Smart Money in Cryptocurrencies*

Vivian W. Fang
Indiana University, ECGI

Preetesh Kantak
Indiana University

Wenwei Lin
Chinese University of Hong Kong, Shenzhen

Wenyu Wang
Indiana University

August 13, 2026

Abstract

Do institutional investors remain the "smart money" in cryptocurrency markets, where private information channels are limited and valuation depends on fundamentally different information from traditional securities? We combine a novel dataset tracking the daily news consumption of U.S. investment firms from 2016 to 2022 with Google search activity by retail investors and estimate a model of information acquisition. We distinguish between headline news, which summarizes Bitcoin and broader crypto market movements, and driver news, which covers developments across the crypto ecosystem and the regulatory environment. We find that while both news types contain information about transitory variation in Bitcoin's expected returns, nearly all information about the persistent component resides in driver news. Institutional investors devote greater attention to driver news than retail investors, allowing them to more precisely estimate the persistent component and retain an informational advantage. This advantage is larger than in equity markets and is more pronounced during market downturns.

Keywords: Bitcoin; Cryptocurrency; Information Acquisition; Institutional Investors; Retail Investors; Asset Pricing
JEL Classification: D83; G11; G12; G14; G23; G24; G41

*Fang can be reached at vfang@iu.edu; Kantak can be reached at pkantak@iu.edu; Lin can be reached at linwenwei@cuhk.edu.cn; and Wang can be reached at wenywang@iu.edu. We thank the Indiana University Pervasive Technology Institute for providing HPC database, storage, and consulting resources. We are also indebted to our data provider, who wishes to remain anonymous, for its generous support. All errors are our own.

1 Introduction

Institutional investors are widely regarded as "smart money" in traditional financial markets, and the literature attributes much of their edge to two informational advantages. First, institutional investors possess superior analytical capabilities through specialized expertise, dedicated research teams, and access to extensive datasets, allowing them to process and trade on new information more quickly and effectively than retail investors (e.g., Cohen, Gompers, and Vuolteenaho, 2002; Ben-Rephael, Da, and Israelsen, 2017). Second, institutional investors benefit from information channels unavailable to most retail investors, such as interactions with corporate managers and policymakers, allowing them to acquire and trade on information beyond public disclosures (e.g., Solomon and Soltes, 2015; Bernile, Hu, and Tang, 2016; Bradley, Jame, and Williams, 2022).

Whether institutional investors enjoy similar advantages in cryptocurrency (or "crypto," in shorthand) markets remains an open question. Cryptocurrencies differ from traditional securities in two important respects. First, because major cryptocurrencies generate no cash flows, valuing them hinges on information about the crypto ecosystem and the regulatory environment rather than on forecasting future cash flows. For Bitcoin in particular, valuation may depend on interpreting drivers such as transactional utility and network activity (Cong, Li, and Wang, 2021; Liu and Tsyvinski, 2021; Biais, Bisiere, Bouvard, Casamatta, and Menkveld, 2023), consensus sustainability and blockchain security (Pagnotta, 2022), competition with fiat currencies (Pagnotta, 2022; Benigno, Schilling, and Uhlig, 2022), and regulatory development (Anderson, Fang, Moon Jr, and Shipman, 2026). More generally, value-relevant information for crypto assets spans developments across the ecosystem, including underlying technology, evolving infrastructure, and growing applications, as well as the regulatory environment.

Second, unlike traditional financial markets, crypto markets lack managers, corporate access events, and insider networks through which institutional investors can systematically acquire private information. Instead, decentralized cryptocurrencies operate through open-source protocols and public blockchains, where consensus rules, transaction records, and network activity are observable to all market participants. Although much of the value-relevant information is publicly available, it is often highly technical, dispersed across multiple sources, and requires specialized expertise to interpret, creating scope for differences in investors' information-processing abilities. As traditional

sources of institutional informational advantage weaken in crypto markets, the challenge instead becomes identifying the most informative public signals and assessing their implications for long-run value. If institutional investors remain smart money in crypto markets, their advantage likely lies in their superior ability to identify, interpret, and integrate these signals.

This paper studies this possibility through the lens of news consumption. We build a model in which investors acquire information to form more accurate beliefs about expected returns, which in turn allow them to make more efficient portfolio allocations.[1] We consider two investor types, institutional and retail, who allocate investment across Bitcoin ("BTC"), the S&P 500 index ("SPY" ETF), and a risk-free asset while facing costly information acquisition. We focus on BTC because it has consistently been the dominant cryptocurrency and the most widely held crypto asset among institutional investors, making it a benchmark asset for the broader digital asset market (e.g., Fidelity Digital Assets, 2021). We use SPY to represent the equity market because it is the most followed index for U.S. equities and provides a natural contrast to BTC for our analysis, as private information channels and valuation frameworks are more developed for equities.[2] For tractability, the model treats the two assets as informationally separate. To the extent that institutional investors are also skilled at extracting information from cross-asset linkages (a channel our model does not capture), our estimates of their informational advantage are conservative.

Expected returns for both BTC and SPY consist of two latent components: a transitory component that captures short-run fluctuations and reverts relatively quickly, and a persistent component that reflects long-run variation and evolves more slowly. While realized returns reveal the noisy sum of these two components, the components themselves are not separately observable. Investors seek to learn about both components, as distinguishing persistent from transitory expected-return variation leads to more efficient portfolio allocation decisions.

Investors learn about these latent expected-return components by consuming public news. We classify news into two broad categories based on the labels assigned to each article by the data vendor's BERT-based classification algorithm. The first category, which we term "headline news," aims to capture summary reporting on recent price movements and market developments, without

[1] We do not observe actual crypto portfolio allocations. Nevertheless, portfolio efficiency provides the natural economic objective underlying institutional due diligence because institutions routinely evaluate prospective investments and update their beliefs about expected returns long before committing capital.

[2] We do not model a separate bond asset and instead follow the standard CAPM framework, in which investors choose between risky assets and a risk-free borrowing/lending asset (e.g., Sharpe, 1964; Lintner, 1965; Mossin, 1966).

delving into their underlying drivers. The second category, which we term "driver news," aims to capture articles analyzing the economic, technological, institutional, and regulatory forces that shape expected returns.[3] For BTC, headline news comprises articles tagged only with broad Bitcoin and cryptocurrency labels, whereas driver news comprises articles tagged with labels spanning the crypto ecosystem—including its technology (e.g., blockchain and cryptography), infrastructure (e.g., tokens, protocols, and exchanges), and applications (e.g., smart contracts and tokenization)—and regulation (e.g., anti-money laundering and know-your-customer requirements). For SPY, headline news comprises articles tagged only with broad equity market index labels, whereas driver news comprises articles tagged with labels related to economic activity, monetary policy, inflation and input costs, financial conditions and risk premia, and the global economy and trade. Tables A1 and A2 present the list of labels used to identify headline news and representative examples of labels used to identify driver news for BTC and SPY, respectively.

Investors observe the quantity of available news for free but must incur costs to consume and process article content. The benefits of information acquisition stem from more precise beliefs about expected returns, which reduce uncertainty and improve portfolio allocation. They depend on three channels: (i) the portfolio value of learning about a particular expected-return component (*value channel*); (ii) the extent to which a news category reveals information about that component (*informativeness channel*); and (iii) the quantity of available news, which determines how much information can be extracted (*quantity channel*). The costs reflect the expertise and resources required to identify and analyze value-relevant information, as well as behavioral biases that impede information processing. These costs may be lower for institutional investors because they possess greater analytical expertise and resources and are less susceptible to behavioral biases due to professional training and organizational oversight (see Barber and Odean (2013) for a review of retail investor trading behavior). Given the available news environment, investors choose how much news to consume and process by weighing the benefits against the costs.

The model combines realized returns with measures of news supply and consumption to disentangle the latent transitory and persistent expected-return components, quantify the informa-

[3]The news production cycle typically begins with headline news, which reports recent price movements and market developments. Because speed is critical, such coverage generally summarizes observable events without examining their causes. Subsequent reporting and analysis may produce driver news that examines the underlying forces and their value implications. As such, we posit that headline news is more likely to reflect short-run fluctuations, whereas driver news is more likely to reveal information relevant to persistent changes in expected returns.

tiveness of different news categories, and infer investors' information-processing costs. Specifically, return data allow us to quantify the value of learning about the transitory and persistent expected-return components (*value channel*), while returns and news-supply data jointly identify the extent to which different news categories reveal information about these components (*informativeness channel*). We model news supply as increasing with the magnitude of underlying information events, so the quantity of published articles provides a publicly observable signal about the amount of value-relevant information generated (*quantity channel*). Conditional on the quantity of available news, investors' news-consumption choices reveal their information-processing costs. The model therefore infers information-processing costs from consumption patterns after controlling for differences in the value, informativeness, and quantity of available information, enabling separate identification of information supply and information demand.

To estimate the model, we combine data from several sources. First, we use a proprietary dataset that tracks daily interactions between anonymous readers and online news content from July 2016 to November 2022 and links readers' article engagement to firms through IP addresses and cookie information. We measure news supply as engagement with the headline and driver news articles defined earlier by readers linked to major news organizations, including *Bloomberg*, *Forbes*, *The New York Times*, *The Wall Street Journal*, *The Washington Post*, and *USA Today*. Because journalists frequently consume related content while developing news coverage, higher engagement proxies for greater news production. We measure institutional news consumption using engagement with the articles by readers linked to U.S. investment firms in NAICS four-digit industry code 5239, including portfolio managers, investment advisers, hedge funds, private equity ("PE"), and venture capital ("VC") firms. We focus on these institutions because they are more likely to consider Bitcoin as an alternative asset given their greater investment flexibility and fewer regulatory constraints. Second, following Da, Engelberg, and Gao (2011), we measure retail news consumption using Google Search Volume Index (SVI) data for the same headline and driver labels. Third, we obtain BTC price data from Coinbase and SPY price data from the CRSP daily files.

The structural estimates yield four main findings. The first speaks to the *value channel*: transitory and persistent expected-return components account for 56.7% and 43.3% of portfolio-relevant expected-return uncertainty in BTC and 64.3% and 35.7% in SPY, respectively. Therefore, investors benefit from learning about both components, with information about persistent expected

returns being particularly valuable in BTC.

The second finding speaks to the *informativeness channel*: headline news and driver news differ sharply in their informational content. Driver news is the dominant source of information about persistent expected returns, accounting for 93.8% of long-run information for BTC and 98.5% for SPY. Headline news, in contrast, primarily reflects transitory return variation, especially for SPY. Therefore, driver news is the principal channel through which investors learn about persistent expected returns in both assets.

The third finding concerns the institutional informational advantage, which the model attributes to lower information-processing costs after controlling for differences in the value, informativeness, and quantity of available information. The estimates suggest that institutional news consumption improves uncertainty resolution by 14.3 percentage points for BTC and 6.6 percentage points for SPY, compared with 6.7 and 3.1 percentage points for retail investors, yielding an informational advantage of 7.6 percentage points in BTC and 3.5 percentage points in SPY. This advantage arises because institutions consume more driver news, the primary source of persistent expected-return information. Because such information is valuable for portfolio allocation, particularly for BTC, greater consumption of driver news translates into larger information gains. Consistent with this mechanism, driver news accounts for 92.2% of institutional information gains in BTC, compared with 81.6% in SPY. This contrast suggests that public-news processing plays a particularly important role in crypto markets, where traditional private information channels are limited.

The fourth finding concerns the state dependence of institutional informational advantage. Using the Relative Strength Index (RSI) as a proxy for market conditions, we find that information-processing costs vary systematically over the market cycle. In BTC, institutional investors devote greater attention to driver news during weaker market conditions and less attention during stronger market conditions, while retail investors exhibit the opposite pattern. This pattern is consistent with evidence from behavioral finance that retail attention is return dependent. For example, Sicherman, Loewenstein, Seppi, and Utkus (2016) document the "ostrich effect," where investors are less likely to seek information following unfavorable trading outcomes. This divergence in attention is particularly pronounced in BTC, where state-dependent processing of driver news plays an important role in generating institutional informational advantage.

Taken together, the results suggest that institutions' informational advantages depend not only

on access to information but also on how they allocate attention across available information sources. Although the information source examined in this study is publicly available, institutional investors acquire substantially more economically valuable information because they devote greater attention to the news categories that are most informative about persistent expected-return variation. These effects appear particularly important in crypto markets, where public information may play a more prominent role in price discovery as traditional private-information channels are relatively limited.

This study contributes to three streams of literature. First, it extends the literature on institutional investment in alternative assets. Prior research shows that institutional investors continually expand their opportunity set by entering markets such as PE and VC in pursuit of diversification and portfolio optimization (e.g., Lerner, Schoar, and Wang, 2008; Ivashina and Lerner, 2018; Lerner, Mao, Schoar, and Zhang, 2022). Crypto markets provide a unique setting because they lack cash-flow-based valuation anchors and offer limited opportunities to acquire private information. Even in such challenging markets, our evidence shows that investment firms began exploring cryptocurrencies as an alternative asset class as early as 2017, well before Bitcoin was widely regarded as an institutional investable asset, and maintained an informational advantage over retail investors.

Second, this study adds to the nascent literature on institutional due diligence. Unlike retail investors, institutional investors bear the responsibility of conducting extensive due diligence on investment opportunities before committing capital. Yet this process is inherently difficult to observe, and existing evidence comes primarily from proprietary data on PE, VC, and hedge funds. Most prominently, Kaplan and Strömberg (2001) and Kaplan and Strömberg (2004) examine investment memoranda, contracts, and internal evaluations from a small number of VC firms, showing that concerns about managerial quality, market opportunities, and execution risk play central roles in investment screening, governance, and contracting decisions. Gompers, Gornall, Kaplan, and Strebulaev (2020) report survey responses from 885 venture capitalists at 681 firms, finding that management quality and market potential are among the most important factors in both screening and due diligence. More recently, Brown, Gredil, and Kantak (2023) draw on proprietary documents and communications with hedge funds to examine how institutional investors allocate attention across due diligence activities. They document substantial differences across funds and over time in both the depth of analysis and breadth of information used in investment decisions.

