When Does Bad News Stick? Topic-Level Asymmetry in Narrative Dynamics

Isha Agarwal* Wentong Chen† Eswar Prasad‡

May 22, 2026

Abstract

When does bad news stick? This paper shows that the persistence of media sentiment depends not only on whether news is good or bad, but also on what the news is about. Using the complete Wall Street Journal archive from 1990 to 2025, we construct monthly positive and negative sentiment indices for twelve economic topics and estimate topic-specific sentiment dynamics. Bad-news innovations generate larger cumulative responses than positive-news innovations in every topic, indicating that negativity asymmetry is a broad feature of media sentiment. But the size of this asymmetry varies sharply across topics. The bad-minus-good persistence gap is largest in security and human rights, health and public safety, real estate and financial stability, and firms and corporate strategy. Bad-news persistence is also sizable in data-rich topics such as financial markets, real estate, and corporate news, suggesting that frequent information arrival does not necessarily lead narratives to revert quickly. These findings show that aggregate sentiment indices mask substantial topic-level heterogeneity in the asymmetric dynamics of narratives, with potential implications for how persistent negative sentiment shapes expectations, asset prices, and real economic outcomes.

JEL Classification Numbers: G14, D83, D84, C55
Keywords: media narratives, textual analysis, natural language processing, bad news persistence, topic asymmetry, media sentiment, expectations.

*Sauder School of Business, University of British Columbia, E-mail: isha.agarwal@sauder.ubc.ca

†Department of Economics, Santa Clara University, Santa Clara, CA, 95053, U.S.A, E-mail: wchen10@scu.edu

‡Dyson School of Applied Economics and Management and Department of Economics, Cornell University, Ithaca, NY, 14853, U.S.A, Brookings Institution, and NBER, E-mail: eswar.prasad@cornell.edu

1 Introduction

News narratives shape how economic agents process information and form expectations. However, the economic effects of media sentiment depend not only on its tone, but also on its persistence. A temporary burst of negative sentiment and a persistently negative narrative can have very different implications for expectation formation. Yet most text-based measures in finance and macroeconomics summarize the news environment with an aggregate sentiment index. These indices are useful, but they hide an important question: does bad news stick differently depending on what the news is about?

This question matters because not all bad news has the same narrative structure. Negative news about financial stability, public health, geopolitics, or human rights may generate follow-up coverage, continuing uncertainty, and repeated reinterpretation. Other types of news may be resolved more quickly. An aggregate sentiment index treats these topics as part of the same object. In doing so, it can mix contemporaneous tone with topic composition and with the persistence of the underlying narrative.

In this paper, we open up the aggregate sentiment index and study the topic-level dynamics of positive and negative news. We use the complete Wall Street Journal archive from 1990 to 2025, covering more than 1.3 million articles. We classify articles into twelve economic topics and construct monthly positive and negative sentiment indices for each topic. Because our analysis is based on the Wall Street Journal, the measures capture U.S.-centered macro-financial narratives as reflected in a major U.S. financial newspaper.

We then estimate topic-specific sentiment dynamics using local projections (Jordà, 2005). The empirical exercise proceeds in two steps. First, we construct positive- and negative-sentiment innovations by removing predictable variation in sentiment using recent topic-level sentiment, aggregate macro-financial conditions, topic-level news intensity, and aggregate sentiment.[1] The resulting shocks capture unexpected changes

[1] In this paper, we use "sentiment innovation", "sentiment shock" and "news shock" interchangeably.

in topic-level sentiment after conditioning on observable drivers of news intensity and aggregate conditions. Critically, this conditioning also absorbs asymmetries in underlying economic fundamentals, such as the well-documented pattern that recessions tend to be sharper and more abrupt than expansions (e.g., Jensen, Petrella, Ravn and Santoro, 2020, Dupraz, Nakamura and Steinsson, 2025), as well as topic-specific variation in the arrival and intensity of underlying events, helping separate that our persistence estimates reflect narrative dynamics rather than the mechanical asymmetry in the persistence of the events themselves. Second, we trace how positive and negative sentiment evolve after positive- and bad-news innovations of comparable size.

Three facts emerge. First, bad-news innovations generate larger cumulative post-impact responses than positive-news innovations in every topic. Negativity asymmetry is therefore not confined to a few crisis-related categories; it is a broad feature of media sentiment dynamics.

Second, the size of this asymmetry varies substantially across topics. The bad-minus-good persistence gap is largest in security and human rights, health and public safety, real estate and financial stability, and firms and corporate strategy. These are areas in which negative news often opens a continuing narrative rather than closing a single episode.

Third, and perhaps most surprisingly, bad-news persistence is sizable even in relatively data-rich topics such as financial markets, real estate, and corporate news. This is useful because one might expect frequent information arrival to accelerate mean reversion. Instead, the evidence suggests that more data does not necessarily mean faster narrative closure. New information can also sustain a narrative by generating follow-up coverage, renewed interpretation, or uncertainty about broader implications.

One possible interpretation is that persistence depends less on the amount of data available than on whether a news shock opens or closes a continuing narrative process. In relatively data-scarce topics such as geopolitics and security, bad-news persistence may reflect the gradual resolution of uncertainty surrounding slow-moving underlying events and fundamentals, while positive news is treated as partial or temporary, generating a large negative-positive persistence gap. In relatively data-rich topics such as financial markets and corporate news, by contrast, both bad and good news may initiate continuing narratives through cascading uncertainty or follow-up coverage, so that bad-news persistence remains high but the asymmetry between negative and positive dynamics is less pronounced.

