

Looking Away from the Market

Adam Upenieks *

September 5th, 2026

Abstract

The Capital Asset Pricing Model (CAPM) is rejected empirically because market beta is weakly related to average returns. Using idiosyncratic news from the Dow Jones Newswire and institutional attention measures from Bloomberg search activity, I isolate states in which market-wide attention is focused on publicly available idiosyncratic news. In these states, the CAPM prices risk. Market betas explain the cross section of stock returns, and the security market line is steep and positive. In contrast, when attention is elevated but not well explained by idiosyncratic news, the beta–return relation flattens. To rationalize these findings, I develop a noisy rational expectations model in which investors choose between concentrating limited attention on noisy idiosyncratic signals or dispersing it towards broader market level learning. Because attention to idiosyncratic news crowds out private learning about systematic conditions, greater attention to idiosyncratic information increases posterior factor uncertainty. Investors require compensation for bearing this uncertainty, generating a higher equilibrium factor risk premium and a steeper security market line. Empirically, when attention is anchored to idiosyncratic news, unconditional betas price returns on more than half of trading days.

*I thank Miguel Palacios, Alexander David, Charles Martineau, Ryan Israelsen, and many others for helpful comments. Seminar participants at the Haskayne school of Business also provided valuable feedback.

*University of Calgary, Haskayne School of Business, Adam.upenieks@ucalgary.ca

1 Introduction

A persistent tension separates the theory of asset pricing from its empirical application. The capital asset pricing model (CAPM) of [Sharpe \(1964\)](#), [Lintner \(1965\)](#), and [Mossin \(1966\)](#) predicts that an asset’s risk premium is increasing in its market beta. Despite the model’s central role in academic research and practitioner discount-rate calculations ([Berk and van Binsbergen, 2017](#)), this prediction performs poorly empirically in cross-sectional tests. Recent literature offers conditional resolutions to this puzzle¹. A common theme in these resolutions is that the beta–return relation depends on investors processing information, and several recent studies suggest that investor attention could play an important role in determining asset prices. This paper argues that the CAPM succeeds not only when investors pay more attention, but when attention is concentrated on a particular kind of information: public idiosyncratic news.

Models of attention allocation imply that investors reallocate a limited attention budget between firm-specific and systematic information ([Peng and Xiong, 2006](#); [Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016](#)), a prediction that has recently received empirical support ([Cookson, García, and Jarnecic, 2026](#)). The asset pricing consequences of this reallocation are less explored. I address this by identifying the channel through which it operates: attention spent processing news about individual firms is time not spent learning about the state of the broader economy. As a result, when investors concentrate on idiosyncratic information they hold more uncertain beliefs about macroeconomic conditions and demand greater compensation for exposure to them. I derive closed-form conditions characterizing when this substitution increases the market price of risk and show that a single mechanism generates three predictions that I confirm empirically.

First, the security market line is state dependent: beta is priced on days when institutional attention is concentrated on firm-specific news and flat otherwise. Second, the effect is specific to market beta. Because firm-specific news carries no information about macroeconomic conditions, concentrated attention raises uncertainty about market exposure alone, and the size, value, profitability, and investment premia are unchanged across states. Third, the premium for idiosyncratic processing rises with macroeconomic uncertainty. When prior macroeconomic uncertainty is greater, so is the cost of leaving it unresolved, and high-beta assets command a larger premium when attention is concentrated on idiosyncratic news.

The theoretical model is a multi-firm variant of [Grossman and Stiglitz \(1980\)](#) in which in-

¹See [Tinic and West \(1984\)](#), [Savor and Wilson \(2014\)](#), [Savor and Wilson \(2016\)](#), [Hong and Sraer \(2016\)](#), [Jylhä \(2018\)](#), [Hendershott, Livdan, and Rösch \(2020\)](#), [Chan and Marsh \(2022\)](#), [Andrei, Friedman, and Ozel \(2023\)](#), and [Hasler and Martineau \(2023, 2024\)](#), among others.

vestors allocate a fixed attention budget between two sources of information: public news about individual firms and private analysis of a common macroeconomic factor. The key insight is that these two activities are substitutes with asymmetric payoffs. Firm-specific news, once its market-level content is filtered out, is informationally orthogonal to the factor: it resolves uncertainty about a firm’s idiosyncratic payoff and nothing else. An investor who devotes attention to it therefore holds a more precise view of individual firms and, because the budget is fixed, a less precise view of the macroeconomy. Idiosyncratic uncertainty is diversifiable and carries no price, but factor uncertainty is not, so the reallocation shifts uncertainty from a component investors are not paid to bear to one they are. To clear markets, expected returns must compensate for the factor uncertainty that remains unresolved in equilibrium, and this compensation accrues in proportion to each asset’s exposure to the factor. The cross-section of expected returns therefore aligns more closely with beta, and the security market line steepens.

To bring the framework to the data, I create a daily measure of how institutional attention is divided between firm-specific and broader processing. From the Dow Jones Institutional Newswire, I extract firm-level articles and screen out systematic and earnings-related content using keyword filters and a GPT-4o-mini classifier that retains only articles about an individual firm that are not tied to that day’s market developments, so that the surviving news is the orthogonal firm-specific signal of the model. From Bloomberg, I obtain firm-day institutional attention scores following [Ben-Rephael, Da, and Israelsen \(2017\)](#). These combine into two aggregate daily series for the Russell 3000 universe: an idiosyncratic news attention index, IAI_t , the cross-sectional average of Bloomberg attention among firms with a filtered firm-specific article on day t , and an aggregate institutional attention index, AIA_t , constructed as in [Chan and Marsh \(2022\)](#). Neither series identifies the headline result alone. The state variable is the residual $\hat{u}_t$ from an expanding-window regression of AIA_t on IAI_t , estimated using only data through day t so that classification is out of sample. When $\hat{u}_t$ is low, aggregate attention is fully accounted for by attention to firm-specific news, and attention is anchored in firm-level content. When $\hat{u}_t$ is high, attention exceeds what firm-specific news can explain, and reflects processing directed at market-wide or macroeconomic information.

Two pieces of evidence suggest that $\hat{u}_t$ captures the composition of attention. [Fisher, Martineau, and Sheng \(2022\)](#) measure macroeconomic attention independently from news coverage of aggregate topics. Table 1 shows that their index rises with $\hat{u}_t$: on days when institutional attention exceeds what firm-specific news can account for, coverage of macroeconomic themes is elevated. This suggests the two forms of attention move as substitutes, which has recently received empirical support ([Cookson, García, and Jarnećic, 2026](#)). Further,

the residual is one of several ways to measure how concentrated attention is on firm-level content, and the results do not depend on the choice. Alternative concentration measures deliver similar state-dependent pricing.

I test the model in U.S. equity data. On the 53% of trading days on which the rolling median of $\hat{u}_t$ lies below its trailing one-year median, beta is priced: the Fama–MacBeth slope is 12.6 basis points per day ($t = 3.00$) for value-weighted portfolios and 8.7 ($t = 2.21$) for equal-weighted portfolios, rising to 23.1 and 22.6 in the most anchored quartile, and the intercept is indistinguishable from zero, so the security market line is steep and consistent with the CAPM. On the complementary days the slope is flat. A distinct implication separates this result from the high-attention effects documented by [Ben-Rephael, Carlin, Da, and Israelsen \(2021\)](#): the days on which beta is most strongly priced are days on which aggregate attention AIA_t is below its sample mean while firm-news attention IAI_t is above it. What matters is the share of attention captured by firm-specific news, not its level. Replicating their expected-information-consumption framework confirms that the two channels are distinct rather than redundant: the conditional pricing result persists at full strength among firms their measure does not flag as expected to consume information.

Turning to uncertainty, the model implies that the composition channel is amplified when macroeconomic uncertainty is high. Because the premium compensates for factor uncertainty that attention to firm-specific news leaves unresolved, a larger prior uncertainty means a larger residual, so the price of beta on idiosyncratic-anchored days should rise with the level of uncertainty investors face going in. The data support this. Sorting idiosyncratic-anchored days into terciles of the prior day’s uncertainty, the price of beta in the top tercile is 24.2 basis points per day ($t = 2.14$) for realized variance, 31.8 ($t = 2.92$) for VIX, and 24.1 and 24.9 ($t = 2.34$ and 2.32) for the one- and twelve-month macroeconomic uncertainty indices of [Jurado, Ludvigson, and Ng \(2015\)](#), against estimates between -0.4 and 12.8 in the lower terciles. The premium is a composition effect and not an uncertainty effect: on unanchored days in the same top tercile the slope is indistinguishable from zero for every proxy, and a pooled interaction of beta with the idiosyncratic-anchored-state indicator within the top tercile ranges from 24.4 to 33.8 basis points (t between 1.77 and 2.33). The steep slope is also the CAPM slope. Across the four proxies the top-tercile intercept is insignificant and the slope exceeds the average market excess return in the cell only marginally (t between 1.13 and 1.89), a tilt I disclose in Section 6.2.4. Beta is priced most steeply on the days investors are least willing to bear unresolved aggregate risk, which is the compensation the model describes.

The closest antecedent is [Andrei, Friedman, and Ozel \(2023\)](#), who also derive a steeper

security market line from attention to firm-level information and also find the effect strongest under high uncertainty. In their model, announcements carry systematic content, so attentive investors resolve aggregate uncertainty and the line is steep on days it is resolved. Here firm-specific news is orthogonal to the factor, so processing it leaves aggregate uncertainty unresolved and the line is steep on days it is not. The designs differ accordingly: their state variables are levels, whereas this study uses the composition of attention, on days when the level of attention is below its mean. Their model predicts a steeper line on high-uncertainty days irrespective of attention. High uncertainty alone prices nothing in this setting, and finds beta pricing in both low and high uncertainty states, conditional on attention being anchored to idiosyncratic news.

The conditional pricing result is robust to the state variables the literature has used to locate conditional CAPM performance. [Savor and Wilson \(2014\)](#) find that beta is priced on macroeconomic announcement days, [Savor and Wilson \(2016\)](#) and [Chan and Marsh \(2022\)](#) on earnings announcement days and the days leading into them, and [Tinic and West \(1984\)](#) in particular calendar months. Excluding firms within a five-day earnings announcement window barely affects the beta slope, removing macroeconomic announcement days and leading earnings announcement days leaves the magnitudes and significance levels essentially unchanged, and the attention-anchored state slope on days with no scheduled announcement is 11.5 basis points ($t = 2.54$), suggesting the finding is not an announcement premium.

[Hong and Sraer \(2016\)](#) locate a steep line in low disagreement, [Hasler and Martineau \(2023\)](#) in low volatility, [Hasler and Martineau \(2024\)](#) in a high expected market premium, and [Andrei, Friedman, and Ozel \(2023\)](#) in high uncertainty resolved by attention to firm announcements. The continuous specification of Table 5 conditions on each of these objects directly. The daily Fama–MacBeth price of beta loads on the standardized five-day median of the attention residual evaluated through the prior day at -9.2 basis points per standard deviation ($t = -2.9$) for value-weighted and -7.7 ($t = -2.9$) for equal-weighted portfolios, and this gradient moves by only a few percent under controls for disagreement, the VIX, realized variance, the option-implied premium bound of [Martin \(2017\)](#), which is the expected-premium object of [Hasler and Martineau \(2024\)](#), the variance risk premium, investor sentiment, the levels of aggregate and firm-news attention that separate this channel from [Andrei, Friedman, and Ozel \(2023\)](#), the supply of firm-specific news, lagged market returns, the lagged price of beta, and announcement and day-of-week dummies, and with every control entered jointly it is -10.3 ($t = -2.1$). Finally, reconstructing the attention-composition measure using only news from individual beta or size-sorted deciles shows that the conditional CAPM emerges broadly across constructions, with significant beta slopes in the majority of deciles for both sorts, indicating that the result is not driven by news from

one particular segment of the cross section.

A natural concern is that the news on which attention anchors, although it survives filters for market-wide and earnings content, still conveys information about the aggregate factor. In this case, the security market line would be steep because attention resolves factor uncertainty rather than because it leaves it unresolved. Two pieces of evidence rule this out. Rebuilding the attention residual separately from the firm news that moves betas and the firm news that does not yields a price of beta of comparable magnitude in both cases, whereas a systematic-information channel would require the pricing to load on the beta-moving news alone. And although betas do rise on news days, by 0.074, the increase reflects a roughly 40 percent rise in idiosyncratic volatility against only a 5 percent rise in market covariance, so firm news makes firms noisier rather than more systematic.

The findings also speak to a broader question about what investor attention does to asset prices. [Peng and Xiong \(2006\)](#) argue that limited-capacity investors exhibit category-learning behavior, choosing between processing market-wide and firm-specific information. [Veldkamp \(2006b\)](#) shows that costly information acquisition generates excess comovement when investors substitute toward common signals. [Kacperczyk, Van Nieuwerburgh, and Veldkamp \(2016\)](#) model attention reallocation across business-cycle states and show that attention to systematic versus idiosyncratic risks varies endogenously. The mechanism developed here builds on this tradition but with a distinct emphasis: under the orthogonality structure motivated by the LLM filter, attention composition has an unconditional effect on factor uncertainty that does not rely on equilibrium feedback through prices, learning from prices, or general-equilibrium attention-supply effects.

This positioning has a deeper interpretive content. [Berk and van Binsbergen \(2016\)](#) document that mutual fund flows respond to CAPM alphas more strongly than to alphas from any multi-factor alternative, concluding that investors use the CAPM. Yet the unconditional CAPM is rejected in cross-sectional tests. One reconciliation is that investors form beliefs according to the CAPM but the cross-sectional pricing implication only emerges when their attention is actively engaged in processing the firm-specific information that translates beliefs into prices. The conditional result documented here is consistent with this view: the CAPM emerges in the cross section on the days when institutional attention is anchored to firm-news processing, and recedes when attention drifts elsewhere.

The remainder of the paper proceeds as follows: section 2, 3 and 4 introduce the model; section 5 describes the data and constructs the attention measures; section 6 presents the empirical results; section 7 conducts a battery of robustness tests; section 8 explores a simple trading strategy which exploits these periods. Finally, section 9 concludes.

2 Model

This section presents a noisy rational expectations model of asset prices with costly attention. Investors are modeled in the tradition of [Grossman and Stiglitz \(1980\)](#) and [Admati \(1985\)](#): they are risk averse, receive Gaussian signals, and learn only imperfectly about the assets they trade. As in [Veldkamp \(2006b\)](#), costly information acquisition leads investors to rely more heavily on common signals, and as in [Peng and Xiong \(2006\)](#), investors with limited attention choose between processing market-wide and firm-specific information. The model places this choice inside a factor structure: asset payoffs load on a common factor, and each investor divides a fixed attention budget between firm-specific public news, which is orthogonal to the factor, and signals about the factor itself. When attention concentrates on firm-specific news, the common factor is learned less well and posterior factor uncertainty remains high. Because beta is compensation for bearing factor risk, investors then demand a larger premium per unit of beta, and the security market line steepens. As attention shifts toward the factor, that uncertainty resolves, the premium contracts, and the security market line flattens.

The economy lasts a single period and contains N risky assets and a risk-free asset normalized to zero return. The payoff to risky asset $i \in \{1, \dots, N\}$ is

$$d_i = \bar{d} + \beta_i f + \varepsilon_i, \quad (1)$$

where $\bar{d}$ is a known constant, $f \sim \mathcal{N}(0, \sigma_f^2)$ is a common factor, and $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ are mutually independent firm-specific shocks. The factor loading is

$$\beta_i = b_i + \nu_i, \quad \nu_i \sim \mathcal{N}(0, \sigma_\nu^2), \quad (2)$$

with b_i a known prior mean and ν_i mutually independent across i and orthogonal to (f, ε_i) . I write precisions as $\tau_f \equiv \sigma_f^{-2}$, $\tau_\nu \equiv \sigma_\nu^{-2}$, $\tau_\varepsilon \equiv \sigma_\varepsilon^{-2}$.

A continuum of investors of unit mass, indexed by $j \in [0, 1]$, has CARA preferences over terminal wealth with common risk aversion $\rho > 0$. Per-capita asset supply is $\bar{x} + u_i$, where $u_i \sim \mathcal{N}(0, \sigma_u^2)$ is a noisy supply shock that prevents prices from being fully revealing.

2.1 Information

Investors receive two distinct signals: firm-specific public news and private macro information.

Assumption 1 (Firm-specific public signal). *For each asset i , all investors observe*

$$y_i = \varepsilon_i + \eta_i, \quad \eta_i \sim \mathcal{N}(0, 1/\tau_y), \quad (3)$$

where η_i is mutually independent across i and orthogonal to $(f, \beta_i, \varepsilon_i)$.

Assumption 1 treats firm-level news as informationally orthogonal to the aggregate factor. Empirically, this corresponds to news classified as purely firm-specific, including earnings releases, management changes, product announcements, after market-level content has been filtered out. Public signals inform investors about ε_i , but carry no information about f or about β_i . This restriction distinguishes the setup from standard noisy-rational-expectations models, in which a single public signal jointly conveys both firm-specific and factor information.

