Board of Governors of the Federal Reserve System

International Finance Discussion Papers

ISSN 1073-2500 (Print)
ISSN 2767-4509 (Online)

Number 1438

May 2026

Attention Allocation and Belief Distortions

Sai Ma

Please cite this paper as:
Ma, Sai (2026). "Attention Allocation and Belief Distortions," International Finance Discussion Papers 1438. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/IFDP.2026.1438.

Attention Allocation and Belief Distortions

Sai Ma*

Federal Reserve Board

This Version: April 17, 2026
Latest version is available here.

Abstract

Using microdata from the Michigan Survey of Consumers, we study how within-household reallocations of attention across news affect inflation expectation *bias*, measured relative to a real-time, machine-learning full-information benchmark. Shifting attention toward unfavorable (favorable) economic news increases (decreases) forecast bias substantially, while dropping attention to an unfavorable topic has little effect. The largest bias increases come not from inflation news itself, but from attention to unfavorable social, political, and geopolitical narratives. Aggregate news sentiment has no effect on bias when a household's reported attention allocation is unchanged. In aggregate, these effects are amplified when the attention network is dominated by an unfavorable focal hub: bias-reducing favorable narratives are crowded out of limited attention sets, and respondents closer to the hub exhibit larger bias increases. We find that past and present attention to news together account for up to 70 percent of observed forecast bias, with the current attention component rising sharply during recessions and large negative news events. Results are robust to a battery of specification checks and external validation.

JEL Classification: E31, E52, D83, D84

Keywords: inflation expectations, limited attention, forecast bias, sentiment, networks

*Federal Reserve Board of Governors, C Ave & 20th Street NW, Washington, DC 20551. (sai.ma@frb.gov). I thank Francesco Bianchi, Anna Cieslak, Yeji Sung, seminar participants at the IF Collaboration Week and Fed Board for comments and suggestions. The analysis and conclusions set forth are those of the author and do not indicate concurrence by other members of the research staff or the Board of Governors. All errors are mine. First Draft: March 5, 2026.

1. Introduction

Households persistently deviate from rational-expectations benchmarks in their inflation forecasts, exhibiting large forecast bias even when publicly observable fundamentals remain little changed. A substantial empirical literature documents this result and traces it to information frictions and limited attention, as shown by Coibion and Gorodnichenko (2015) and Bianchi et al. (2022). One leading interpretation is that information-processing constraints affect how households form beliefs.[1] Empirically, many papers study how the level and sentiment of news coverage shape households' inflation expectations and expectation revisions (for example, Larsen et al. 2021; Schmidt et al. 2023; Chahrour et al. 2025). What remains less understood is whether attention reallocations drive *forecast bias* per se and whether good and bad news contribute symmetrically to those belief distortions.

To bridge this gap, we need direct measures of both forecast bias and attention allocation, and constructing each poses a distinct challenge. On one hand, measuring attention requires a survey that records both household expectations and the specific news recently heard. The Michigan Survey of Consumers (MSC) does exactly this, eliciting one-year inflation expectations alongside open-ended descriptions of recent news heard on business conditions, at most two items per interview, each classified by topic family and sentiment. Because respondents can report at most two news items, these reported topics proxy for the news that occupied their limited attention between interviews (Sahm and Sockin, 2016). A respondent who mentioned labor market conditions in one interview and war or geopolitical tensions six months later has, in all likelihood, reallocated attention toward geopolitical risks over the intervening period. By linking respondents across consecutive interviews from January 1990 through December 2025, we observe directly how attention composition changes within households and how those changes relate to shifts in forecast bias over the same horizon.

On the other hand, forecast bias is not directly observed and depends on a benchmark

[1] Mankiw and Reis (2002) introduce sticky information and Sims (2003) rational inattention as frameworks in which processing constraints shape belief formation, and Maćkowiak et al. (2023) provide a comprehensive review. Nimark and Pitschner (2019) show that editorial selection is itself informative: what media outlets choose to report conveys information beyond any single story's content, and concentrated coverage shifts which events become common knowledge.

of what households could have known when forming expectations. To quantify it, we follow Bianchi et al. (2022) and construct a machine-learning inflation forecast built from a rich set of observable real-time information available at the interview date. Because the machine-learning forecast proxies for a full-information benchmark, the deviation of survey expectations from this forecast directly measures belief distortions.