The paper contributes to three strands of literature. The first is the media's impact on financial decisions and asset pricing (e.g., Tetlock, 2007; Engelberg and Parsons, 2011). This literature treats sentiment as a scalar aggregate. We show that doing so conflates topics with very different propagation structures. The second is the literature on narrative economics, heterogeneous beliefs, and disagreement (e.g., Sims, 2003; Shiller, 2017, 2020; Cookson and Niessner, 2020; Cookson, Engelberg and Mullins, 2020, 2023; Agarwal, Chen and Prasad, 2026). We provide systematic evidence that narrative salience, measured by post-impact persistence, is itself topic-dependent. The third is the literature on asymmetric news coverage and negativity bias. Soroka and McAdams (2015) and the related "man bites dog" literature (Nimark, 2014) document that negative and unusual events receive disproportionate media attention. We complement this by showing that the asymmetry extends beyond coverage intensity to narrative persistence: bad news not only attracts more attention, it also persists longer, and does so differentially across topics.

This distinction matters for the interpretation of text-based measures in finance and macroeconomics. A rise in aggregate negativity may reflect a more negative tone within topics, a shift in attention toward topics where negative narratives are especially persistent, or both. Treating sentiment as a single aggregate index can therefore conflate contemporaneous tone, topic composition, and dynamic propagation. Our results suggest that separating positive from negative sentiment and conditioning on topic composition can sharpen the measurement of narrative dynamics and improve the identification of how news affects expectations and economic outcomes.

The paper proceeds as follows. Section 2 describes the data, topic classification, and sentiment measures. Section 3 presents descriptive evidence on aggregate and topic-level sentiment. Section 4 estimates topic-specific local projections and documents asymmetry in the persistence of positive and negative sentiment shocks. Section 5 concludes.

2 Data, Topic Classification, and Sentiment Measurement

We use the complete Wall Street Journal archive from 1990 to 2025, comprising 1,321,905 articles. Sentiment is measured at the article level using ProsusAI/FinBERT, a transformer-based classifier trained on financial text, which assigns each article a label (positive, neutral, or negative) and a confidence score.[2] We set label values of +1, 0, and −1 for positive, neutral, and negative articles respectively, and define the article-level sentiment score as the product of the label value and the confidence score.

Aggregate Index. We construct the aggregate sentiment index (sen_t) as the average sentiment score across all articles in month t. We also construct a positive sentiment index (pos_t) and a negative sentiment index (neg_t), defined as the sum of sentiment scores for positively and negatively labeled articles, respectively, each normalized by the total number of articles in that month. These measures jointly capture both the prevalence and intensity of positive and negative narratives.

Topic-level Index. We classify each article into one of $K = 12$ topics using a large language model (`gpt-4.1-mini`) applied to article titles: (1) Trade & Supply Chain; (2) Domestic Economy & Growth; (3) Financial Markets, Banking & FX; (4) Real Estate, Debt & Financial Stability; (5) Technology, Innovation & Industry Policy; (6) Governance, Regulation & Institutions; (7) Environment, Climate & Energy; (8) Diplomacy & Geopolitics; (9)

[2]Classification is applied to the headline and first 1,500 characters of each article, reflecting both the input-length constraints of BERT-based models and the concentration of sentiment-relevant content near the top of news articles.

Society, Labor & Demographics; (10) Security, Surveillance & Human Rights; (11) Health, Pandemics & Public Safety; and (12) Firms, Industries & Corporate Strategy.[3] We then construct topic-specific sentiment indices (sen_t^k, pos_t^k, and neg_t^k) by restricting attention to articles within each topic k.

3 Stylized Facts

Time Series Patterns. Figure 1 plots the monthly aggregate, positive, and negative sentiment indices. Aggregate sentiment is persistently negative throughout the sample, with pronounced deteriorations around the early 1990s recession, the dot-com collapse, the Global Financial Crisis, and the COVID-19 pandemic. These fluctuations are driven primarily by negative sentiment, which spikes sharply and remains elevated around each episode.[4] Positive sentiment also varies over time and comoves with major macroeconomic episodes.

Cross-sectional Patterns. Figure 2 reports time-series averages of sentiment indices across the twelve topics. Negative sentiment is particularly pronounced in Trade & Supply Chain, Domestic Economy & Growth, Security, Surveillance & Human Rights, Financial Markets, Banking & FX, and Diplomacy & Geopolitics. Positive sentiment is smaller in magnitude than negative sentiment, but is not uniform across topics: it is highest in Trade & Supply Chain and Firms, Industries & Corporate Strategy, with relatively high values also in Domestic Economy & Growth and Financial Markets, Banking & FX. These patterns suggest that cross-sectional variation in sentiment is substantial, with average negativity concentrated in a subset of topics rather than evenly distributed across the news space.

[3] Title-based classification substantially reduces computational cost. Validation on random subsamples confirms that title-based assignments closely match full-text classifications.

[4] Note that vertical scales differ across panels.

4 Topic-Level Asymmetry in Narrative Persistence

4.1 Methodology

Our empirical approach estimates topic-specific local projections of sentiment on conditional sentiment shocks using a two-stage method.

Stage 1: Construction of Topic-Specific Sentiment Shocks. We study the persistence of *unexpected* changes in topic-level sentiment. The key identification challenge is that raw negative sentiment may be persistent either because narratives themselves persist or because the underlying events do. To separate these, we construct topic-specific sentiment shocks by removing predictable variation from three sources: recent sentiment dynamics, aggregate macro-financial conditions, and topic-level news intensity.

For notational clarity, define

$$bad_{k,t} \equiv -neg_{k,t},$$

so that larger values of $bad_{k,t}$ correspond to more negative sentiment.