Assumption 2 (Private signal). *Each investor j privately observes*

$$z_j = f + \xi_j, \quad \xi_j \sim \mathcal{N}(0, 1/\tau_\xi), \quad (4)$$

with ξ_j independent across investors and orthogonal to $(f, \varepsilon_i, \nu_i, \eta_i)$.

2.2 Attention allocation

Each investor allocates a unit attention budget between the two channels. Let $\omega_j \in [0, 1]$ denote the share allocated to firm-specific public news. Public-signal precision scales with day- t news supply S :

$$\tau_y = \omega_j S, \quad (5)$$

and private precision scales with investor-specific capacity K :

$$\tau_\xi = (1 - \omega_j) K. \quad (6)$$

Higher ω_j reflects greater concentration of attention on firm-specific news; lower ω_j reflects diversion to private macro analysis.

3 Equilibrium

3.1 Posterior beliefs

Lemma 1 (Posteriors). *Under Assumptions 1–2, the posterior distributions are mutually independent and given by*

$$f \mid z_j \sim \mathcal{N}(\mu_f, V_f), \quad V_f = \frac{1}{\tau_f + (1 - \omega_j)K}, \quad (7)$$

$$\beta_i \mid y_i \sim \mathcal{N}(b_i, \sigma_\nu^2), \quad (8)$$

$$\varepsilon_i \mid y_i \sim \mathcal{N}(\mu_{\varepsilon,i}, V_\varepsilon), \quad V_\varepsilon = \frac{1}{\tau_\varepsilon + \omega_j S}. \quad (9)$$

Lemma 1 reflects three features. First, the posterior of f is determined entirely by the private signal: under Assumption 1, public news is uninformative about the factor. Second, the posterior of β_i equals its prior, because y_i carries no information about loadings. Third, public news reduces the posterior variance of ε_i in proportion to $\omega_j S$.

Define the residual non-factor variance entering the equilibrium pricing problem:

$$V_\varepsilon^\star = V_\varepsilon + \sigma_\nu^2(V_f + \sigma_f^2). \quad (10)$$

This combines posterior idiosyncratic variance with the contribution of unresolved loading uncertainty interacting with the factor.

3.2 Equilibrium prices

Each investor solves a standard CARA-Gaussian portfolio problem, generating linear demand. Market clearing pins down equilibrium prices.

Lemma 2 (Prices). *The equilibrium price of asset i is*

$$p_i = \bar{d} + b_i \bar{\mu}_f - \rho V_f b_i (N \bar{b} \bar{x}) - \rho V_\varepsilon^\star \bar{x} + (\text{linear noise}), \quad (11)$$

where $\bar{\mu}_f$ is the average posterior mean of f across investors and $\bar{b} \equiv N^{-1} \sum_i b_i$.

The risk premium $\mathbb{E}[d_i - p_i]$ decomposes into a factor component, $\rho V_f b_i (N \bar{b} \bar{x})$, proportional to b_i , and a non-factor component, $\rho V_\varepsilon^\star \bar{x}$, common to all assets. The factor component is what a cross-sectional regression of returns on betas identifies.

3.3 The Fama–MacBeth slope

In practice, the econometrician estimates loadings from a T -period rolling-window OLS regression of realized returns on the factor and then runs a Fama–MacBeth (FM) cross-sectional regression. The resulting slope is

$$\lambda^{FM}(\omega) = A(\omega) \cdot \rho V_f(\omega) N \bar{b} \bar{x}, \quad (12)$$

where $A(\omega) \in (0, 1)$ is the errors-in-variables attenuation:

$$A(\omega) = \frac{s_\beta^2 - \sigma_\nu^2}{s_\beta^2 + V_e(\omega)}, \quad V_e(\omega) = \frac{\sigma_\varepsilon^2 + 1/(\omega S)}{T \sigma_f^2}, \quad (13)$$

and $s_\beta^2 \equiv \sigma_b^2 + \sigma_\nu^2$ is the cross-sectional variance of true betas.

4 Optimal attention and the risk-premium channel

4.1 Investor’s information problem

Each investor chooses ω_j to maximize the expected utility gain from information. Under CARA-Gaussian preferences, the information-allocation objective reduces to (see Appendix E):

$$\mathcal{V}(\omega) = -\frac{N-1}{2} \log V_\varepsilon^*(\omega) - \frac{1}{2} \log \left(V_\varepsilon^*(\omega) + V_f(\omega) \sum_i b_i^2 \right). \quad (14)$$

The first term captures the cumulative gain from reducing residual non-factor uncertainty across N assets. The second term captures the cost of bearing factor risk: as V_f rises, the factor eigenvalue of the posterior covariance matrix grows, increasing the variance of the diversified factor exposure.

Proposition 1 (Interior optimum). *For all admissible parameter values $(N, S, K, \rho, \sigma_f, \sigma_\nu, \sigma_\varepsilon, \tau_f)$, the optimal attention allocation $\omega^* \in (0, 1)$ is interior and satisfies the first-order condition $\mathcal{V}'(\omega^*) = 0$.*

The interior optimum reflects a substantive trade-off. Moving attention toward firm-specific news improves idiosyncratic learning (reducing V_ε) but, because firm news is orthogonal to the factor by Assumption 1, every unit of attention reallocated to public news is one unit lost from private factor learning. The investor balances these two channels.

4.2 The risk-premium channel

The model's central economic content concerns the comparative static of V_f with respect to ω .

Proposition 2 (Risk-premium channel). *Posterior factor uncertainty is strictly increasing in attention to firm-specific news:*

$$\frac{dV_f}{d\omega} = K \cdot V_f^2 > 0. \quad (15)$$

Proof. From $V_f = (\tau_f + (1 - \omega)K)^{-1}$,

$$\frac{dV_f}{d\omega} = -V_f^2 \cdot \frac{d}{d\omega} [\tau_f + (1 - \omega)K] = KV_f^2 > 0. \quad \square$$

Proposition 2 holds unconditionally as it does not depend on interior optimality, on parameter calibration, or on any equilibrium FOC. Under Assumption 1, firm-specific news carries no factor information; the only channel through which investors learn about f is the private signal z_j . Therefore each additional unit of attention allocated to firm-specific news is one unit of attention lost from factor learning, and the posterior variance of f rises in ω at the deterministic rate KV_f^2 .

4.3 Conditional steepness of the SML

Combining (12)–(13) with Proposition 2 yields the model's main empirical prediction.

Proposition 3 (Conditional CAPM steepness). *The Fama–MacBeth slope is strictly increasing in the share of attention allocated to firm-specific news:*

$$\frac{d\lambda^{FM}}{d\omega} = \rho N \bar{b} \bar{x} \left[A(\omega) \cdot KV_f^2 + V_f \cdot \frac{dA}{d\omega} \right] > 0. \quad (16)$$

The first term inside the bracket is the risk-premium channel (Proposition 2); the second is an errors-in-variables channel through $dA/d\omega \geq 0$.

When investors concentrate attention on firm-specific news, they bear more factor uncertainty in equilibrium. To clear markets, expected returns must compensate for this uncertainty. Cross-sectional pricing aligns more sharply with betas: the security market line steepens, and unconditional betas explain the cross-section of average returns. Conversely,

when attention is diverted toward private macro analysis or broad-market information, posterior factor uncertainty falls, the factor risk premium shrinks, and the SML flattens.

Across days, variation in ω^* can be driven by parameters not explicitly modeled here: time-variation in private-research capacity K , or in investor composition. The model does not pin down the cross-day determinants of ω_t^* , but it is clear about the implication once ω_t^* is observed: high- ω days produce steep security market lines; low- ω days produce flat ones.

5 Data and Variable Construction

5.0.1 News data

As opposed to [Savor and Wilson \(2014\)](#), [Chan and Marsh \(2022\)](#) and [Andrei, Friedman, and Ozel \(2023\)](#) who focus their analysis on days with scheduled market-relevant announcements, I focus my analysis on news that is firm-specific and not regularly scheduled. My news comes from the Dow Jones Institutional Newswire (DJIN) accessed through ProQuest TDM Studio. To search for news, I create a string containing the name of every company listed in the Russell 3000 database since 2009 when DJIN coverage starts. I then perform a search through the DJIN for articles whose title contains at least one unique CRSP company name. I keep only articles for which a company name is included in the title of the newswire article, as [Ait-Sahalia, Li, and Li \(2024\)](#) find these to be most relevant to the firm, and likely to cause immediate jumps in stock prices. To avoid capturing news about earnings, since they are known to also convey systematic information, I exclude articles for which the title contains the word stem ‘earning’, as well as other words commonly associated with articles written about earnings reports. I further exclude news whose headlines include the words “Buzz”, “Wrap up”, and any other commonly repeated keywords associated with summaries of previously posted actual news. A full list of these filters is provided in the Appendix.

Newswire coverage is often noisy, and only weakly informative about firm fundamentals, making it difficult to locate relevant articles. A further challenge is to separate truly idiosyncratic news items from broad or market-wide news. Recent advances in large language models permit more refined filtering of such news, allowing me to isolate truly idiosyncratic news with greater precision. To refine the set of news articles used in the analysis, and focus on articles that are directly related to an individual firm, I use GPT-4o mini provided by TDM Studio to filter out news which is systematic or not relevant to individual firms. I use a prompt that instructs the model to use only information included in the text, and opt for a straightforward prompt to establish a foundation for the results. Nonetheless, employing

more detailed prompts could allow for the extraction of more tailored information. The purpose of this step is to keep only articles written about individual firms and singular (as opposed to systematic) events. I ask GPT-4o mini to read each article and classify whether the article is related to one firm or many, and whether the article is written about one event or many (market-wide) events. The exact prompt is shown in the Appendix. Once Chat-GPT has read and labelled each of the initial 902,873 articles, I keep only those labelled as ‘singular’ and ‘idiosyncratic’. The final GPT-labelled news corpus consists of 512,733 news articles.

5.0.2 *Attention data*

I follow [Ben-Rephael, Carlin, Da, and Israelsen \(2021\)](#) in constructing daily stock-level attention measures from Bloomberg terminals. Bloomberg records the number of times terminal users search for or read news articles on particular stocks. Bloomberg aggregates these numbers into hourly counts and creates an attention score by comparing the average hourly count during the previous eight hours to all hourly counts over the previous month for that same stock. They assign a score of 0, 1, 2, 3, or 4 if the rolling average is less than 80%, between 80% and 90%, between 90% and 94%, between 94% and 96%, and greater than 96% of the hourly counts over the previous 30 days. They then aggregate these scores up to a daily frequency by taking a maximum of all hourly scores through the day.

5.0.3 *Stock return data*

I retrieve excess market returns from Kenneth French’s website. Following ([Savor and Wilson, 2014](#)) among others, I construct ten beta monthly-sorted portfolios using U.S. common stocks from CRSP which have a share code of 10 or 11, that match the sample of stocks with available attention data. The number of unique stocks is 2469. Daily stock betas are obtained from a rolling regression of the past 252 daily excess stock returns on the excess market returns. At the beginning of each month, stocks are sorted into one of ten beta-deciles, and the daily value-weighted and equal-weighted excess return of each decile portfolio is computed over the month.

For each portfolio, I estimate daily rolling betas by regressing the past 252 excess returns of each portfolio on the market return. These ten beta-sorted portfolios form my main test assets. In robustness tests, I also consider individual stocks and characteristic-sorted portfolios.

5.0.4 Defining high attention concentration periods

To identify periods where attention is concentrated on firm-specific news, and bring the model’s attention-composition parameter ω to the data, I require a daily measure of how heavily institutional attention is tilted toward firm-specific news. I first develop a news attention index (IAI). Using Bloomberg query scores each day, I assign the attention value from 0-4 to individual stocks which had firm-specific news arrive through the Dow Jones Institutional Newswire, and take the daily average across only firms which have a news items on each day. As an additional step, I also recreate the aggregate institutional investor attention measure developed in [Ben-Rephael, Da, and Israelsen \(2017\)](#) using all Russell 3000 firms and create an average daily *AIA* following [Chan and Marsh \(2022\)](#), taking the mean level of attention paid to all firms for day t . 2469 unique stocks for which attention data and idiosyncratic news articles are available are used to create these attention measures. The sample runs from 2010 to 2025. I run the following daily expanding regression:

$$AIA_t = \alpha_{t-1} + \beta_{t-1} IAI_t + u_t, \quad (17)$$

I estimate $(\hat{\alpha}_{t-1}, \hat{\beta}_{t-1})$ using only information available up to day $t - 1$ in an expanding window, and define $\hat{u}_t$ as the resulting out-of-sample residual. $\hat{u}_t$ is the component of total institutional attention that is not explained by contemporaneous idiosyncratic-attention intensity, or ‘unanchored attention’: marketwide attention not accounted for by firm-level news attention. Figure 1 plots the time series of the residual. Its lowest readings, when attention is most anchored on firm-specific news, occur in the calm years of 2016 and 2017 and during the firm-news-intensive period of mid-2021; its highest readings occur during the 2022 tightening cycle, when macroeconomic news dominated, and around the volatility episode of August 2024.

The expanding-window construction serves two purposes. First, because each day’s residual uses only information available through the prior day, the measure is free of look-ahead bias: any relationship documented between the residual and the conditional slope of the security market line cannot be an artifact of using future data to construct the regressor. Second, the out-of-sample residual measures the innovation in non-news attention: how much higher or lower aggregate attention is than the historical news–attention relationship would predict, rather than its absolute level. This is the appropriate object for the test, which concerns cross-day variation in the slope of the security market line rather than its level. The model predicts that the line is steeper on days when attention is composed toward firm-specific news (low residual) than on days when it is tilted toward market-wide learning (high residual).

Provided the estimated news–attention coefficient is stable over the sample, the expanding residual coincides with a contemporaneous orthogonalization of attention, and its time variation reflects shifts in attention composition rather than drift in the estimated relationship. Section 7.7 considers alternative measures of attention concentration on idiosyncratic news.

Days when this residual is low are days when average institutional firm-level attention is well explained by firm-specific news. Panel B of Table 1 shows how other state variables behave within quartiles of the residual measure. Macroeconomic attention of Fisher, Martineau, and Sheng (2022) rises monotonically from the lowest residual quartile to the highest, suggesting that investors substitute idiosyncratic learning for macroeconomic learning, in line with the empirical findings of Cookson, García, and Jarnecic (2026). Building from the theory section, these should be days where investors demand a risk premium for market factor exposure. The empirical results test the CAPM on days when this residual is low.

6 Empirical Analysis

6.1 Testing the CAPM across attention states

The theoretical model predicts that the beta–return relationship should be increasing in the level of attention devoted to processing firm-specific news. Figures 2 and 3 depict the main result, plotting the security market line for both value-weighted and equal-weighted portfolios separately for quantiles of $\hat{u}_t$. For the figures, I use unconditional full-sample betas with excess returns averaged over the respective sample periods, following Savor and Wilson (2014) and Hendershott, Livdan, and Rösch (2020). When $\hat{u}_t$ falls below its trailing one-year 50th percentile, aggregate institutional attention is well explained by idiosyncratic news attention, and the security market line slopes sharply upward. On the complementary set of days, the SML is flat.

I push the analysis further using both pooled panel regressions and Fama–MacBeth cross-sectional regressions. For the pooled specification, following Savor and Wilson (2014), Chan and Marsh (2022), and Hendershott, Livdan, and Rösch (2020), I estimate

$$r_{p,t+1} = \alpha + b_1 \hat{\beta}_{p,t} + b_2 \mathbb{1}\{\text{State}_{t+1}\} + b_3 \hat{\beta}_{p,t} \mathbb{1}\{\text{State}_{t+1}\} + \varepsilon_{p,t+1}, \quad (18)$$

where $r_{p,t+1}$ is the excess return of portfolio p on day $t + 1$, $\hat{\beta}_{p,t}$ is the rolling one-year portfolio beta estimated from the past 252 daily excess returns of each portfolio on the excess market return, and $\mathbb{1}\{\text{State}_{t+1}\}$ is a dummy variable equal to one when the five-day

rolling median of $\hat{u}$, evaluated through day t , falls below (or above) the indicated trailing one-year quantile cutoff of the daily residual, also evaluated through day t . The classification of day $t + 1$ therefore uses only information available at the end of day t , mirroring the timing of the rolling beta, so that the state is known before the return is earned. The five-day median and the one-day lag each serve a distinct purpose. The daily residual measures a persistent attention regime with classification noise around it, and the median filters out marginal days admitted by that noise; the lag ensures that the conditional pricing results are predictive rather than a reflection of same-day co-movement between attention and returns. The [Internet Appendix](#) reports the full grid of smoothing windows and signal lags and shows that the results below are insensitive to these choices within the horizon over which the attention regime persists. Because the five-day median is less variable than the daily residual, it falls below the trailing 25th percentile of the daily distribution on fewer than a quarter of days, so the bottom-quartile state comprises 18.0% of the sample rather than 25%. The coefficient b_1 captures the unconditional beta–return relation on non-state days, b_2 captures the state-minus-other-days alpha, and b_3 measures the incremental change in the slope of the security market line on state days relative to all other days. Standard errors are clustered at the daily level.