From regressing within-household changes in forecast bias on attention reallocations between consecutive interviews, we find a sharp asymmetry. Households that switch attention toward unfavorable economic news exhibit substantially higher forecast bias, while those that switch toward favorable economic news exhibit lower forecast bias of comparable magnitude. Both effects are economically large relative to the average change in forecast bias between interviews. The sign of the bias also depends on news sentiment: unfavorable attention drives household forecasts above the full-information benchmark, reflecting pessimistic overreaction, while favorable attention drives them below it, producing underreaction. In contrast, dropping mentions of an unfavorable topic between interviews has no material effect on forecast bias, statistically indistinguishable from zero in most specifications. This asymmetry is consistent with evidence that households allocate attention asymmetrically toward unfavorable signals and that unfavorable news leaves a more persistent imprint on beliefs than favorable news.[2]

Does only inflation-related news matter? We find that the bias-increasing effect of unfavorable news is not confined to news about prices or inflation. In a joint specification that conditions on attention switches across all topic families and sentiments simultaneously, the largest increases in forecast bias arise from unfavorable attention to social issues, political developments, and geopolitical risk, not from inflation narratives directly. As a result, broad

[2] Gennaioli and Shleifer (2010) formalize the availability heuristic, showing that easily recalled examples distort probability judgments. Bordalo et al. (2023) extend this framework, modeling how selective memory shaped by recency and similarity determines which past experiences are retrieved, providing a micro-foundation for why recently processed unfavorable signals exert outsized influence on beliefs. Bordalo et al. (2020) apply diagnostic expectations to macro forecasters and find individual-level overreaction that aggregates to consensus underreaction. Bordalo et al. (2022) provide a broader review of overreaction evidence in macroeconomics and finance. Related evidence from credit markets shows that salience of climate disasters drives belief distortions in loan pricing even for unaffected borrowers, with effects amplified by media attention (Correa et al., 2025).

unfavorable sentiment distorts inflation forecasts regardless of whether the topic bears directly on prices. The narratives most associated with bias reduction are favorable topics in sentiment and confidence, financial conditions, and market performance, yet these offsets are weaker in months when one dominant unfavorable topic accounts for most reported mentions.

Does the attention asymmetry simply recapture the well-established finding that news sentiment shapes inflation expectations?[3] We find the connection is more nuanced. Section 4 shows that fluctuations in aggregate news sentiment have no significant effect on forecast bias when a household's attention allocation is unchanged. This comparison matters because it suggests that attention reallocation, rather than the ambient news environment by itself, is what moves bias. The category-level evidence based on Economic Policy Uncertainty (EPU) spikes points in the same direction. Categories that shift attention toward matched unfavorable topics also raise forecast bias, whereas categories that are rarely mentioned move neither attention nor bias. Taken together, this evidence distinguishes between the news environment households face and the news they actually incorporate into their limited attention sets.

The results are robust to a battery of checks, including alternative sample restrictions, demographic subgroups, different inflation regimes, and permutation tests that randomly reassign topic sentiment labels. A central concern is that the estimates reflect attention *intensity* rather than composition, as Bracha and Tang (2025) show that households pay more attention to inflation and respond more strongly to inflation information when inflation is high. We find instead that directly controlling for the total number of reported items leaves the estimates essentially unchanged, while a placebo that replaces sentiment-specific switches with generic add-or-drop indicators yields insignificant results. The composition of news in the attention set, not its volume, drives forecast bias.

The household-level results have meaningful aggregate implications. We first construct a monthly *attention network* in which each node combines a topic family and a sentiment

[3] Carroll (2003), Coibion et al. (2018), Larsen et al. (2021), Binder et al. (2025), and Chahrour et al. (2025) document such effects through variation in news coverage or supply.

label, for example unfavorable price news, and each link records how frequently households report two narratives together in the same interview. We show that some months feature an *unfavorable focal hub*, in which reported attention concentrates on a single unfavorable narrative. Following the terrorist attacks in September 2001, for instance, aggregate attention concentrated on unfavorable war and national security narratives, producing a pronounced hub-and-spoke network. In months featuring such a hub, favorable narratives on confidence, financial conditions, and markets are less able to offset the effect of unfavorable attention on bias. Proximity to the unfavorable hub, defined by how frequently topics are mentioned together in the network, also matters within month: households mentioning topics closer to the hub exhibit larger forecast bias, but only when the hub is itself unfavorable.