Our study advances this literature in two ways. First, we construct one of the first high-

frequency, large-scale measures of institutional due diligence based on investors' engagement with public news. Our approach is most closely related to Kwan, Liu, and Matthies (2026), who use institutional news-reading data to measure fund-level attention to public information and examine how attention allocation relates to portfolio decisions and investment performance. In contrast, we compare institutional and retail engagement with crypto news, providing new evidence on investor learning in an emerging asset class. Second, we show that institutional investors systematically allocate attention to the public information most relevant for long-run value in both equity and crypto markets, providing large-scale evidence that their informational advantages arise in part from superior public-information acquisition and processing. This finding complements recent evidence that institutions optimally allocate due diligence effort (Fu and Taylor, 2025) and that information acquisition improves capital allocation outcomes (Brown et al., 2023).

Finally, this study broadens the empirical literature on crypto markets, which has largely focused on retail investors' trading behavior, attention dynamics, and speculative activity (e.g., Benetton and Compiani, 2020; Hackethal, Hanspal, Lammer, and Rink, 2022; Pursiainen and Toczynski, 2022; Aiello, Baker, Balyuk, Di Maggio, Johnson, and Kotter, 2023; Bertomeu, Martin, and Zhang, 2023; Didisheim and Somoza, 2024; Divakaruni and Zimmerman, 2024; Kogan, Makarov, Niessner, and Schoar, 2024; Auer, Cornelli, Doerr, Frost, and Gambacorta, 2025). In contrast, evidence on institutional participation remains limited due to data constraints and a common perception that institutions largely remained on the sidelines during the formative years of the crypto market (Duchin, Solomon, Tu, and Wang, 2025). Because trading and portfolio allocations are generally unobservable for most investment firms, our use of news-consumption data provides a novel window into institutional participation in this emerging market. Our evidence points to a much earlier and more sustained institutional presence in crypto markets than previously recognized, showing that investment firms actively engaged with crypto-related news beginning in 2017 and remaining engaged throughout subsequent market cycles, including the prolonged 2018–2019 crypto winter, and systematically focused on the information most relevant for long-run value.[4] These patterns are consistent with an extended period of institutional learning and due diligence preceding the accelerating institutional adoption of digital assets in later years.

[4]This result aligns with anecdotal evidence. For example, Fidelity, one of the largest U.S. asset managers, launched Fidelity Digital Assets in 2018. Initially focused on researching cryptocurrencies and blockchain applications, the division soon began publishing annual institutional surveys and, by 2019, expanded into trading and custody services.

2 Data

We obtain a proprietary dataset from a data vendor that tracks interactions between online readers and news articles published by thousands of its member publishers, including the major news outlets discussed earlier as well as more specialized industry sources. The dataset's broad coverage and detailed article-level engagement records make it particularly useful for studying information acquisition and attention allocation (e.g., Baba Yara, Grigoris, and Kantak, 2024; Kwan et al., 2026). The data vendor observes the article URL, the reader's external IP address, and cookie information. It uses the URL to retrieve article content and applies a BERT-based classification algorithm to assign each article to a set of predefined labels. It then uses the IP address and cookie information to map each interaction to an entity-level domain. The resulting dataset links label-level article engagement to the firms whose employees engage with the content. The data are aggregated to the entity-date-label level.

Because our analysis focuses on BTC, the dominant cryptocurrency, and SPY, a widely used proxy for the U.S. equity market, we extract engagement data associated with news relevant to these two assets. Our sample spans July 5, 2016, through November 28, 2022, the last date for which engagement data are available to us. We obtain daily BTC return data from Coinbase and daily SPY return data from the CRSP daily stock files over the same period.

2.1 Selecting Labels

To construct measures of news supply and news consumption for BTC and SPY, we first identify the vendor-supplied labels used to classify headline and driver news. The vendor's taxonomy is extensive, containing roughly 10,000 fine-grained labels. Two challenges are noteworthy. First, no taxonomy can perfectly capture every aspect of a complex and evolving market. Economic concepts may be represented at different levels of granularity, while closely related concepts may appear under separate labels. Second, we observe only the vendor-supplied labels rather than the proprietary algorithm used to assign them or the underlying article text. Ideally, we would retrieve the original articles and apply a single time-invariant classification algorithm ourselves. However, because of copyright restrictions, the vendor retains article URLs rather than the underlying text, and many older URLs are no longer accessible.

Importantly, these challenges are less limiting for our attention measures than they may initially appear because our approach does not rely on keyword-based searches. Instead, the vendor-supplied labels are used only to identify relevant news articles rather than to characterize their complete content. Once an article is identified, our measure captures readers' actual engagement with the article itself, not merely the presence of a particular label. For example, an article tagged with a given crypto-related label may discuss additional aspects of a technology development beyond those explicitly reflected in the taxonomy, and readers' engagement with that article is fully captured by our measure. In other words, the taxonomy need not exhaustively enumerate every economic concept. To the extent that it adequately identifies the relevant articles, our attention measures capture readers' engagement with the full content of those articles.

For BTC, we identify headline news using articles tagged with only umbrella labels including "Bitcoin," "cryptocurrency," and related variants. We identify driver news using articles tagged with labels spanning the crypto ecosystem, including its technological foundations (e.g., "distributed ledgers," "private keys," and "SHA-256"), infrastructure (e.g., "Ethereum," "initial coin offering," and "Coinbase"), and applications (e.g., "smart contracts," "tokenization," and "Hyperledger"), as well as the regulatory environment (e.g., "anti-money laundering," "know-your-business," and "know-your-customer"). Table A1 reports the complete list of labels used to identify BTC headline news together with representative examples of the labels used to identify BTC driver news.

We apply the same approach to SPY. We identify headline news using articles tagged only with broad equity market index labels, including "stock market index," "S&P 500 index," "Nasdaq," "Dow Jones Industrial Average," and related variants. We identify driver news using articles tagged with labels related to economic activity (e.g., "gross domestic product," "Bureau of Labor Statistics," and "supply chain"), monetary policy (e.g., "Federal Open Market Committee," "bond yield," and "LIBOR"), inflation and commodity markets (e.g., "consumer price index," "inflation," and "Brent crude"), financial conditions and risk premia (e.g., "credit risk," "liquidity risk," and "market volatility index"), and the global economy and trade (e.g., "exchange rate," "global supply chain," and "free trade"). Table A2 reports the complete list of labels used to identify SPY headline news together with representative examples of the labels used to identify SPY driver news.

2.2 Measures of News Consumption

We next construct measures of retail and institutional attention using the vendor-supplied labels to identify relevant news articles. Following prior literature (e.g., Da et al., 2011; Kwan et al., 2026), we also refer to these measures as retail and institutional attention.

For retail investors, we build on Da et al. (2011) and measure attention using Google SVI. Because Google Trends normalizes search volume within each query, SVI series obtained from separate queries are not directly comparable. To place labels on a common scale, we download daily Google Trends data by querying each label together with the asset's headline-news labels, which serve as a common benchmark. As a result, all label-level SVI measures for a given asset are scaled relative to the same benchmark and therefore become directly comparable. We measure retail attention for each asset-news category by aggregating the corresponding label-level SVI measures and normalizing the resulting series by the historical maximum level of headline-news attention for that asset.

We measure institutional attention using data from the proprietary news engagement dataset. We define institutional investors as investment firms in NAICS four-digit industry code 5239, which includes portfolio management companies and investment advisers. The vendor provides three measures of label-level engagement: user counts, record counts, and label-weighted interaction counts. We focus on user counts because they are conceptually closest to Google SVI. Although Google SVI measures relative search volume rather than unique users, both metrics capture the breadth of attention devoted to a label rather than the intensity of repeated engagement by the same individual.

For each asset-news category, we aggregate user counts across the corresponding labels and scale the resulting measure by total news engagement across all labels among 5239 firms. We then normalize the series by the historical maximum level of headline-news attention for that asset, mirroring the normalization used for the retail measure. The resulting measure is unit-free and directly comparable across assets, news categories, and investor groups.

Figure 1 plots the BTC attention series. Panels 1a and 1b show institutional attention to headline and driver news, while Panels 1c and 1d show the corresponding retail Google SVI measures. Two patterns stand out. First, institutional attention to BTC news emerges as early as

retail attention, with visible activity beginning in early 2017, but evolves more gradually over time. While both groups pay greater attention during Bitcoin price booms, retail attention rises and falls much more sharply and is concentrated around bull markets. Second, compared to retail investors, institutional attention exhibits a sharper distinction between headline news and driver news. Institutional attention to driver news remains relatively stable throughout the sample and declines much less during market downturns, particularly during the two-year crypto winter following the 2017 boom. In contrast, retail attention to driver news largely disappears during these periods. These patterns start to reveal that institutional investors engage more continuously with driver news, whereas retail investors' attention is more episodic and tied to market conditions.

2.3 News Supply and Volatility

The attention measures defined above may reflect both the supply of news and investors' demand for that news. To separate these forces, we also construct a measure of news supply using the same label-based classifications and aggregation procedure as the institutional attention measure, but restrict the engagement data to six major media outlets: *Bloomberg*, *Forbes*, *The New York Times*, *The Wall Street Journal*, *The Washington Post*, and *USA Today*.

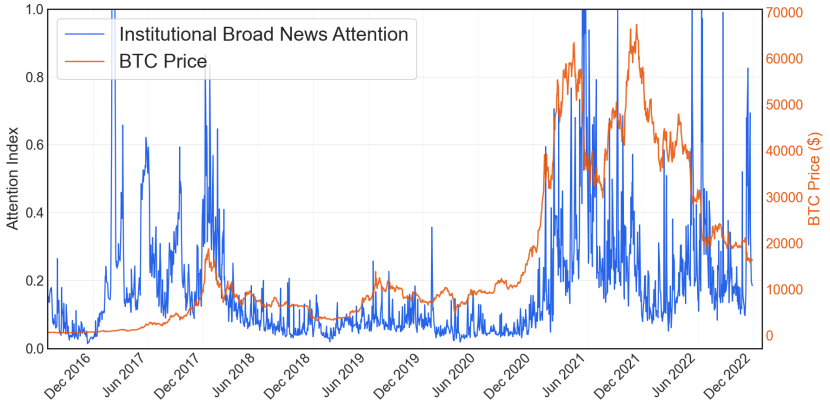
The idea is that when newsworthy events occur, journalists research, update, and expand coverage around the relevant labels. Engagement by these organizations therefore reflects the intensity of news production rather than different investors' demand for the content. Specifically, for each asset-news category, we aggregate engagement across the corresponding labels and scale the resulting series relative to overall activity by these outlets. We then apply the same normalization used for the institutional attention measure. The resulting measure captures changes in the supply of coverage devoted to a given asset-news category over time, allowing us to distinguish variation in investor attention arising from changes in news supply from variation arising from different investors' information demand.

In the model, news supply is rational and increasing in the variance of latent information events, implying that the quantity of news articles serves as a publicly observable signal about the amount of value-relevant information being generated. As a result, the relation between news supply and volatility would shed light on whether different news categories are more closely associated with transitory or persistent information events.

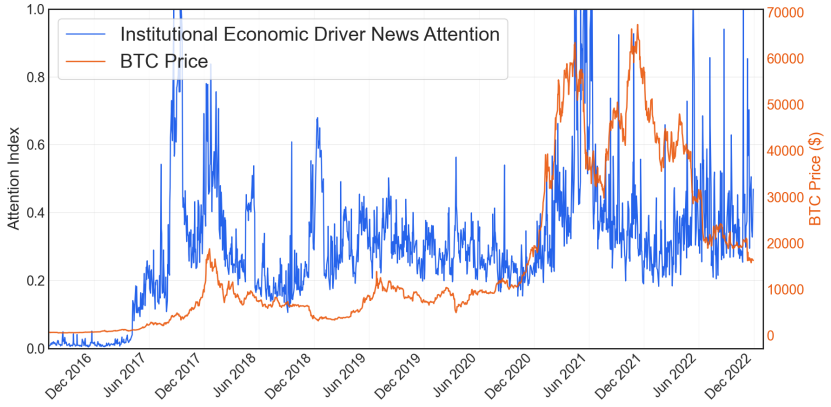
Figure 1: News Attention Timeseries

This figure plots the time-series of BTC price and our headline news and driver news consumption across both Institutional and Retail Investors.