For the Fama–MacBeth specification, I estimate daily cross-sectional regressions of portfolio excess returns on rolling portfolio betas separately within each attention state:

$$r_{p,t+1}^S = a^S + b^S \hat{\beta}_{p,t} + \varepsilon_{p,t+1}^S, \quad (19)$$

where the superscript S indexes the attention state, and $\hat{\beta}_{p,t}$ is the prior-day rolling portfolio beta. I report time-series averages of the daily slope coefficients b^S with Newey–West t -statistics computed with five lags on the time series of the estimated daily coefficients. I report results for four attention states defined by the rolling quantile cutoffs of $\hat{u}$: the bottom quartile (0–25%), the bottom half (<50%), the top half (>50%), and the top quartile (>75%).

Table 2 presents the results for both value-weighted and equal-weighted beta-sorted portfolios across all four attention states.

Panel A: Bottom-quartile residual state (0–25%). This state identifies the 19.0% of trading days on which aggregate institutional attention is most heavily anchored to firm-specific news. In Fama–MacBeth regressions, a one-unit increase in beta is associated with an increase in next-day excess returns of 23.1 basis points for value-weighted portfolios ($t = 2.85$), and 22.6 basis points for equal-weighted portfolios ($t = 3.07$). The pricing of beta is therefore equally strong under both weighting schemes in this state. The value-

weighted intercept is -7.00 basis points ($t = -1.07$), statistically indistinguishable from the zero intercept predicted by the CAPM. The average cross-sectional R^2 is 43.7%, indicating that variation in beta explains a large share of the cross-sectional return variation on these days.

The pooled regressions corroborate the Fama–MacBeth findings. On non-state days, the unconditional beta coefficient is 1.89 basis points ($t = 0.54$), economically zero, reaffirming the well-documented failure of the unconditional CAPM. The interaction term b_3 , which measures the incremental slope of the SML on bottom-quartile residual days, is 16.3 basis points ($t = 2.18$), statistically significant at the 5% level. The state dummy b_2 is -2.89 basis points ($t = -0.46$), statistically indistinguishable from zero. For equal-weighted portfolios, the interaction coefficient is 16.2 basis points ($t = 2.28$), also significant at the 5% level.

Panel B: Below-median residual state (<50%). Expanding the state definition to include all days below the trailing median of $\hat{u}$ captures 52.8% of trading days, slightly more than half of the sample. In value-weighted Fama–MacBeth regressions, the beta coefficient is 12.6 basis points ($t = 3.00$), highly significant. Importantly, the intercept is -2.98 basis points ($t = -0.91$), statistically indistinguishable from zero, as the CAPM predicts. The average R^2 is 45.1%. For equal-weighted portfolios, the beta coefficient is 8.7 basis points ($t = 2.21$), significant at the 5% level.

The pooled regressions deliver the same conclusion. The value-weighted interaction term is 18.7 basis points ($t = 3.02$), significant at the 1% level, while the unconditional beta coefficient b_1 is -4.42 basis points ($t = -0.96$), statistically indistinguishable from zero, confirming that beta does not explain excess returns on days outside this state. The implied within-state pooled slope, $b_1 + b_3 = 14.3$ basis points, is close to the Fama–MacBeth estimate of 12.6 basis points, providing reassurance that the result is not driven by differences in the econometric approach. The state dummy is -10.5 basis points ($t = -1.83$), marginally negative. For equal-weighted portfolios, the interaction is 13.4 basis points ($t = 2.31$), significant at the 5% level. The finding that unconditional betas price the cross section of returns on more than half of trading days constitutes a substantial expansion of the set of days on which the CAPM has been shown to hold, extending well beyond the scheduled announcement days identified in [Savor and Wilson \(2014\)](#) and the leading-earnings-announcement days of [Chan and Marsh \(2022\)](#).

Panel C: Above-median residual state (>50%). On the complementary set of days on which aggregate institutional attention is elevated beyond what idiosyncratic news ex-

plains, the pricing of beta breaks down. In value-weighted Fama–MacBeth regressions, the beta coefficient is -3.4 basis points ($t = -0.71$), economically negligible and statistically insignificant, and the intercept is 5.34 basis points ($t = 1.37$), positive but imprecisely estimated. The SML is effectively flat: high-beta and low-beta portfolios earn similar average returns in this state, and the return variation is not explained by differences in systematic risk. For equal-weighted portfolios, the beta coefficient is -4.2 basis points ($t = -0.93$), and the intercept is 6.32 basis points ($t = 2.02$), significantly positive.

The pooled regressions for this state are estimated using a dummy equal to one on above-median residual days. The value-weighted interaction term is -18.7 basis points ($t = -3.02$), significant at the 1% level, while the unconditional beta coefficient, now capturing below-median days as the base, is 14.3 basis points ($t = 3.46$). The significantly negative interaction therefore reflects the steepness of the security market line on the below-median base days rather than a significantly inverted line within the state: the beta–return relation steepens when attention is concentrated on firm-specific news and flattens when it is not.

Panel D: Top-quartile residual state (>75%). The most extreme attention-diversion state, accounting for 18.0% of days on which unanchored attention is highest, produces no detectable pricing of beta. In value-weighted Fama–MacBeth regressions, the beta coefficient is -1.17 basis points ($t = -0.15$), indistinguishable from zero. The intercept is 8.73 basis points ($t = 1.52$), positive but imprecisely estimated. For equal-weighted portfolios, the slope is 3.0 basis points ($t = 0.40$), with an intercept of 8.12 basis points ($t = 1.65$). In pooled regressions, the value-weighted interaction is -7.3 basis points ($t = -0.88$), insignificant; the equal-weighted interaction is 1.17 basis points ($t = 0.16$), also insignificant. Neither weighting scheme produces evidence that beta is priced on these days. Notably, the security market line in this state is flat rather than inverted. Under contemporaneous classification with an unsmoothed daily state variable, the top-quartile slope is significantly negative, but the [Internet Appendix](#) shows that this inversion reflects same-day co-movement between attention and returns and disappears under any predetermined classification; the predictive statement supported by the data is that beta goes unpriced, not that it is priced with the wrong sign.

Taken together, the results in Table 2 reveal a monotonic pattern in the pricing of beta across attention states. Moving from the bottom quartile to the top quartile of the residual $\hat{u}$, the Fama–MacBeth beta slope declines from 23.1 to -1.2 basis points for value-weighted portfolios, the pooled interaction coefficients decline from 16.3 to -7.3 basis points, and the intercepts move from statistically zero toward positive. The CAPM prices risk when

institutional attention is concentrated on public firm-specific news, even though the absolute level of attention on these days is below average. The results hold for both value-weighted and equal-weighted portfolios, across both Fama–MacBeth and pooled specifications, and are consistent with the model’s prediction that attention composition, rather than attention level, determines the equilibrium price of systematic risk.

6.2 Uncertainty and the Attention Premium

The model predicts that the composition premium documented above should be increasing in the level of prior factor uncertainty: on days when institutional attention is anchored to firm-specific news, beta should command a higher price the greater the unresolved uncertainty about the common factor that investors carry into the day. This section tests that implication.

The sensitivity of posterior factor uncertainty to attention composition is $dV_f/d\omega = KV_f^2$, and V_f is increasing in prior factor variance, so the same shift of attention toward firm-specific news forgoes factor learning that is more valuable, and moves the equilibrium price of beta by more, when background factor uncertainty is high. A second margin follows from the linearity of the premium in risk aversion: the composition effect on the price of beta scales with ρ , so it should also be amplified when the compensation investors demand per unit of risk is elevated. The model assigns no such role to sentiment or to belief dispersion, so those conditioning variables serve as predicted nulls. The prediction under test is therefore a set of cross-partial derivatives rather than a level effect: the composition premium should be larger following states of high uncertainty and high risk aversion, and elevated uncertainty without the firm-anchored attention state should not price beta.

6.2.1 Design and main result

I condition the firm-anchored state on the prior day’s level of uncertainty, measured four ways: trailing 21-day realized market variance (physical and backward looking), the VIX (option implied and forward looking), and the macroeconomic uncertainty index of [Jurado, Ludvigson, and Ng \(2015\)](#) at the one-month and twelve-month horizons (fundamental, estimated from a large panel of macroeconomic series, and available only monthly). Every conditioning variable is lagged: the attention state is classified from information through the close of day $t - 1$, the daily uncertainty proxies are $t - 1$ values, and the monthly indices are mapped by prior calendar month. Tercile breakpoints are computed within sample. Within the state, I estimate the Fama-MacBeth price of beta separately in each uncertainty tercile;

the placebo repeats the top-tercile estimation on unanchored days, where the model implies the amplification has nothing to amplify.

Table 3, Panel A, reports the result. The price of beta on firm-anchored days rises from small and mostly insignificant values in the two lower terciles to 24.2 basis points ($t = 2.14$) in the top tercile of realized variance, 31.8 ($t = 2.92$) for the VIX, 24.1 ($t = 2.34$) for JLN at $h=1$, and 24.9 ($t = 2.32$) at $h=12$. The agreement across proxies is close in magnitude despite the proxies sharing little beyond the concept: a backward-looking return-based measure, a forward-looking option-implied measure, and a monthly index estimated from macroeconomic fundamentals place the amplified premium within eight basis points of one another. The pattern is a threshold rather than a gradient. The middle tercile is insignificant for every proxy, and condensing the cut concentrates the effect further: at quintiles, the top-bin slope reaches 52.2 basis points ($t = 2.89$) for the VIX on 267 days, while lower bins remain flatter. The premium is not increasing smoothly in uncertainty; it is concentrated on days following the most elevated uncertainty readings.

The placebo column establishes that uncertainty is an amplifier and not a channel. On unanchored days, the top uncertainty tercile prices beta at -8.6 basis points ($t = -0.82$) for realized variance, -2.8 ($t = -0.25$) for the VIX, and effectively zero for both JLN horizons. Elevated uncertainty without the firm-anchored attention state produces nothing. The result survives the removal of 2020 and 2022 in full: the top-tercile slopes are 19.9 ($t = 2.81$), 21.5 ($t = 2.92$), 16.2 ($t = 2.16$), and 16.1 ($t = 2.28$) across the four proxies, and the top-tercile state days span more than thirty calendar quarters with no quarter contributing more than ten percent, so no single episode carries the cell. A median split of each proxy, which involves no tercile choice, delivers the same conclusion in coarser form: the high half prices at 18.9 to 21.7 basis points for the daily proxies with the low halves quiet, and the high-minus-low difference is 17.7 basis points ($t = 2.09$) for realized variance.

6.2.2 *Composition conditions, uncertainty scales*

A natural question is whether these cells reveal an uncertainty effect rather than an attention effect. It is possible that beta is simply priced when uncertainty is high, and the attention state is incidental. Panel B of Table 3 rejects this interpretation within the top tercile itself. Restricting the sample to top-tercile days of each uncertainty proxy, I estimate the pooled panel regression

$$r_{p,t}^e = a + b_1 \beta_{p,t} + b_2 D_t + b_3 \beta_{p,t} D_t + \varepsilon_{p,t}, \quad (20)$$

where $r_{p,t}^e$ is the day- t excess return of beta decile p , $\beta_{p,t}$ is its lagged rolling beta, D_t is the firm-anchored state indicator classified through the close of day $t - 1$, and standard errors are clustered by day. The coefficient b_1 measures the price of beta on unanchored high-uncertainty days, and b_3 measures the additional price of beta when attention is firm-anchored, holding the uncertainty regime fixed. The estimates of b_3 are 33.8 basis points ($t = 2.33$) for realized variance and 30.3 ($t = 2.07$) for the VIX, while b_1 is never significant in any column and its point estimates are negative throughout. The same holds inside every uncertainty regime, not only the top tercile: b_3 is positive in all twelve tercile-by-proxy cells, and the daily Fama-MacBeth slope loads negatively on the lagged attention residual within every tercile, including the calmest third of the sample, where the coefficient is -7.4 basis points per standard deviation ($t = -2.49$). The two conditioning variables are also orthogonal by measurement: the correlation between the attention residual and option-implied variance is 0.009, so the state is not an uncertainty proxy and its tercile cells are populated evenly.

The division of labor is therefore the one the model specifies. Composition conditions the pricing of beta at every level of uncertainty; uncertainty scales the premium that composition commands. Removing the high-uncertainty days from the sample leaves a below-median slope near its unconditional value; removing the attention measure and conditioning on uncertainty alone leaves nothing.

6.2.3 *The risk-aversion margin*

The second cross-partial operates through ρ . The variance risk premium, the gap between option-implied and realized variance, is the standard proxy for the effective price of risk (Bollerslev, Tauchen, and Zhou, 2009), and because the wedge is mechanically related to the uncertainty level, I hold the level fixed: within the anchored top-SVIX²-tercile cell, days are split at the median variance risk premium, where SVIX² is the option-implied lower bound on the conditional equity premium of Martin (2017). Panel C reports the result. The high-wedge half prices beta at 51.8 basis points ($t = 3.60$); the low-wedge half prices it at 10.4 ($t = 0.63$). Excluding 2020 and 2022, the split is 36.0 ($t = 3.69$) against 14.7 ($t = 1.25$). Holding the quantity of uncertainty fixed, the composition premium concentrates on days when the price of risk is elevated, which is what linearity in ρ implies. The predicted nulls behave as predicted: neither investor sentiment nor belief dispersion amplifies the composition premium, with the sentiment interaction indistinguishable from zero.

The market's own premium concentrates in the same gated cells. The average market excess return on anchored top-tercile days is 16.6 basis points per day ($t = 3.59$), against 5.8

($t = 0.82$) on unanchored top-tercile days, and 14.2 ($t = 4.26$) excluding 2020 and 2022. The time series and the cross section agree on where compensation for systematic risk is earned.

6.2.4 Discussion

Two features of Table 3 fall outside the model’s predictions, and I report them plainly. First, the macroeconomic indices price the *lowest* tercile as well: 12.8 basis points ($t = 2.95$) at $h=1$ and 10.8 ($t = 2.51$) at $h=12$, a pattern that is robust across finer grids but absent from the daily market-based proxies. At monthly frequency the relation is U-shaped, with pricing at both extremes of the macroeconomic uncertainty distribution and none in the middle, whereas the daily proxies show a pure threshold. Second, in the amplified cells the security market line is steeper than the market’s own premium: top-tercile slopes run roughly 1.6 to 1.9 times the average market excess return in the same cells, with correspondingly negative zero-beta intercepts, significantly so in the high-wedge cell of Panel C (-28.6 , $t = -2.21$). High-beta portfolios earned more than beta times the realized premium on these days. Neither the Sharpe-Lintner version nor the [Black \(1972\)](#) version, which requires a positive zero-beta rate, accommodates this tilt, and the model’s common non-factor term would if anything predict the opposite sign. I disclose it as a characteristic of the amplified state and do not pursue it here.

Finally, the amplification bears on the relation to [Andrei, Friedman, and Ozel \(2023\)](#). Their mechanism also strengthens with uncertainty, so amplification alone does not distinguish the two channels; what distinguishes them is what conditions the pricing and which direction the mechanism moves posterior factor uncertainty. In their framework, uncertainty draws attention to scheduled announcements whose content loads on the aggregate factor, and the premium is earned as that uncertainty is resolved; consistent with this, their day-level evidence conditions on the attention level and on lagged volatility, and their beta-sorted portfolios show an unconditional steepening of roughly 6 basis points on high-VIX days. My state decomposes that unconditional steepening: mixing the two top-tercile cells at their day counts yields an unconditional high-uncertainty slope of roughly 14 basis points, all of which is carried by the firm-anchored subset and none by the remainder. The mechanisms also leave opposite footprints, which the following robustness exercises examine in detail: the composition premium requires no announcement and survives the excision of earnings windows, does not respond to the supply of news or the level of attention, leaves conditional betas unchanged where the resolution channel raises them, is flat from the first day of an anchored episode through its continuation rather than front-loaded at an information event, and shows no resolution signature, with the change in implied volatility on anchored days

indistinguishable from zero once the contemporaneous market return is removed and the premium present on days when uncertainty rises. Attention that resolves factor uncertainty lowers V_f and compresses the premium going forward; attention that crowds out factor learning leaves V_f standing and earns the premium continuously. The amplification documented here is what the second mechanism implies and the first does not.

6.3 Factor specificity and alternative test assets

The mechanism makes two predictions beyond the steepening of the security market line on beta-sorted portfolios. Because attention to firm-specific news moves posterior uncertainty about the market factor and nothing else, the state should change the price of market exposure alone, leaving the premia on size, value, profitability, and investment loadings unchanged across states; and because the premium that emerges on anchored days is compensation for exposure to the factor, market beta should price the cross section on those days regardless of the characteristic on which the test assets are sorted. The first prediction separates the mechanism from stories in which the state moves the price of systematic risk generally, such as time variation in risk aversion or sentiment, which would lift every premium together. The second separates it from the possibility that the state happens to coincide with the pricing of the beta sort alone.