For each topic k, we estimate

$$bad_{k,t} = \alpha_k^- + \sum_{\ell=1}^{p} \pi_{k,\ell}^- bad_{k,t-\ell} + \sum_{\ell=1}^{p} \theta_{k,\ell}^- pos_{k,t-\ell} + \sum_{\ell=1}^{p} \Phi_{k,\ell}^{-'} X_{t-\ell} + \sum_{\ell=1}^{p} \Psi_{k,\ell}^{-'} Q_{k,t-\ell} + u_{k,t}^-,$$

and

$$pos_{k,t} = \alpha_k^+ + \sum_{\ell=1}^{p} \pi_{k,\ell}^+ pos_{k,t-\ell} + \sum_{\ell=1}^{p} \theta_{k,\ell}^+ bad_{k,t-\ell} + \sum_{\ell=1}^{p} \Phi_{k,\ell}^{+'} X_{t-\ell} + \sum_{\ell=1}^{p} \Psi_{k,\ell}^{+'} Q_{k,t-\ell} + u_{k,t}^+,$$

where X_t contains the total U.S. stock market return, the VIX, industrial production growth, and the change in the two-year Treasury yield, and $Q_{k,t}$ contains the topic's article share and leave-one-topic-out aggregate sentiment.

We set a parsimonious lag length of $p = 3$ and use heteroskedasticity- and autocorrelationconsistent standard errors throughout. Residuals are standardized within topic to yield one-standard-deviation shocks $\varepsilon^{+}_{k,t}$ and $\varepsilon^{-}_{k,t}$. These are not raw news shocks but conditional sentiment shocks that absorb both macro-financial conditions and topic-specific news intensity.

Stage 2: Main Dynamic Test. We estimate topic-specific local projections to trace the propagation of these conditional sentiment shocks. For each horizon $h = 0, 1, \ldots, H$, we estimate

$$pos_{k,t+h} = \alpha^{+}_{k,h} + \beta^{++}_{k,h}\varepsilon^{+}_{k,t} + \beta^{+-}_{k,h}\varepsilon^{-}_{k,t} + \Gamma^{+'}_{k,h}W_{k,t-1} + \eta^{+}_{k,t+h},$$

and

$$bad_{k,t+h} = \alpha^{-}_{k,h} + \beta^{-+}_{k,h}\varepsilon^{+}_{k,t} + \beta^{--}_{k,h}\varepsilon^{-}_{k,t} + \Gamma^{-'}_{k,h}W_{k,t-1} + \eta^{-}_{k,t+h},$$

where $W_{k,t-1}$ contains the same controls used in the shock construction step. Including both positive and bad news shocks in each projection ensures that the estimated response to one type of sentiment shock is identified conditional on contemporaneous shocks of the opposite sign.

Our main objects of interest are the own-sentiment impulse responses:

$$\left\{\beta^{++}_{k,h}\right\}_{h=0}^{H} \quad \text{and} \quad \left\{\beta^{--}_{k,h}\right\}_{h=0}^{H}.$$

The first traces the response of positive sentiment to a positive sentiment innovation, while the second traces the response of bad news sentiment to a bad news sentiment innovation.

We summarize persistence by computing cumulative effects beyond the impact month:

$$CP^{+}_{k}(H) = \sum_{h=1}^{H}\hat{\beta}^{++}_{k,h}, \qquad CP^{-}_{k}(H) = \sum_{h=1}^{H}\hat{\beta}^{--}_{k,h}.$$

We further summarize asymmetry using the difference in cumulative responses:

$$ASYM_k^{CP}(H) = CP_k^-(H) - CP_k^+(H).$$

A positive value of $ASYM_k^{CP}(H)$ indicates that, following one-standard-deviation shocks, bad news shocks generate larger cumulative propagation than positive-news shocks within topic k.

4.2 Results

Bad news persistence is pervasive. Figure 3 shows that bad news shocks generate substantially larger cumulative post-impact responses than positive shocks across all topics. This pattern is broad-based: in every topic, $CP_k^-(12)$ exceeds $CP_k^+(12)$, indicating that negative sentiment shocks propagate more strongly over time than positive sentiment shocks of comparable size. The largest bad news persistence appears in Real Estate, Debt & Financial Stability; Financial Markets, Banking & FX; Firms, Industries & Corporate Strategy; Domestic Economy & Growth; and Health, Pandemics & Public Safety. Positive shocks, by contrast, exhibit weaker cumulative propagation and are close to zero in several topics, most notably Governance, Regulation & Institutions and Security, Surveillance & Human Rights. These results provide direct evidence that within-topic sentiment persistence is asymmetric, and that this asymmetry is broad-based rather than confined to any single topic.

Within-topic asymmetry is heterogeneous. Figure 4 ranks topics by the 12-month bad news-minus-positive cumulative persistence statistic, $ASYM_k^{CP}(12)$. The ranking shows that within-topic asymmetry is positive for every topic: bad news sentiment shocks generate larger cumulative post-impact propagation than positive sentiment shocks of comparable size. The largest gaps appear in Security, Surveillance & Human Rights; Health, Pandemics & Public Safety; Real Estate, Debt & Financial Stability; and Firms, Industries & Corporate Strategy. These are topics in which adverse events often generate follow-up coverage, uncertainty, and repeated interpretation. By contrast, Technology, Innovation & Industry Policy; Domestic Economy & Growth; and Governance, Regulation & Institutions display smaller asymmetries, suggesting more balanced or faster-decaying positive and negative sentiment dynamics.

