the consensus. We match each Morgan Stanley report to the most recent IBES consensus price target for the same firm within 90 days, yielding a matched sample of 26,858 observations.

Table 9 reports the results. Columns (1)–(2) report the two indicators separately. When the base-case target is above consensus ($\mathbf{1}\{P_{MS}^{\text{base}} > P^{\text{cons}}\} = 1$), MacroBearGap is 1.05 percentage points above the sample mean ($t = 4.47$); when it is below consensus, MacroBearGap is 1.09 percentage points below the mean ($t = -4.66$). Columns (3)–(4) replace the discrete indicators with their continuous analogs—the signed consensus deviation $(P_{MS}^{\text{base}} - P^{\text{cons}})/P^{\text{cons}}$ and its absolute value. Neither continuous measure is statistically significant, and the within-R^2 collapses to essentially zero: the relationship is not a smooth slope through consensus but a discrete contrast across regimes of the base-case forecast.

Table 9 When Macro Bear Bias Is Realized

Note: This table reports panel regressions of MacroBearGap (bear minus base macro attention, 0–100 scale) on indicators and continuous measures of the Morgan Stanley base-case target's position relative to the IBES consensus. Columns (1)–(2) use the discrete indicators $\mathbf{1}\{P_{MS}^{\text{base}} > P^{\text{cons}}\}$ and $\mathbf{1}\{P_{MS}^{\text{base}} < P^{\text{cons}}\}$ separately. Columns (3)–(4) use the signed continuous deviation $(P_{MS}^{\text{base}} - P^{\text{cons}})/P^{\text{cons}}$ and its absolute value.The sample is restricted to observations with an available IBES consensus match within 90 days. All specifications include firm and year fixed effects. Standard errors are clustered by firm. t-statistics in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	MacroGap			
	(1)	(2)	(3)	(4)
$\mathbf{1}\{P_{MS}^{\text{base}} > P^{\text{cons}}\}$	1.0491***			
	(4.47)			
$\mathbf{1}\{P_{MS}^{\text{base}} < P^{\text{cons}}\}$		−1.0938***		
		(−4.66)		
$(P_{MS}^{\text{base}} - P^{\text{cons}})/P^{\text{cons}}$			0.0865	
			(1.34)	
$\lvert(P_{MS}^{\text{base}} - P^{\text{cons}})/P^{\text{cons}}\rvert$				0.0428
				(0.61)
Year FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Observations	26,858	26,858	26,380	26,380
R^2	0.0012	0.0013	0.0000	0.0000

Taken together, the four columns identify *when* Macro Bear Bias is realized. When the base case is already bearish, firm-specific downside is supplied internally and the bear case is less macro-heavy: MacroBearGap falls below the sample mean. When the base case is at or above consensus, no firm-specific downside is supplied, macro narratives carry the bear case, and MacroBearGap stays at or above the mean. The continuous specifications confirm that the cross-section is a regime contrast, not a slope: averaging across the at- and above-consensus regimes washes out the contrast that the discrete indicators recover. The bias is therefore state-contingent—realized in the base case in which the availability heuristic has no firm-specific competitor for the downside narrative.

One might object that this pattern is mechanical: a bullish analyst has "used up" the firm-specific upside and must turn to macro factors for the downside. But firm-specific downside is always available—execution risk, competitive entry, regulatory changes, management turnover, product delays—regardless of the base-case outlook. The systematic tendency to emphasize macro over these alternatives, conditional on base-case direction, is itself the bias. Moreover, if the pattern were purely mechanical, bearish analysts should exhibit the mirror image: they should need macro factors in their *bull* case. But as documented in Section 2.4, the macro attention gap in bull narratives does not predict forecast errors. The asymmetry is bear-case-specific, consistent with a psychological mechanism rather than a mechanical artifact.

4.2 Career Concerns

A natural competing explanation for Macro Bear Bias is strategic rather than cognitive. Sell-side analysts face asymmetric reputational costs: firm-specific errors are falsifiable and attributable, while consensus-correlated errors provide reputational cover (Hong and Kubik, 2003). This asymmetry could generate a "cover your asset" (CYA) heuristic: when constructing downside scenarios, analysts may emphasize macro risks because macro-based bear narratives are difficult to falsify, relationally neutral toward management, and consensus-aligned. Under this view, Macro Bear Bias reflects an incentive-compatible communication strategy rather than a cognitive distortion.

If career incentives drive macro emphasis, three predictions follow that our proxies can test: (i) junior analysts, who face greater reputational risk from differentiated misses, should exhibit larger MacroBearGap; (ii) firms with investment-banking relationships with Morgan Stanley should show larger MacroBearGap, since firm-specific bear narratives are relationally costly; and (iii) analyst identity should explain a substantial share of MacroBearGap variation. The data provide little support for (i) and (ii) and only partial support for (iii).

Appendix Table A.5 reports regressions of MacroBearGap on analyst experience, measured both as Morgan Stanley tenure and as years of industry experience across all brokerages (using the full IBES price-target detail file to identify each analyst's career start at any firm). Career concerns predict that junior analysts exhibit larger MacroBearGap; the data contradict this. The MS-tenure junior indicator enters with a *negative* coefficient (-0.69, $t = -1.66$): analysts new to Morgan Stanley show *less* macro emphasis in bear narratives, not more. The profession-wide junior indicator is statistically indistinguishable from zero. The negative MS-tenure result is more naturally read through the availability lens, with the caveat that availability in our setting operates through both personal recall of crisis episodes and repeated exposure to firm-specific narrative templates. Analysts new to Morgan Stanley have not yet been steeped in the firm's macro-heavy bear-case template; as they accumulate tenure, the template becomes cognitively more accessible. This is the institutional-repetition leg of availability, not personal-memory availability in the original Tversky–Kahneman sense. Our tests cannot discriminate between the two within a single-brokerage sample, and we do not claim to: the GFC-at-MS finding in Table 8 is consistent with both personal crisis memory and institution-specific socialization into the firm's crisis-era framing. We read both legs as instances of the availability channel and discuss the limits of this identification in Section 5.

If career concerns operate through relational incentives, then analysts covering firms with active Morgan Stanley investment-banking relationships should show larger MacroBearGap, since firm-specific bear narratives risk antagonizing management during the advisory relationship. We construct an IB-relationship indicator equal to one if Morgan Stanley served as bookrunner or lead manager on a new issuance for the covered firm within a specified lookback window, using the LSEG SDC New Issues database. Table 10, columns (1)–(4), reports the results across four lookback windows (6, 12, 24, and 36 months). All coefficients are positive, consistent with the predicted direction, but none approaches statistical significance (t-statistics range from 0.46 to 1.00).[22]

Columns (5)–(7) of Table 10 test three additional proxies motivated by the career-concerns literature. If analysts who cover more firms rely on generic templates, coverage breadth should predict higher MacroBearGap (column 5: -0.04, $t = -0.80$). If analysts who write more reports entrench house templates, report frequency should predict higher MacroBearGap (column 6: -0.01, $t = -0.52$). If analysts with poor historical forecast accuracy use macro framing as reputational insurance, average forecast error should predict higher MacroBearGap (column 7: 0.48, $t = 0.78$). None of these proxies is statistically significant, and the coverage and frequency measures carry signs opposite to the career-concerns prediction.[23]

[22]The SDC data match 34.5% of deals to our panel via CUSIP-to-PERMNO mapping, yielding IB-relationship rates of 12% (6-month window) to 39% (36-month window). The null result is unlikely to be driven by insufficient variation in the IB indicator.