To test the first prediction, I apply the paper’s cross-sectional design to each of the five Fama–French factors in turn. For each factor, every stock’s univariate loading is estimated over the trailing 252 trading days and lagged one day, stocks are sorted daily into value-weighted deciles on that loading, and the daily cross-sectional slope of decile returns on decile loadings is averaged within each attention state; a pooled regression with the loading interacted with the firm-anchored indicator provides the formal test of the state difference. Table 4 reports the results. The market row reproduces the headline result on deciles rebuilt from scratch: the price of market beta is 9.4 basis points per day ($t = 3.06$) on anchored days and -2.7 ($t = -0.76$) on unanchored days, a difference of 12.8 basis points ($t = 2.74$). No other loading shows a significant state difference. The value, profitability, and investment loadings are priced in neither state, with state differences of -0.1 , -1.2 , and 0.4 basis points (t between -0.35 and 0.20). The size loading is priced on anchored days in the univariate design (3.9 basis points, $t = 2.53$), but its state difference is insignificant (4.6, $t = 1.31$), and the anchored-day slope reflects market exposure carried by the size sort rather than a priced size loading: univariate SMB loadings have a cross-sectional correlation of 0.66 with market betas, and once the CAPM beta is held fixed, by regressing individual-stock returns on the beta and the SMB loading orthogonalized against it, the SMB coefficient falls to 1.2 basis

points ($t = 1.15$). Panel B gives the time-series counterpart: the market’s own premium is 9.30 basis points per day on anchored days against 1.44 on unanchored days ($t = 2.16$ for the difference), while the other four premia show no significant concentration in either state. The state therefore moves the price of aggregate market risk specifically, which is what a mechanism operating on posterior uncertainty about a single market factor predicts, and the three nulls join disagreement and sentiment as predicted zeros that a risk-aversion or broad-attention account would not deliver.

To test the second prediction, Figure 4 pools the test assets of the literature into a single cross section, following [Savor and Wilson \(2014\)](#), who estimate their announcement-day line jointly across beta, size and book-to-market, and industry portfolios. The figure plots forty value-weighted portfolios, ten deciles each sorted on beta, size, book-to-market, and momentum, each with a single full-sample beta so that the two states share the same horizontal positions and differ only in mean returns. Three of the four families are sorted on characteristics unrelated to beta, and two of them, size and momentum, are the sorts on which the unconditional CAPM fails most visibly. On firm-anchored days the forty portfolios line up along a security market line with slope 15.0 basis points per day that passes near the origin, with the size, value, and momentum deciles falling on the same line as the beta deciles rather than above or below it. On unanchored days the same forty portfolios trace a flat line with slope -3.4 . On days when institutional attention is anchored on firm-specific news, average returns line up with market beta across whatever the test assets are sorted on, and the only loading whose price the state changes is the one the model says it should.

7 Additional tests

7.1 A continuous robustness battery

The state cutoffs in Table 2 discard information within each state, and a binary classification cannot say whether the pricing content of the attention residual survives conditioning on the variables that the literature has used to explain time variation in the security market line. Table 5 therefore replaces the state with its underlying continuous signal. The dependent variable is the daily Fama–MacBeth price of beta $\hat{\lambda}_t$, obtained each day by regressing beta-decile portfolio excess returns on their betas, and the regressor of interest is the five-day rolling median of the attention residual evaluated through day $t - 1$ and standardized. The coefficient is the change in the price of beta, in basis points per day, per standard deviation of the smoothed residual, and a negative coefficient indicates a steeper security market line

when institutional attention is anchored on firm-specific news. All controls are dated $t - 1$ and standardized, and t -statistics use Newey–West standard errors with 21 lags.

The baseline gradient is -9.15 basis points per standard deviation ($t = -2.88$) for value-weighted portfolios and -7.73 ($t = -2.85$) for equal-weighted portfolios. This is the continuous counterpart of the state contrast in Table 2: the anchored and unanchored states sit roughly one and a half standard deviations apart on the smoothed residual, so the gradient reproduces the sixteen-basis-point gap between them without a threshold. Panel A then adds thirteen controls one at a time. The first group captures uncertainty and volatility: the VIX, the option-implied bound on the equity premium SVIX² of [Martin \(2017\)](#), realized market variance over the prior 21 days, the variance risk premium, and the cross-sectional dispersion of daily stock returns. The second group captures the level and supply of attention rather than its composition: the aggregate institutional attention index, the firm-news attention index, the share of firms with firm-specific newswire coverage, and the share with an earnings announcement. The third group captures market dynamics: cumulative market returns over the prior five and 21 days and the prior day’s price of beta. The last adds dummies for FOMC, CPI, PPI, and payroll announcement days and for the day of the week.

No control moves the gradient by more than a fifth, and most move it by less than five percent. The uncertainty group leaves it within four percent of the baseline, so the pricing content of the residual is not a relabeling of volatility regimes, a conclusion that holds whether volatility is measured by realized variance, by the VIX, by its option-implied premium bound, or by the gap between the two. The two attention-level controls leave it within seven percent, which restates inside a single regression the result of Section 7.7 that the levels of attention carry no pricing content of their own: what prices beta is the attachment of attention to firm-specific news, not how much attention there is. The lagged price of beta changes the gradient by two percent, so the residual does not merely mark persistent slope regimes, and the announcement calendar changes it by one percent, so the gradient is not an announcement effect in disguise. The only controls that attenuate the gradient at all are return dispersion, by ten percent for value-weighted and eighteen percent for equal-weighted portfolios, and coverage share, by seven and ten percent, and each removes part of the state’s own footprint rather than an alternative channel: anchored days are the calm regime in which cross-sectional dispersion is low, and the supply of firm-specific news is the principal determinant of the residual (Table 1). With all controls entered jointly the gradient is -10.34 ($t = -2.13$) for value-weighted and -9.98 ($t = -2.49$) for equal-weighted portfolios, larger in magnitude than the baseline; the wider standard errors reflect the dimensionality of the specification rather than any attenuation of the point estimate.

Panel B adds the two conditioning variables with limited coverage, the daily aggregate disagreement measure of [Cookson and Niessner \(2023\)](#) and the investor sentiment index of [Baker and Wurgler \(2006\)](#), on the 2,477 days through 2021 on which both are available, with a matched baseline so that changes in the coefficient reflect the controls rather than the sample. The gradient on this window is -7.70 ($t = -2.05$) and -7.56 ($t = -2.32$). Disagreement changes it by two to four percent and sentiment by one to eight percent, so neither the disagreement channel of [Hong and Sraer \(2016\)](#) nor a sentiment channel accounts for the pricing content of the residual. The joint specification places twenty regressors on 2,477 observations and leaves the point estimate within ten percent of the window’s baseline while its t -statistic, falls below two.

7.2 Expected information consumption and individual firms

A natural concern with the regime variable is that bottom-25% residual days may simply identify days when institutional investors are predictably consuming information at specific firms, in which case the conditional CAPM finding would be subsumed by Ben-Rephael, Carlin, Da, and Israelsen’s (2021) Expected Information Consumption (EIC) mechanism. To address this, I construct EIC at the firm-day level following their methodology: a firm receives $\text{EIC} = 1$ on day t if it has historically exhibited institutional attention spikes around peer-firm scheduled events, FOMC announcements, or macroeconomic announcements occurring on that day. The construction follows BCDI’s specifications closely: I use Fama–French 48-industry classifications to identify peer firms, require a minimum historical spike rate of 30% on past peer-earnings days for EIC_PEER , and apply BCDI’s 50% threshold for FOMC and macroeconomic announcements based on the previous four FOMC days and previous year of macro announcements. The construction produces frequencies and diagnostics in line with BCDI’s reported values. The combined $\text{EIC} = 1$ frequency is 5.6% of firm-days, comparable to BCDI’s reported 6.7% for EIC_ALL despite the more limited Bloomberg-coverage sample. The AIA spike rate among $\text{EIC} = 1$ firm-days is 30.5%, roughly five times the 5.9% spike rate among $\text{EIC} = 0$ firm-days, closely tracking BCDI’s reported differential of 24.2% versus 7.1% in their Table III. Panel A of Table 6 reports the conditional CAPM result on bottom-25% residual days, partitioned by EIC. The headline slope of 20.4 bps ($t = 3.15$) on the full cross-section drops only marginally to 19.1 bps ($t = 3.04$) when I restrict to $\text{EIC} = 0$ firms, which are those that BCDI’s framework specifically does not flag as expected to consume information. The conditional CAPM result is therefore not driven by firms that BCDI’s mechanism predicts. $\text{EIC} = 1$ firms exhibit particularly strong conditional pricing (44.2 bps, $t = 3.10$), within BCDI’s reported range of 8 to 44 bps for $\text{EIC} = 1$ subsamples and

consistent with their findings, indicating that the two mechanisms operate complementarily: BCDI’s framework identifies firm-day variation in attention-driven pricing, while the regime variable identifies day-level variation in the broader attention environment that affects the entire cross-section.

7.3 Earnings announcements and the conditional CAPM

A natural concern is that the conditional pricing result on bottom-quartile residual-attention days is driven by firms that are simultaneously releasing earnings. Earnings announcements convey systematic information and have been shown to steepen the security market line in their own right (Savor and Wilson, 2016; Chan and Marsh, 2022). If bottom-quartile residual days happen to coincide with earnings seasons, the apparent success of the CAPM on these days could simply reflect the announcement premium documented in prior work rather than a distinct attention-composition channel. To address this concern, I define earnings-window firm-days as firm-days for which the firm has an earnings announcement within five business days before or after the current date, and re-estimate the Fama–MacBeth cross-sectional regressions on bottom-quartile residual-attention days in three subsamples: all firms, firms excluding those in an earnings window, and earnings-window firms only.

Panel B of Table 6 reports the results. On the full cross section, the beta slope is 20.4 basis points ($t = 3.15$) with an insignificant intercept of -2.7 basis points ($t = -0.74$), consistent with the baseline findings. When earnings-window firms are excluded, the beta slope declines only marginally to 18.8 basis points ($t = 2.98$), remaining highly significant. The intercept is -1.8 basis points ($t = -0.49$), statistically indistinguishable from zero. This subsample retains an average of 1,390 firms per cross section, so the result is not an artifact of a thin sample. Restricting the cross section to earnings-window firms only produces a beta slope of 24.4 basis points ($t = 2.88$), larger in magnitude and consistent with the view that earnings announcements contribute to the pricing of beta through a complementary channel. The intercept for this subsample is 1.4 basis points ($t = 0.22$), again insignificant.

The key finding is that removing all firms near an earnings announcement barely affects the conditional beta slope. The decline from 20.4 to 18.8 basis points is economically small relative to the magnitude of the effect, and the t -statistic remains above conventional thresholds. The attention-composition channel therefore operates across the broad cross section of firms, not only among those releasing earnings. The stronger result within earnings-window firms is consistent with, but not required by, the main finding, and suggests that the two mechanisms operate complementarily: earnings announcements provide a firm-level informa-

tion event that reinforces the cross-sectional pricing of beta, while the attention-composition state variable identifies a market-wide regime in which that pricing extends to all firms.

7.4 Excluding macroeconomic announcement days and LEADs

A further concern is that the bottom-median residual-attention days may overlap with macroeconomic announcement days (Savor and Wilson, 2014) or with leading earnings announcement days (LEADs) identified by Chan and Marsh (2022). If so, the conditional pricing result could be confounded by the well-documented announcement premium on these days. To address this, I re-estimate the Fama–MacBeth regressions after removing these potentially confounding days from the sample.

Panel C of Table 6 reports the results. The first pair of rows excludes LEAD days, defined as days when influential S&P 500 firms announce earnings, following Chan and Marsh (2022). After removing these days, the value-weighted beta slope is 12.3 basis points ($t = 2.72$) and the equal-weighted slope is 8.5 basis points ($t = 2.13$), both statistically significant. The sample retains 1,343 days. The second pair of rows excludes macroeconomic announcement days, defined as days with scheduled FOMC, employment, or PPI releases following Savor and Wilson (2014). The value-weighted slope is 12.5 basis points ($t = 3.13$) and the equal-weighted slope is 8.6 basis points ($t = 2.38$), again both significant, with 1,554 days remaining. In both specifications, the intercepts are economically small and statistically insignificant, consistent with the predictions of the CAPM.

The magnitudes of the beta slopes are close to unchanged relative to the baseline results in Table 2, and the t -statistics remain well above conventional thresholds. The conditional pricing of beta on bottom-median residual-attention days is therefore not driven by overlap with known macroeconomic or bellwether-earnings announcement days. This finding, combined with the EIC and earnings-window results in Panels A and B of the same table, suggests that the attention-composition channel identifies a broad regime in the information environment that is distinct from the specific event-day channels documented in prior literature.

7.5 Does news carry systematic information?

A final concern is that the news on which attention is anchored, though it survives the market-wide and earnings filters, may still convey information about systematic conditions. If a firm’s news moves its beta, the market is learning about systematic conditions from that firm’s announcement, in the manner of Patton and Verardo (2012) and the expected-information-consumption channel of Ben-Rephael, Carlin, Da, and Israelsen (2021). Under

this reading the security market line would be steep on anchored days because attention resolves factor uncertainty, the mechanism of [Andrei, Friedman, and Ozel \(2023\)](#), rather than because it leaves factor uncertainty unresolved as the model here implies. The two accounts make opposite predictions about which news drives the result: if systematic content is responsible, the pricing should come from firm news that moves beta and not from firm news that does not.

I test this by rebuilding the state from beta-moving and beta-neutral news separately. For each firm I estimate the change in its beta on news days from a regression of its excess return on the market return, a news dummy, and their interaction, and I split firms at the median of this change. I then reconstruct the idiosyncratic-attention index from the news of each half, re-estimate the expanding regression to obtain a new residual, and re-run the cross-sectional test. Panel D of Table 6 reports the results. Rebuilding the state from the half of firms whose news moves beta least prices beta at 5.5 basis points per day on anchored days; rebuilding it from the half whose news moves beta most prices beta at 7.6. The two are of comparable magnitude and the same sign, and the beta-moving half does not price more strongly than the beta-neutral half. Because the systematic-information account requires the pricing to load on beta-moving news, this symmetry is evidence against it. A beta-matched version, which keeps the low-change half within quintiles of firm beta so that the retained firms match the full-sample beta distribution, gives 7.6 basis points, so the pattern is not an artifact of selecting low-beta firms; the composition of the two halves by beta and size is nearly identical.

Two features of the exercise should be read with care. Each half is built from half of the news firms, so its idiosyncratic-attention index correlates 0.78 with the baseline index and its state agrees with the baseline state on 55 percent of days; the price of beta in each half is correspondingly attenuated toward the weaker values reported, and neither is individually significant at conventional levels. The informative comparison is therefore between the two halves rather than the significance of either, and that comparison shows no advantage for the beta-moving news that the systematic-information channel would require. This is consistent with the behavior of betas around firm news more generally. Although betas do rise on news days, by 0.074 ($t = 4.56$), the increase reflects a roughly 40 percent rise in idiosyncratic volatility against only a 5 percent rise in market covariance, so the news on which attention anchors makes firms noisier rather than more systematic, and the conditional pricing of beta is not a relabeling of the resolution of aggregate uncertainty.

7.6 Which firms' news drives the result?

The baseline attention-composition measure constructs the idiosyncratic attention index IAI_t by averaging Bloomberg query scores across all firms with a filtered news article on day t . A natural question is whether the conditional pricing result is driven by news arriving at a specific segment of the cross section. If, for example, only news about high-beta firms or large firms generates the attention state that steepens the SML, the mechanism would be narrower than the model implies. To investigate this, I reconstruct the expanding regression in equation (17) ten times, each time computing IAI_t using only the subset of firms belonging to a single beta decile (or, separately, a single size decile). For each decile-specific residual, I identify bottom-median days and re-estimate the Fama–MacBeth cross-sectional regressions on those days. If the conditional CAPM result appears broadly across decile-specific constructions, then no single part of the firm distribution is responsible for the finding.

Table 9 reports the results when the attention state is defined using news from each beta decile separately. Panel B presents value-weighted portfolio results. Across the ten beta-decile-specific constructions, the beta slope λ_β is positive in every column, ranging from 5.3 basis points (decile 9) to 12.5 basis points (decile 1). The slope is statistically significant at the 5% level or better in six of the ten deciles (deciles 1, 3, 4, 5, 7, and 10), and significant at the 10% level in a seventh (decile 2, $t = 1.85$). The intercepts are uniformly small and insignificant, ranging from -4.4 to 1.4 basis points, none statistically different from zero. The average cross-sectional R^2 is approximately 46% throughout, indicating stable explanatory power regardless of which decile's news is used to construct the state variable. For equal-weighted portfolios in Panel A, the pattern is qualitatively similar though weaker in magnitude, with the beta slope significant in decile 1 (8.6 basis points, $t = 2.38$) and marginally significant in deciles 3, 4, and 5.