Persistence versus asymmetry. The results highlight an important distinction between the persistence of bad news and the bad-minus-positive cumulative persistence. Bad-news persistence is high in a broad range of topics. Several of the topics with the highest bad-news persistence — financial markets, real estate, and corporate news — are also among the most data-rich topics, where frequent updating might be expected to accelerate mean reversion and crowd out stale narratives. The results suggest otherwise. Bad news remains persistent even in environments with frequent information arrival.

A unifying way to interpret these findings is that persistence depends less on the amount of data available than on whether a news shock opens or closes a continuing narrative process. Bad news tends to persist when it is interpreted not as a single resolvable event, but as a signal of an unresolved underlying state. In complex financial and corporate environments, bad news may trigger cascading uncertainty about valuations, credit conditions, and strategic viability, so that each new data release reopens rather than closes the narrative, and may generate sustained follow-up coverage through secondary events such as earnings revisions, credit rating actions, and litigation.

This mechanism also helps explain why the degree of negativity asymmetry differs across topics, though these interpretations are suggestive and not directly tested in the data. In data-rich topics such as financial markets and corporate news, positive shocks may also initiate continuing narratives through earnings upgrades, analyst revisions, investment responses, or valuation reassessments. This is consistent with the notably high positive sentiment half-life in the Firms topic (5.51 months, the largest across all topics) in Figure A1, suggesting that good news in corporate settings can itself become a sustained narrative rather than a transient signal. As a result, both positive and negative sentiment can be persistent, and the negative-positive persistence gap is less pronounced. By contrast, in data-scarce topics such as geopolitics, security, and human rights, bad news is more likely to signal unresolved and slow-moving underlying risks, while good news may be treated as partial, temporary, or less informative about the broader state of the world. These topics therefore exhibit a larger negative-positive persistence gap. More data does not necessarily imply faster resolution of uncertainty; what matters is whether new information brings narrative closure or instead sustains repeated updating and reinterpretation.

The asymmetry builds over time. Figure 5 further shows how this asymmetry evolves across cumulative horizons. The bad news advantage is not an endpoint artifact: it accumulates gradually in most topics, with the steepest buildup in security and human rights. Firms, financial stability, health, environment, and financial markets also show widening gaps at longer horizons. This pattern confirms that negative narratives do not simply have larger initial impacts. Rather, they propagate more persistently than positive narratives of the same initial magnitude.

5 Conclusion

This paper documents that bad news sticks, but not equally across topics. Using monthly topic-level sentiment indices constructed from the complete Wall Street Journal archive, we show that negativity asymmetry, the tendency for bad news shocks to generate larger and more persistent cumulative responses than positive shocks of comparable magnitude, is a broad feature of sentiment dynamics, present in every one of twelve economic topics. Most strikingly, topic-level heterogeneity is sharper in the degree of this asymmetry than in the level of bad-news persistence itself: the asymmetry gap is largest in security and human rights, health and public safety, real estate and financial stability, and firms and corporate strategy, while bad-news persistence remains sizable across a much broader range of topics, including data-rich areas such as financial markets and corporate news. These patterns are consistent with the interpretation that persistence depends not on the frequency of data arrival, but on whether shocks open continuing narratives of uncertainty, with bad news more likely to trigger sustained reinterpretation across both data-scarce and data-rich environments.

These findings carry direct implications for text-based measurement in finance and macroeconomics. Aggregate sentiment indices obscure substantial heterogeneity: a rise in aggregate negativity may reflect not only a more negative tone but a shift toward topics in which negative narratives are intrinsically more persistent. Treating sentiment as a scalar measure of news mood therefore conflates level effects with dynamic propagation. Our results suggest that text-based measures should distinguish positive from negative sentiment, condition on topic composition, and account for the persistence structure of the underlying narrative environment rather than relying solely on contemporaneous tone.

More broadly, the findings establish that narrative asymmetry is a topic-dependent phenomenon. Where bad news lands matters as much as how bad it is. This opens a natural agenda for future work: tracing how topic-specific negative narratives shape expectations formation, asset prices, and real economic outcomes over longer horizons, and understanding the mechanisms, whether behavioral, institutional, or informational, that make negative sentiment in some topics so much harder to dislodge than in others.