[23]We also examine whether analyst identity explains a meaningful share of MacroBearGap variation by adding 232 analyst fixed effects to a specification with firm and year fixed effects (Appendix Table A.6). The within-R^2 increases

Table 10 Career Concern Proxies: IB Relationships and Analyst Characteristics

Note: This table reports panel regressions of MacroBearGap (bear minus base macro attention, 0–100 scale) on career-concern proxies. Columns (1)–(4) test whether Morgan Stanley investment-banking relationships predict MacroBearGap, using lookback windows of 6, 12, 24, and 36 months. IB relationship equals one if MS served as bookrunner or lead manager on a new issuance for the covered firm within the indicated window (source: SDC New Issues). Coverage equals the number of firms the analyst covers in the same calendar year. Reports equals the analyst's annual report count. Accuracy equals the analyst's mean absolute forecast error. Columns (5)–(7) are restricted to the analyst-matched subsample. All specifications include firm and year fixed effects. Standard errors are clustered by firm. t-statistics in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	MacroGap						
	IB 6m	IB 12m	IB 24m	IB 36m	Coverage	Reports	Accuracy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
MS IB Relationship	0.1484	0.3363	0.2481	0.3474			
	(0.46)	(1.00)	(0.59)	(0.70)			
Firms Covered					−0.0351		
					(−0.80)		
Reports / Year						−0.0055	
						(−0.52)	
Avg. Forecast Error							0.4776
							(0.78)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	31,445	31,445	31,445	31,445	17,257	17,257	17,255
R^2	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000

On balance, none of the career-concern proxies—analyst experience, IB relationships, coverage breadth, report frequency, forecast accuracy—predicts MacroBearGap in the direction the channel implies. The only marginally significant coefficient, junior MS tenure, carries the wrong sign. We read this as little support for the career-concerns channel through the proxies available to us; we do not claim the channel is conclusively ruled out, since institutional-client demand, sector-head review, internal compliance, and management-access pressure are not captured by these measures and could in principle generate Macro Bear Bias through routes our data do not observe.[24] Availability, by

by 2.05 percentage points (from 0.28% to 2.33%), indicating persistent heterogeneity in macro-narrative style across analysts. This is consistent with either career incentives or cognitive differences in scenario construction, and does not uniquely identify career concerns.

[24] The proxies in Table 10 have within-R^2 values close to zero, so the tests have limited power against small but economically meaningful career-concern effects. A formal power analysis would help bound the magnitude of effect that the IB-relationship coefficient is consistent with; we leave this calculation to a future revision.

contrast, is supported across the industry–crisis interaction, GFC-experience, sector-herding, and directional-boldness tests.

5 Conclusion

We document a systematic pattern in how analysts narrate downside risk. Using scenario-based analyst reports, we show that narratives contain information that is not visible from numerical forecasts alone. Analysts do not merely forecast prices; they also explain why different outcomes may occur. These explanations reveal a systematic distortion in subjective risk attribution: bear-case scenarios are framed as disproportionately macro-driven, even when realized outcomes do not support a larger macro component of downside risk.

We call this pattern Macro Bear Bias. The key insight is not that macro risk is unimportant, nor that analysts are wrong to discuss macro conditions in downside states. Rather, the bias lies in the failure to condition correctly. Conditional on the same subsequent market environment, realized CAPM R^2 is similar across realized bear, base, and bull outcomes, yet bear narratives remain persistently more macro-heavy than base or bull narratives. Downside risk is therefore not only perceived as bad; it is narrated as systematically macro in nature.

This narrative distortion matters economically. The bear–base macro-attention gap predicts subsequent errors in analysts' base-case forecasts, suggesting that macro-heavy downside narratives spill over into central valuation judgments and make analysts too pessimistic ex ante. The predictive relation is especially strong in nonlinear specifications, consistent with the idea that the cost of a fixed macro-downside template depends on the market environment in which it is applied.

Mechanism tests are consistent with an availability-heuristic channel for Macro Bear Bias, best understood as a refinement of canonical memory-based accounts (Bordalo, Gennaioli, and Shleifer, 2020; Bordalo et al., 2023): the heuristic anchors downside narratives on a small number of vivid macro-crisis episodes, but the resulting templates persist beyond individual recall through institutional and peer-network reinforcement. As predicted by this mechanism, macro-heavy bear narratives respond to negative macro-news shocks, concentrate in crisis-exposed settings, and converge among analysts facing shared sector-level narrative environments. The economic realization is consistent with a discount-rate contamination channel—salient macro-crisis templates inflate perceived systematic risk in bear scenarios, raising the discount rate used in valuation and depressing base-case targets—though our forecast-error evidence does not decompose the bias into discount-rate and cash-flow components, and a cash-flow-revision channel cannot be ruled out without that decomposition. By contrast, the simple career-concerns proxies we observe carry little explanatory power; we read this as no support for that channel through these proxies rather than as a wholesale rejection of strategic explanations.

The broader impact of this paper is to show that narratives are not merely decorative explanations attached to forecasts; they are an economically important part of belief formation in financial markets. By revealing how sophisticated intermediaries systematically misattribute downside risk to macro forces, our findings highlight a new channel through which narrative templates can shape valuations, forecast errors, and ultimately market beliefs. The paper therefore contributes to a broader understanding of price formation: distortions need not arise only from biased numerical expectations, but can also emerge from the way agents organize and communicate uncertainty. Methodologically, the paper shows how combining structured scenario reports, LLM-based narrative measurement, and realized return benchmarks can open a new empirical window into subjective risk perception.

References

Ang, Andrew, and Joseph Chen, 2002, Asymmetric correlations of equity portfolios, *Journal of financial Economics* 63, 443–494.

Ang, Andrew, Joseph Chen, and Yuhang Xing, 2006, Downside risk, *The Review of Financial Studies* 19, 1191–1239.

Bastianello, Francesca, Paul H Décaire, and Marius Guenzel, 2024, Mental models and financial forecasts .

Bohren, J Aislinn, Josh Hascher, Alex Imas, Michael Ungeheuer, and Martin Weber, 2024, A cognitive foundation for perceiving uncertainty, Technical report, National Bureau of Economic Research.

Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer, 2023, Memory and probability, *The Quarterly Journal of Economics* 138, 265–311.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2019, Diagnostic expectations and stock returns, *Journal of Finance* 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2024, Belief overreaction and stock market puzzles, *Journal of Political Economy* 132, 1350–1396.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, attention, and choice, *The Quarterly journal of economics* 135, 1399–1442.

Bouchaud, Jean-Philippe, Philipp Krueger, Augustin Landier, and David Thesmar, 2019, Sticky expectations and the profitability anomaly, *Journal of Finance* 74, 639–674.

Bradshaw, Mark, Ziwei Qiao, Hongping Tan, and Changqiu Yu, 2025, Unlocking analysts' risk insights .

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu, 2025, Forest through the trees: Building cross-sections of stock returns, *The Journal of Finance* 80, 2447–2506.

Chen, Luyang, Markus Pelger, and Jason Zhu, 2024, Deep learning in asset pricing, *Management Science* 70, 714–750.

Chen, Shuaiyu, Shuaiqi Li, and Yucheng John Yang, 2025, Subjective expectations for variance and skewness: Evidence from analyst forecasts, *Available at SSRN* .