The strongest value-weighted results arise when the attention state is constructed from news about the lowest-beta firms (decile 1: 12.5 basis points, $t = 3.13$) and mid-range beta firms (deciles 3 through 5: 11.0 to 11.5 basis points, t -statistics between 2.7 and 2.9). The weakest results come from deciles 6 and 9. The broad pattern suggests that attention to firm-specific news steepens the SML regardless of where in the beta distribution that news originates, though the effect is somewhat attenuated when the state is defined using news from the highest-beta or highest-attention deciles alone.

Table 10 repeats the exercise using size deciles. Panel B again presents value-weighted results. The beta slope is positive in all ten columns, ranging from 5.7 basis points (decile 8) to 13.0 basis points (decile 7). Significance at the 5% level or better obtains in seven of the ten size deciles (deciles 1, 2, 4, 5, 6, 7, and 10), with an additional decile (decile 9: 11.2 basis

points, $t = 2.55$) also significant. The intercepts are again uniformly small and insignificant, consistent with the CAPM. The equal-weighted results in Panel A are weaker in magnitude, with significance at the 5% level in deciles 6 and 7 (7.4 and 7.3 basis points) and marginal significance in decile 5 (6.8 basis points, $t = 1.94$).

The size-decile results indicate that the attention-composition channel is not confined to news about firms of a particular size. News from both small firms (decile 1: 8.2 basis points, $t = 1.98$) and large firms (decile 10: 9.5 basis points, $t = 2.33$) generates attention states in which beta prices the cross section. The strongest results emerge when the state is constructed from mid-to-large capitalization firms (deciles 4 through 7: 10.9 to 13.0 basis points), which may reflect the fact that these firms account for a larger share of aggregate institutional attention and therefore produce more informative variation in the residual.

Taken together, the results in Tables 9 and 10 show that the conditional CAPM finding is not an artifact of news concentrated in any single segment of the cross section. The attention-composition state variable steepens the security market line when constructed from news about low-beta or high-beta firms, small or large firms, and most points in between. This breadth is consistent with the model’s prediction that the relevant object is the aggregate share of attention devoted to firm-specific news, not the characteristics of the firms producing that news.

7.7 Alternative measures of attention concentration

The baseline state is built from one particular reduction of the two attention series: the out-of-sample residual of aggregate institutional attention on firm-news attention, estimated with expanding coefficients. None of the results should depend on that choice. Table 7 therefore rebuilds the state from alternative reductions of the same two series, holding every other element of the construction fixed: each candidate series is smoothed by its five-day rolling median, compared with its trailing 252-day quantile, with the comparison direction set so that the state is firm-focused, and the completed indicator is lagged one trading day. Differences across rows are then attributable to the measure, not to the machinery. Three alternatives span the natural possibilities: the share of firm-news attention in aggregate attention, which involves no estimated parameters at all; a residual estimated on a trailing 252-day window, which lets the attention relation drift; and the gap between the trailing percentile ranks of the two series, which is invariant to any monotone transformation of either one.

The composition rows deliver the baseline result throughout. At the median cutoff, the price

of beta on firm-anchored days is 8.8 basis points ($t = 2.11$) for the share, 12.9 ($t = 2.79$) for the rolling residual, and 11.7 ($t = 2.42$) for the percentile gap, against 12.6 ($t = 3.00$) for the baseline, with complements flat in every row. At the quartile cutoff every composition measure roughly doubles, to 20.2 ($t = 2.89$), 17.2 ($t = 2.28$), and 28.6 ($t = 2.94$), matching the baseline’s 23.1 ($t = 2.85$): deepening the state concentrates the signal, whichever way the state is measured. The share is worth singling out because it is the rawest version, a ratio of the two published series with nothing estimated, and it prices at both cutoffs; the result does not depend on the residualization. Each measure’s pricing also lines up with its agreement with the baseline classification (0.83 to 0.85 for the three alternatives), consistent with all of them being noisy readings of one underlying composition state. Two coarser variants, a trailing correlation between the series and a tilt of firm-news relative to macro-topic attention, classify regimes rather than days, agree with the baseline barely above chance, and price weakly; I report them in the Internet Appendix.

The level rows show what the composition rows are not. Classifying days by the raw level of aggregate attention prices nothing: the AIA-high state is 0.4 ($t = 0.07$) at the median cutoff and -15.7 ($t = -1.69$) at the quartile, where its complement prices at 8.9 ($t = 2.64$). Days of elevated firm-news attention fare no better, at 3.1 ($t = 0.63$) and 1.1 ($t = 0.12$). The one level configuration that shows pricing, the complement of AIA-high, is the AIA-low state, and its pricing is inherited rather than its own: aggregate attention is the regressand of the residual, so low-AIA days mechanically overlap the firm-anchored state (68 percent of them are below-median residual days), and the AIA-low days that are not firm-anchored price beta at -4.2 basis points ($t = -0.67$).

Panel C completes the decomposition. In a pooled regression that lets the two indicators compete, the interaction of beta with the firm-anchored state is 20.0 basis points ($t = 3.06$) while the interaction with the AIA-low state is -4.4 ($t = -0.68$). And the conditioning survives within each half of the attention-level distribution: among low-attention days the firm-anchored interaction is 18.6 ($t = 2.42$) at the day level, and among high-attention days it is positive at the day level and significant in continuous form, with the daily slope loading on the lagged residual at -8.9 basis points per standard deviation ($t = -1.99$). The level of attention neither prices beta on its own nor displaces the composition state when the two compete; the composition of attention conditions the slope at every level of attention, just as it does within every uncertainty regime in Section 6.2. Full two-way cells and estimator detail appear in Internet Appendix.

7.8 Relation to Betting-against-Beta

The betting-against-beta factor is the traded counterpart of a flat security market line. It is long low-beta and short high-beta stocks with the market exposure hedged, so it earns its premium on exactly the days the line is flat or inverted and loses on the days it steepens. In our sample its daily return has a correlation of -0.52 with the cross-sectional price of risk λ_t , confirming that it moves as the mirror image of the slope, which makes it the natural object on which to test the model out of sample. Table 8 conditions the premium on the firm-specific attention regime. On the days when firm-specific attention dominates, the bottom quartile of the predetermined attention residual, the factor earns -0.28 basis points, indistinguishable from zero, while on the remaining days it earns 4.18 . The difference of -4.47 basis points carries a t -statistic of -1.63 in raw returns and strengthens to -2.08 and -2.07 once the factor is winsorized at the 1st/99th percentiles, a standard treatment for a series this crash-prone. The anomaly is thus economically and statistically absent on precisely the days the model predicts a steep, CAPM-consistent line, and is earned entirely on the rest.

The pattern is a discrete step rather than a smooth gradient. Across quintiles of the attention residual the most firm-specific quintile is the single low cell and the other four form a flat plateau, which is why a linear specification understates the effect: there is no monotone trend to fit, only a step at the firm-specific extreme. The apparent dip in the most macro-dominated quintile in raw returns, 0.81 basis points, is an artifact of a few fat-tailed days. It disappears under winsorization, rising to 3.26 and 3.77 , and reverses to the highest cell, 5.09 , when 2020 is excluded. We read this as conditioning and locating the anomaly rather than displacing the funding-constraint account of it. The crashes that give the factor its fat left tail, the days on which high-beta assets rally and the line steepens violently, fall on the non-firm-specific days and are orthogonal to the attention measure, which is why winsorizing leaves the low- U cell almost unchanged while lifting the others. The firm-specific effect is instead a steady difference that sharpens when those tails are removed. In this sense the flat unconditional line is a composition average: quiet where firm-specific attention dominates, and carrying both the premium of the anomaly and its crash risk on the days where it does not.

Panel C views the same factor from the betting-on-beta side, which is where the economic content is sharpest. I build a beta-neutral low-minus-high factor from the beta-sorted decile portfolios, equal-weighted, hedge it to zero market beta, and report its mirror, the return to betting on beta. Because the factor carries no market exposure, its mean return is its CAPM alpha, so a zero entry is the literal statement that the CAPM holds. On firm-specific

days betting on beta earns -0.19 basis points, indistinguishable from zero at $t = -0.06$: once beta is priced there is nothing left to capture. On the remaining days it earns -2.98 basis points at $t = -1.66$, the losing side of the same anomaly that betting against beta harvests in Panel A. The underlying factor is sound, with a market beta near zero and a 0.69 correlation with the traded AQR factor.

The absence of alpha is the result, not its presence, and it is what separates a restored CAPM from a reversed anomaly. If high beta were underpriced on firm-specific days, betting on beta would earn positive alpha there; instead the alpha is zero, so on those days beta is priced correctly rather than mispriced in the other direction. This is the cleanest form of the paper’s claim: not that a new strategy pays when firm-specific attention dominates, but that the deviation from the CAPM disappears. Consistent with that, the regime-conditioned factor, betting against beta on macro days and with beta on firm-specific days, earns an alpha of 0.39 basis points against the Fama-French five factors plus the AQR factor, at $t = 0.30$, adding nothing to the existing set. The contribution is a location rather than a new factor: the flat-line anomaly that betting against beta monetizes lives on macro-attention days and vanishes precisely where firm-specific attention concentrates, which is where the CAPM is restored.

8 Trading strategy

The empirical results documented above establish that the security market line steepens when institutional attention is concentrated on firm-specific news. I now explore the economic significance of this finding through a set of trading strategies that exploit the attention-composition state variable. The strategies are designed to be easily implementable and to assess whether the conditional pricing of beta translates into economically meaningful abnormal returns after standard risk adjustment.

The baseline strategy takes a long position in the highest-beta decile portfolio and a short position in the lowest-beta decile portfolio on days identified as bottom-quartile residual-attention days, and holds cash on all other days. The regime signal is days when the 5-day rolling median of $\hat{u}_t$ is below its rolling 1-year 25th percentile, lagged by one trading day so that positions are entered or exited the day after the signal switches, ensuring that the strategy uses only information available at the time of trading. Because the residual $\hat{u}_t$ is constructed from a persistent, slowly moving rolling median, the signal switches states infrequently, keeping turnover low. In practice, an investor could approximate the long and short legs using liquid high-beta and low-beta ETFs, which would keep implementation costs

manageable. The portfolio returns I report are gross of transaction costs, financing frictions, and shorting fees. The sample runs from 2012 to 2025 and contains 3,222 trading days, of which 448 are classified as regime days by the lagged signal.

Table 11 reports the results. Panel A provides two benchmarks. The first is an always-on long-short strategy that holds the high-minus-low beta spread portfolio every day regardless of the attention state. This strategy earns an average daily return of 4.3 basis points with an annualized Sharpe ratio of 0.34, and generates no significant alpha under either the CAPM (-4.35% annualized, $t = -0.63$) or the Fama-French five-factor model (2.58% , $t = 0.42$). The failure of the unconditional long-short strategy to generate alpha is consistent with the well-documented flatness of the unconditional security market line. The second benchmark is the buy-and-hold market portfolio, which earns 5.4 basis points per day with an annualized Sharpe ratio of 0.79.

Panel B reports the headline conditional strategy. By restricting the long-short position to bottom-quartile residual-attention days and holding cash otherwise, the strategy earns 3.4 basis points per day with an annualized Sharpe ratio of 0.67. The strategy generates a CAPM alpha of 6.5% per year ($t = 1.86$) and a five-factor alpha of 7.4% per year ($t = 2.11$), both statistically significant. The market beta of the strategy is 0.13, reflecting the fact that the strategy is in the market on only a fraction of trading days and takes a long-short position that partially hedges market exposure when active. The key comparison is with the always-on long-short benchmark in Panel A: timing the beta spread using the attention-composition signal converts an unconditionally unprofitable strategy into one that generates significant abnormal returns, even after controlling for size, value, profitability, and investment factors.

Panel C considers an alternative approach to managing the off-regime position. Instead of holding cash on non-regime days, the strategy reverses the long-short position to bet against beta, going long the lowest-beta portfolio and short the highest-beta portfolio, following the logic of Frazzini and Pedersen (2014). This hybrid strategy earns a CAPM alpha of 17.4% per year ($t = 2.24$) and a five-factor alpha of 12.3% per year ($t = 1.67$), the latter significant at the 10% level. The market beta is -0.59 , reflecting the net short-market exposure that accumulates from betting against beta on the majority of trading days. The substantially larger CAPM alpha relative to Panel B arises because the strategy profits from the steep SML on regime days and from the flat or inverted SML on non-regime days, capturing both tails of the conditional beta-return relationship documented in Table 2. The annualized Sharpe ratio of 0.20 is lower than that of the headline strategy, however, reflecting the higher volatility inherent in maintaining an active position at all times.

Savor and Wilson (2014) and subsequent work document that the CAPM also prices risk on

macroeconomic announcement days. Panel D of Table 11 explores whether the attention-composition signal and the FOMC announcement calendar provide complementary sources of conditional beta pricing. A long-short strategy restricted to FOMC days alone earns 1.5 basis points per day with a Sharpe ratio of 0.73, and generates a CAPM alpha of 3.3% annualized ($t = 2.43$) and a five-factor alpha of 3.5% ($t = 2.55$), both significant at the 5% level. These magnitudes are consistent with the FOMC premium documented in prior work. Combining the two signals by taking the long-short position on either bottom-quartile residual days or FOMC days, and holding cash otherwise, produces a strategy that earns 4.8 basis points per day with an annualized Sharpe ratio of 0.88. The CAPM alpha is 9.5% per year ($t = 2.58$) and the five-factor alpha is 10.6% per year ($t = 2.87$), both significant at the 1% level. The combined strategy outperforms either signal in isolation, suggesting that the attention-composition regime and the FOMC calendar identify distinct sets of days on which beta is priced. This interpretation is supported by the minimal overlap between the two signals: the regime and FOMC indicators coincide on only 0.16% of trading days.

Finally, the most aggressive variant replaces cash with a betting-against-beta position on non-signal days. This strategy earns a CAPM alpha of 23.3% per year ($t = 2.97$) and a five-factor alpha of 18.6% per year ($t = 2.49$). The alphas are large and significant, reflecting the combined gains from timing beta in both directions. The market beta of -0.54 again reflects the net short exposure during the majority of non-signal days.

The trading results reinforce the central finding of this paper. The conditional price of systematic risk is higher when institutional attention is concentrated on firm-specific news, steepening the SML so that high-beta assets earn higher returns. The unconditional long-short beta strategy earns no alpha, but timing it with the attention-composition residual generates economically and statistically significant abnormal returns that are not subsumed by standard factor models. The complementarity with FOMC announcement days further suggests that the attention-composition channel is distinct from the macroeconomic announcement channel documented in prior literature: both identify states in which beta is priced, but they do so through different features of the information environment.

9 Conclusion

This paper shows that the empirical performance of the CAPM depends on how institutional investors allocate attention between firm-specific news and broader market-wide information. Using firm-specific news from the Dow Jones Institutional Newswire, filtered by a large language model to remove systematic content, and institutional attention data from Bloomberg

terminal activity, I construct a daily measure of attention composition: the residual from an expanding regression of aggregate institutional attention on idiosyncratic news attention. When this residual is low, aggregate attention is well explained by firm-specific news processing. When it is high, attention is elevated but directed elsewhere.

On the 52.8% of trading days when the attention-composition residual falls below its trailing median, beta prices the cross section of stock returns. The Fama–MacBeth slope is 12.6 basis points per unit of beta for value-weighted portfolios ($t = 3.00$) and 8.7 for equal-weighted portfolios ($t = 2.21$), rising to 23.1 and 22.6 in the most anchored quartile, with intercepts indistinguishable from zero. On the complementary set of days, the security market line is flat. These findings hold in both Fama–MacBeth and pooled panel specifications, for value-weighted and equal-weighted portfolios, and survive controls for disagreement, the VIX, macroeconomic announcement days, leading earnings announcement days, and the expected information consumption framework of [Ben-Rephael, Carlin, Da, and Israelsen \(2021\)](#). Importantly, this is not a high-attention result: the days on which the CAPM succeeds are days of below-average aggregate attention, but above-average concentration of that attention on firm-specific news.

To rationalize these findings, I develop a noisy rational expectations model in which investors allocate a fixed attention budget between processing firm-specific public news and acquiring private signals about a common macroeconomic factor. Because filtered firm-specific news is informationally orthogonal to the aggregate factor, attention devoted to firm-level news crowds out private factor learning. Posterior factor uncertainty rises with the share of attention directed toward firm-specific news, the equilibrium factor risk premium increases, and the security market line steepens. This mechanism is the mirror image of the resolution channel in [Andrei, Friedman, and Ozel \(2023\)](#), in which attention to earnings accelerates the resolution of factor uncertainty and reduces posterior factor variance. The two channels operate through opposite effects on the same theoretical object and leave opposite footprints: the composition premium requires no announcement, is amplified rather than displaced by prior uncertainty, and is earned where factor uncertainty persists rather than where it is resolved.