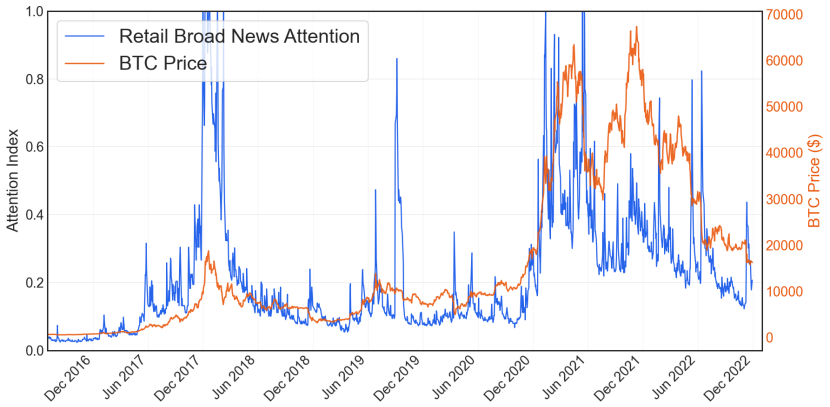
(a) *Institutional attention to BTC headline news*



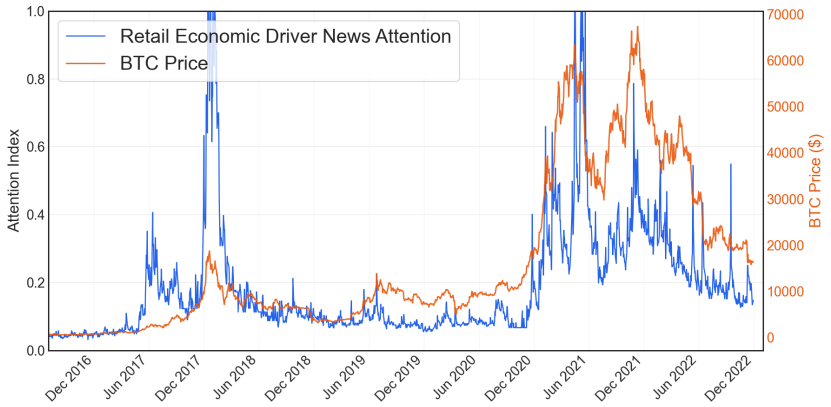
(b) *Institutional attention to BTC driver news*



(c) *Retail attention to BTC headline news*



(d) *Retail attention to BTC driver news*



To examine how news supply relates to different types of market uncertainty, we first separate daily returns into high-frequency and low-frequency components using a band-pass filter with a 30-trading-day cutoff following Christiano and Fitzgerald (2003). The high-frequency component captures return fluctuations occurring at horizons shorter than roughly one month, while the low-frequency component captures more persistent movements. For each asset and month, we then compute annualized volatility separately for the two components using the corresponding daily filtered returns. We also aggregate news supply to the monthly level by averaging daily news supply within each month for each asset-news category. Finally, we regress monthly high-frequency and low-frequency volatility on monthly news supply. The estimated coefficient therefore reflects how strongly variation in headline or driver news supply is associated with each volatility component and can be interpreted as a reduced-form measure of whether a news category is more closely linked to transitory or persistent information events.

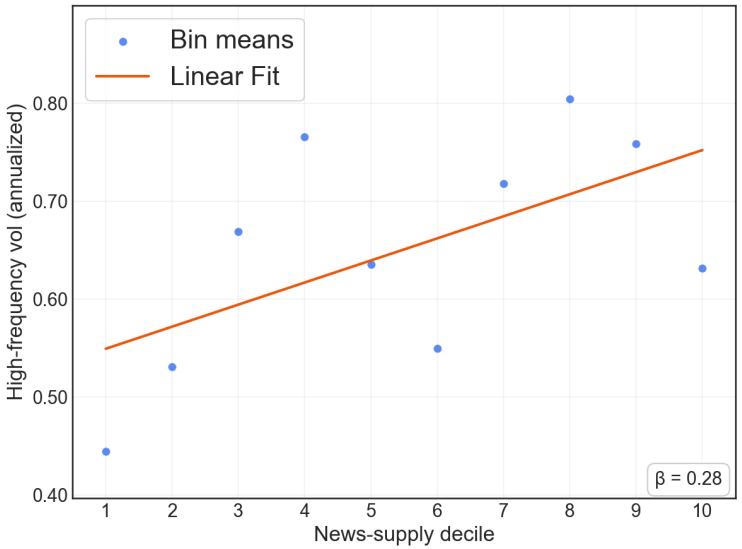
Figure 2 illustrates how BTC news production varies with different sources of return volatility. The figure plots average monthly high-frequency and low-frequency volatility across deciles of headline and driver news supply. Each point represents the average volatility within a news-supply decile, while the reported β is estimated using the underlying monthly observations. Two patterns emerge. First, high-frequency BTC volatility is positively associated with both headline and driver news supply. Second, low-frequency BTC volatility is positively associated with driver news supply but shows little relation to headline news supply. These results suggest that headline news is primarily linked to short-run volatility, whereas driver news is linked to both transitory and persistent volatility. Figure A1 reports the corresponding results for SPY and reveals similar patterns.

Taken together, these patterns suggest that investors differ not only in how much attention they pay, but also in the type of information they consume. They also suggest that news production responds systematically to underlying market conditions. Because observed news consumption reflects a combination of news supply, news informativeness, and information-processing costs, reduced-form evidence alone cannot disentangle these channels. We therefore develop a structural model to identify these forces and assess whether institutional investors have an informational advantage in learning about persistent expected-return variation.

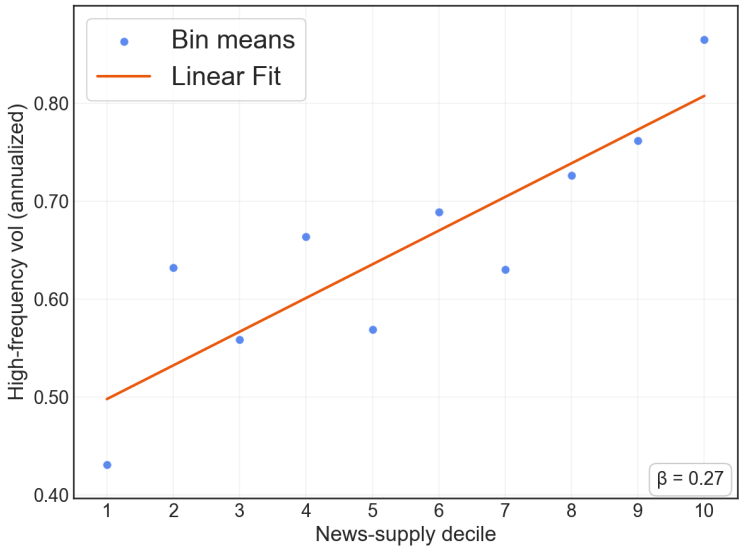
Figure 2: BTC News Supply

This figure plots the association between news supply and BTC volatility. News supply comes in two categories—headline and driver-related news. Likewise, BTC volatility is decomposed into a low and high-frequency component using a Christiano and Fitzgerald (2003) band-pass filter.

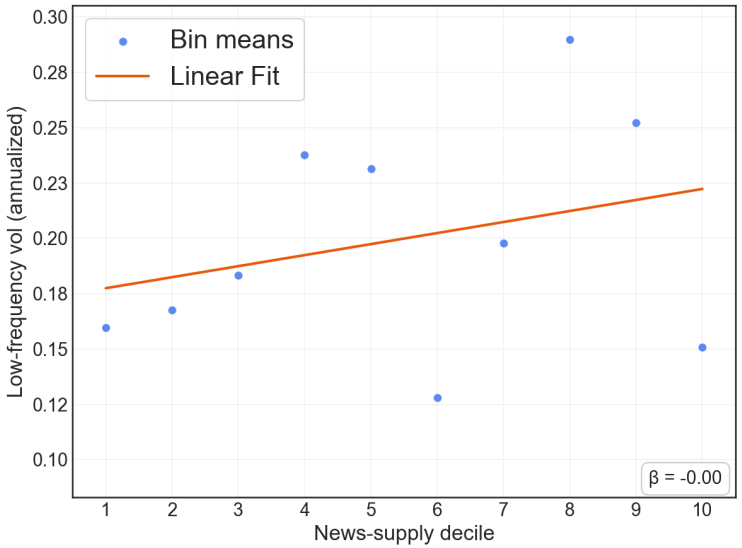
(a) *Headline news supply and BTC high-frequency volatility*



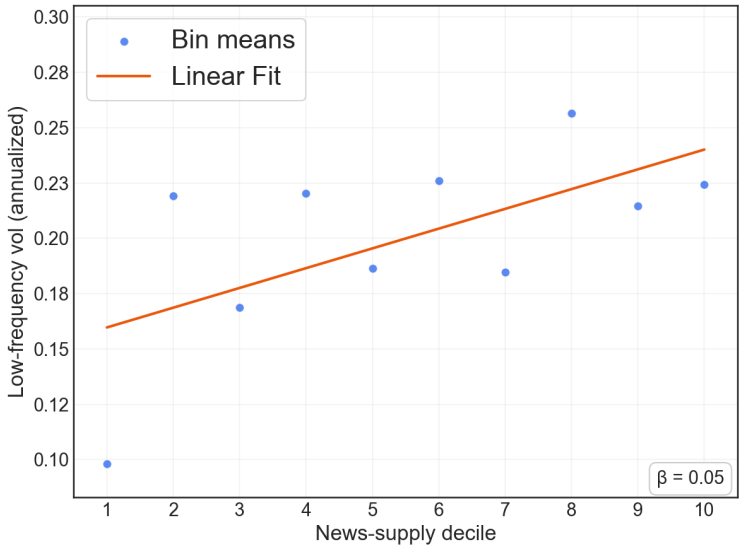
(b) *Driver news supply and BTC high-frequency volatility*



(c) *Headline news supply and BTC low-frequency volatility*



(d) *Driver news supply and BTC low-frequency volatility*



3 Model

We study a discrete-time economy in which two types of investors—institutional and retail—allocate wealth across Bitcoin (BTC), the S&P 500 index (SPY), and a risk-free asset. The conditional expected return on each risky asset varies over time and is the sum of two latent components: a persistent component that captures slow-moving forces shaping the asset's value over long horizons, and a transitory component that captures short-lived fluctuations. Realized returns reveal only a noisy sum of the two components. Telling them apart matters for portfolio choice because the two components carry different implications for future returns: a persistent movement in expected returns forecasts returns far into the future, whereas a transitory movement dies out within days. Forming accurate beliefs about the two components is therefore valuable, and this is the motive for information acquisition in the model.

Investors learn about expected returns by consuming public news. For each asset, news articles fall into two categories: headline news, which summarizes recent price movements and market developments, and driver news, which examines the underlying economic, technological, institutional, and regulatory forces. The two categories may reveal different amounts of information about the persistent and transitory components. Importantly, we do not impose which category reveals which component; instead, we estimate this mapping from the data through a news-loading matrix. News enters investors' decisions through two distinct margins. The *quantity* of published articles is observed by everyone at no cost, and—because more articles are published when more is happening—quantity itself is informative about how uncertain the environment currently is. The *content* of articles, in contrast, must be actively processed at a cost. Each period, each investor type therefore makes two linked decisions for each asset: how many articles of each news category to process, and how much wealth to allocate to the asset. The cost of processing news is investor-type specific and time varying; these costs are the model's measure of differences in information-processing ability between institutional and retail investors.

To keep the model tractable, we maintain a few assumptions that make the news-consumption and portfolio-choice problems separable across the two risky assets, so that each asset can be analyzed on its own. First, investors have CARA preferences, which eliminate wealth effects.[5]

[5]Three reasons favor CARA over CRRA. (i) *Wealth independence.* The certainty-equivalent gain from holding a belief depends only on that belief, not on current wealth, so investors with the same beliefs take the same dollar

Second, returns are conditionally independent across the two risky assets. Third, news about one asset carries no information about the other asset's expected returns. Together, these assumptions imply that the portfolio and news-consumption decisions for BTC and SPY separate into two independent problems (Section OA.1 of the Online Appendix presents the argument), allowing us to describe the model for a generic risky asset $a \in \{BTC, SPY\}$ throughout. We discuss how these assumptions may affect our estimates and results in Section 4.2.

Within each period t, events unfold in four steps:

1. *Public signals.* All investors freely observe the realized return and the quantity of published articles in each news category.

2. *Attention choice.* Before seeing any article content, each investor type chooses how many articles of each news category to process.

3. *Belief update.* Each investor observes the content signals generated by her reading choice and updates beliefs about the expected-return components.

4. *Portfolio choice.* Each investor chooses her position in the risky asset; returns are realized at the start of period $t + 1$.

The rest of this section fills in the four steps. Section 3.1 specifies the return process, Section 3.2 the news environment together with the benefit and cost of processing news, and Section 3.3 the portfolio and information-acquisition choices, including how news consumption shapes beliefs. Technical details are provided in the Online Appendix, and Appendix A collects the model notation.

3.1 Expected Returns and Information Events

Time is discrete and indexed by $t = 0, 1, 2, \ldots$, with one period corresponding to one trading day. Investors maximize expected discounted lifetime utility with discount factor $\beta \in (0, 1)$. The risk-free return is normalized to zero, so all returns below are interpreted as excess returns.[6]

position regardless of wealth. (ii) *Tractable dynamic programming.* Per-period payoffs are independent of wealth, so the Bellman problem has a belief-state representation without wealth as a state variable. (iii) *Comparison with CRRA.* Under CRRA, the optimal portfolio *fraction* is wealth-independent, but the attention problem then depends on wealth, making the dynamic program intractable once expected returns are predictable.

[6] We do not model a separate bond asset and instead follow the standard CAPM framework, in which investors choose between risky assets and risk-free borrowing/lending (e.g., Sharpe, 1964; Lintner, 1965; Mossin, 1966).

For each risky asset $a \in \{BTC, SPY\}$, the realized return is

$$r_t^a \;=\; \bar{\mu}^a + \mu_t^{a,L} + \mu_t^{a,S} + \varepsilon_t^a, \qquad \varepsilon_t^a \sim \mathcal{N}(0, \sigma_{\varepsilon,a}^2), \tag{1}$$

where $\bar{\mu}^a$ is the unconditional average return, $\mu_t^{a,L}$ and $\mu_t^{a,S}$ are the two latent expected-return components, and ε_t^a is unpredictable noise. We refer to $\tilde{r}_t^a \equiv r_t^a - \bar{\mu}^a = \mu_t^{a,L} + \mu_t^{a,S} + \varepsilon_t^a$ as the centered return; it is the noisy sum through which returns reveal the two components.