Chevalier, Judith, and Glenn Ellison, 1999, Career concerns of mutual fund managers, *The Quarterly Journal of Economics* 114, 389–432.

De la O, Ricardo, and Sean Myers, 2021, Subjective cash flow and discount rate expectations, *Journal of Finance* 76, 1339–1387.

Dew-Becker, Ian, Stefano Giglio, and Pooya Molavi, 2025, Learning and the emergence of nonlinearity in financial markets, Technical report, National Bureau of Economic Research.

Didisheim, Antoine, Shikun Barry Ke, Bryan T Kelly, and Semyon Malamud, 2024, Apt or "aipt"? the surprising dominance of large factor models, Technical report, National Bureau of Economic Research.

Feng, Guanhao, Jingyu He, Nicholas G Polson, and Jianeng Xu, 2024, Deep learning in characteristics-sorted factor models, *Journal of Financial and Quantitative Analysis* 59, 3001–3036.

Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, 2020, Dissecting characteristics nonparametrically, *The Review of Financial Studies* 33, 2326–2377.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481–1522.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *The Review of Financial Studies* 27, 714–746.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2020, Empirical asset pricing via machine learning, *The Review of Financial Studies* 33, 2223–2273.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2021, Autoencoder asset pricing models, *Journal of Econometrics* 222, 429–450.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2018, Time varying risk aversion, *Journal of Financial Economics* 128, 403–421.

Holmström, Bengt, 1999, Managerial incentive problems: A dynamic perspective, *The Review of Economic Studies* 66, 169–182.

Hong, Harrison, and Jeffrey D Kubik, 2003, Analyzing the analysts: Career concerns and biased earnings forecasts, *Journal of Finance* 58, 313–351.

Hong, Harrison, Jeffrey D Kubik, and Amit Solomon, 2000, Security analysts' career concerns and herding of earnings forecasts, *The Rand journal of economics* 121–144.

Jensen, Theis Ingerslev, 2022, Subjective risk and return .

Joos, Peter, Joseph D Piotroski, and Suraj Srinivasan, 2016, Can analysts assess fundamental risk and valuation uncertainty? an empirical analysis of scenario-based value estimates, *Journal of Financial Economics* 121, 645–663.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2014, Time-varying fund manager skill, *Journal of Finance* 69, 1455–1484.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2016, A rational theory of mutual funds' attention allocation, *Econometrica* 84, 571–626.

Ke, Shikun, 2024, Analysts' belief formation in their own words, *Available at SSRN 5025830* .

Kelly, Bryan, Semyon Malamud, and Kangying Zhou, 2024, The virtue of complexity in return prediction, *Journal of Finance* 79, 459–503.

Kelly, Bryan, and Dacheng Xiu, 2023, Financial machine learning, *Foundations and Trends® in Finance* 13, 205–363.

Kelly, Bryan T, and Semyon Malamud, 2025, Understanding the virtue of complexity, *Available at SSRN 5346842* .

Kelly, Bryan T, Semyon Malamud, and Kangying Zhou, 2022, The virtue of complexity everywhere, *Available at SSRN 4166368* .

La Porta, Rafael, 1996, Expectations and the cross-section of stock returns, *Journal of Finance* 51, 1715–1742.

Lettau, Martin, Matteo Maggiori, and Michael Weber, 2014, Conditional risk premia in currency markets and other asset classes, *Journal of Financial Economics* 114, 197–225.

Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289–314.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken, 2010, A skeptical appraisal of asset pricing tests, *Journal of Financial Economics* 96, 175–194.

Lochstoer, Lars A, and Tyler Muir, 2022, Volatility expectations and returns, *The Journal of Finance* 77, 1055–1096.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: do macroeconomic experiences affect risk taking?, *The quarterly journal of economics* 126, 373–416.

Rahimi, Ali, and Benjamin Recht, 2007, Random features for large-scale kernel machines, *Advances in neural information processing systems* 20.

Rahimi, Ali, and Benjamin Recht, 2008, Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning, *Advances in neural information processing systems* 21.

Roll, Richard, 1988, R^2, *The Journal of Finance* 43, 541–566.

Scharfstein, David S, and Jeremy C Stein, 1990, Herd behavior and investment, *The American Economic Review* 80, 465–479.

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

Srinivasan, Suraj, and David Lane, 2011, The risk-reward framework at morgan stanley research, *Harvard Business School Accounting & Management Unit case* .

Tversky, Amos, and Daniel Kahneman, 1973, Availability: A heuristic for judging frequency and probability, *Cognitive Psychology* 5, 207–232.

Tversky, Amos, and Daniel Kahneman, 1974, Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty., *science* 185, 1124–1131.

Weyns, Guy, Juan-Luis Perez, Barry Hurewitz, and Vlad Jenkins, 2011, Morgan stanley's risk-reward views: Unlocking the full potential of fundamental analysis, *Journal of Applied Corporate Finance* 23, 59–68.

Zhou, Kangying, 2025, Professional investors and media narratives, Technical report, Working Paper.

INTERNET APPENDIX

A Morgan Stanley Risk-Reward Framework Examples

A.1 Example 1: Autodesk (ADSK.O), September 3, 2021

Risk Reward – Autodesk (ADSK.O)

New Products and Pricing Power Should Support Durable Growth at Autodesk

PRICE TARGET $324.00

Based on a multiple of 28x EV/FCF FY26e FCF/share discounted back at 11% WACC

Consensus Price Target Distribution $265.00 $350.02 MS PT $400.00

Source: Thomson Reuters, Morgan Stanley Research

◆ Mean ◆ Morgan Stanley Estimates

EQUAL-WEIGHT THESIS

- The subscription transition (now largely complete) has significantly reduced volatility in earnings, which has driven the multiple higher. Additionally, Autodesk has found multiple avenues to better monetize a sticky customer base. However, long-standing debates on the longer-term sustainable growth, higher macro sensitivity and quality of near-term FCF likely limits further multiple expansion NT, leaving us EW on ADSK.
- We forecast ~$11 in FCF/share by FY23 & ~$16 by FY26 & apply a 28X EV/FCF multiple against this FCF and discounting back at 11.3% to drive our price target. Our multiple implies 1.6x growth-adjusted which is largely inline with large cap software and design software peers.

RISK REWARD CHART AND OPTIONS IMPLIED PROBABILITIES (12M)

Source: Thomson Reuters, Morgan Stanley Research, Morgan Stanley Institutional Equities Division. The probabilities of our Bull, Base, and Bear case scenarios playing out were estimated with implied volatility data from the options market as of 02 Sep, 2021. All figures are approximate risk-neutral probabilities of the stock reaching beyond the scenario price in either three-months' or one-years' time. View explanation of Options Probabilities methodology here

Consensus Rating Distribution

68%	Overweight
27%	Equal-weight
5%	Underweight

● MS Rating

Source: Thomson Reuters, Morgan Stanley Research

Risk Reward Themes

Disruption:	*Positive*
Technology Diffusion:	*Positive*

View descriptions of Risk Rewards Themes here

BULL CASE $435.00

Discount of 34x EV/FCF FY26e FCF/share of $17.62

Secular Growth Opportunities Come Into Fruition. Autodesk is able to successfully expand into a large greenfield opportunity within Field Construction, as total subscriptions reach 9M+ by FY26, while expenses grow at ~11% CAGR, resulting in FCF of $17.62/share in FY26. Applying a 34x EV/FCF multiple to the rapidly growing cash generation profile and a 11.3% discount rate yields our bull case target.