A simple trading strategy that conditions on the attention-composition signal converts an unconditionally unprofitable long-short beta spread into a strategy that earns significant five-factor alpha. Combining the signal with the FOMC announcement calendar further improves performance, consistent with the view that the attention-composition regime and macroeconomic announcement days identify distinct sources of conditional beta pricing.

There are several directions for future work. First, the model treats the attention allocation

ω as exogenous and does not pin down its cross-day determinants. Endogenizing the supply of firm-specific news and the resulting equilibrium allocation of attention would provide a richer account of when and why the CAPM succeeds. Second, the empirical construction relies on Bloomberg terminal data, which limits the sample to institutional investors and to the period from 2010 onward. Extending the analysis to longer samples using alternative attention proxies would test the generality of the finding. Third, the analysis measures the attention of institutional investors. Whether the composition of retail attention carries pricing content of its own, or instead operates through the price-pressure channel that [Da \(2025\)](#) document for its level, is an open question.

These findings have implications for both researchers and practitioners. For researchers, they identify attention composition as a new state variable governing the empirical performance of the CAPM, distinct from the level of attention, disagreement, volatility, or the expected market return. For practitioners, they suggest that the reliability of CAPM-based discount rates and portfolio decisions depends on the prevailing attention environment: the model's cross-sectional predictions are most accurate when institutional attention is anchored to firm-specific information processing, and least accurate when attention drifts toward market-wide or macroeconomic concerns.

References

- Admati, Anat R. (1985). A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets. *Econometrica*, 53(3), 629–657.
- Aït-Sahalia, Yacine, Chen Xu Li, and Chenxu Li (2024). So Many Jumps, So Few News. *NBER Working Paper No. 32746*, National Bureau of Economic Research.
- Andrei, Daniel, Henry Friedman, and N. Bugra Ozel (2023). Economic Uncertainty and Investor Attention. *Journal of Financial Economics*, 149(2), 179–217.
- Baker, Malcolm, and Jeffrey Wurgler (2006). Investor sentiment and the cross-section of stock returns. *Journal of Finance*, 61(4), 1645–1680.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4), 1593–1636.
- Ben-Rephael, Azi, Bruce I. Carlin, Zhi Da, and Ryan D. Israelsen (2021). Information Consumption and Asset Pricing. *The Journal of Finance*, 76(1), 357–394.
- Ben-Rephael, Azi, Zhi Da, and Ryan D. Israelsen (2017). It depends on where you search: Institutional investor attention and underreaction to news. *The Review of Financial Studies*, 30(9), 3009–3047.
- Berk, Jonathan B., and Jules H. van Binsbergen (2016). Assessing Asset Pricing Models Using Revealed Preference. *Journal of Financial Economics*, 119(1), 1–23.
- Berk, Jonathan B., and Jules H. van Binsbergen (2017). How Do Investors Compute the Discount Rate? They Use the CAPM (Corrected June 2017). *Financial Analysts Journal*, 73(2), 25–32.
- Black, Fischer (1972). Capital Market Equilibrium with Restricted Borrowing. *The Journal of Business*, 45(3), 444–455.
- Black, Fischer, Michael C. Jensen, and Myron Scholes (1972). The Capital Asset Pricing Model: Some Empirical Tests. In Michael C. Jensen (ed.), *Studies in the Theory of Capital Markets*, New York: Praeger, pp. 79–121.
- Bollerslev, Tim, George Tauchen, and Hao Zhou (2009). Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11), 4463–4492.

- Caldara, Dario, and Matteo Iacoviello (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225.
- Chan, Kam Fong, and Terry Marsh (2022). Asset Pricing on Earnings Announcement Days. *Journal of Financial Economics*, 144(3), 1022–1042.
- Cookson, J. Anthony, Diego García, and Elvis Jarnećic (2026). News consumption in the wild. Working paper, University of Colorado Boulder.
- Cookson, J. Anthony, and Marina Niessner (2023). Investor disagreement: Daily measures from social media. Available at SSRN 4529594.
- Da, Zhi (2025). Market Returns and a Tale of Two Types of Attention. *Management Science*, forthcoming.
- Fama, Eugene F., and James D. MacBeth (1973). Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy*, 81(3), 607–636.
- Fisher, Adlai, Charles Martineau, and Jinfei Sheng (2022). Macroeconomic Attention and Announcement Risk Premia. *The Review of Financial Studies*, 35(11), 5057–5093.
- Frazzini, Andrea, and Lasse Heje Pedersen (2014). Betting Against Beta. *Journal of Financial Economics*, 111(1), 1–25.
- Grossman, Sanford J., and Joseph E. Stiglitz (1980). On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70(3), 393–408.
- Hasler, Michael, and Charles Martineau (2023). Explaining the Failure of the Unconditional CAPM with the Conditional CAPM. *Management Science*, 69(3), 1835–1855.
- Hasler, Michael, and Charles Martineau (2024). Equity Return Predictability with the ICAPM. *The Review of Asset Pricing Studies*, 14(3), 481–512.
- Hendershott, Terrence, Dmitry Livdan, and Dominik Rösch (2020). Asset Pricing: A Tale of Night and Day. *Journal of Financial Economics*, 138(3), 635–662.
- Hong, Harrison, and David A. Sraer (2016). Speculative Betas. *The Journal of Finance*, 71(5), 2095–2144.
- Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng (2015). Measuring uncertainty. *American Economic Review*, 105(3), 1177–1216.

- Jylhä, Petri (2018). Margin Constraints and the Security Market Line. *Journal of Finance*, 73(2), 1281–1321.
- Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp (2016). A Rational Theory of Mutual Funds’ Attention Allocation. *Econometrica*, 84(2), 571–626.
- Lintner, John (1965). Security Prices, Risk, and Maximal Gains from Diversification. *The Journal of Finance*, 20(4), 587–615.
- Martin, Ian (2017). What is the expected return on the market? *The Quarterly Journal of Economics*, 132(1), 367–433.
- Mossin, Jan (1966). Equilibrium in a Capital Asset Market. *Econometrica*, 34(4), 768–783.
- Patton, Andrew J., and Michela Verardo (2012). Does Beta Move with News? Firm-Specific Information Flows and Learning about Profitability. *The Review of Financial Studies*, 25(9), 2789–2839.
- Peng, Lin, and Wei Xiong (2006). Investor Attention, Overconfidence and Category Learning. *Journal of Financial Economics*, 80(3), 563–602.
- Savor, Pavel, and Mungo Wilson (2014). Asset Pricing: A Tale of Two Days. *Journal of Financial Economics*, 113(2), 171–201.
- Savor, Pavel, and Mungo Wilson (2016). Earnings Announcements and Systematic Risk. *The Journal of Finance*, 71(1), 83–138.
- Sharpe, William F. (1964). Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *The Journal of Finance*, 19(3), 425–442.
- Tinic, Seha M., and Richard R. West (1984). Risk and Return: January vs. the Rest of the Year. *Journal of Financial Economics*, 13(4), 561–574.
- Upenieks, Adam (2026). The Elusive CAPM: Idiosyncratic News and the Tilt of the Security Market Line. *Working Paper*.
- Veldkamp, Laura L. (2006). Information Markets and the Comovement of Asset Prices. *The Review of Economic Studies*, 73(3), 823–845.

Table 1: Beta-Decile Characteristics, Attention Regimes, and Determinants of Residual Attention

Panel A: Firm characteristics by beta decile										
	Beta decile									
	1	2	3	4	5	6	7	8	9	10
	Low									High
Average beta	0.488	0.710	0.838	0.947	1.050	1.149	1.260	1.397	1.582	1.939
Average log market capitalization	14.456	14.636	14.653	14.622	14.574	14.545	14.495	14.432	14.318	14.182
News days	28,867	26,944	26,669	27,013	27,209	28,092	29,813	30,346	31,922	32,848
Median attention score	0.360	0.362	0.364	0.351	0.337	0.336	0.330	0.339	0.353	0.377

Panel B: Regime characteristics across residual-attention quartiles					
	Q1	Q2	Q3	Q4	Q1–Q4
	Low	25–50	50–75	High	Difference
Firm-news attention, standardized	0.544	−0.217	−0.223	−0.401	0.945
Macro attention, standardized	−0.030	−0.041	0.056	0.154	−0.184
IAI	0.075	0.051	0.050	0.045	0.030
AIA	0.405	0.346	0.397	0.444	−0.039
Firm-news supply	0.063	0.050	0.047	0.042	0.021
VIX	17.465	17.003	17.618	18.248	−0.783
Disagreement	0.249	0.242	0.244	0.250	−0.001
Market return	0.100	0.084	0.079	−0.041	0.141

Panel C: Determinants of residual attention			
<i>Continuous drivers (single regressions, drivers at $t - 1$)</i>			
	Coefficient	t	R^2
Firm-news coverage share	−0.22	−7.28	0.046
Earnings-announcement share	−0.14	−4.77	0.018
Macro-topic attention	+0.06	1.39	0.003
EPU innovation (5-day)	+0.02	1.30	0.000
GPR innovation (5-day) ^a	−0.06	−4.50	0.004
<i>Scheduled announcement days (day-of-week controls; standardized daily objects)</i>			
	FOMC	CPI	PPI
Aggregate attention (AIA)	+0.38 (4.77)	+0.11 (1.99)	+0.13 (2.40)
Firm-news attention (IAI)	+0.43 (4.66)	−0.02 (−0.31)	−0.06 (−1.02)
Firms covered	−0.17 (−2.06)	−0.29 (−4.29)	−0.34 (−4.76)
Attention residual $\hat{u}$	−0.04 (−0.40)	+0.15 (1.98)	+0.27 (3.45)
Joint R^2 , all drivers and announcement calendar			0.056

Notes: Panel A reports firm characteristics across beta deciles. Average beta is the average estimated 252-day market beta; news days is the number of firm-days with newswire coverage within each decile; attention score is the Bloomberg institutional attention measure. Panel B reports averages across quartiles of residual attention, where Q1 denotes the lowest quartile (attention most anchored on firm-specific news) and Q4 the highest. Firm-news attention and macro attention are the standardized daily firm-news attention index and the standardized macroeconomic attention index of [Fisher, Martineau, and Sheng \(2022\)](#); disagreement is from J. Anthony Cookson’s website. Panel C reports determinants of the standardized daily residual. The first block reports single regressions on each standardized driver dated $t - 1$ with Newey–West standard errors (21 lags); coverage share is the fraction of panel firms with firm-specific news, earnings share the fraction with an earnings announcement, macro-topic attention the Bloomberg macroeconomic reading index, and EPU and GPR innovations five-day changes in the log daily indices of [Baker, Bloom, and Davis \(2016\)](#) and [Caldara and Iacoviello \(2022\)](#), lagged one trading day. The second block reports coefficients on announcement-day dummies from regressions of each standardized daily object on all four dummies and day-of-week fixed effects, with t -statistics in parentheses.

Table 2: CAPM tests across attention-state cutoffs.

The table reports estimates from daily Fama–MacBeth regressions of portfolio excess returns on CAPM betas and from pooled panel regressions with a state dummy and a beta–state interaction. Day t is assigned to an attention state by comparing the five-day rolling median of the residual $\hat{u}$, evaluated through day $t - 1$, to trailing 252-day quantiles of the daily residual, also evaluated through day $t - 1$, so that each day’s classification uses only information available before the return is earned. The sample is split into four attention states defined by rolling quantile cutoffs of the residual-based state variable: (i) 0–25%, (ii) $< 50\%$, (iii) $> 50\%$, and (iv) $> 75\%$. For the pooled regressions, I estimate $r_{p,t} = \alpha + \beta \hat{\beta}_p + \gamma D_t^{state} + \delta(\hat{\beta}_p \times D_t^{state}) + \varepsilon_{p,t}$. Coefficients are reported in basis points per day and t -statistics are in parentheses. Fama–MacBeth t -statistics use Newey–West standard errors with five lags on the time series of estimated daily slope coefficients. Pooled regression standard errors are clustered by trading day. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

State	Port weights	Fama–MacBeth			Pooled regression					
		Intercept	β	Avg R^2	Intercept	β	D^{state}	$\beta \times D^{state}$	R^2	% days
Panel A: (0–25% state)										
	Value-weighted	-7.00 (-1.07)	23.13*** (2.85)	0.437	1.68 (0.51)	1.89 (0.54)	-2.89 (-0.46)	16.29** (2.18)	0.0024	19.0%
	Equal-weighted	-2.55 (-0.53)	22.63*** (3.07)	0.551	4.55 (1.48)	-0.88 (-0.27)	0.89 (0.15)	16.18** (2.28)	0.0029	19.0%
Panel B: (<50% state)										
	Value-weighted	-2.98 (-0.91)	12.63*** (3.00)	0.451	6.23 (1.38)	-4.42 (-0.96)	-10.54* (-1.83)	18.68*** (3.02)	0.0021	52.8%
	Equal-weighted	2.36 (0.94)	8.73** (2.21)	0.551	6.36 (1.47)	-4.32 (-1.00)	-4.13 (-0.76)	13.36** (2.31)	0.0016	52.8%
Panel C: (>50% state)										
	Value-weighted	5.34 (1.37)	-3.43 (-0.71)	0.450	-4.31 (-1.20)	14.27*** (3.46)	10.54* (1.83)	-18.68*** (-3.02)	0.0021	47.2%
	Equal-weighted	6.32** (2.02)	-4.20 (-0.93)	0.567	2.23 (0.68)	9.04** (2.34)	4.13 (0.76)	-13.36** (-2.31)	0.0016	47.2%
Panel D: (>75% state)										
	Value-weighted	8.73 (1.52)	-1.17 (-0.15)	0.418	-1.24 (-0.41)	6.96** (2.07)	9.23 (1.10)	-7.27 (-0.88)	0.0003	18.0%
	Equal-weighted	8.12* (1.65)	2.96 (0.40)	0.556	3.45 (1.22)	2.51 (0.78)	3.98 (0.50)	1.17 (0.16)	0.0002	18.0%

Table 3: Uncertainty and the Composition Premium

Panel A: Price of beta by prior-day uncertainty tercile				
	Firm-anchored state			Unanchored
	T1	T2	T3	T3
	Low		High	Placebo
Realized variance	−0.4 (−0.08)	14.7*** (2.69)	24.2** (2.14)	−8.6 (−0.82)
VIX	6.9* (1.84)	1.4 (0.26)	31.8*** (2.92)	−2.8 (−0.25)
JLN $h = 1$	12.8*** (2.95)	1.7 (0.30)	24.1** (2.34)	−0.4 (−0.03)
JLN $h = 12$	10.8** (2.51)	3.5 (0.63)	24.9** (2.32)	−0.3 (−0.03)
Panel B: Pooled interaction within the top uncertainty tercile				
	β	$\beta \times D^{state}$	Days	
Realized variance	−14.1 (−1.40)	33.8** (2.33)	1,054	
VIX	−7.0 (−0.65)	30.3** (2.07)	1,054	
JLN $h = 1$	−9.1 (−0.86)	26.1* (1.80)	1,046	
JLN $h = 12$	−7.1 (−0.73)	24.4* (1.77)	1,047	
Panel C: Variance risk premium within the anchored high-uncertainty cell				
	λ_β	Intercept	Days	
Low variance risk premium	10.4 (0.63)	1.7 (0.14)	261	
High variance risk premium	51.8*** (3.60)	−28.6** (−2.21)	261	

Notes: Panel A reports Fama–MacBeth prices of beta in basis points per day on firm-anchored days, sorted into terciles of the prior trading day’s uncertainty proxy, with breakpoints computed within sample; the final column repeats the top-tercile estimation on unanchored days. The attention state is the five-day rolling median of the residual $\hat{u}$ relative to its trailing 252-day median, lagged one trading day, so all conditioning information precedes the return day. Anchored cells contain 495 to 623 days and span more than thirty calendar quarters, with no quarter contributing more than ten percent. Panel B estimates $r_{p,t} = \alpha + \beta \hat{\beta}_p + \gamma D_t^{state} + \delta(\hat{\beta}_p \times D_t^{state}) + \varepsilon_{p,t}$ on top-tercile days only, with standard errors clustered by trading day. Panel C sorts the anchored top-SVIX²-tercile days by the variance risk premium, SVIX² minus trailing 21-day realized variance, where SVIX² is the option-implied lower bound on the conditional equity premium of [Martin \(2017\)](#); the attention residual is uncorrelated with SVIX² (correlation 0.009). JLN denotes the macroeconomic uncertainty index of [Jurado, Ludvigson, and Ng \(2015\)](#), mapped by prior calendar month. Newey–West t -statistics with five lags are in parentheses in Panels A and C. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 4: Factor Specificity: Own-Factor-Beta Pricing and Factor-Premium Timing by Attention State

Panel A: Fama–MacBeth price of each factor loading on own-loading-sorted deciles			
Factor	Firm-anchored	Unanchored	Difference (pooled)
Market	9.4*** (3.06)	−2.7 (−0.76)	12.8*** (2.74)
SMB ^a	3.9** (2.53)	−1.5 (−0.87)	4.6 (1.31)
HML	0.9 (0.50)	−0.1 (−0.05)	−0.1 (−0.02)
RMW	−0.5 (−0.37)	0.3 (0.20)	−1.2 (−0.35)
CMA	−1.7* (−1.67)	−0.1 (−0.08)	0.4 (0.20)
Days	1,668	1,493	

Panel B: Mean daily factor return by attention state			
Factor	Firm-anchored	Unanchored	Difference t
Market	9.30	1.44	2.16
SMB	1.07	−2.42	1.59
HML	−0.79	−0.05	−0.25
RMW	0.38	1.76	−0.79
CMA	−0.74	0.39	−0.72

Notes: Panel A applies the paper’s cross-sectional design to each Fama–French factor in turn. For each factor, every stock’s univariate loading is estimated over the trailing 252 trading days and lagged one day; stocks are sorted daily into deciles on that loading; decile excess returns are value-weighted using prior-day market capitalization, and the decile’s average loading is the regressor. The first two columns report time-series averages of daily cross-sectional slopes of decile returns on decile loadings within each attention state, in basis points per day, with Newey–West t -statistics (five lags) in parentheses. The third column reports the coefficient on the loading interacted with the firm-anchored indicator from a pooled regression of decile returns on the loading, the indicator, and their interaction, with standard errors clustered by day. Panel B reports the mean daily return of each factor, in basis points, within each state, and the t -statistic on the state indicator from a regression of the factor return on the indicator with Newey–West standard errors (five lags). The state is the five-day rolling median of $\hat{u}$ compared with its trailing 252-day median, classified at the prior day’s close. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

^a Univariate SMB loadings inherit market exposure: their average daily cross-sectional correlation with market betas is 0.66. In a daily individual-stock regression on the CAPM beta and the multivariate SMB loading orthogonalized against it, the SMB coefficient is 1.2 basis points ($t = 1.15$) on anchored days, so the anchored-state SMB slope reflects market covariance carried by the size sort rather than a priced size loading, consistent with the insignificant state difference.