The two components differ in persistence. Each follows an AR(1) process,

$$\mu_{t+1}^{a,L} = \rho_L^a \mu_t^{a,L} + \eta_{t+1}^{a,L}, \qquad \eta_{t+1}^{a,L} \mid V_t^{a,L} \sim \mathcal{N}(0, V_t^{a,L}), \tag{2}$$

$$\mu_{t+1}^{a,S} = \rho_S^a \mu_t^{a,S} + \eta_{t+1}^{a,S}, \qquad \eta_{t+1}^{a,S} \mid V_t^{a,S} \sim \mathcal{N}(0, V_t^{a,S}), \tag{3}$$

with $0 < \rho_S^a < \rho_L^a < 1$: the long-run component $\mu_t^{a,L}$ mean-reverts slowly, so movements in it forecast returns over extended horizons, whereas the short-run component $\mu_t^{a,S}$ decays within days.

The innovation variances $V_t^{a,L}$ and $V_t^{a,S}$ are themselves time varying. We interpret them as measuring the size of current *information events* for each component: periods of high $V_t^{a,j}$ are periods in which large shocks are hitting component j—for example, regulatory or adoption developments moving the long-run component, or shifts in short-run trading conditions moving the transitory one. In such periods the trajectory of expected returns is especially uncertain, so forming accurate beliefs is especially valuable. Each log variance follows an AR(1) process,

$$\log V_{t+1}^{a,j} - \mu_{V,a,j} = \phi_{V,a,j}\left(\log V_t^{a,j} - \mu_{V,a,j}\right) + \omega_{t+1}^{a,j}, \qquad \omega_{t+1}^{a,j} \sim \mathcal{N}(0, \sigma_{\omega,a,j}^2), \quad j \in \{L, S\}, \tag{4}$$

a specification that guarantees strictly positive variances while allowing volatility to be persistent.

Finally, consistent with the independence assumptions discussed above, the state processes $(\mu_t^{a,L}, \mu_t^{a,S}, V_t^{a,L}, V_t^{a,S})$ and the return innovations ε_t^a are independent across assets, so returns exhibit zero conditional covariance:

$$\mathrm{Cov}\big(r_{t+1}^{BTC}, r_{t+1}^{SPY} \mid \mathcal{F}_t^i\big) = 0, \tag{5}$$

where $\mathcal{F}_t^i$ denotes investor i's information set at time t.

3.2 News and Information

News articles are the vehicle through which investors learn about the latent expected-return components. For each asset a, news falls into two categories indexed by

$$k \in \{H, D\},$$

where H denotes *headline* news (mnemonically, *market*-summary coverage), such as articles summarizing recent price movements and market developments, and D denotes *driver* news (*economic*-driver coverage), such as articles examining the technological, institutional, and regulatory developments behind Bitcoin's long-run viability and adoption, or the macroeconomic conditions shaping corporate cash flows, discount rates, and risk premia. The two categories may well differ in what they reveal about the two expected-return components, but the model imposes no restriction on this mapping; instead, we let the data reveal each category's information content through the news loadings introduced below.

News content as signals. We formalize the idea that news is the source of information by modeling article content as a noisy signal of the expected-return components. Recall from Section 3.1 that the centered return reveals only the noisy sum $\mu_t^{a,L} + \mu_t^{a,S}$; news content is what helps investors tell the two components further apart. An individual category-k article carries the signal

$$\tilde{y}_{k,t}^a = b_L^{a,k} \mu_t^{a,L} + b_S^{a,k} \mu_t^{a,S} + \tilde{\nu}_{k,t}^a, \tag{6}$$

where the loadings $b_L^{a,k}$ and $b_S^{a,k}$ capture how strongly category-k content reflects the long-run and short-run components, and the article-level noise $\tilde{\nu}_{k,t}^a$, specified below, is drawn independently across articles. For example, if driver news is mainly about long-run forces, its signal should load heavily on $\mu_t^{a,L}$ and little on $\mu_t^{a,S}$. We collect the loadings in an asset-specific *news-loading matrix*, which we estimate from the data:

$$B^a \equiv \begin{pmatrix} b_L^{a,H} & b_S^{a,H} \\ b_L^{a,D} & b_S^{a,D} \end{pmatrix}, \tag{7}$$

where row $k \in \{H, D\}$ gives the loadings of news category k on the long-run and short-run components.

The information value of a category-k article depends on both its loadings and the current variance states: even a category that loads heavily on the long-run component teaches little if the long-run environment is quiet. We summarize this as the *news value* of category-k articles at time t:

$$U_{k,t}^{a} \equiv \left(b_{L}^{a,k}\right)^{2} V_{t}^{a,L} + \left(b_{S}^{a,k}\right)^{2} V_{t}^{a,S}. \tag{8}$$

News value is high when the expected-return components that category-k articles track are currently subject to large information events, so that a single article resolves substantial uncertainty. Accordingly, we specify the article-level noise in equation 6 as

$$\tilde{\nu}_{k,t}^{a} = \frac{\sigma_{\nu,a}}{\sqrt{U_{k,t}^{a}}}\,\tilde{\epsilon}_{k,t}^{a}, \qquad \tilde{\epsilon}_{k,t}^{a} \sim \mathcal{N}(0,1), \tag{9}$$

where $\sigma_{\nu,a} > 0$ governs the baseline content-noise level for asset a: each article is more informative precisely when its news value is high, that is, when there is more to learn.

News supply. Before an investor can read anything, articles must exist. Let $n_{k,t}^{a}$ denote the number of category-k articles about asset a published in period t: this is the menu available to investors, observed freely by all. We model the supply of news as rational: publishers produce more category-k articles precisely when there is more valuable information to report, that is, when $U_{k,t}^{a}$ is high. Log news supply is

$$\log n_{k,t}^{a} = \mu_{k}^{a} + U_{k,t}^{a} + u_{k,t}^{a}, \qquad u_{k,t}^{a} \sim \mathcal{N}(0, \sigma_{u,a,k}^{2}), \qquad k \in \{H, D\}, \tag{10}$$

where the intercept μ_{k}^{a} captures the baseline supply level and $u_{k,t}^{a}$ absorbs supply fluctuations unrelated to news value. This specification has an important implication: because article counts are freely observed, the *quantity* of news serves as a public signal of the size of current information events. Even an investor who reads nothing learns from a burst of category-k coverage that a large shock has hit the component(s) that category tracks—she learns how uncertain the environment is, though not the direction in which it is moving. The supply equation is also what generates the

co-movement between news volume and the magnitude of price movements documented in Section 2 (Figure 2), without news quantities carrying any directional information.

News consumption. How many of the available articles to process is the investor's central information-acquisition decision. Let $c_{k,t}^{i,a} \leq n_{k,t}^{a}$ denote *news consumption*: the number of category-k articles about asset a that investor type i processes in period t, chosen each period before article content is observed. Consumption governs how much information the investor extracts: because article-level noises are independent, averaging $c_{k,t}^{i,a}$ articles' worth of signals (6) scales the per-article noise down by $1/\sqrt{c_{k,t}^{i,a}}$, so the investor's aggregated content signal is

$$y_{k,t}^{i,a} = b_L^{a,k}\mu_t^{a,L} + b_S^{a,k}\mu_t^{a,S} + \frac{1}{\sqrt{c_{k,t}^{i,a}}}\,\tilde{\nu}_{k,t}^{a}, \tag{11}$$

where $\tilde{\nu}_{k,t}^{a}$ is the per-article noise defined in eq. (9). Signal precision therefore accumulates with the investor's own consumption choice ($c_{k,t}^{i,a}$, through the $1/\sqrt{c_{k,t}^{i,a}}$ scaling) and with per-article informativeness ($U_{k,t}^{a}$, through $\tilde{\nu}_{k,t}^{a}$). When economically important developments occur, each article carries more information, so investors who choose to read can form precise beliefs exactly when precision matters most.

Reading more news, however, is not free: processing article content takes time, analytical skill, and resources, and is subject to behavioral frictions that may impede information processing. We refer to this cost interchangeably as an attention cost or an information-processing cost; it is what keeps optimal consumption finite and is the other side of the information-acquisition tradeoff. Processing $\boldsymbol{c}_t^{i,a} = (c_{H,t}^{i,a}, c_{D,t}^{i,a})$ articles costs investor type i

$$\mathcal{C}_t^{i,a} = \sum_{k\in\{H,D\}} \frac{\lambda_{k,t}^{i,a}}{2}\left(c_{k,t}^{i,a}\right)^2\left(1+\tau_k^a\, q_{k,t}(n_{k,t}^{a})\right), \tag{12}$$

a traditional convex (quadratic) cost scaled by a congestion factor that rises when news is scarce relative to typical consumption levels, capturing that relevant content is harder to locate and process when few articles are available. Central to our analysis is the cost scale $\lambda_{k,t}^{i,a}$: it is investor-type specific and time varying, with a low $\lambda_{k,t}^{i,a}$ meaning that processing category-k news is cheap for investor type i. We choose functional forms for the congestion factor and the $\lambda_{k,t}^{i,a}$ process that let

the model fit the data closely, deferring the details to Section OA.1 of the Online Appendix. The path of $\lambda_{k,t}^{i,a}$ is the latent object that the estimation recovers from observed news consumption, and it is the model's measure of each investor type's information-processing capacity.

Table 1 summarizes the information available to each investor type: returns and news quantities are free public signals, whereas content signals are private and costly. Note that each content signal is itself a noisy signal of a combination of the two components; it is because the two news categories and the return weight the components *differently* that combining these signals helps the investor tell the components apart.

Table 1: Information set of investor type i for risky asset a.

Observable	Who receives	Cost	Informs about
Centered return $\tilde{r}_t^a$	All investors	Free	Noisy sum $\mu_t^{a,L} + \mu_t^{a,S}$
Log supply $\log n_{k,t}^a$	All investors	Free	Variance states $V_t^{a,L}$, $V_t^{a,S}$
Content signal $y_{k,t}^{i,a}$	Investor i only	Costly	Noisy weighted sum $b_L^{a,k}\mu_t^{a,L} + b_S^{a,k}\mu_t^{a,S}$

3.3 Portfolio Choice and Information Acquisition

With the return process, the news environment, and the cost of attention in place, we can now state the investor's problem: each period, investor type i chooses how much of asset a to hold and how many articles to process, with the two decisions tightly linked through beliefs—a more precise belief translates into a better-calibrated position, and it is precisely this portfolio value that makes processing news worthwhile. Investor type i has CARA utility with risk-aversion coefficient $\gamma > 0$, common across investor types.[7]

Portfolio choice. Each period, investor i chooses a dollar position $\theta_t^{i,a}$ in asset a. The certainty-equivalent payoff from holding this position depends on two objects: the posterior expected return $m_t^{i,a}$ and the total conditional variance of the next-period return, $\Sigma_{\text{total},t}^{i,a}$. Given these two objects, the payoff is

$$\Pi_t^{i,a} = \theta_t^{i,a}\, m_t^{i,a} - \frac{\gamma}{2}\left(\theta_t^{i,a}\right)^2 \Sigma_{\text{total},t}^{i,a}, \tag{13}$$

[7]Under the three assumptions in Section 3, the portfolio-choice and news-consumption problems for BTC and SPY are informationally and financially independent of each other (Section OA.1 of the Online Appendix gives the argument), so they can be solved separately and analogously for each asset; we state the problem below for a generic asset a.

and the first-order condition yields the optimal position and the resulting indirect utility:

$$\theta_t^{i,a,*} = \frac{m_t^{i,a}}{\gamma \Sigma_{\text{total},t}^{i,a}}, \qquad \Pi_t^{i,a,*} = \frac{\left(m_t^{i,a}\right)^2}{2\gamma \Sigma_{\text{total},t}^{i,a}}. \tag{14}$$

Mean and variance from news consumption. These two objects—the expected return and the total variance in eq. (13)—are exactly what news consumption shapes, though not uniformly: as the two equations below make clear, only part of each object responds to how much investor i reads. Let $\hat{\boldsymbol{\mu}}_{t|t}^{i,a}$ denote investor i's conjectured value of the two expected-return components, and $\Sigma_{t|t}^{\mu,i,a}$ her uncertainty around that conjecture—the precise specification is given in Section OA.1 of the Online Appendix. Let $\hat{\boldsymbol{V}}_t^a = (\hat{V}_t^{a,L}, \hat{V}_t^{a,S})'$ denote the size of current information events, known to all investors. Together, these beliefs deliver the two objects in eq. (13):

$$m_t^{i,a} \equiv \bar{\mu}^a + \boldsymbol{\rho}^{a\prime} \hat{\boldsymbol{\mu}}_{t|t}^{i,a}, \tag{15}$$

$$\Sigma_{\text{total},t}^{i,a} \equiv \boldsymbol{\rho}^{a\prime} \Sigma_{t|t}^{\mu,i,a} \boldsymbol{\rho}^a + \hat{V}_t^{a,L} + \hat{V}_t^{a,S} + \sigma_{\varepsilon,a}^2, \tag{16}$$

where $\boldsymbol{\rho}^a \equiv (\rho_L^a, \rho_S^a)'$ are the forecast weights on the two components—rather than one, because each component decays toward zero at its own rate before the next-period return is realized.