Figures and Tables

(a) Aggregate sentiment

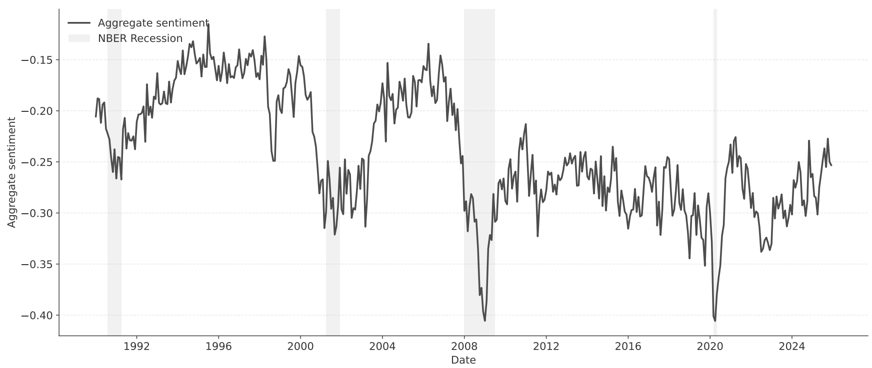

(b) Positive sentiment

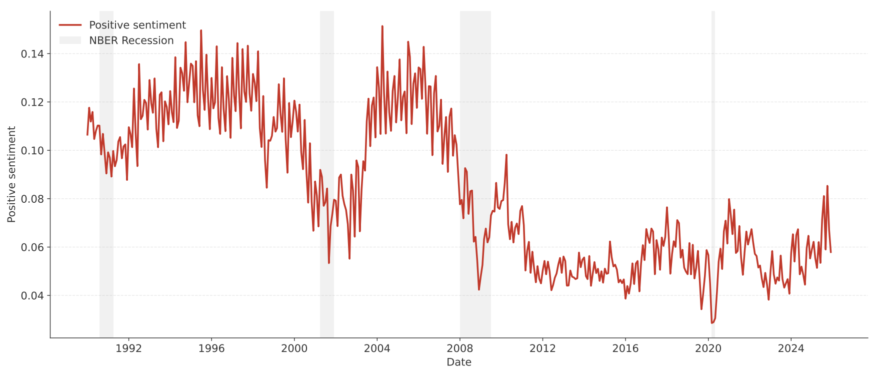

(c) Negative sentiment

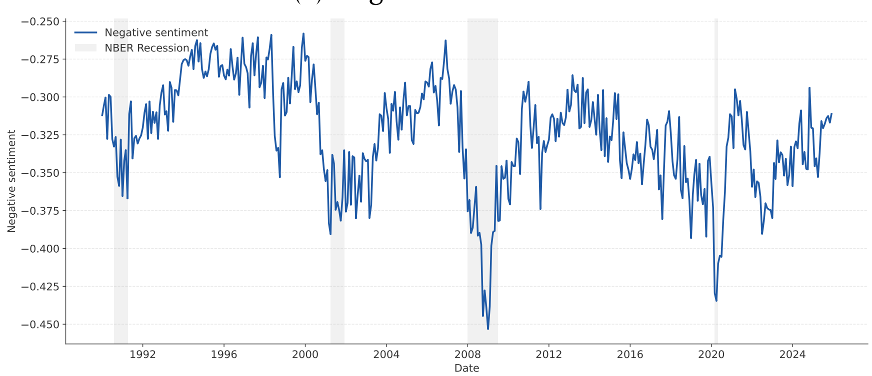

Figure 1: Monthly aggregate, positive, and negative sentiment indices

Notes: This figure presents the time-series of the aggregate, positive, and negative sentiment indices.

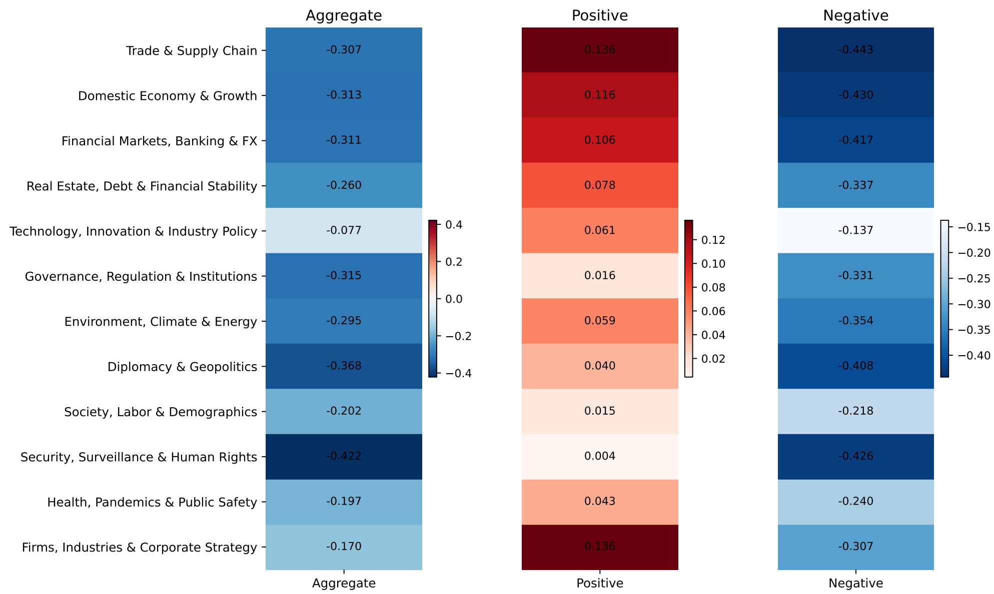

Figure 2: Topic-level average sentiment indices

Notes: This figure presents the time-series average of the aggregate, positive, and negative sentiment indices for each of the 12 topics.