Table 5: Continuous Robustness of the Attention Gradient in the Price of Beta

Panel A: Full sample							
Specification	Value-weighted			Equal-weighted			Days
	b	(t)	R^2	b	(t)	R^2	
Baseline	−9.15***	(−2.88)	0.002	−7.73***	(−2.85)	0.002	3,220
+ VIX	−9.30***	(−2.92)	0.004	−7.89***	(−2.93)	0.005	3,220
+ SVIX ²	−9.53***	(−2.98)	0.005	−8.08***	(−2.98)	0.005	3,220
+ Realized variance (21 days)	−9.24***	(−2.90)	0.002	−7.85***	(−2.89)	0.002	3,220
+ Variance risk premium	−9.11***	(−2.87)	0.002	−7.72***	(−2.84)	0.002	3,220
+ Return dispersion	−8.27***	(−2.59)	0.003	−6.36**	(−2.38)	0.006	3,220
+ Aggregate attention (AIA)	−8.53***	(−2.65)	0.002	−7.82***	(−2.82)	0.002	3,220
+ Firm-news attention (IAI)	−9.24***	(−2.86)	0.002	−7.75***	(−2.83)	0.002	3,220
+ Firm-news coverage share	−8.50**	(−2.43)	0.002	−6.93**	(−2.34)	0.002	3,220
+ Earnings-announcement share	−9.02***	(−2.73)	0.002	−7.52***	(−2.66)	0.002	3,220
+ Lagged 5-day market return	−10.22***	(−3.02)	0.008	−8.58***	(−3.01)	0.006	3,220
+ Lagged 21-day market return	−9.60***	(−2.97)	0.003	−7.99***	(−2.93)	0.002	3,220
+ Lagged price of beta	−8.94***	(−2.90)	0.003	−7.32***	(−2.85)	0.004	3,219
+ Announcement and day-of-week dummies	−9.09***	(−2.89)	0.006	−7.59***	(−2.83)	0.005	3,220
+ All controls	−10.34**	(−2.13)	0.017	−9.98**	(−2.49)	0.017	3,219

Panel B: Common sample with disagreement and sentiment							
Specification	Value-weighted			Equal-weighted			Days
	b	(t)	R^2	b	(t)	R^2	
Baseline	−7.70**	(−2.05)	0.002	−7.56**	(−2.32)	0.002	2,477
+ Disagreement	−7.57**	(−1.96)	0.002	−7.26**	(−2.19)	0.002	2,477
+ Sentiment	−7.05**	(−1.97)	0.002	−7.46**	(−2.36)	0.002	2,477
+ Disagreement and sentiment	−7.09*	(−1.93)	0.002	−7.26**	(−2.25)	0.002	2,477
+ All controls	−6.91	(−1.23)	0.021	−7.20	(−1.61)	0.020	2,476

Notes: The table reports time-series regressions of the daily Fama–MacBeth price of beta $\hat{\lambda}_t$, in basis points per day, on the standardized five-day rolling median of the attention residual $\hat{u}$ evaluated through day $t - 1$, with controls added as indicated. Each day, beta-decile portfolio excess returns are regressed cross-sectionally on their betas to obtain $\hat{\lambda}_t$; the left block uses value-weighted and the right block equal-weighted portfolios. The reported coefficient b is in basis points per one-standard-deviation increase in the smoothed residual, so a negative value indicates a steeper security market line when institutional attention is anchored on firm-specific news. All controls are standardized and dated $t - 1$: VIX; SVIX², the option-implied bound on the equity premium of [Martin \(2017\)](#); realized variance, the sum of squared daily market excess returns over the prior 21 trading days; the variance risk premium, SVIX² minus annualized realized variance; return dispersion, the cross-sectional standard deviation of daily stock returns; AIA and IAI, the aggregate institutional attention and firm-news attention indices; coverage share, the fraction of panel firms with firm-specific newswire coverage; earnings-announcement share, the fraction with an earnings announcement; cumulative market excess returns over the prior 5 and 21 trading days; the prior day’s price of beta; and dummies for FOMC, CPI, PPI, and payroll announcement days and for day of the week. The all-controls specification includes every control except SVIX², whose correlation with VIX is 0.998. Panel B restricts the sample to days on which the disagreement measure of J. Anthony Cookson and the sentiment index of [Baker and Wurgler \(2006\)](#) (orthogonalized to macroeconomic conditions and mapped by prior calendar month) are both available, and reports a matched baseline so that changes in the coefficient reflect the controls rather than the sample; the all-controls row in Panel B adds twenty regressors to 2,477 observations, and its point estimate rather than its t -statistic is the informative quantity. Newey–West t -statistics with 21 lags are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 6: Robustness to Information-Event Channels

<i>Panel A: Expected information consumption (Ben-Rephael, Carlin, Da, and Israelsen, 2021)</i>			
	All firms	EIC = 0 firms	EIC = 1 firms
α	−2.72 (−0.74)	−2.12 (−0.60)	−18.77 (−1.54)
λ_β	20.44*** (3.15)	19.11*** (3.04)	44.16*** (3.10)
Mean firms per day	1,750	1,635	168
Days	449	449	282
<i>Panel B: Earnings announcement windows (Savor and Wilson, 2016)</i>			
	All firms	Excl. earnings window	Earnings window only
α	−2.72 (−0.74)	−1.75 (−0.49)	1.36 (0.22)
λ_β	20.44*** (3.15)	18.81*** (2.98)	24.43*** (2.88)
Mean firms per day	1,750	1,390	396
Days	449	449	403
<i>Panel C: Macro-news days (Chan and Marsh, 2022; Savor and Wilson, 2013)</i>			
	α	λ_β	Days
Excluding LEAD days, EW	0.55 (0.20)	8.52** (2.13)	1343
Excluding LEAD days, VW	−5.25 (−1.39)	12.33*** (2.72)	1343
Excluding macro-announcement days, EW	0.73 (0.30)	8.58** (2.38)	1554
Excluding macro-announcement days, VW	−4.87 (−1.48)	12.47*** (3.13)	1554
<i>Panel D: State rebuilt from beta-moving versus beta-neutral news</i>			
	λ_β (anchored)	λ_β (other)	Firms
All firm news (baseline)	12.63*** (3.00)	−3.43 (−0.71)	1,627
Beta-neutral news only	5.50 (1.14)	4.65 (1.03)	814
Beta-moving news only	7.62* (1.70)	2.79 (0.63)	813
Beta-neutral, beta-matched	7.63* (1.72)	2.71 (0.58)	815

Notes: Fama–MacBeth regressions of excess returns on CAPM betas (estimated over the previous 252 days); Newey–West t -statistics in parentheses (10 lags in Panels A–B, 5 in C–D). Coefficients are in basis points per day. Panels A and B use individual firms on bottom-quartile residual-attention days. In Panel A, EIC = 1 firms are those whose institutional attention is predicted to spike around peer-firm, FOMC, and macroeconomic announcements, following Ben-Rephael, Carlin, Da, and Israelsen (2021); EIC = 0 firms are the rest. In Panel B, earnings-window firm-days are those within five business days of an earnings announcement. Panel C uses beta-decile portfolios on below-median residual-attention days, excluding LEAD days (Chan and Marsh) or macro-announcement days (Savor and Wilson); equal-weighted (EW) and value-weighted (VW). Panel D rebuilds the attention residual from a subset of firms and reports the value-weighted price of beta by state. Each firm’s news-day beta change is estimated from a regression of its excess return on the market return, a news dummy, and their interaction (Andrei, Friedman, and Ozel, 2023); firms are split at the median. “Beta-neutral” and “beta-moving” rebuild the residual from the halves whose news moves beta least and most; “beta-matched” keeps the low-change half within firm-beta quintiles. Sample 2011–2025. ***, **, *: 1%, 5%, 10% significance.

Table 7: Alternative Measures of Attention Concentration

	Firm-anchored days				Complement		
Measure	λ_β (VW)	λ_β (EW)	α (VW)	Days	λ_β (VW)	Days	Agreement
<i>Panel A: Median cutoff</i>							
Expanding residual (baseline)	12.6*** (3.00)	8.7** (2.21)	−3.0 (−0.91)	1,668	−3.4 (−0.71)	1,493	1.00
Share IAI/AIA	8.8** (2.11)	6.1 (1.55)	0.1 (0.02)	1,791	−1.0 (−0.22)	1,430	0.85
Rolling-252 residual	12.9*** (2.79)	8.4* (1.92)	−4.7 (−1.28)	1,402	−2.2 (−0.48)	1,567	0.83
Percentile gap	11.7** (2.42)	7.9* (1.75)	−2.8 (−0.70)	1,469	−1.8 (−0.37)	1,501	0.85
Level: AIA high	0.4 (0.07)	0.1 (0.02)	4.4 (1.06)	1,486	8.0** (2.24)	1,735	0.33
Level: IAI high	3.1 (0.63)	3.6 (0.80)	3.5 (0.88)	1,637	5.9* (1.65)	1,584	0.51
<i>Panel B: Bottom-quartile cutoff</i>							
Expanding residual (baseline)	23.1*** (2.85)	22.6*** (3.07)	−7.0 (−1.07)	601	0.8 (0.23)	2,560	
Share IAI/AIA	20.2*** (2.89)	18.2*** (2.98)	−6.8 (−1.19)	673	0.3 (0.09)	2,548	
Rolling-252 residual	17.2** (2.28)	14.2** (1.98)	−6.3 (−1.03)	567	2.0 (0.53)	2,402	
Percentile gap	28.6*** (2.94)	23.0*** (2.58)	−15.3* (−1.96)	530	−0.2 (−0.07)	2,440	
Level: AIA high	−15.7* (−1.69)	−13.4* (−1.67)	17.4** (2.28)	577	8.9*** (2.64)	2,644	
Level: IAI high	1.1 (0.12)	2.6 (0.28)	2.6 (0.31)	610	5.3* (1.70)	2,611	
<i>Panel C: Attention level versus attention composition</i>							
	<i>Joint pooled regression, full panel (day-clustered)</i>						
	$\beta \times$ firm-anchored				$\beta \times$ AIA-low		
	20.0*** (3.06)				−4.4 (−0.68)		
	<i>Composition interactions within attention-level halves</i>						
	$\beta \times$ anchored (day-level)		$\beta \times$ anchored (pooled)		$\beta \times z(\hat{u})$		Days
Within AIA-low	18.6**	(2.42)	16.2*	(1.91)	−5.4	(−1.62)	1,726
Within AIA-high	11.5	(0.97)	24.0**	(2.40)	−8.9**	(−1.99)	1,435

Notes: Each measure is pushed through the baseline state construction: five-day rolling median compared with the trailing 252-day quantile (median in Panel A, quartile in Panel B, direction set so the state is firm-focused), lagged one trading day. Fama–MacBeth prices of beta in basis points per day; Newey–West t -statistics (five lags) in parentheses. Agreement is the share of days classified identically to the baseline. The complement of the AIA-high rows is the AIA-low state; its pricing reflects mechanical overlap with the residual state, of which AIA is the regressand: AIA-low days that are not firm-anchored price beta at -4.2 ($t = -0.67$). Panel C interacts beta with the firm-anchored and AIA-low indicators jointly (day-clustered), and, within each half of the attention-level distribution, reports the firm-anchored interaction from daily slopes regressed on the state (day-level, equal-weighted across days) and from the pooled panel (dispersion-weighted), plus the interaction with the standardized lagged residual $z(\hat{u})$. Day-level estimates equal differences of the corresponding Fama–MacBeth cells. ***, **, *: 1%, 5%, 10%.

Table 8: The Betting-Against-Beta Premium by Firm-Specific Attention Regime

Panel A: Daily BAB return by attention regime				
	Firm-anchored	Other	Difference	<i>t</i>
Raw	−0.28	4.18	−4.47	(−1.63)
Winsorized 1/99	−0.09	4.65	−4.74**	(−2.08)
Excluding 2020	0.63	3.94	−3.31	(−1.44)

Panel B: Daily BAB return by attention quintile					
	Q1 Most anchored	Q2	Q3	Q4	Q5 Least anchored
Raw	1.21	3.72	3.96	5.62	0.81
Winsorized 1/99	1.72	4.51	3.36	4.47	3.26
Excluding 2020	0.68	3.38	2.99	3.47	5.09

Panel C: Betting on beta (beta-neutral) by regime		
	Firm-anchored	Other
Mean return (= CAPM α)	−0.19 (−0.06)	−2.98 (−1.66)

Notes: Panels A and B report the mean daily return of the AQR U.S. betting-against-beta (BAB) factor, in basis points, conditional on the firm-specific attention regime; Panel C reports its mirror, the return to betting on beta. The attention regime is based on the out-of-sample residual $\hat{u}$; low values indicate that aggregate attention is well explained by firm-specific news, and “Other” denotes all days outside the firm-anchored quartile. Days are classified by the bottom quartile of the five-day rolling median of $\hat{u}$ relative to its trailing 252-day quantiles in Panels A and C, and by quintiles of $\hat{u}$ in Panel B, with Q1 the most firm-anchored and Q5 the least; all classifications use information through day $t - 1$. In Panel A, t -statistics in parentheses use Newey–West standard errors with 21 lags on the regime difference, and winsorization is applied to the daily BAB return at the stated percentiles. The Panel C factor is the negative of a beta-neutral low-minus-high factor constructed from the beta-sorted decile portfolios, equal-weighted and hedged to zero market beta (realized market beta near zero, full-sample correlation 0.69 with the AQR factor); because it carries no market exposure, its mean return equals its CAPM alpha, and its t -statistics are computed against zero within each regime. The regime-conditioned version of this factor, betting against beta on unanchored days and with beta on firm-anchored days, earns an alpha of 0.39 basis points per day ($t = 0.30$) against the Fama–French five factors plus the AQR BAB factor. The AQR-based panels span February 2012 to December 2024 (3,221 trading days; 2,968 excluding 2020); Panel C begins after a 252-day beta-estimation window. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 9: Fama–MacBeth Regressions: Attention Explained by Firm News Across Beta Deciles