BASE CASE $324.00

Discount of 28x EV/FCF FY26e FCF/share of $16.94

Double Digit FCF CAGR. Autodesk is successful in ramping to a base of ~8M subscriptions by FY26, while expenses grow at a ~9% CAGR, yielding FCF/share of $15.91 and double digit CAGR in FY23-26. 28x that FCF, discounted back at 11.3%, yields our price target. Our 28x FCF multiple is a premium to Autodesk's historical FCF multiple, but roughly inline with its growth adjusted peer group average (1.6x).

BEAR CASE $224.00

Discount of 20x EV/FCF FY26e FCF/share of $15.41

Prolonged Macro Downturn. In our bear case, we include the risk of prolonged recession with modest recovery thereafter, which drives more modest gross subscription adds. End result would be $15.41/share in FY26 and 20x EV/FCF with 11.3% discount rate, arriving at our bear case value.

A.2 Example 2: Gartner Inc. (IT.N), April 26, 2024

Risk Reward – Gartner Inc. (IT.N)

Strong Growth Story, but Balanced Risk/Reward

PRICE TARGET $447.00

Derived from our DCF model. Our DCF includes a 8.6% WACC, 1.06 beta, and a 3.0% terminal growth rate.

Consensus Price Target Distribution

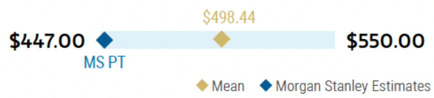

Source: Refinitiv, Morgan Stanley Research

RISK REWARD CHART

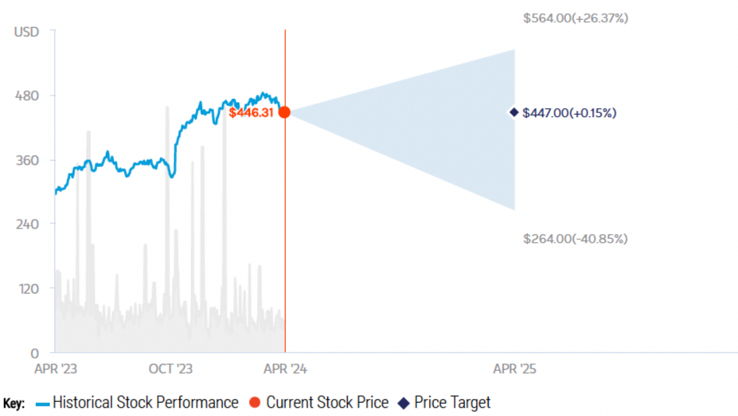

Source: Refinitiv, Morgan Stanley Research

EQUAL-WEIGHT THESIS

- Gartner is the market leader in technology research, and is expanding into new functional areas (marketing, supply chain, finance, legal, sales, and HR)
- IT tends to have one of the highest top line growth rates in the Info Services peer group
- Global Business Sales' (GBS) growth has been strong and we expect it to continue due to its large TAM
- Appetite for technology research and secular growth in IT continue to provide solid backdrop
- We find valuation fair at current levels vs. peers, and would be more constructive if there were signs of approaching a positive CV inflection.

Consensus Rating Distribution

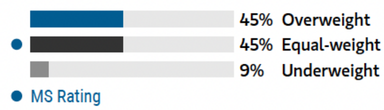

Source: Refinitiv, Morgan Stanley Research

Risk Reward Themes

Secular Growth: *Positive*
Pricing Power: *Positive*

View descriptions of Risk Rewards Themes here

BULL CASE $564.00

27.0x Bull Case '25e EV/MS Adj. EBITDA

Salesperson productivity accelerates, driving contract value growth and adj. EBITDA margin expansion. Revenue and adj. EBITDA CAGRs of 12% and 13% through '27; adj. EBITDA margins expand from 16% in '19 to 26% in '27.

BASE CASE $447.00

23.8x Base Case '25e EV/MS Adj. EBITDA

Sales force expansion drives top-line growth and modest margin expansion. Revenue and adj. EBITDA CAGRs of 9% through '27; adj. EBITDA margins expand to 25% in '27.

BEAR CASE $264.00

17.8x Bear Case '25e EV/MS Adj. EBITDA

Revenue growth decelerates as macro environment becomes more challenged and margin expansion is limited. Revenue and adj. EBITDA grow at 7% and 3% CAGRs through '27; adj. EBITDA margins of 22% in '27.

B LLM Prompt

B.1 Subjective Risks Identification Prompt

```
You are a financial analyst trained in equity valuation and empirical asset
    pricing. You will analyze the Bull, Base, and Bear case texts from a
    Morgan Stanley equity research report. Your task is to extract reasoning,
    classify it as firm-level, industry-level, or macro-level, identify the
    economic channel, and quantify attention shares across cases. Write in a
    concise, research-style tone used in equity research reports.

Return your full answer **strictly in JSON format**, with one JSON object per
    report (suitable for .jsonl storage). Each JSON object must include keys:
    "filename", "bull_case", "base_case", "bear_case", and "cross_case_summary
    ". Do not include explanations outside the JSON.

Tasks:
1.Price Targets: List the bull case, base case, and bear case price targets
    stated in the Risk Reward section.
2.Reason Extraction: Identify all distinct reasons for the bull, base, and
    bear cases. Requirements:
2.1 Include any explicitly stated upside risks as part of the bull-case
    reasons.
2.2 Include any explicitly stated downside risks as part of the bear-case
    reasons.
2.3 Preserve the original wording as closely as possible (avoid paraphrasing).
2.4 Number each reason sequentially (e.g., 1., 2., 3.).
2.5 Label each reason as one of the following categories: Firm, Industry,
    Macro, or Mixed. If labeled Mixed, specify which component (Firm, Industry
    , or Macro) is dominant.
2.6 Assign each reason to a single economic channel that best represents its
    underlying mechanism, such as expected cash flow, discount rate, risk
    aversion or asset demand.
```

3. Attention Quantification: Quantify the analyst's relative attention across Firm, Industry, and Macro news as: (Firm %, Industry %, Macro %), ensuring Firm % + Industry % + Macro % <= 100.
4. Attention Evolution Summary: Summarize how the analyst's attention distribution across Firm, Industry, and Macro news (from Step 2.5) evolves across the bull, base, and bear cases.
5. Channel Evolution Summary: Summarize how the dominant economic channels (from Step 2.6) evolve across the bull, base, and bear cases.