The mean $m_t^{i,a}$ is simply a forecast built from the conjecture $\hat{\boldsymbol{\mu}}_{t|t}^{i,a}$, and it is entirely shaped by news consumption: more reading lets the investor better disentangle the persistent and transitory dynamics underlying her conjecture, sharpening it and leading to wiser portfolio choices. The variance $\Sigma_{\text{total},t}^{i,a}$ has three pieces, only the first of which responds to news consumption. Belief uncertainty, $\boldsymbol{\rho}^{a\prime} \Sigma_{t|t}^{\mu,i,a} \boldsymbol{\rho}^a$, shrinks with reading: more news consumption reduces the aggregate noise in the investor's content signals, yielding a more accurate conjecture—especially when she reads both news categories, since their different loadings on the two components help her separate persistent from transitory movements, something the return alone can never do. The other two pieces do not respond to reading: $\hat{V}_t^{a,L} + \hat{V}_t^{a,S}$ is the size of next-period shocks, known to everyone for free, and $\sigma_{\varepsilon,a}^2$ is irreducible return noise. Indirect utility in eq. (14) is therefore the payoff to precision: more informative content lowers belief uncertainty, which both increases the position the investor is willing to take and raises the certainty-equivalent payoff from taking it.

Information acquisition as a dynamic decision. The expected-return components and their volatility both follow persistent AR(1) processes, so investor i's beliefs are persistent too: today's conjecture and its precision carry over into tomorrow's prior. Reading more news today therefore does not just sharpen today's mean and variance—it also raises the precision of beliefs, and hence portfolio value, in future periods. The benefits of information acquisition thus extend beyond the current period, making news consumption an inherently dynamic decision, solved by dynamic programming. Let $J^{i,a}$ denote the value function for investor type i in asset a. The Bellman problem is

$$J^{i,a}\Big(\hat{\boldsymbol{\mu}}_t^{i,a}, \hat{\boldsymbol{V}}_t^a, \boldsymbol{\lambda}_t^{i,a}\Big) = \max_{\boldsymbol{c}_t^{i,a} \geq \mathbf{0}} \Big[\Pi_t^{i,a,*}(\boldsymbol{c}_t^{i,a}) - \mathcal{C}_t^{i,a} + \beta\, \mathbb{E}_t\Big[J^{i,a}\Big(\hat{\boldsymbol{\mu}}_{t+1}^{i,a}, \hat{\boldsymbol{V}}_{t+1}^a, \boldsymbol{\lambda}_{t+1}^{i,a}\Big)\Big]\Big], \tag{17}$$

where $\boldsymbol{c}_t^{i,a} = (c_{H,t}^{i,a}, c_{D,t}^{i,a})$ and the state collects the investor's posterior beliefs, the public variance path, and the attention-cost processes.[8] The first-order condition for category-k consumption makes this contemporaneous-versus-future tradeoff explicit: it equates the marginal benefit of information—today's payoff plus the discounted effect on tomorrow's continuation value—to its marginal cost:

$$\underbrace{\frac{\partial \Pi_t^{i,a,*}}{\partial c_{k,t}^{i,a}} + \beta\, \frac{\partial \mathbb{E}_t\left[J_{t+1}^{i,a}\right]}{\partial c_{k,t}^{i,a}}}_{\text{marginal benefit of information}} = \underbrace{\lambda_{k,t}^{i,a}\left(1 + \tau_k^a\, q_{k,t}(n_{k,t}^a)\right) c_{k,t}^{i,a}}_{\text{marginal attention cost}}. \tag{18}$$

4 Identification and Estimation

The empirical design follows the structure of the model. For each asset, returns discipline the latent expected-return process, news supply disciplines the arrival and composition of information events, and investor-specific news consumption disciplines attention costs. BTC and the SPY are treated as two parallel applications of the same framework: the state variables, parameters, and likelihood blocks are asset-specific, but the identification logic is the same.

[8] The posterior covariance $\Sigma_{t|t}^{\mu,i,a}$ is also part of the belief state; it evolves deterministically through the filtering recursions in Section OA.1 of the Online Appendix given the variance path and consumption choices, so we suppress it in the notation.

4.1 Parameters to Be Estimated

The model contains four economically distinct parameter blocks. First, the return block describes how expected returns move over time: the persistence parameters ρ_L^a and ρ_S^a govern long- and short-run mean reversion, while $\sigma_{\varepsilon,a}^2$ captures residual return risk. Second, the stochastic-volatility block governs the size and persistence of information events through the log-variance parameters $(\mu_{V,a,j}, \phi_{V,a,j}, \sigma_{\omega,a,j}^2)$, for $j \in \{L, S\}$. Third, the news-supply block determines how the headline and driver news articles load on these information events through the matrix B^a, together with supply intercepts and supply-noise variances. Finally, the attention-cost block governs how costly it is for each investor type to process available articles: the parameters $(\bar{\lambda}_k^{i,a}, \rho_{\lambda,k}^{i,a}, \sigma_{\eta,i,a,k}^2)$ describe the latent cost process, while τ_k^a captures the role of news supply in shifting effective reading costs.

Table 2: Parameter blocks and identifying variation

Block	Key parameters	Main identifying variation
Expected returns	$\rho_L^a, \rho_S^a, \sigma_{\varepsilon,a}^2$	Frequency decomposition of returns and the Kalman fit of the short- and long-run expected-return components
Information-event volatility	$\mu_{V,a,j}, \phi_{V,a,j}, \sigma_{\omega,a,j}^2$	Time variation in return innovations and in the quantity of news published about the asset
News informativeness	$B^a, \mu_H^a, \mu_D^a, \sigma_{u,a,H}^2, \sigma_{u,a,D}^2$	Relative co-movement of the headline and driver news supply with the short- and long-run volatility states
Attention costs	$\bar{\lambda}_k^{i,a}, \rho_{\lambda,k}^{i,a}, \sigma_{\eta,i,a,k}^2, \tau_k^a, \sigma_{u,i,a}^2$	Institutional and retail consumption of headline and driver news, conditional on the estimated information value of available articles

4.2 Identification

The first source of identification is returns. Realized excess returns reveal a noisy sum of the persistent and transitory expected-return components. A frequency decomposition of returns (Section OA.4 of the Online Appendix) provides empirical proxies for the different horizons in this sum, allowing the expected-return block to separate $\mu_t^{a,L}$ from $\mu_t^{a,S}$ and to estimate their persistence.

The second source of identification is news supply. In the model, news supply is high when the latent information-event variance is high. The headline and driver news may load differently on the short- and long-run components, so their relative co-movement with return volatility identifies

the news-loading matrix B^a. This is the point at which the data decide whether headline news is primarily short-horizon information, whether the driver news contains more persistent information, or whether the two types are empirically similar.

The third source of identification is news consumption. Once the return process, volatility states, and news-loading matrix are fixed, the marginal benefit of reading each type of article is pinned down by the model. Observed consumption then reveals attention costs, following the logic of revealed preference: the benefit of reading is pinned down by observables—returns discipline the portfolio value of learning about each expected-return component, and news supply disciplines the informativeness and quantity of available articles—so once these benefit components are fixed, high consumption when information is valuable implies low costs, while low consumption in the same information environment implies high costs. Institutional and retail investors are observed separately, so the model can recover investor-type-specific cost dynamics while holding the asset environment fixed.

This last comparison is central to the paper's interpretation. Institutional and retail investors face the same asset returns, the same public news supply, and the same estimated news informativeness. Differences in their consumption therefore map into differences in posterior precision and, ultimately, in the portfolio value of their information. The post-estimation analysis in Section 6 turns these estimated cost and belief objects into information-advantage measures.

Finally, the independence assumptions that separate the two assets in the model (Section 3) also matter for how the estimates should be read. The identification argument treats BTC and SPY as informationally self-contained: news about one asset is assumed to carry no information about the other. In practice, investors may process information jointly across asset classes, and institutional investors may be particularly skilled at extracting information from cross-asset linkages. Any such cross-asset learning operates through a channel the model does not capture, so to the extent that it exists, our estimates understate institutions' effective informational advantage. The estimated institutional advantage should therefore be interpreted as a lower bound.

4.3 Estimation Procedure

We estimate the model through a four-step sequential procedure that mirrors the identification logic above.[9] Steps 1–2 estimate the expected-return dynamics from returns; Step 3 estimates the volatility states and the news-loading matrix from news supply, holding the return block fixed; and Step 4 estimates the attention-cost processes from news consumption, holding the return, volatility, and news-loading objects fixed. Section OA.4 of the Online Appendix gives the technical details of the filters, parameter vectors, and likelihood evaluation used at each step.

5 Parameter Estimates

This section reports the structural estimates that feed into the information advantage analysis. Both assets are estimated over a matched daily sample running from July 5, 2016 to November 28, 2022 (T = 1,587 trading days after lag construction). Attention measures are defined in Section 2.2. The estimates reveal three empirical patterns. First, BTC and SPY have similar persistence in expected returns, but BTC has substantially larger residual return volatility. Second, headline news primarily tracks short-run variation, while driver news is more connected to persistent or broader-market information. Third, institutional investors face much lower driver attention costs than retail investors, especially in BTC. These patterns discipline the post-estimation analysis in Section 6.

5.1 Expected-Return Dynamics

The expected-return estimates pin down the persistence of the short- and long-run components that investors are trying to learn about (eqs. (2)–(3)), as well as the residual return variance (eq. (1)). Table 3 reports the two-state AR(1) estimates for BTC and SPY. The filtered return components are shown in Section OA.4 of the Online Appendix.

[9] An alternative would be to estimate all parameters and latent states jointly in a single step, exploiting the full likelihood and all cross-equation restrictions at once. An early version of this paper pursued that approach; we found the joint likelihood surface to be quite flat over large regions of the parameter space, so that the estimator frequently stalled well short of a well-identified optimum. The sequential procedure used here trades some of this theoretical efficiency for considerably greater practical stability: by estimating each parameter block in the order implied by the identification argument above, holding previously estimated blocks fixed, it imposes a sharper, more direct mapping between each parameter block and the data that identifies it, and we find that it performs substantially better than the joint alternative in practice.

Table 3: Step 2: expected-return dynamics.
MLE estimates of the two-component AR(1) expected-return process. ρ_S and ρ_L are the daily AR(1) coefficients of the short- and long-run components; σ_ε^2 is the idiosyncratic return variance. Both assets are estimated on the same July 2016 – November 2022 sample ($T = 1{,}587$).

Parameter	Description	BTC	SPY
ρ_S	Short-run AR(1) persistence	0.680	0.688
ρ_L	Long-run AR(1) persistence	0.985	0.984
σ_ε^2	Idiosyncratic return variance	1.121×10^{-3}	6.870×10^{-5}

The AR(1) persistence parameters are nearly identical across assets: $\rho_S \approx 0.68$ and $\rho_L \approx 0.985$ for both BTC and SPY. This implies that shocks to the short-run expected return have a half-life of about two days, while the long-run component has a half-life of roughly 46 trading days. The idiosyncratic return variance σ_ε^2 is about 16 times larger for BTC than for SPY, reflecting the well-known higher idiosyncratic volatility of crypto.

5.2 News-Loading Matrix and Stochastic Volatility

The news-loading estimates describe which type of news is informative about which return component, through the news-supply equation (10). Table 4 reports the estimated loading matrix and two complementary informativeness decompositions. The first decomposition is news-type-centered: conditional on reading the headline or driver news, it asks whether that news mainly tracks short-run or long-run return uncertainty. The second decomposition is component-centered: conditional on wanting to learn about the short-run or long-run component, it asks which type of news is the more informative source. Section OA.4 of the Online Appendix gives the exact construction of both decompositions from the fitted news-value component. The filtered volatility and news-supply paths are reported in Section OA.4 of the Online Appendix.

Two features stand out in Table 4. First, the within-news decomposition shows that the headline news is strongly short-run in both markets: nearly all of the fitted broad-market-news value is associated with V_t^S. Driver news is less one-dimensional. For BTC, driver news still covaries substantially with the short-run component because the filtered short-run variance state is much larger in the sample. For SPY, driver news places meaningful weight on both components.

Second, the within-news decomposition already shows that the persistent component explains only a small fraction of either news type's value: for BTC, long-run information accounts for

Table 4: News loadings and informativeness decompositions.
The **B** matrix is 2×2 with rows = news types (headline H, driver D) and columns = return components (S, L). Within-news shares ask which component each news type tracks; rows within each news type sum to 100%. Within-component shares ask which news type is more informative about each component; rows within each component sum to 100%. ϕ_j and $\sigma^2_{\omega,j}$ are the log-volatility AR(1) coefficient and innovation variance for component $j \in \{S, L\}$.

	BTC	SPY
B matrix elements		
b_{HS} (headline loads on S)	46.88	153.02
b_{HL} (headline loads on L)	37.76	62.79
b_{DS} (driver loads on S)	52.65	83.79
b_{DL} (driver loads on L)	146.63	513.30
Within-news shares: what component does each news type track?		
Headline news		
Share from V^S	98.7%	≈100%
Share from V^L	0.1%	≈0%
Covariance share	1.2%	≈0%
Driver news		
Share from V^S	81.5%	55.6%
Share from V^L	6.4%	32.5%
Covariance share	12.1%	11.9%
Within-component shares: which news type tracks each component?		
Short-run component V^S		
Headline news	44.2%	76.9%
Driver news	55.8%	23.1%
Long-run component V^L		
Headline news	6.2%	1.5%
Driver news	93.8%	98.5%
Volatility dynamics		
ϕ_S	0.993	0.989
$\sigma^2_{\omega,S}$	0.0062	0.0535
ϕ_L	0.935	0.943
$\sigma^2_{\omega,L}$	0.265	0.372

just 0.1% of headline-news value and 6.4% of driver-news value, versus almost 0% and 32.5% for SPY. Learning about the persistent component from news is therefore inherently more difficult than learning about the transitory component, and substantially more so for BTC than for SPY. Given this challenge, the within-component decomposition shows which news type is the better—if still limited—source: driver news is by far the dominant channel for long-run information in both markets, accounting for 93.8% of BTC long-run news value and 98.5% of SPY long-run news value, against just 6.2% and 1.5% from headline news. For the short-run component, the two markets differ in composition rather than magnitude: BTC short-run news value is split between headline and driver news, whereas SPY short-run news value is concentrated in headline news. This composition difference matters for interpretation: driver BTC news can look short-run in the within-news decomposition because V_t^S is large and volatile, but from the component-centered perspective it remains the primary source of long-run information.