	Beta decile used to define the attention-state condition									
	1	2	3	4	5	6	7	8	9	10
<i>Panel A: Equal-weighted returns</i>										
Intercept	1.23 (0.51)	2.84 (1.10)	1.47 (0.60)	2.37 (0.99)	1.87 (0.71)	5.47** (2.25)	2.00 (0.83)	3.29 (1.33)	5.39** (2.05)	3.47 (1.49)
λ_β	8.58** (2.38)	4.50 (1.20)	6.31* (1.75)	5.84* (1.66)	6.61* (1.73)	2.09 (0.56)	3.98 (1.12)	4.68 (1.26)	1.61 (0.41)	5.14 (1.46)
Avg. R^2	0.547	0.562	0.555	0.558	0.548	0.560	0.557	0.557	0.551	0.558
Days	1,554	1,461	1,415	1,482	1,437	1,500	1,531	1,568	1,528	1,635
<i>Panel B: Value-weighted returns</i>										
Intercept	−4.37 (−1.33)	−0.54 (−0.16)	−4.10 (−1.29)	−4.24 (−1.42)	−3.96 (−1.15)	0.92 (0.29)	−1.77 (−0.57)	−0.87 (−0.27)	1.43 (0.42)	0.06 (0.02)
λ_β	12.47*** (3.13)	7.62* (1.85)	11.47*** (2.91)	11.01*** (2.93)	11.48*** (2.72)	6.11 (1.52)	7.51** (1.99)	8.40** (2.08)	5.31 (1.27)	8.52** (2.22)
Avg. R^2	0.454	0.471	0.461	0.461	0.459	0.459	0.462	0.463	0.462	0.464
Days	1,554	1,461	1,415	1,482	1,437	1,500	1,531	1,568	1,528	1,635

Notes: This table reports Fama–MacBeth regressions of daily beta-decile portfolio returns on portfolio betas, restricted to days on which aggregate attention is well explained by firm-specific news among firms in the indicated beta decile. Each day, portfolio returns are regressed cross-sectionally on beta, and the table reports time-series averages of the daily intercepts and slopes in basis points per day, with Newey–West t -statistics (five lags) in parentheses. Avg. R^2 is the average daily cross-sectional R^2 . ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 10: Fama–MacBeth Regressions: Attention Explained by Firm News Across Size Deciles

	Size decile used to define the attention-state condition									
	1 Small	2	3	4	5	6	7	8	9	10 Large
<i>Panel A: Equal-weighted returns</i>										
Intercept	4.32* (1.66)	3.93 (1.63)	4.38* (1.85)	3.16 (1.24)	3.43 (1.43)	3.22 (1.32)	3.09 (1.29)	3.89 (1.54)	2.28 (0.84)	1.79 (0.71)
λ_β	3.23 (0.86)	4.16 (1.16)	2.54 (0.69)	5.68 (1.52)	6.82* (1.94)	7.42** (2.03)	7.33** (2.08)	1.70 (0.44)	6.28 (1.54)	5.11 (1.34)
Avg. R^2	0.553	0.558	0.557	0.558	0.553	0.553	0.549	0.555	0.555	0.558
Days	1,495	1,526	1,479	1,529	1,532	1,492	1,521	1,472	1,534	1,489
<i>Panel B: Value-weighted returns</i>										
Intercept	−0.71 (−0.21)	−0.80 (−0.25)	0.71 (0.22)	−2.90 (−0.89)	−1.66 (−0.54)	−2.22 (−0.69)	−3.26 (−1.04)	0.29 (0.09)	−3.39 (−0.97)	−2.80 (−0.85)
λ_β	8.19** (1.98)	8.67** (2.18)	6.10 (1.48)	10.90*** (2.67)	10.92*** (2.88)	11.61*** (2.89)	12.96*** (3.36)	5.67 (1.37)	11.18** (2.55)	9.53** (2.33)
Avg. R^2	0.461	0.461	0.463	0.464	0.463	0.462	0.452	0.468	0.459	0.464
Days	1,495	1,526	1,479	1,529	1,532	1,492	1,521	1,472	1,534	1,489

Notes: This table reports Fama–MacBeth regressions of daily beta-decile portfolio returns on portfolio betas, restricted to days on which aggregate attention is well explained by firm-specific news among firms in the indicated size decile; decile 1 contains the smallest firms and decile 10 the largest. Each day, portfolio returns are regressed cross-sectionally on beta, and the table reports time-series averages of the daily intercepts and slopes in basis points per day, with Newey–West t -statistics (five lags) in parentheses. Avg. R^2 is the average daily cross-sectional R^2 . ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 11: Trading Strategy Performance and Risk-Adjusted Returns

Strategy	Performance		CAPM		Fama–French Five-Factor		
	Mean (bps/day)	Sharpe (ann.)	Alpha (% ann.)	<i>t</i> -stat	Alpha (% ann.)	<i>t</i> -stat	Mkt beta
<i>Panel A: Benchmark Strategies</i>							
Always-on long–short	4.30	0.34	-4.35	-0.63	2.58	0.42	0.84
Buy-and-hold market	5.37	0.79	0.00	–	0.00	–	1.00
<i>Panel B: Headline Conditional Strategy</i>							
Long–short on regime days; cash otherwise	3.43	0.67	6.50*	1.86	7.42**	2.11	0.13
<i>Panel C: Alternative Cash Management</i>							
Long–short on regime days; BAB otherwise	2.57	0.20	17.35**	2.24	12.26*	1.67	-0.59
<i>Panel D: Combination with FOMC Announcement Days</i>							
FOMC days only	1.47	0.73	3.25**	2.43	3.46**	2.55	0.02
Long–short on regime or FOMC days; cash otherwise	4.80	0.88	9.47***	2.58	10.58***	2.87	0.15
Long–short on regime or FOMC days; BAB otherwise	5.29	0.42	23.30***	2.97	18.58**	2.49	-0.54




Figure 1: The attention residual over the sample. The figure plots the 63-day rolling median of $\hat{u}_t$, the residual from an expanding (ex ante) regression of aggregate institutional attention AIA_t on firm-news attention IAI_t , from the first classified day in May 2012 through December 2024. Low values indicate that aggregate attention is well explained by attention to firm-specific news. The 63-day window is a presentation choice; the attention state is classified on the five-day median compared with trailing 252-day quantiles of the daily residual.




Figure 2: SMLs across residual attention states - Value weighted portfolios. This figure plots average portfolio excess returns against unconditional CAPM betas (Security Market Lines) for beta-sorted portfolios, separately by quantile cutoffs of the residual-based attention state. Panels correspond to state definitions based on rolling quantiles of the residual (0–25%, 0–50%, 50–100%, and 75–100%). Points show mean excess returns within each beta portfolio in the given state, and the solid line is the fitted linear relation between returns and betas within the state. The SML is steepest in low-residual states (0–25% and 0–50%), indicating a stronger positive risk–return tradeoff; it flattens or turns downward in high-residual states, consistent with weaker pricing of market risk when idiosyncratic attention explains less of market-wide attention.




Figure 3: SMLs across residual attention states - Equal weighted portfolios. This figure plots average portfolio excess returns against unconditional CAPM betas (Security Market Lines) for beta-sorted portfolios, separately by quantile cutoffs of the residual-based attention state. Panels correspond to state definitions based on rolling quantiles of the residual (0–25%, 0–50%, 50–100%, and 75–100%). Points show mean excess returns within each beta portfolio in the given state, and the solid line is the fitted linear relation between returns and betas within the state. The SML is steepest in low-residual states (0–25% and 0–50%), indicating a stronger positive risk–return tradeoff; it flattens or turns downward in high-residual states, consistent with weaker pricing of market risk when idiosyncratic attention explains less of market-wide attention.




Figure 4: The security market line by attention state, all test-asset families. The figure plots mean daily excess returns (basis points per day) against CAPM betas for forty value-weighted portfolios: ten beta-sorted deciles (circles), ten size deciles (triangles), ten book-to-market deciles (diamonds), and ten momentum deciles (squares). Each portfolio's beta is estimated once over the full sample, so both states share the same horizontal positions and differ only in mean returns. Filled markers and the solid line correspond to days on which institutional attention is anchored on firm-specific news (the attention residual $\hat{u}_t$, smoothed over five days, lies below its trailing 252-day median, classified at the prior day's close); hollow markers and the dashed line correspond to the remaining days. Lines are ordinary least squares fits through the forty portfolio means within each state, with slopes reported in the legend. The sample covers May 2010 through January 2025.

A Proof of Lemma 1

The joint distribution of primitives and signals is multivariate Gaussian under Assumptions 1–2. Posterior distributions are obtained by standard conjugate Gaussian updating.

Posterior of f . The only signal that depends on f is $z_j = f + \xi_j$, with ξ_j orthogonal to f . Conjugate Gaussian updating gives

$$\begin{aligned}\text{Var}(f \mid z_j) &= (\tau_f + \tau_\xi)^{-1} = (\tau_f + (1 - \omega_j)K)^{-1}, \\ \mathbb{E}[f \mid z_j] &= V_f \cdot \tau_\xi \cdot z_j.\end{aligned}$$

Posterior of β_i . The only public signal carrying any potential dependence on β_i is y_i . Under Assumption 1, $y_i = \varepsilon_i + \eta_i$, which by construction is independent of β_i . Hence the posterior equals the prior:

$$\beta_i \mid y_i \sim \mathcal{N}(b_i, \sigma_\nu^2).$$

Posterior of ε_i . Analogous Gaussian updating from $y_i = \varepsilon_i + \eta_i$ gives

$$\begin{aligned}\text{Var}(\varepsilon_i \mid y_i) &= (\tau_\varepsilon + \tau_y)^{-1} = V_\varepsilon, \\ \mathbb{E}[\varepsilon_i \mid y_i] &= V_\varepsilon \cdot \tau_y \cdot y_i.\end{aligned}$$

Independence. Because $(f, \beta_i, \varepsilon_i)$ are mutually independent under the prior, and each posterior updates only through its dedicated signal channel under Assumptions 1–2, posterior independence is preserved. $\square$

B Derivation of $V_\varepsilon^\star$

The ex-ante residual variance of the non-factor component of d_i , conditional on information but excluding the systematic factor loading, is

$$V_\varepsilon^\star = \text{Var}(\beta_i f + \varepsilon_i \mid \text{signals}) - b_i^2 V_f.$$

Using $\beta_i = b_i + \nu_i$ and the independence structure of posteriors,

$$\begin{aligned}\text{Var}(\beta_i f + \varepsilon_i \mid \text{signals}) &= \text{Var}(b_i f + \nu_i f + \varepsilon_i \mid \text{signals}) \\ &= b_i^2 V_f + \mathbb{E}[\nu_i^2 f^2 \mid \text{signals}] + V_\varepsilon,\end{aligned}$$

where cross-terms vanish by mean-zero ν_i and ε_i . Taking expectations over signal realizations,

$$\mathbb{E}[\nu_i^2 f^2] = \mathbb{E}[\nu_i^2] \mathbb{E}[f^2] = \sigma_\nu^2 (V_f + \sigma_f^2),$$

where the second factor uses $\mathbb{E}[f^2] = \text{Var}(\mu_f) + \mathbb{E}[V_f] = (\sigma_f^2 - V_f) + V_f = \sigma_f^2$ when measured ex ante, or equivalently $\mathbb{E}[\mathbb{E}[f^2 \mid \text{signals}]] = V_f + \sigma_f^2 - V_f$ when measured in the appropriate posterior-times-prior product space. Either way the residual variance is

$$V_\varepsilon^\star = V_\varepsilon + \sigma_\nu^2 (V_f + \sigma_f^2). \quad \square$$

C Proof of Lemma 2

Investor j chooses portfolio q_j to maximize $\mathbb{E}_j[-\exp(-\rho W_j)]$ where $W_j = q_j^\top (d - p)$. Under CARA-Gaussian preferences, optimal demand is

$$q_j = \frac{1}{\rho} \Sigma^{-1} (\mu_j - p),$$

with $\mu_j = \mathbb{E}_j[d]$ and $\Sigma = \text{Var}_j(d - p)$.

By Lemma 1 and (10), the posterior covariance has the rank-one-plus-diagonal structure

$$\Sigma = V_f b b^\top + V_\varepsilon^\star I,$$

where $b = (b_1, \dots, b_N)^\top$. The Sherman–Morrison identity gives

$$\Sigma^{-1} = \frac{1}{V_\varepsilon^\star} \left(I - \frac{V_f}{V_\varepsilon^\star + V_f b^\top b} b b^\top \right).$$

Market clearing requires $\int q_j dj = \bar{x} + u$. Averaging investor demands and solving for p yields, component-wise,

$$p_i = \bar{d} + b_i \bar{\mu}_f - \rho V_f b_i (N \bar{b} \bar{x}) - \rho V_\varepsilon^\star \bar{x} + (\text{linear functions of } z \text{ and } u),$$

which is (11). □

D Derivation of the FM slope and attenuation factor

From Lemma 2, the equilibrium expected excess return of asset i is

$$\mathbb{E}[d_i - p_i] = b_i \rho V_f N \bar{b} \bar{x} + \rho V_\varepsilon^* \bar{x}.$$

A cross-sectional regression of expected returns on true b_i would identify

$$\lambda^A \equiv \rho V_f N \bar{b} \bar{x}.$$

In practice, the econometrician estimates loadings from a T -period rolling-window OLS regression of realized returns r_i on the factor:

$$\hat{\beta}_i^{OLS} = \beta_i + e_i,$$

where e_i is mean-zero sampling noise. From the model, realized return variance has a factor and a non-factor component; standard OLS algebra gives

$$\text{Var}(e_i) \approx \frac{\sigma_\varepsilon^2 + 1/\tau_y}{T\sigma_f^2} = V_e(\omega).$$

The cross-sectional FM slope of expected returns on $\hat{\beta}^{OLS}$ is

$$\hat{\lambda}^{FM} = \frac{\text{Cov}_i(\mathbb{E}[d_i - p_i], \hat{\beta}_i^{OLS})}{\text{Var}_i(\hat{\beta}_i^{OLS})}.$$

Since investors price using b_i (the prior mean), and $\hat{\beta}_i^{OLS} = b_i + \nu_i + e_i$, computing covariances cross-sectionally (treating b_i as fixed but interpretable as a random draw with variance σ_b^2),

$$\text{Cov}_i(\mathbb{E}[d_i - p_i], \hat{\beta}_i^{OLS}) = \lambda^A \cdot \sigma_b^2, \quad \text{Var}_i(\hat{\beta}_i^{OLS}) = \sigma_b^2 + \sigma_\nu^2 + V_e = s_\beta^2 + V_e.$$

Therefore

$$\hat{\lambda}^{FM} = \lambda^A \cdot \frac{\sigma_b^2}{s_\beta^2 + V_e} = \lambda^A \cdot \frac{s_\beta^2 - \sigma_\nu^2}{s_\beta^2 + V_e} = A(\omega) \cdot \lambda^A,$$

where $A(\omega)$ is defined in (13). □

E Investor's value function and proof of Proposition 1

Value function. Under CARA-Gaussian preferences with information set $\mathcal{I}_j(\omega_j)$, the certainty-equivalent gain from information acquisition is proportional to $\frac{1}{2} \log \det(\Sigma_0 \Sigma^{-1}(\omega_j))$, where Σ_0 is the prior covariance of d and $\Sigma(\omega_j)$ the posterior covariance under investor attention ω_j . The first term is constant in ω_j .

From the rank-one-plus-diagonal structure of Σ derived in Appendix C,

$$\det \Sigma(\omega) = \left(V_\varepsilon^*(\omega) \right)^{N-1} \cdot \left(V_\varepsilon^*(\omega) + V_f(\omega) \sum_i b_i^2 \right),$$

where the second factor comes from the matrix-determinant lemma applied to the rank-one update. Substituting into the value-function objective and dropping the constant yields (14).

Existence of interior optimum. The value function $\mathcal{V}(\omega)$ is continuous on $[0, 1]$ and differentiable on $(0, 1)$, with

$$\mathcal{V}'(\omega) = -\frac{N-1}{2V_\varepsilon^*} \cdot \frac{dV_\varepsilon^*}{d\omega} - \frac{1}{2} \frac{1}{V_\varepsilon^* + V_f \sum_i b_i^2} \cdot \left(\frac{dV_\varepsilon^*}{d\omega} + \sum_i b_i^2 \cdot \frac{dV_f}{d\omega} \right),$$

where $dV_f/d\omega = KV_f^2 > 0$ and

$$\frac{dV_\varepsilon^*}{d\omega} = -SV_\varepsilon^2 + \sigma_\nu^2 KV_f^2.$$

At $\omega \rightarrow 0^+$: $V_\varepsilon \rightarrow 1/\tau_\varepsilon$, $V_f \rightarrow 1/(\tau_f + K)$; the first term of $\mathcal{V}'$ dominates and is large positive, so $\mathcal{V}'(0^+) > 0$.

At $\omega \rightarrow 1^-$: $V_\varepsilon \rightarrow 1/(\tau_\varepsilon + S)$, $V_f \rightarrow 1/\tau_f = \sigma_f^2$; the dominant negative contribution comes from $\sum_i b_i^2 \cdot dV_f/d\omega$ in the second term, which is large when $\sum_i b_i^2$ is large. For sufficiently informative private signals (any $K > 0$) and any non-degenerate cross-section, $\mathcal{V}'(1^-) < 0$.

By the intermediate value theorem, $\mathcal{V}'(\omega^*) = 0$ for some $\omega^* \in (0, 1)$. Strict concavity of $\mathcal{V}$ over the interior (verified numerically across the relevant parameter region) yields uniqueness. $\square$