Output format (strict JSON):

```
{
  "filename": "<report_name>",
  "bull_case": {
    "bull_price_target": "<numerical price target or range>",
    "reasons": [
      {
        "id": 1,
        "text": "<verbatim reason>",
        "category": "Firm | Industry | Macro | Mixed (dominant: Firm/Industry/
            Macro)",
        "economic_channel": "Expected Cash Flow | Discount Rate | Risk
            Aversion | Asset Demand"
      }
    ],
    "attention_distribution": { "Firm": X, "Industry": Y, "Macro": Z },
    "analytical_summary": "[one concise sentence summarizing key focus of bull
        case]"
  },

  "base_case": {
    "base_price_target": "<numerical price target or range>",
    "reasons": [
      {
```

```
      "id": 1,
      "text": "<verbatim reason>",
      "category": "Firm | Industry | Macro | Mixed (dominant: ...)",
      "economic_channel": "Expected Cash Flow | Discount Rate | Risk
         Aversion | Asset Demand"
    }
  ],
  "attention_distribution": { "Firm": X, "Industry": Y, "Macro": Z },
  "analytical_summary": "[one concise sentence summarizing key focus of base
      case]"
},

"bear_case": {
  "bear_price_target": "<numerical price target or range>",
  "reasons": [
    {
      "id": 1,
      "text": "<verbatim reason>",
      "category": "Firm | Industry | Macro | Mixed (dominant: ...)",
      "economic_channel": "Expected Cash Flow | Discount Rate | Risk
         Aversion | Asset Demand"
    }
  ],
  "attention_distribution": { "Firm": X, "Industry": Y, "Macro": Z },
  "analytical_summary": "[one concise sentence summarizing key focus of bear
      case]"
},

"cross_case_summary": {
  "attention_evolution": "[summary of how Firm/Industry/Macro attention
     shifts across bull -> base -> bear]",
  "channel_evolution": "[summary of how dominant economic channels shift
     across bull -> base -> bear]"
```

```
  }
}

Rules:
- All percentages must be integers (no decimals).
- Maintain an analytical and academic tone throughout.
- Do NOT include any text outside the JSON structure.
- Ensure structural consistency and valid JSON formatting.
```

B.2 Topic Classification Prompt

```
You are a financial research assistant.

The input summaries already contain ONLY macro-related content.
Do NOT extract or filter anything.

Tasks:

1. Classify each scenario (Bull, Base, Bear) using the taxonomy.
   - Multiple topics allowed.
   - If broad -> "Mixed"
   - If not covered -> "Other: <specific macro theme>"
   - If empty -> "No explicit macro content"
   - Also assign a percentage weight to each topic based on its importance in
     the summary:
      - The percentages must sum to 100 for each scenario
      - Use integers (e.g., 20, 30, 50)
      - The order of percentages must match the order of topics exactly

1. Compare macro differences:

Bear vs Base:
```

```
- If both Bear and Base are empty -> "No explicit macro content"

- If Base is empty but Bear is not:
  -> return ALL Bear topics

- If Bear is empty but Base is not:
  -> "No meaningful difference"

- Otherwise:
  - Identify ONLY macro topics that appear in Bear but NOT in Base
  - Do NOT include topics already in Base
  - If no additional topics -> "No meaningful difference"

- For any returned topics:
  - Assign percentage weights that sum to 100
  - Reflect relative importance of differences
  - Order must match topics exactly


Bull vs Base:

- If both Bull and Base are empty -> "No explicit macro content"

- If Base is empty but Bull is not:
  -> return ALL Bull topics

- If Bull is empty but Base is not:
  -> "No meaningful difference"

- Otherwise:
  - Identify ONLY macro topics that appear in Bull but NOT in Base
  - Do NOT include topics already in Base
  - If no additional topics -> "No meaningful difference"
```

```
- For any returned topics:
  - Assign percentage weights that sum to 100
  - Reflect relative importance of differences
  - Order must match topics exactly


Rules:
- Use only provided text
- No inference beyond text
- Topics must come from taxonomy (+ Mixed / Other)
- "Other" must specify a real macro concept
- Output valid JSON only

Taxonomy:
{TOPIC_TAXONOMY}

Input:

[BULL]
{bull_summary}

[BASE]
{base_summary}

[BEAR]
{bear_summary}

Output JSON:

{{
  "bull": {{
    "topics": ["..."],
```

```
    "percentages": []
  }},
  "base": {{
    "topics": ["..."],
    "percentages": []
  }},
  "bear": {{
    "topics": ["..."],
    "percentages": []
  }},
  "bear_vs_base": {{
    "topics": ["..."],
    "percentages": []
  }},
  "bull_vs_base": {{
    "topics": ["..."],
    "percentages": []
  }}
}}""".strip()
```

B.3 Systematic vs. Idiosyncratic Prompt

```
You are a financial research analyst.

Your task is to classify each statement based only on its text content.

Determine whether the primary driver described in the statement is:

- Systematic
- Idiosyncratic
- Mixed

Judge only from the text itself.
```

DEFINITIONS

Systematic:

Primarily driven by broad common factors,
affecting multiple firms or the overall market,
and generally not diversifiable.

Idiosyncratic:

Primarily driven by firm-specific factors,
mainly affecting one firm or a small number of firms,
and generally diversifiable.

Mixed:

Use only when both systematic and idiosyncratic factors
are explicitly stated and equally important.

CRITICAL RULES

----

1. Primary Driver Rule

Classify based on the main economic driver.

2. Mixed Rule

Use Mixed only when both types are explicitly present
and neither is clearly dominant.

3. Forced Choice Rule

Every statement must be classified as:

- Systematic
- Idiosyncratic
- Mixed

4. Background vs Driver Rule

If broad market or environmental conditions are only background context,
but the company itself drives the outcome,
classify as Idiosyncratic.

INSTRUCTIONS

For each item:

1. Read the statement carefully.
2. Determine the true primary driver.
3. Return one final label.

OUTPUT FORMAT

Return ONLY valid JSON:

```
{
  "items": [
    {
      "scenario": "bull/base/bear",
      "segment_number": "1",
      "original_text": "...",
      "correct_label": "Systematic",
      "explanation": "short explanation"
    }
  ]
}
```