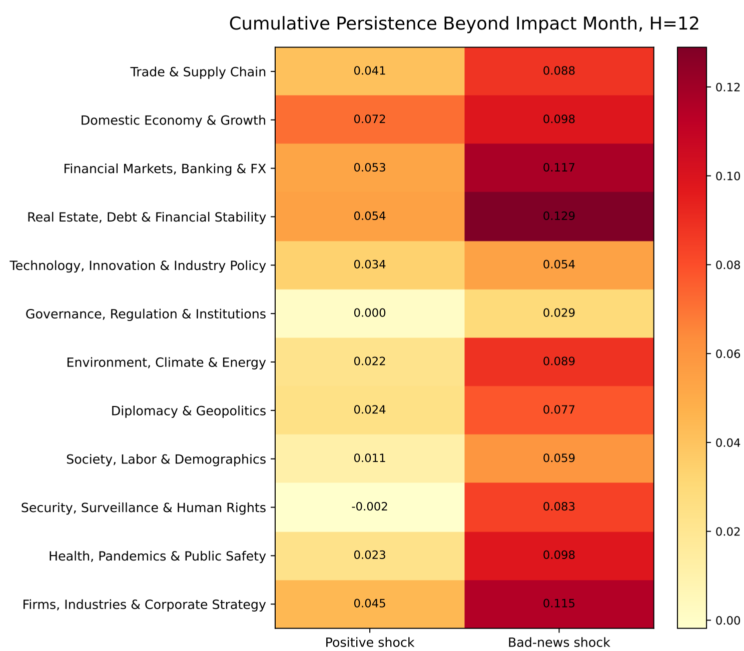

Figure 3: Cumulative persistence of positive and bad news sentiment shocks

Notes: This figure reports cumulative post-impact responses over a 12-month horizon. Rows correspond to topics. The first column reports $CP_k^+(12) = \sum_{h=1}^{12} \hat{\beta}_{k,h}^{++}$, the cumulative response of positive sentiment to a one-standard-deviation positive sentiment innovation. The second column reports $CP_k^-(12) = \sum_{h=1}^{12} \hat{\beta}_{k,h}^{--}$, the cumulative response of bad news sentiment to a one-standard-deviation bad news sentiment innovation. Larger values indicate more persistent post-impact propagation. Individual impulse response functions are in Figure A2 in Appendix.

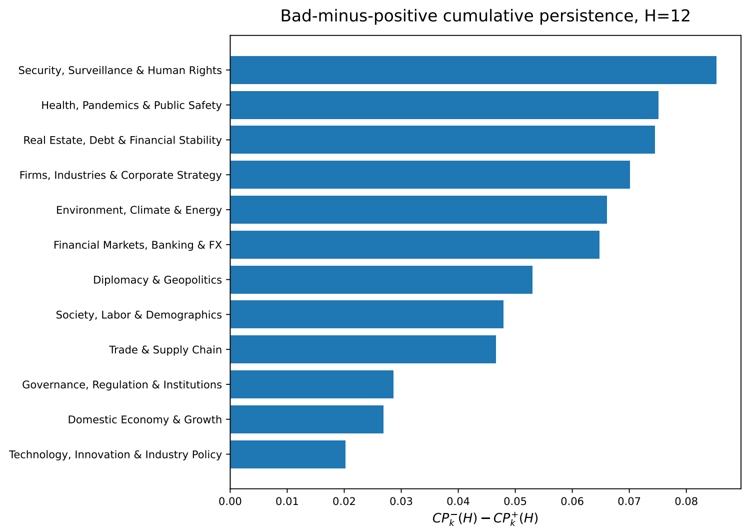

Figure 4: Ranking of within-topic asymmetry in cumulative sentiment persistence

Notes: This figure ranks topics by bad news-minus-positive cumulative persistence over a 12-month horizon. Each bar reports $ASYM_k^{CP}(12) = CP_k^-(12) - CP_k^+(12)$, where $CP_k^-(12) = \sum_{h=1}^{12} \hat{\beta}_{k,h}^{--}$ and $CP_k^+(12) = \sum_{h=1}^{12} \hat{\beta}_{k,h}^{++}$. Positive values indicate that bad news sentiment shocks generate larger cumulative post-impact propagation than positive sentiment shocks within the same topic.

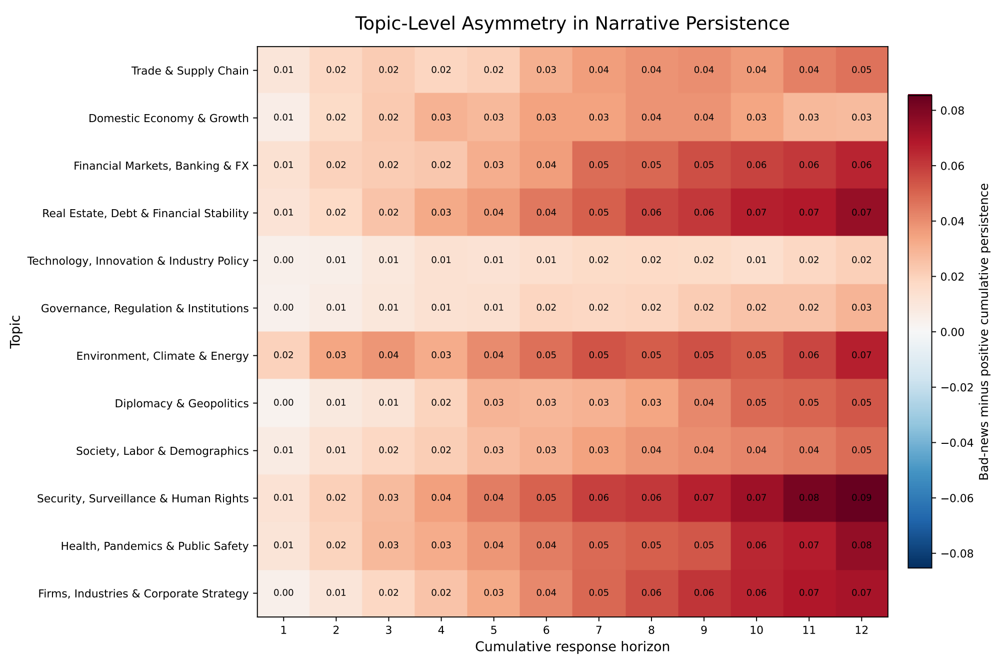

Figure 5: Within-topic asymmetry in cumulative sentiment persistence

Notes: Rows correspond to topics and columns correspond to cumulative response horizons. Each cell reports the bad news-minus-positive cumulative persistence statistic, $ASYM_k^{CP}(H) = CP_k^-(H) - CP_k^+(H)$, where $CP_k^-(H) = \sum_{h=1}^{H} \hat{\beta}_{k,h}^{--}$ and $CP_k^+(H) = \sum_{h=1}^{H} \hat{\beta}_{k,h}^{++}$. Positive values indicate that bad news sentiment shocks generate larger cumulative post-impact propagation than positive sentiment shocks within the same topic. A normalized version of this figure by the initial response can be found in Figure A3 in Appendix.