The volatility-dynamics rows show that both assets have highly persistent short-run variance states, with ϕ_S close to one. The smaller BTC $\sigma^2_{\omega,S}$ implies a smoother short-run log-variance state, whereas SPY's larger innovation variance implies sharper revisions to that state. Because the scale of these states is jointly determined with the news loadings in the supply equation, the economically more transparent quantities are the fitted news-value shares reported above.

5.3 Attention-Cost Estimates

The attention-cost estimates quantify how costly it is for each investor type to turn available news into consumed information. Section 3 left the cost-scale process $\lambda^{i,a}_{k,t}$ unspecified; it follows a log-linear AR(1) process,

$$\log \lambda^{i,a}_{k,t+1} = (1-\rho^{i,a}_{\lambda,k}) \log \bar{\lambda}^{i,a}_k + \rho^{i,a}_{\lambda,k} \log \lambda^{i,a}_{k,t} + \eta^{i,a}_{k,t+1}, \qquad \eta^{i,a}_{k,t+1} \sim \mathcal{N}(0, \sigma^2_{\eta,i,a,k}),$$

with long-run mean $\bar{\lambda}^{i,a}_k$, persistence $\rho^{i,a}_{\lambda,k}$, and innovation variance $\sigma^2_{\eta,i,a,k}$. The log specification keeps attention cost positive; high persistence implies that episodes of unusually cheap or costly reading are long-lived. Observed news consumption links to this latent cost through the attention first-order condition (eq. (18)). The supply-scaling parameters τ_H and τ_D, which enter the congestion adjustment in eq. (12), are shared across investor types.

Table 5: Attention-cost estimates.
Entries are posterior means with posterior standard errors in parentheses. $\bar{\lambda}_k^i$ is the long-run mean attention cost. The cost ratio is retail divided by institutional. Supply-wedge parameters τ_H and τ_D are shared across investor types.

Parameter	BTC	SPY
Headline news (H)		
$\bar{\lambda}_H^{\text{inst}}$	0.343 (0.061)	0.053 (0.008)
$\bar{\lambda}_H^{\text{retail}}$	0.379 (0.126)	0.078 (0.011)
Retail/institutional headline cost ratio	1.1	1.5
$\rho_{\lambda H}^{\text{inst}}$	0.894 (0.023)	0.959 (0.006)
$\rho_{\lambda H}^{\text{retail}}$	0.936 (0.016)	0.956 (0.005)
τ_H	1.073 (0.325)	0.240 (0.041)
Driver news (D)		
$\bar{\lambda}_D^{\text{inst}}$	1.057 (0.550)	0.118 (0.023)
$\bar{\lambda}_D^{\text{retail}}$	78.596 (149.956)	3.329 (0.384)
Retail/institutional driver cost ratio	74.4	28.2
$\rho_{\lambda D}^{\text{inst}}$	0.983 (0.010)	0.945 (0.009)
$\rho_{\lambda D}^{\text{retail}}$	0.969 (0.013)	0.964 (0.009)
τ_D	8.801 (4.343)	0.225 (0.126)

Several patterns in Table 5 deserve attention. *Mean attention costs.* For BTC, institutional investors have a lower mean driver attention cost ($\bar{\lambda}_D^{\text{inst}} = 1.057$) than retail ($\bar{\lambda}_D^{\text{ret}} = 78.60$): retail investors face a 74× higher average cost barrier for driver consumption. For SPY, the same ordering holds but the ratio is more moderate (3.329 vs. 0.118, roughly 28×). Headline attention costs are more comparable across types in both markets.

Persistence and volatility. Driver cost persistence $\rho_{\lambda D}$ is high for both types and both assets (0.945–0.983), implying that periods of elevated or suppressed driver consumption are prolonged. Headline cost persistence $\rho_{\lambda H}$ is also high (0.894–0.959), though slightly lower than driver persistence. The supply-wedge parameter τ_D is much larger for BTC (8.8) than for SPY (0.225), consistent with the higher driver supply elasticity in crypto markets.

6 Post-Estimation Analysis

Section 5 estimated the model's structural parameters. We now put these estimates to work: using the fitted model, we quantify the informational advantage that institutional investors hold over retail investors in each market and explore where that advantage comes from. In other words, we ask whether institutional investors remain smart money in BTC and SPY markets, and if so,

in what specific sense they are smart—which components of expected returns they are better at learning about, and through which margins of news consumption.

Our key metric is the *information resolution ratio*, the share of prior expected-return uncertainty that investor i's beliefs resolve:

$$\mathrm{IR}^{i,a} = 1 - \frac{\Sigma^{i,a}_{\mathrm{proj},ss}}{\Sigma^{a}_{\mu,\mathrm{proj}}}, \tag{19}$$

where $\Sigma^{a}_{\mu,\mathrm{proj}}$ is the unconditional variance of the expected-return forecast absent any signals, and $\Sigma^{i,a}_{\mathrm{proj},ss}$ is its posterior counterpart under investor i's steady-state beliefs. For each asset, the no-news benchmark computes $\mathrm{IR}^{i,a}$ when the investor still observes the return process and the public state variables embedded in the fitted model, but receives no headline or driver content signals. The news increment for investor type i is the difference between $\mathrm{IR}^{i,a}$ under observed steady-state consumption and under this no-news benchmark, and the *institutional information advantage* is the institutional news increment minus the retail news increment. We complement this information-based metric with the expected squared Sharpe ratio, which captures portfolio performance rather than information itself—a different dimension of the same underlying precision gain. Section OA.5 of the Online Appendix gives the exact formulas underlying both metrics and the headline/driver Shapley decomposition.

6.1 Magnitude of the Institutional Advantage

Table 6 reports the institutional information advantage in percentage points: a 14.3 ppt increment means that costly news consumption raises resolved expected-return uncertainty by 14.3 percentage points relative to the no-news benchmark.

The main finding is unambiguous: in both BTC and SPY, institutional investors resolve a strictly larger fraction of prior expected-return uncertainty through news consumption. For BTC, the institutional news gain is 14.3 ppt versus 6.7 ppt for retail, implying a 7.6 ppt institutional advantage. For SPY, the corresponding figures are 6.6 ppt versus 3.1 ppt, yielding a 3.5 ppt advantage.

The advantage is nearly twice as large in BTC as in SPY. This difference need not mean that institutions are inherently more skilled at processing BTC news than SPY news—the model measures informational outcomes, not skill, and offers no direct comparison of ability across assets.

Table 6: Institutional information advantage.
Panel A reports news increments: percentage-point (ppt) increases in information resolution relative to the no-news benchmark. The institutional advantage is the institutional minus retail news increment. Panel B reports the steady-state expected squared Sharpe ratio.

	BTC	SPY
Panel A: Information resolution (ppt increment)		
Institutional news increment	14.3	6.6
Retail news increment	6.7	3.1
Institutional advantage	7.6	3.5
Panel B: Expected squared Sharpe ratio		
Institutional	0.233	0.200
Retail	0.205	0.187
Institutional/retail ratio	1.137	1.067

A more natural reading, consistent with the motivation in the Introduction, is that public news consumption carries more of institutions' overall informational edge in BTC precisely because other channels of advantage—analyst coverage, corporate access, insider networks, and established valuation frameworks—are comparatively underdeveloped in crypto markets. Where these channels are richer, as in equities, the public-news channel we measure here is only one among several sources of institutional advantage, so it need not loom as large even if institutions are equally adept at processing news in both markets.

6.2 Source of the Institutional Informational Advantage

The previous subsection shows that institutional investors extract more information from news than retail investors, and that the gap is larger for BTC than for SPY. We now ask where this advantage comes from. One possibility is that institutions benefit because they learn about a uniquely valuable component of expected returns. Another is that they read news categories that contain more return-relevant information. A third is that they simply consume more of the news that matters.

We organize the answer with a three-channel decomposition. The news gain for investor type i in asset a can be read through

$$\Delta \mathrm{IR}^{i,a}_{\mathrm{news}} \approx \sum_{k\in\{H,D\}} \underbrace{c^{i,a}_{k}}_{\text{quantity}} \sum_{h\in\{S,L\}} \underbrace{\omega^{a}_{k,h}}_{\text{informativeness}} \underbrace{v^{a}_{h}}_{\text{portfolio value}}, \tag{20}$$

where k indexes headline and driver news and h indexes the transitory and persistent expected-return components: $c_k^{i,a}$ is average news consumption, $\omega_{k,h}^a$ summarizes how much news type k reveals about component h, and v_h^a measures how valuable learning about component h is for portfolio choice. Equation (20) is an economic accounting device rather than an exact identity of the nonlinear filtering problem; after describing the three channels, we report the exact model-implied headline/driver attribution using Shapley values.

Table 7 reports the three channels in three panels. Panel A is the *value channel*: which component of expected returns is most valuable to learn for portfolio choice? Panel B is the *informativeness channel*: which news type reveals each component? Panel C is the *quantity channel*: given which news types are useful, what do institutions and retail investors actually read?

Table 7: Three-channel decomposition of the institutional information advantage.
Panel A reports the share of portfolio-relevant expected-return uncertainty associated with the transitory and persistent components. Panel B repeats the within-component shares from Table 4. Panel C reports average steady-state news consumption for institutional and retail investors.

	BTC		SPY	
Panel A: Value channel — share of portfolio-relevant uncertainty				
	Transitory	Persistent	Transitory	Persistent
Uncertainty share	56.7%	43.3%	64.3%	35.7%
Panel B: Informativeness channel — which news type reveals each component				
	Headline	Driver	Headline	Driver
Transitory component	44.2%	55.8%	76.9%	23.1%
Persistent component	6.2%	93.8%	1.5%	98.5%
Panel C: Quantity channel — average news consumption				
	Inst.	Retail	Inst.	Retail
Headline news	0.074	0.093	0.157	0.172
Driver news	0.229	0.042	0.568	0.085

Value channel. Panel A of Table 7 asks what investors should want to learn. The answer is not obvious ex ante: if portfolio choice depended almost entirely on one component, the institutional advantage would reduce to a story about that component alone. The estimates imply neither extreme. For BTC, portfolio-relevant uncertainty is 56.7% transitory and 43.3% persistent; for SPY, the corresponding shares are 64.3% and 35.7%. The transitory component is somewhat larger, but the persistent component is economically important in both markets, so investors should care about learning both.

Informativeness channel. Because both components matter, the next question is how investors

can learn about each one. Panel B shows a sharp asymmetry. Transitory information can be learned from both news types, in proportions that differ across markets. Persistent information cannot: it comes almost entirely from driver news, which accounts for 93.8% of BTC persistent news value and 98.5% of SPY persistent news value. The type of news investors read therefore matters, because persistent expected-return information is available almost only through driver news.

Quantity channel. Given that news type matters, the final channel is what investors actually read. Panel C shows that institutional and retail investors consume similar amounts of headline news in both markets. The sharp difference is driver news: institutions consume roughly five times as much as retail investors in BTC (0.229 versus 0.042) and nearly seven times as much in SPY (0.568 versus 0.085). The quantity advantage is therefore concentrated precisely in the news category needed to learn about the persistent component.

The estimated attention costs in Table 5 explain why this quantity gap opens on the driver margin: retail investors face driver attention costs roughly 74 times higher than institutions in BTC and 28 times higher in SPY, while headline costs are comparable across investor types. This pattern is consistent with an expertise or processing-cost interpretation: institutions are not simply reading more of every type of news; they have a lower cost of processing the driver information that is essential for learning the persistent component.

The three channels thus point to driver news as the locus of the institutional advantage. Table 8 quantifies this attribution exactly in the full model, using Shapley values to split each investor type's total news gain between the headline and driver categories.

Table 8: Headline/driver decomposition of the news gain.
Entries are percentage-point (ppt) contributions to information resolution. Contributions are Shapley values and sum to the total news increment; percentages in parentheses are each news category's share of the total news gain for that investor type.

	Total	Headline	Driver
Panel A: BTC			
Institutional	14.3	1.1 (7.8%)	13.2 (92.2%)
Retail	6.7	1.7 (24.8%)	5.0 (75.2%)
Panel B: SPY			
Institutional	6.6	1.2 (18.4%)	5.4 (81.6%)
Retail	3.1	1.4 (43.4%)	1.8 (56.6%)

Driver news indeed accounts for the bulk of the institutional gain: 92.2% of the 14.3-ppt BTC

gain and 81.6% of the 6.6-ppt SPY gain. Retail investors also benefit from driver news, but headline news accounts for a much larger share of their total gain (24.8% for BTC retail and 43.4% for SPY retail).

Putting the three channels together, the information advantage does not arise from institutions reading more news of every kind. It comes from a particular alignment: driver news is the main route to persistent expected-return information, and institutions have the strongest consumption advantage exactly on that margin. The same mechanism operates in both markets; what makes the advantage larger in BTC is not a uniquely large driver-consumption ratio—institutional driver consumption is high relative to retail in both markets—but the combination of a larger persistent share of portfolio-relevant uncertainty (43.3% versus 35.7%), a much larger driver attention-cost gap, and a smaller retail offset from headline news. The root cause is a driver processing advantage, and it is especially consequential in BTC.