C Additional Exhibits

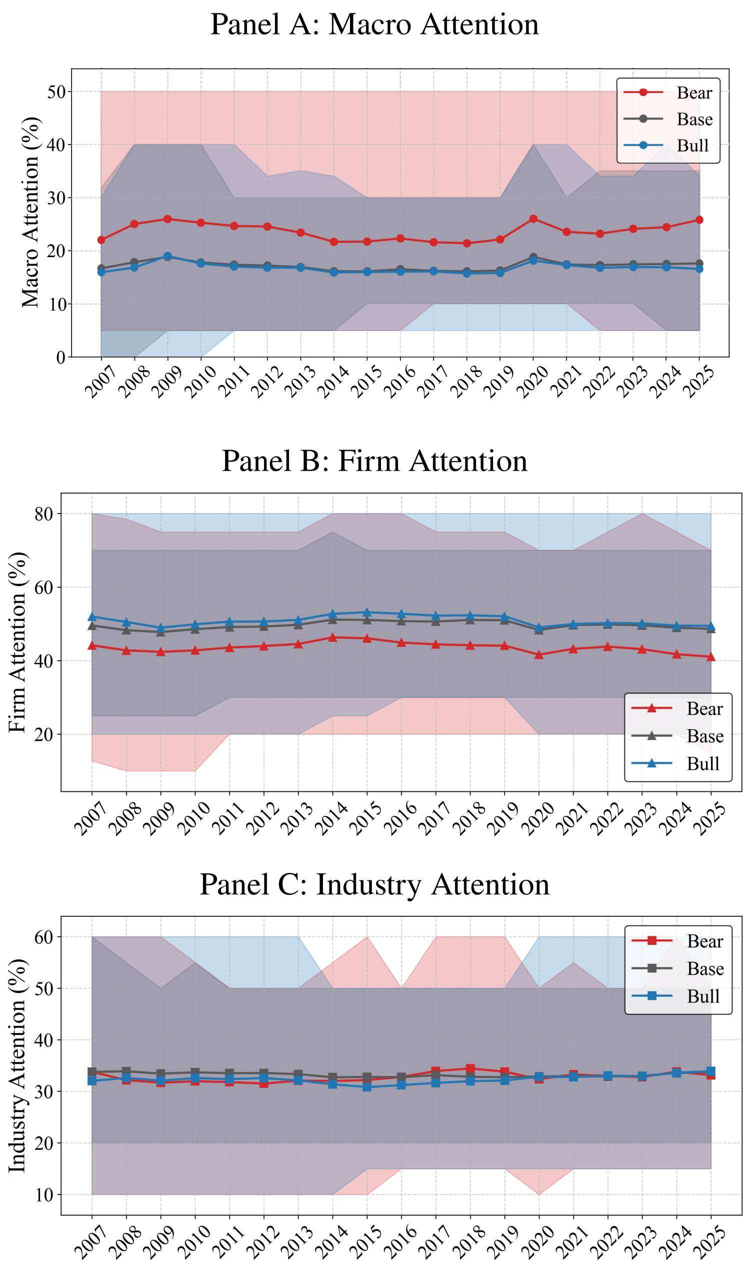

Figure A.1 Time Series of Attention to Firm, Industry, and Macro News.

Note: Time series of annual average attention shares to firm-, industry-, and macro-level news. The solid lines represent yearly means. The shaded areas correspond to the empirical 5th to 95th percentile range of the cross-sectional distribution in each year.

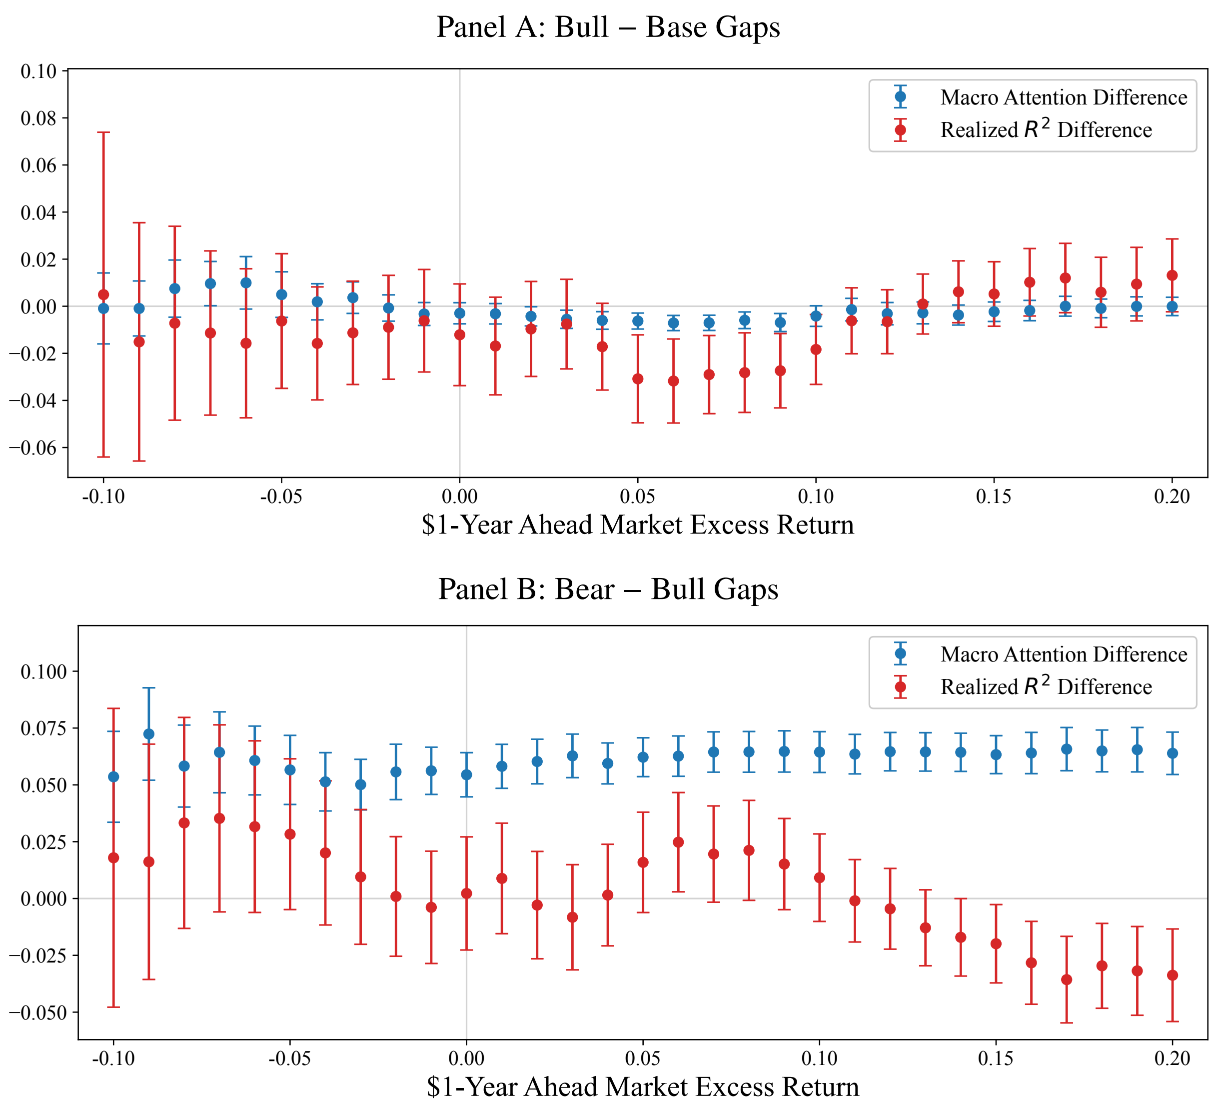

Figure A.2 Gaps in Macro Attention and Realized CAPM R^2 Conditional on Market States.

Note: This figure plots firm-level bull–base and bear–bull differences in macro attention and realized CAPM R^2 across subsamples defined by 10% bins of forward 252-day market excess returns following the analyst report release. Each point includes reports with $x - 5\% \leq (R_m - R_f) \leq x + 5\%$ and applies the same-firms rule within each bin, requiring firms to have realized-bear, realized-base, and realized-bull reports. Macro attention is measured as the firm-level average macro attention in the bull, base, and bear cases. Realized CAPM R^2 is computed as the firm-level average R^2 over the 252 trading days following the report release for reports ex post classified as realized bull, realized base, and realized bear cases. The plotted values represent the bull-minus-base and bear-minus-bull differences. Error bars show 95% confidence intervals from constant-only regressions.