References

Agarwal, Isha, Wentong Chen, and Eswar S Prasad (2026) "Beyond the Fundamentals: Cross-Country Media Narrative Dispersion and International Capital Flows," Working Paper 33159, National Bureau of Economic Research.

Cookson, J Anthony, Joseph E Engelberg, and William Mullins (2020) "Does partisanship shape investor beliefs? Evidence from the COVID-19 pandemic," *The Review of Asset Pricing Studies*, 10 (4), 863–893.

——— (2023) "Echo chambers," *The Review of Financial Studies*, 36 (2), 450–500.

Cookson, J Anthony and Marina Niessner (2020) "Why don't we agree? Evidence from a social network of investors," *The Journal of Finance*, 75 (1), 173–228.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson (2025) "A plucking model of business cycles," *Journal of Monetary Economics*, 152, 103766.

Engelberg, Joseph E. and Christopher A. Parsons (2011) "The Causal Impact of Media in Financial Markets," *Journal of Finance*, 66 (1), 67–97.

Jensen, Henrik, Ivan Petrella, Søren Hove Ravn, and Emiliano Santoro (2020) "Leverage and deepening business-cycle skewness," *American Economic Journal: Macroeconomics*, 12 (1), 245–281.

Jordà, Òscar (2005) "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95 (1), 161–182.

Nimark, Kristoffer P (2014) "Man-bites-dog business cycles," *American Economic Review*, 104 (8), 2320–2367.

Shiller, Robert J. (2017) "Narrative Economics," *American Economic Review*, 107 (4), 967–1004.

Shiller, Robert J (2020) *Narrative economics: How stories go viral and drive major economic events*: Princeton University Press.

Sims, Christopher A (2003) "Implications of rational inattention," *Journal of monetary Economics*, 50 (3), 665–690.

Soroka, Stuart and Stephen McAdams (2015) "News, politics, and negativity," *Political communication*, 32 (1), 1–22.

Tetlock, Paul C. (2007) "Giving Content to Investor Sentiment: The Role of Media in the Stock Market," *Journal of Finance*, 62 (3), 1139–1168.

Appendix for Online Publication

Table of Contents

A Additional Results for Section 3

Topic Persistence. As a simple summary statistic, we estimate topic-specific autoregressive models for positive and negative sentiment:

$$pos_{k,t} = a_k^+ + \rho_k^+ pos_{k,t-1} + e_{k,t}^+, \quad neg_{k,t} = a_k^- + \rho_k^- neg_{k,t-1} + e_{k,t}^-.$$

We then compute the implied half-life of each process as

$$HL_k^+ = \frac{\log(0.5)}{\log(\hat{\rho}_k^+)}, \quad HL_k^- = \frac{\log(0.5)}{\log(\hat{\rho}_k^-)},$$

for topics with $\hat{\rho}_k \in (0,1)$. These measures provide a transparent first-pass characterization of the persistence of positive and negative sentiment within each topic.

Figure A1 reports AR(1)-implied half-lives of sentiment across topics. A clear asymmetry emerges: in most topics, negative sentiment is more persistent than positive sentiment, indicating that bad news decays more slowly than good news. At the same time, persistence varies substantially across topics. Interpretation-heavy topics such as financial markets and macroeconomic conditions exhibit relatively long half-lives–particularly for negative sentiment–while technology and policy-related topics display much faster decay. Interestingly, some data-rich categories such as financial markets and firms also show high persistence, especially on the negative side. This is somewhat counterintuitive given the high frequency of information updates in these topics, and suggests that negative narratives in these areas may reflect sustained uncertainty and repeated reinterpretation rather than purely transient news. Finally, asymmetry itself differs across topics: some topics exhibit large gaps between negative and positive persistence, while others display more symmetric dynamics, highlighting substantial heterogeneity in how narratives evolve across economic contexts. These estimates capture reduced-form persistence in topic-level sentiment, reflecting both the evolution of underlying events and the dynamics of narrative processing.