6.3 Attention-Cost Dynamics over the Market Cycle

Section 6.2 traces the institutional information advantage to a single critical factor: the low cost at which institutions process driver news. The estimates in Table 5 summarize the *average level* of these costs, but attention costs are persistent stochastic processes that move over time. This subsection examines that time variation: when, over the market cycle, are attention costs high or low, and do institutional and retail costs move together or apart?

Table 9 projects the filtered log attention-cost series on the standardized RSI (Relative Strength Index), a momentum-based indicator of market strength that we use as a proxy for market conditions. A negative coefficient means that high-RSI periods, when prices have recently been rising, are associated with lower attention costs and therefore higher attention appetite; a positive coefficient means the opposite.

Most loadings in Table 9 are negative: attention costs fall, and attention appetite therefore rises, when markets are strong. This procyclical pattern is the expected one—attention chases rising prices—and it describes headline attention in both markets as well as retail driver attention. The retail side is consistent with the return-dependent attention documented in behavioral finance, such as the "ostrich effect" of Sicherman et al. (2016), whereby investors avoid information after unfavorable outcomes.

Table 9: RSI loadings on latent attention cost.
OLS coefficient on standardized RSI in a projection of the filtered log attention cost $\log \lambda_{k,t}^{i,a}$ on RSI_t. p-values from heteroskedasticity-robust standard errors are in parentheses.

	BTC	SPY
Inst. headline cost	−0.094 (0.032)	−0.366 (<0.001)
Retail headline cost	−0.058 (0.192)	−0.260 (0.009)
Inst. driver cost	+0.192 (0.023)	−0.090 (0.159)
Retail driver cost	−0.179 (0.017)	−0.368 (<0.001)

The exception is the most important cell in the table: institutional driver cost in BTC carries the only positive loading (+0.192, $p = 0.023$). Institutional appetite for BTC driver news is *counter*-cyclical: institutions step up their processing of driver news precisely when the market cools, at the same time as retail investors pull back. This counter-cyclical attention falls exactly on the news category that carries persistent expected-return information—the category Section 6.2 identifies as the source of the institutional advantage—so the advantage widens in weak markets, when persistent expected-return information is most valuable. In SPY, by contrast, no counter-cyclical margin appears: institutional driver cost is unrelated to RSI (−0.090, $p = 0.159$), and attention dynamics are concentrated on the procyclical headline margin and less sharply differentiated across investor types.

Taken together, the time-series evidence completes the picture from Section 6.2: institutions not only face lower driver attention costs on average, they also time this margin differently, leaning into driver news when markets weaken. Section OA.6 of the Online Appendix reports the same post-estimation exercises over alternative content-noise calibrations; the qualitative patterns are unchanged.

7 Conclusion

In this paper, we develop and estimate a model of news consumption and portfolio choice to study how investors acquire and process information in cryptocurrency and equity markets. The model distinguishes between headline and driver news and allows investors to allocate costly attention across these news categories when forming beliefs about expected returns.

We document four findings. First, both persistent and transitory expected-return components are economically important, suggesting that investors benefit from learning about both. Second,

driver news is the primary source of information about persistent expected returns, whereas both headline and driver news contain information about transitory fluctuations. Third, institutional investors derive substantially greater benefits from news consumption than retail investors because they devote more attention to driver news and face lower information-processing costs. This informational advantage is larger in crypto than in the equity market, where opportunities to acquire private information are greater. Fourth, this institutional advantage is state dependent. For BTC, institutions increase their relative attention to driver news during weaker market conditions, whereas retail investors do not, making the institutional advantage especially pronounced when persistent-return information is most valuable. Overall, these findings suggest that institutional investors retain their traditional informational advantage even in a market with limited traditional valuation anchors and few opportunities to acquire private information.

More broadly, our findings suggest that investors differ substantially in their ability to identify, process, and synthesize value-relevant public information. In crypto markets, where traditional private information channels are relatively limited, such differences in public information processing appear to be an especially important source of informational advantage. Our analysis does not rule out the possibility that public news prompts investors to collect additional information, and institutional investors may also enjoy advantages in accessing or processing such information. Exploring the interaction between public information-processing and private information acquisition, particularly across different types of digital assets and market environments, is a promising direction for future research.

Bibliography

Aiello, D., Baker, S. R., Balyuk, T., Di Maggio, M., Johnson, M. J., Kotter, J. D., 2023. Who invests in crypto? Wealth, financial constraints, and risk attitudes. Unpublished Manuscript .

Anderson, C. M., Fang, V. W., Moon Jr, J. R., Shipman, J. E., 2026. Accounting for cryptocurrencies. Journal of Accounting Research 64, 45–79.

Auer, R., Cornelli, G., Doerr, S., Frost, J., Gambacorta, L., 2025. Crypto trading and bitcoin prices: Evidence from a new database of retail adoption. IMF Economic Review pp. 1–36.

Baba Yara, F., Grigoris, F., Kantak, P., 2024. Risk from the inside out: understanding firm risk through employee news consumption. Kelley School of Business Research Paper Forthcoming .

Barber, B. M., Odean, T., 2013. The behavior of individual investors. In: *Handbook of the Economics of Finance*, Elsevier, vol. 2, pp. 1533–1570.

Ben-Rephael, A., Da, Z., Israelsen, R. D., 2017. It depends on where you search: Institutional investor attention and underreaction to news. Review of Financial Studies 30, 3009–3047.

Benetton, M., Compiani, G., 2020. Investors' beliefs and asset prices: A structural model of cryptocurrency demand. Unpublished Manuscript .

Benigno, P., Schilling, L. M., Uhlig, H., 2022. Cryptocurrencies, currency competition, and the impossible trinity. Journal of International Economics 136, 103601.

Bernile, G., Hu, J., Tang, Y., 2016. Can information be locked up? Informed trading ahead of macro-news announcements. Journal of Financial Economics 121, 496–520.

Bertomeu, J., Martin, X., Zhang, S., 2023. Uncle Sam's stimulus and crypto boom. Unpublished Manuscript .

Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., Menkveld, A. J., 2023. Equilibrium bitcoin pricing. Journal of Finance 78, 967–1014.

Bradley, D., Jame, R., Williams, J., 2022. Non-deal roadshows, informed trading, and analyst conflicts of interest. Journal of Finance 77, 265–315.

Brown, G. W., Gredil, O. R., Kantak, P., 2023. Finding fortune: How do institutional investors pick asset managers? Review of Financial Studies 36, 3071–3121.

Christiano, L. J., Fitzgerald, T. J., 2003. The band pass filter. International Economic Review 44, 435–465.

Cohen, R. B., Gompers, P. A., Vuolteenaho, T., 2002. Who underreacts to cash-flow news? Evidence from trading between individuals and institutions. Journal of Financial Economics 66, 409–462.

Cong, L. W., Li, Y., Wang, N., 2021. Tokenomics: Dynamic adoption and valuation. Review of Financial Studies 34, 1105–1155.

Da, Z., Engelberg, J., Gao, P., 2011. In search of attention. Journal of Finance 66, 1461–1499.

Didisheim, A., Somoza, L., 2024. The end of the crypto-diversification myth. Unpublished Manuscript .

Divakaruni, A., Zimmerman, P., 2024. Uncovering retail trading in bitcoin: The impact of covid-19 stimulus checks. Management Science 70, 2066–2085.

Duchin, R., Solomon, D. H., Tu, J., Wang, X., 2025. The cryptocurrency elephant in the room. Review of Finance 29, 1721–1767.

Fidelity Digital Assets, 2021. Institutional investor digital assets study. Available at `https://www.fidelitydigitalassets.com/sites/g/files/djuvja3256/files/acquiadam/993307.2.0_2021%20Institutional%20Investor%20Digital%20Assets%20Study%20Report.pdf`.

Fu, X. J., Taylor, L. A., 2025. Due diligence and the allocation of venture capital. Unpublished Manuscript .

Gompers, P. A., Gornall, W., Kaplan, S. N., Strebulaev, I. A., 2020. How do venture capitalists make decisions? Journal of Financial Economics 135, 169–190.

Hackethal, A., Hanspal, T., Lammer, D. M., Rink, K., 2022. The characteristics and portfolio behavior of bitcoin investors: Evidence from indirect cryptocurrency investments. Review of Finance 26, 855–898.

Ivashina, V., Lerner, J., 2018. Looking for alternatives: Pension investments around the world, 2008 to 2017. Unpublished Manuscript .

Kaplan, S. N., Strömberg, P., 2001. Venture capitalists as principals: Contracting, screening, and monitoring. American Economic Review 91, 426–430.

Kaplan, S. N., Strömberg, P. E., 2004. Characteristics, contracts, and actions: Evidence from venture capitalist analyses. Journal of Finance 59, 2177–2210.

Kogan, S., Makarov, I., Niessner, M., Schoar, A., 2024. Are cryptos different? Evidence from retail trading. Journal of Financial Economics 159, 103897.

Kwan, A., Liu, Y., Matthies, B., 2026. Institutional investor attention. Journal of Finance 81, 791–827.

Lerner, J., Mao, J., Schoar, A., Zhang, N. R., 2022. Investing outside the box: Evidence from alternative vehicles in private equity. Journal of Financial Economics 143, 359–380.

Lerner, J., Schoar, A., Wang, J., 2008. Secrets of the academy: The drivers of university endowment success. Journal of Economic Perspectives 22, 207–222.

Lintner, J., 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. Review of Economics and Statistics 47, 13–37.

Liu, Y., Tsyvinski, A., 2021. Risks and returns of cryptocurrency. Review of Financial Studies 34, 2689–2727.

Mossin, J., 1966. Equilibrium in a capital asset market. Econometrica 34, 768–783.

Pagnotta, E. S., 2022. Decentralizing money: Bitcoin prices and blockchain security. Review of Financial Studies 35, 866–907.

Pursiainen, V., Toczynski, J., 2022. Retail investors' cryptocurrency investments. Unpublished Manuscript .

Sharpe, W. F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. Journal of Finance 19, 425–442.

Sicherman, N., Loewenstein, G., Seppi, D. J., Utkus, S. P., 2016. Financial attention. Review of Financial Studies 29, 863–897.

Solomon, D., Soltes, E., 2015. What are we meeting for? The consequences of private meetings with investors. Journal of Law and Economics 58, 325–355.

Table A1: Vendor-supplied labels used to identify BTC-related news

This appendix summarizes the vendor-supplied labels used to identify BTC-related news and whether the associated articles are classified as headline or driver news. For ease of presentation, we report both labels associated with headline news and five representative labels associated with driver news within each of four broad categories. These categories are used for illustration only, may not be mutually exclusive, and are not part of our empirical classification.

News Category	Label Category	Representative Label: Vendor Description
Headline	N/A	Bitcoin (BTC): A cryptocurrency and a payment system invented by an unidentified programmer(s).
		Cryptocurrency: A digital asset designed to work as a medium of exchange using cryptography to secure the transactions and to control the creation of additional units of the currency.
Driver	**Ecosystem: Technology** (e.g., blockchain, cryptography, and consensus)	Blockchain: A digital ledger in which transactions made in Bitcoin or another cryptocurrency are recorded chronologically and publicly.
		Distributed Ledgers: A consensus of replicated, shared, and synchronized digital data geographically spread across multiple sites, countries, or institutions.
		Private Keys: A cryptographic key that can be obtained and used by anyone to encrypt messages intended for a particular recipient, such that the encrypted messages can be deciphered only by using a second key that is known only to the recipient.
		SHA-256 (Secure Hash Algorithm): A set of cryptographic hash functions designed by the United States National Security Agency.
		Strong Encryption: An encryption method that uses a very large number as its cryptographic key.
Driver	**Ecosystem: Infrastructure** (e.g., tokens, protocols, and exchanges)	Ethereum (ETH): A decentralized software platform that enables smart contracts and distributed applications to be built and run without any downtime, fraud, control or interference from a third party.
		Ripple (XRP): A real-time gross settlement system, currency exchange and remittance network.
		Initial Coin Offering (ICO): A type of crowdfunding that uses cryptocurrencies as a means of raising capital and has been associated with scams and securities law violations.
		Coinbase: A digital asset exchange company headquartered in San Francisco, California.
		Interoperability: The ability of computer systems or software to exchange and make use of information.
Driver	**Ecosystem: Application** (e.g., blockchain applications and enterprise solutions)	Smart Contracts: Self-executing contracts with the terms of the agreement between buyer and seller being directly written into lines of code.
		Tokenization: The process of replacing sensitive data with unique identification symbols that retain all the essential information about the data without compromising its security.
		Hyperledger: An umbrella project of open source blockchains and related tools, started in December 2015 by the Linux Foundation and supported by big industry players like IBM, Intel and SAP to support the collaborative development of blockchain-based distributed ledgers.
		IBM Blockchain: A shared, immutable ledger for recording the history of transactions.
		Oracle Blockchain Service: A network consisting of validating nodes (peers) that update the ledger and respond to queries by executing smart contract code the business logic that runs on the blockchain.
Driver	**Regulation**	Anti-Money Laundering: The set of policies, regulations, and laws designed to prevent the generation of income through illegal activities.
		Know Your Business (KYB): Know Your Business (KYB) verification is a company's Anti-Money Laundering compliance. Companies must protect their interests before doing business with another business. Companies need to know if their income is misused by corrupt business owners, shareholders, and money launderers.