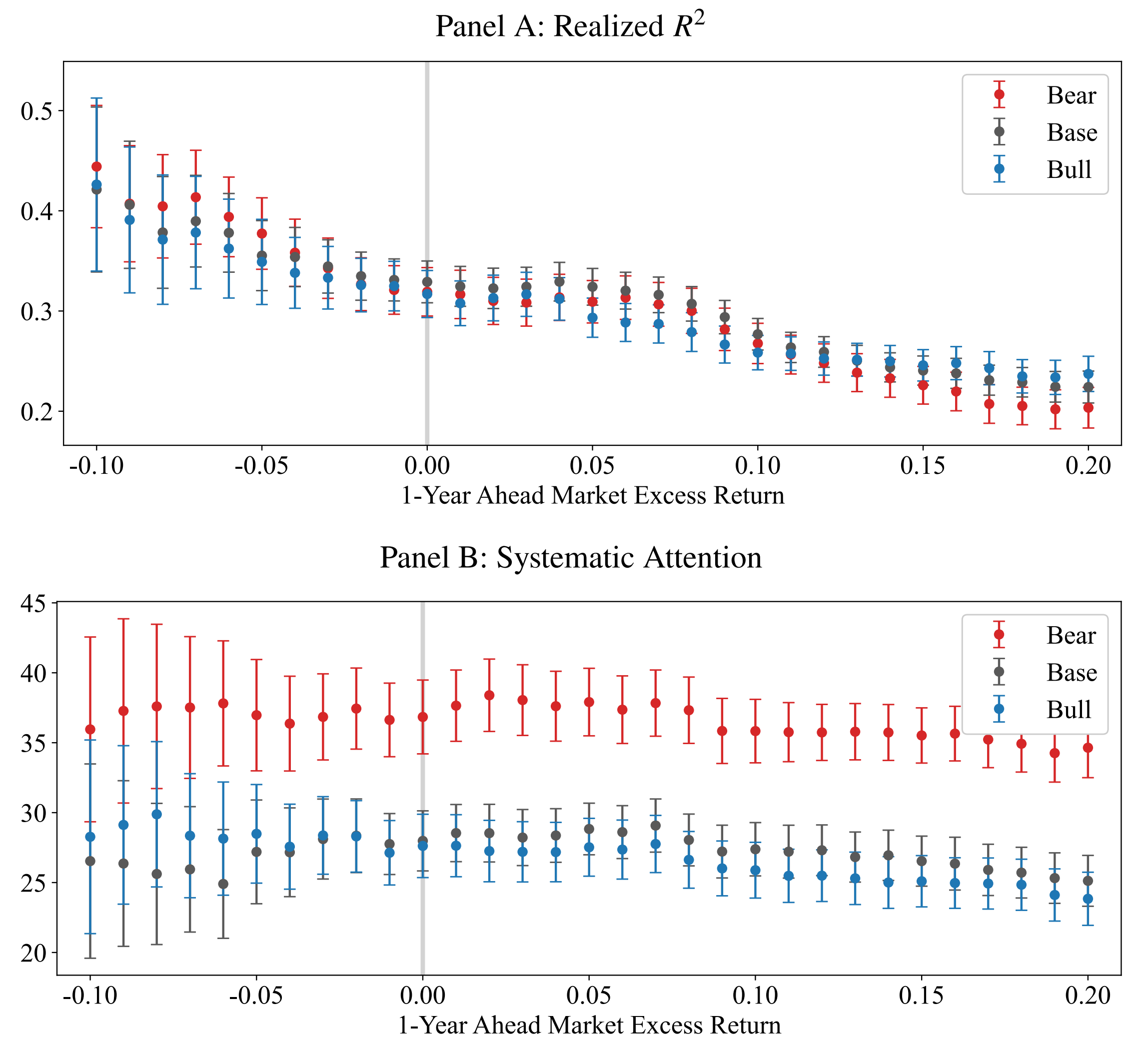

Figure A.3 Realized R^2 and Systematic Attention Conditional on Market States.

Note: This figure shows the realized R^2 and systematic attention conditional on different market states. For each point x on the horizontal axis, we consider the reports within 10% bins ($[x - 5\%, x + 5\%]$) of forward 252-day market excess returns following the analyst report release. We apply the same-firms rule within each bin, requiring firms to have realized-bear, realized-base, and realized-bull reports. Systematic attention is measured as the firm-level average systematic attention within each realized scenario (bear, base, and bull). Realized CAPM R^2 is computed as the firm-level average R^2 over the 252 trading days following the report release for reports ex post classified as realized bear/base/bull cases. Error bars show 95% confidence intervals from constant-only regressions.

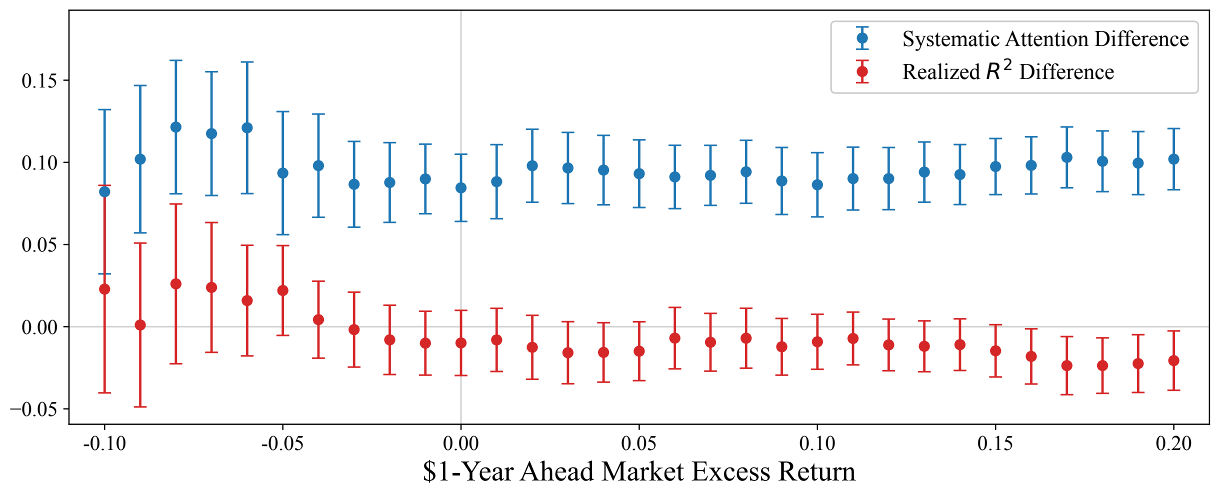

Figure A.4 Bear – Base Gaps in Systematic Attention and Realized CAPM R^2 Conditional on Market States.

Note: This figure plots firm-level bear–base differences in systematic attention and realized CAPM R^2 on different market states. For each point x on the horizontal axis, we consider the reports within 10% bins ($[x - 5\%, x + 5\%]$) of forward 252-day market excess returns following the analyst report release. We apply the same-firms rule within each bin, requiring firms to have realized-bear, realized-base, and realized-bull reports. Systematic attention is measured as the firm-level average systematic attention in the bear and base cases. Realized CAPM R^2 is computed as the firm-level average R^2 over the 252 trading days following the report release for reports ex post classified as realized bear or realized base cases. The plotted values represent the bear-minus-base differences. Error bars show 95% confidence intervals from constant-only regressions.