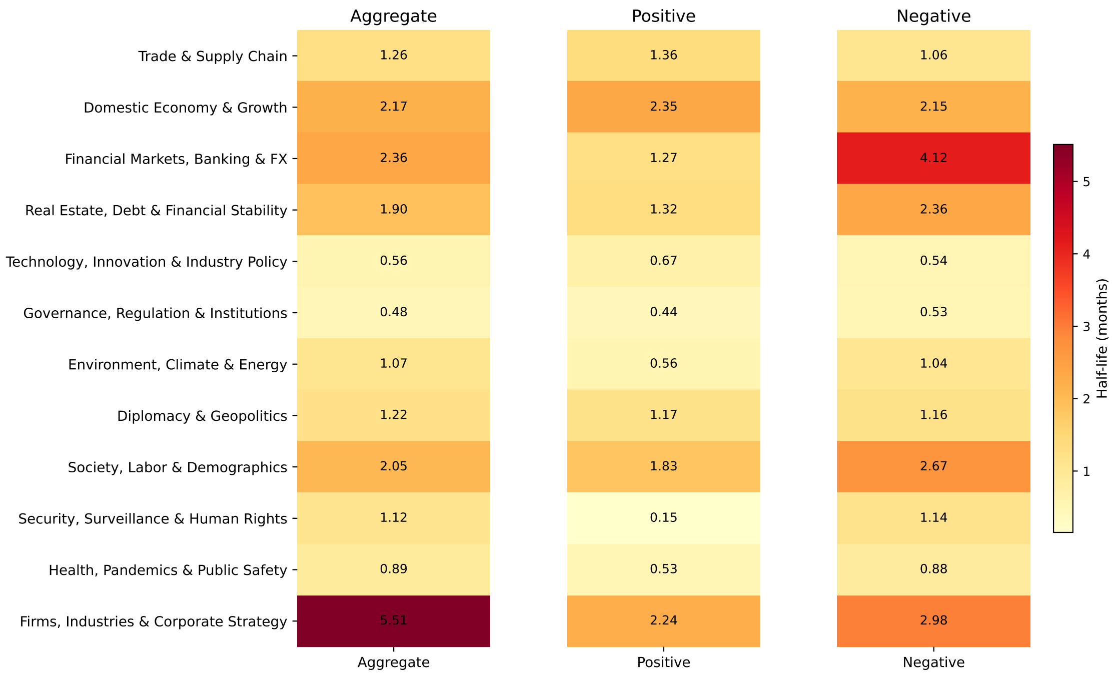

Figure A1: Topic-level half-lives of sentiment indices

Notes: For each topic, the figure reports the AR(1)-implied half-life, in months, of the aggregate sentiment index, the positive sentiment index, and the negative sentiment index. Higher values indicate more persistent sentiment dynamics.

B Additional Results for Section 4

B.1 Impulse Response Functions

Figure A2: Topic-specific impulse responses

Notes: Each panel reports topic-specific local-projection impulse responses over monthly horizons. The positive-sentiment response traces the own response of positive sentiment to a one-standard-deviation positive-sentiment innovation, while the bad-news response traces the own response of bad-news sentiment to a one-standard-deviation bad-news-sentiment innovation. Innovations are standardized within topic. Shaded bands denote 95% pointwise Newey–West/HAC confidence intervals.

B.2 Normalized Persistence

Because cumulative responses may reflect both the size of the initial response and the speed of decay, we also report a normalized persistence measure:

$$NCP_k^+(H) = \sum_{h=1}^{H} \frac{\hat{\beta}_{k,h}^{++}}{\hat{\beta}_{k,0}^{++}}, \qquad NCP_k^-(H) = \sum_{h=1}^{H} \frac{\hat{\beta}_{k,h}^{--}}{\hat{\beta}_{k,0}^{--}}.$$

We then define normalized asymmetry as

$$NASYM_k(H) = NCP_k^-(H) - NCP_k^+(H).$$

This statistic compares the persistence of positive and bad news shocks after scaling each response by its own impact effect. A positive value therefore indicates that bad news responses decay more slowly than positive-news responses within the same topic.

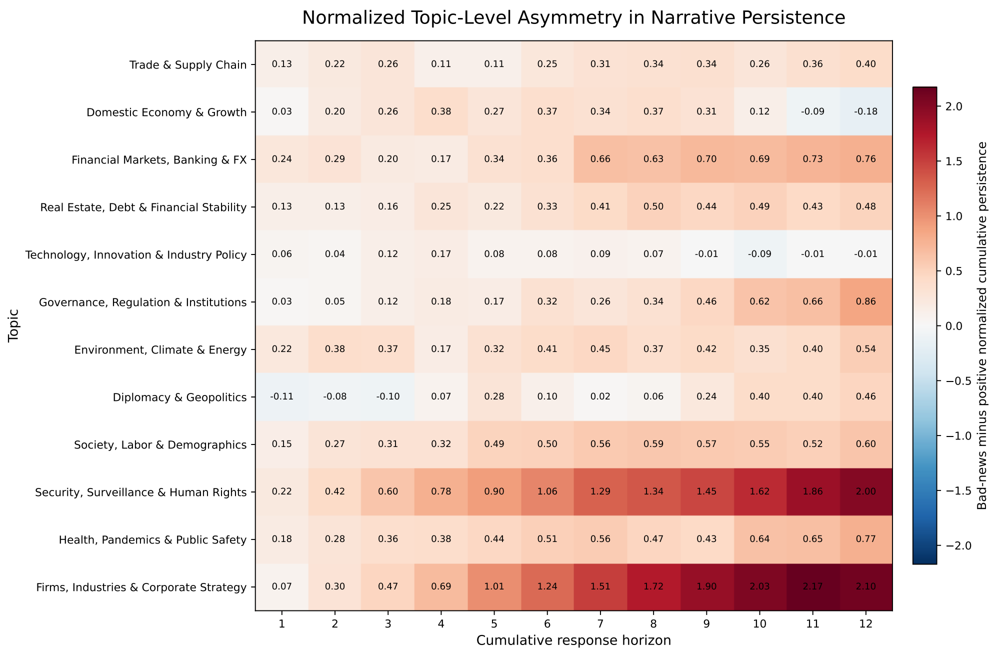

Figure A3: Normalized within-topic asymmetry in sentiment persistence

Notes: Rows correspond to topics and columns correspond to cumulative response horizons. Each cell reports the normalized bad news-minus-positive persistence statistic, $NASYM_k(H) = NCP_k^-(H) - NCP_k^+(H)$, where $NCP_k^-(H) = \sum_{h=1}^{H} \frac{\hat{\beta}_{k,h}^{--}}{\hat{\beta}_{k,0}^{--}}$ and $NCP_k^+(H) = \sum_{h=1}^{H} \frac{\hat{\beta}_{k,h}^{++}}{\hat{\beta}_{k,0}^{++}}$. This statistic compares the persistence of bad news and positive sentiment shocks after scaling each response by its own impact effect. Positive values indicate that bad news responses decay more slowly than positive-news responses within the same topic. Because normalized responses can be sensitive to small impact effects, this figure is reported as an appendix robustness check.