News Category	Label Category	Representative Label: Vendor Description
		Know Your Customer (KYC): The identification of a business's target clients and customers.
		Fourth EU Anti-Money Laundering Directive: EU Government legislation to combat disruptive money laundering and terrorist funding.
		RegTech: Regulatory technology, in short regtech, is a new technology that uses information technology to enhance regulatory processes. With its main application in the Financial sector, it is expanding into any regulated business with a particular appeal for the Consumer Goods Industry.

Table A2: Vendor-supplied labels used to identify SPY-related news

This appendix summarizes the vendor-supplied labels used to identify SPY-related news and whether the associated articles are classified as headline or driver news. For ease of presentation, we report labels associated with headline news and representative labels associated with driver news within each of five broad categories. These categories are used for illustration only, may not be mutually exclusive, and are not part of our empirical classification.

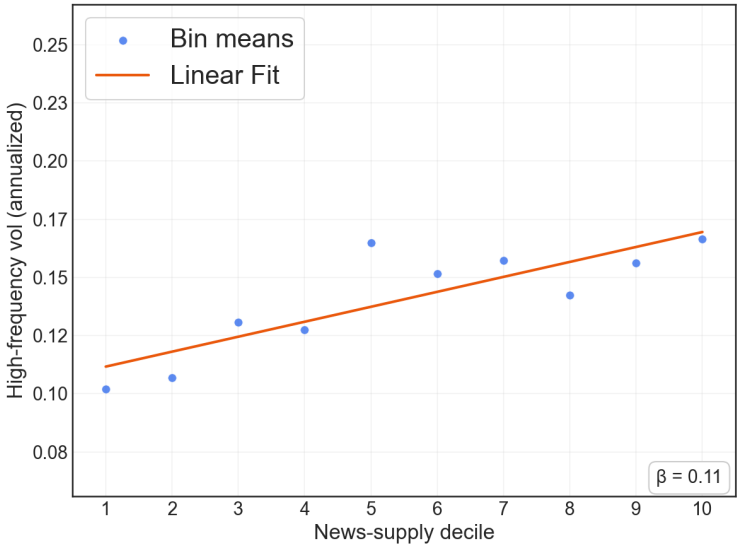
News Category	Label Category	Representative Label: Vendor Description
Headline	N/A	Stock Market Index: A measurement of the value of a section of the stock market.
		S&P 500 Index (SPX): An American stock market index based on the market capitalizations of 500 large companies having common stock listed on the NYSE or NASDAQ.
		Nasdaq: An American stock exchange that operates in New York City.
		Dow Jones Industrial Average (DJIA): A stock market index that tracks 30 large, publicly-owned companies trading on the New York Stock Exchange and the NASDAQ.
Driver	**Economic Activity** (e.g., economic output, labor markets, and production)	Gross Domestic Product (GDP): Valuation of goods produced and services provided by a country during one year.
		Bureau of Labor Statistics (BLS): The Bureau of Labor Statistics (BLS) is a unit of the United States Department of Labor. It is the principal fact-finding agency for the U.S. government in the broad field of labor economics and statistics and serves as a principal agency of the U.S. Federal Statistical System.
		Unemployment Insurance: Unemployment benefits (depending on the jurisdiction also called unemployment insurance or unemployment compensation) are payments made by the state or other authorized bodies to unemployed people. In the United States, benefits are funded by a compulsory governmental insurance system, not taxes on individual citizens.
		Supply Chain: Sequence of processes involved in the production and distribution of a commodity.
Driver	**Monetary Policy** (e.g., central banks and interest rates)	Federal Open Market Committee (FOMC): A committee within the Federal Reserve System (the Fed) is charged under the United States law with overseeing the nation's open market operations (e.g., the Fed's buying and selling of United States Treasury securities).
		Bond Yield: The amount of return an investor realizes on a bond.
		Treasury Notes: Marketable debt securities issued by the U.S. Department of the Treasury.
		LIBOR: A benchmark rate that some of the world's leading banks charge each other for short-term loans.
		Interest Rate Risk: The risk that an investment's value will change due to a change in the absolute level of interest rates, in the spread between two rates, in the shape of the yield curve, or in any other interest rate relationship.
Driver	**Inflation and Commodity Markets** (e.g., inflation, inflation expectations, and energy prices)	Consumer Price Index (CPI): An index of the variation in prices paid by typical consumers for retail goods and other items.
		Inflation: General increase in prices and the decrease in the buying power of money.
		Treasury Inflation Protected Securities (TIPS): Treasury inflation-protected securities (TIPS) refer to a treasury security that is indexed to inflation in order to protect investors from the negative effects of inflation.
		Gold Spot: Gold spot price is used as the basis for most bullion dealers to determine the exact price to charge for a specific coin bar.
		Brent Crude: A major trading classification of sweet light crude oil that serves as a major benchmark price for purchases of oil worldwide.

News Category	Label Category	Representative Label: Vendor Description
Driver	**Financial Conditions and Risk Premia** (e.g., credit conditions and market risk)	Credit Risk: The probable risk of loss resulting from a borrower's failure to repay a loan or meet contractual obligations.
		Liquidity Risk: A financial risk that for a certain period of time a given financial asset, security or commodity cannot be traded quickly enough in the market without impacting the market price.
		Commercial Lending: Process of funding an organization for short-term expenditures or capital equipment acquisitions.
		High Yield Bonds: A high paying bond with a lower credit rating than investment-grade corporate bonds, treasury bonds, and municipal bonds, with a high risk of default.
		Market Volatility Index (VIX): A popular measure of the implied volatility of S&P 500 index options, calculated by the Chicago Board Options Exchange (CBOE).
Driver	**Global Economy and Trade** (e.g., foreign exchange, global production, and international trade)	Exchange Rate: The exchange rate is the value of one currency for the purpose of conversion to another.
		Currency Risk: The possibility that currency depreciation will negatively affect the value of one's assets, investments, and their related interest and dividend payment streams, especially those securities denominated in foreign currency.
		Global Supply Chain: International network of companies participating in a dynamic supply and demand system.
		Free Trade: Free trade is international trade left to its natural course without tariffs, quotas, or other restrictions.
		Barrier to Trade: A barrier to trade is a government-imposed restraint on the flow of international goods or services.

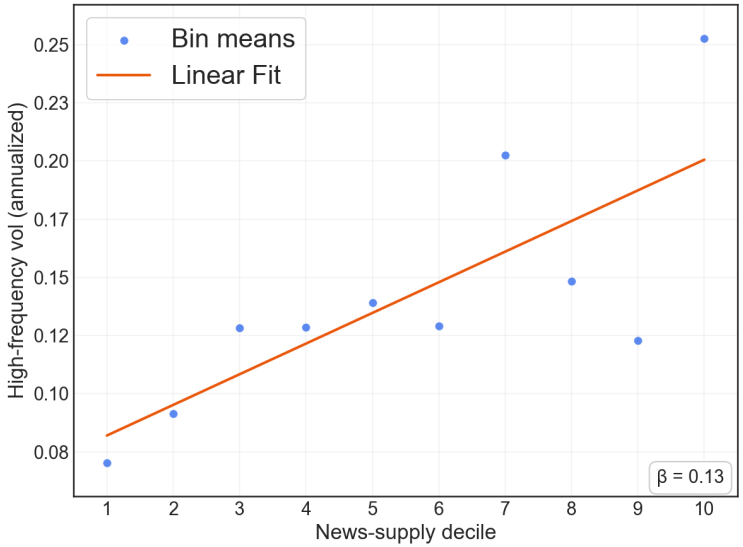
Figure A1: SPY News Supply

This figure plots the association of news supply to SPY volatility. News supply comes in two categories—headline and driver news. Likewise, SPY volatility is decomposed into a low and high-frequency component using a Christiano and Fitzgerald (2003) band-pass filter.

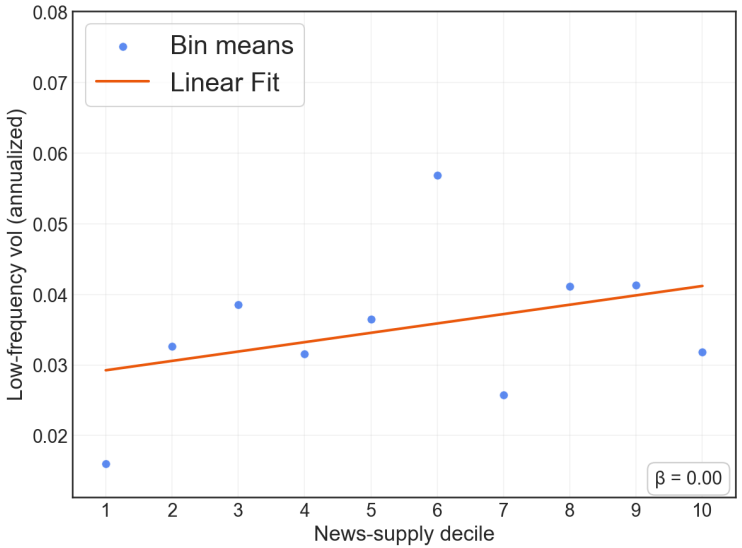
(a) *Headline news supply and SPY high-frequency volatility*



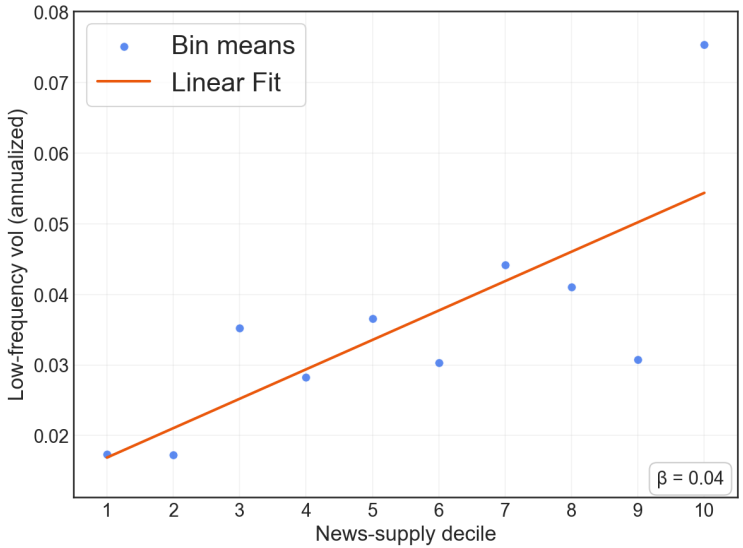
(b) *Driver news supply and SPY high-frequency volatility*



(c) *Headline news supply and SPY low-frequency volatility*



(d) *Driver news supply and SPY low-frequency volatility*



A Notation

Table A3 summarizes the notation of the model in Section 3. Full technical details—the separability of the two-asset problem, the belief-updating recursions, the attention-cost specification, the four-step estimation procedure, and the post-estimation measurement—are provided in the Online Appendix (Sections OA.1–OA.6).

Table A3: Model notation.

Group	Symbol	Description
Indices	$a \in \{BTC, SPY\}$	Risky asset
	$i \in \{\text{inst}, \text{retail}\}$	Investor type
	$k \in \{H, D\}$	News category (headline, driver)
	$j \in \{L, S\}$	Expected-return component (long-run, short-run)
	t	Trading day
Returns	r_t^a, $\tilde{r}_t^a$	Realized return; centered return $r_t^a - \bar{\mu}^a$
	$\bar{\mu}^a$	Unconditional average return
	$\mu_t^{a,L}$, $\mu_t^{a,S}$	Persistent and transitory expected-return components
	ρ_L^a, ρ_S^a	AR(1) persistence of the two components; $\boldsymbol{\rho}^a = (\rho_L^a, \rho_S^a)'$
	ε_t^a, $\sigma_{\varepsilon,a}^2$	Return noise and its variance
Information events	$V_t^{a,j}$	Innovation variance of component j (size of information events)
	$\mu_{V,a,j}$, $\phi_{V,a,j}$, $\sigma_{\omega,a,j}^2$	Mean, persistence, and innovation variance of $\log V_t^{a,j}$
	$\hat{V}_t^{a,j}$	Publicly filtered variance state
News	B^a, $b_j^{a,k}$	News-loading matrix and its elements
	$U_{k,t}^a$	News value of category-k articles
	$n_{k,t}^a$	Quantity of published category-k articles
	μ_k^a, $u_{k,t}^a$, $\sigma_{u,a,k}^2$	Supply intercept, supply shock, and its variance
	$\tilde{y}_{k,t}^a$, $\tilde{\nu}_{k,t}^a$	Article-level content signal and its noise
	$y_{k,t}^{i,a}$	Aggregated content signal received by investor i
	$\sigma_{\nu,a}$	Baseline content-noise level
Choices	$c_{k,t}^{i,a}$	Number of category-k articles processed
	$\theta_t^{i,a}$	Dollar position in asset a
Beliefs	$\hat{\boldsymbol{\mu}}_{t\|t}^{i,a}$, $\Sigma_{t\|t}^{\mu,i,a}$	Posterior mean and covariance of the expected-return components
	$m_t^{i,a}$	Posterior expected next-period return
	$\Sigma_{\text{total},t}^{i,a}$	Total conditional variance of the next-period return
Costs and preferences	$\lambda_{k,t}^{i,a}$, $\bar{\lambda}_k^{i,a}$	Attention-cost scale and its long-run mean
	$\rho_{\lambda,k}^{i,a}$, $\sigma_{\eta,i,a,k}^2$	Persistence and innovation variance of $\log \lambda_{k,t}^{i,a}$
	τ_k^a, $q_{k,t}(\cdot)$	Congestion-adjustment scale and congestion function
	$\mathcal{C}_t^{i,a}$	Total attention cost
	γ, β	Risk aversion and discount factor
	$\Pi_t^{i,a}$, $J^{i,a}$	Certainty-equivalent payoff and value function