Table A.1 Descriptions of Key Topics Used in Attention and Correlation Analysis

Topic	Description	Bear Attn. (%)	News Attn. (%)	News Sent.
Economic Growth and Recovery Outlook	Indicators suggesting potential economic recovery emerge from consumer confidence, job growth, and corporate earnings. Positive signals improve investor outlook, although inflation and labor market weakness may remain concerns.	10.66	2.36	0.48
COVID-19 Pandemic and Vaccine Developments	The COVID-19 pandemic significantly affected global health and economic conditions. Vaccine development, distribution logistics, variants, and public health responses shaped market expectations and business conditions.	5.36	5.34	0.41
Market Activity and Financial Performance	Stock market activity reacts to earnings releases, firm performance, and sector-wide developments. Strong results support rallies, while weak performance increases caution and volatility.	4.75	5.44	0.39
Consumer Price Index and Inflation Trends	Rising consumer prices and inflation pressures influence monetary policy, investor sentiment, and broader economic strategies. Central bank responses are especially important for markets.	2.40	0.86	0.08
International Trade Relations and Economic Policies	Trade tensions, tariff policy, and negotiations among major economies affect global commerce, supply chains, and investor confidence.	2.31	1.60	0.12
Consumer Behavior Amid Economic Strains	Economic pressures change household spending patterns, especially for discretionary purchases. Firms adjust pricing and marketing strategies amid inflation and fuel cost concerns.	1.91	0.24	-0.05
Geopolitical Tensions and Economic Impacts	Political instability, wars, and international conflicts affect stock prices, commodity markets, and investor sentiment through heightened uncertainty.	0.46	4.40	-0.11
Economic Stimulus and Government Interventions	Fiscal stimulus packages, government support programs, and intervention policies are used to stabilize growth, employment, and financial markets during downturns.	0.40	0.46	0.07

Table A.3 Forecast Error Regressions on Attention Gaps
Note: This table reports empirical results from the following forecast-error prediction regression:

$$\text{Forecast Error}_{i,t+1} = a + b_1 \text{AttentionGap}_{i,t} + b_2 M_t + b_3 \hat{\beta}_{i,t} + b_4 R_{i,t} + \gamma_i + u_{i,t+1},$$

where each narrative attention-gap variable is NOT interacted with stock i's CAPM beta. Stock i's CAPM beta $\hat{\beta}_{i,t}$ is estimated using daily returns over the 252 trading days prior to the report release. M_t and $R_{i,t}$ denote the market excess return and stock i's excess return over the same 252-day pre-release window. Columns (1)–(4) include each attention gap individually. Column (5) includes the two macro attention gaps jointly. Column (6) includes all four attention-gap variables simultaneously. Firm fixed effects are included in all specifications. Standard errors are clustered by firm and month. t-statistics are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
Attention Gaps						
Macro Bear − Base	0.0012**				0.0011**	0.0023***
	(2.38)				(2.09)	(3.26)
Macro Bull − Base		0.0009			0.0005	0.0002
		(1.03)			(0.54)	(0.24)
Firm Bear − Base			0.0006			0.0018**
			(1.11)			(2.28)
Firm Bull − Base				-0.0002		-0.0004
				(-0.23)		(-0.55)
Control Variables						
M_t	-0.1451	-0.1462	-0.1459	-0.1458	-0.1454	-0.1454
	(-1.55)	(-1.56)	(-1.56)	(-1.55)	(-1.55)	(-1.55)
$\hat{\beta}_{i,t}$	-0.0323	-0.0317	-0.0315	-0.0316	-0.0323	-0.0322
	(-0.58)	(-0.56)	(-0.56)	(-0.56)	(-0.57)	(-0.57)
$R_{i,t}$	-2.4847	-2.4249	-2.2599	-2.3870	-2.5042	-2.2778
	(-0.15)	(-0.15)	(-0.14)	(-0.15)	(-0.15)	(-0.14)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43,306	43,306	43,306	43,306	43,306	43,306
Within R^2	0.0004	0.0004	0.0004	0.0003	0.0004	0.0006

Table A.4 Bull-Case Macro Attention and Topic Salience

Note: This table reports panel regressions of bull-case analyst attention to topic k at time t on the autoregressive (AR) residual of WSJ news coverage (attention shock), its interaction with WSJ news sentiment, the sentiment level, and three lags of WSJ topic attention; see Equation (18). Columns correspond to lag $\tau \in \{1, 2, 3, 4\}$. All specifications include time fixed effects. t-statistics in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$
WSJ Attention Shock$_{t-\tau,k}$	0.0315*	0.0250*	0.0103	-0.0025
	(1.85)	(1.70)	(0.97)	(-0.35)
WSJ Attention Shock$_{t-\tau,k}$ $\times$ WSJ Sentiment$_{t-\tau,k}$	-0.0783	-0.0876	-0.0540	-0.0338
	(-1.35)	(-1.52)	(-1.08)	(-0.83)
WSJ Sentiment$_{t-\tau,k}$	0.0003**	0.0003*	0.0002	0.0003**
	(2.08)	(1.85)	(1.29)	(1.97)
WSJ Attention$_{t-1,k}$	0.2889***	0.2849***	0.2842***	0.2849***
	(11.22)	(10.95)	(10.86)	(10.85)
WSJ Attention$_{t-2,k}$	0.1844***	0.1864***	0.1818***	0.1804***
	(6.65)	(6.70)	(6.51)	(6.45)
WSJ Attention$_{t-3,k}$	0.4919***	0.4938***	0.4987***	0.5004***
	(16.22)	(16.30)	(16.39)	(16.37)
Time FE	Yes	Yes	Yes	Yes
Observations	11,440	11,388	11,336	11,284
R^2	85.05%	85.10%	85.16%	85.27%

Table A.5 Analyst Experience and Macro Bear Bias

Note: This table reports panel regressions of MacroGap (bear minus base macro attention, 0–100 scale) on analyst experience measures. Junior (MS) equals one if the analyst's tenure at Morgan Stanley is at most 36 months. Junior (Profession) equals one if the analyst's industry experience across all brokerages is at most 3 years (identified from the full IBES price-target detail file). Columns (3)–(4) use continuous experience measures in years. Column (5) includes both MS and profession experience simultaneously. All columns are restricted to the analyst-matched subsample. All specifications include firm and year fixed effects. Standard errors are clustered by firm. *t*-statistics in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	MacroGap				
	(1) Junior (MS)	(2) Junior (Prof.)	(3) Exp. MS	(4) Exp. Prof.	(5) Horse Race
Junior (MS)	−0.6924*				
	(−1.66)				
Junior (Profession)		0.0759			
		(0.15)			
Experience MS (years)			0.0750		0.0914
			(0.94)		(0.79)
Experience Prof. (years)				0.0223	−0.0142
				(0.48)	(−0.21)
Firm FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Observations	17,257	17,257	17,257	17,257	17,257
R^2	0.0005	0.0000	0.0003	0.0001	0.0003

C.1 Additional Evidence on Mechanisms

Table A.5 reports panel regressions of MacroBearGap on analyst experience, measured as Morgan Stanley tenure (columns 1 and 3), industry-wide experience across all brokerages (columns 2 and 4), and a horse-race specification including both measures (column 5). See Section 4.2 for discussion.

Table A.6 compares the within-R^2 of a baseline specification (firm and year fixed effects) with a specification that adds 232 analyst fixed effects. See Section 4.2 for discussion.

Table A.6 Analyst Fixed Effects and Macro Bear Bias

Note: This table reports the incremental R^2 from adding analyst fixed effects to a specification with firm and year fixed effects. The dependent variable is MacroGap (bear minus base macro attention, 0–100 scale). Firm effects are absorbed by within-firm demeaning. Year effects are captured by year dummies. Model (B) adds 232 analyst dummies. Both models are estimated on the analyst-matched subsample. ΔR^2 is the incremental R^2 from adding analyst fixed effects.

	(A) Baseline	(B) + Analyst FE
R^2	0.0028	0.0233
ΔR^2		0.0205
Firm FE	✓	✓
Year FE	✓	✓
Analyst FE		✓
Observations	17,257	17,257
Firms	1,274	1,274
Analysts		232