We then use this network structure to construct a quantitative decomposition of aggregate forecast bias. Calibrating a parsimonious framework of attention draws from the monthly attention network to the reduced-form estimates, we recover three components: persistence, contemporaneous attention, and a residual. Persistence and contemporaneous attention together account for 58 percent of average absolute bias in the full sample and 70 percent in the dual-response sample. The attention share is not constant: it averages 31 percent in NBER recession months and rises to 48 percent post-2020, when unfavorable narratives dominate the network. This decomposition links the household-level evidence to aggregate fluctuations in forecast bias over time.

Related Literature. This paper relates to work on news selection, expectation formation, and belief distortions. A growing literature studies how media selection and narrative transmission shape beliefs and aggregate fluctuations, as shown by Nimark (2014), Chahrour et al. (2021), and Nimark and Pitschner (2019). Empirically, news coverage is often tilted toward negative developments, as documented by Soroka (2012), and media outlets may slant coverage toward audience priors, as argued by Gentzkow and Shapiro (2006). We study how household-level attention reallocations within a common monthly news environment map into forecast bias.

The paper is also related to evidence that media exposure, news content, and attention affect inflation expectations, as documented by Carroll (2003), Coibion et al. (2018), Larsen et al. (2021), Pfajfar and Santoro (2013), Schmidt et al. (2023), and Binder et al. (2025). Most closely related, Chahrour et al. (2025) use MSC panel data to study asymmetric responses to higher-price versus lower-price inflation news in a rational-expectations framework. Our paper differs in two ways. Rather than expectation revisions, we study forecast bias, the persistent gap between survey expectations and a full-information benchmark. Rather than levels of reported mentions, we use Add/Drop switches in limited attention sets, which allows separate identification of adding, dropping, and persistence effects not available in level-based tests. We also condition on the full set of fourteen MSC topic families rather than only inflation or price-related news. This matters empirically because the largest bias-increasing effects come from unfavorable attention to social conditions, political developments, and geopolitical risk rather than from inflation narratives directly.[4]

Our results also connect to research on informational rigidity and overreaction (Coibion and Gorodnichenko, 2015; Gennaioli and Shleifer, 2010; Bordalo et al., 2020, 2018, 2022; Kohlhas and Walther, 2021). The asymmetry between adding and dropping unfavorable news is consistent with the idea that unfavorable signals receive disproportionate weight and continue to matter even after they are no longer explicitly reported. In that sense, the paper is close in spirit to Gennaioli et al. (2024), who emphasize selective memory in household inflation expectations, and to work showing that personal inflation experiences and shopping exposure shape beliefs (Malmendier and Nagel, 2016; D'Acunto et al., 2021). It also complements survey information-treatment studies (Armantier et al., 2016; Cavallo et al., 2016; Coibion et al., 2022; Weber et al., 2025; Kim and Binder, 2023) by using naturally occurring attention reallocations rather than researcher-provided signals. This literature has largely focused on the content and format of central-bank messages, but our results add a distinct dimension: which competing narratives dominate household attention sets at the

[4]The broader literature on news and inflation expectations identifies effects through variation in news coverage or supply, as shown by Carroll (2003), Larsen et al. (2021), Binder et al. (2025), and Chahrour et al. (2025). Our interaction-IV result adds a further distinction: changes in the aggregate news environment that do not shift a household's reported attention allocation have no significant effect on forecast bias, which is consistent with reallocation of limited attention as the operative channel.

time shapes how far those messages pass through to expectations.

Finally, the paper connects to the narrative-network and inflation-attention literatures. Work on narratives emphasizes how economic stories spread and shape macroeconomic outcomes (Shiller, 2017), while recent work shows that attention to inflation rises when inflation becomes more salient (Link et al., 2023; Bracha and Tang, 2025; Pfäuti, 2025). We complement these papers by showing that the composition of attention matters alongside its intensity, and by linking the attention network to both household-level bias changes and aggregate fluctuations in forecast bias.

Section 2 describes the data. Section 3 documents the core Add/Drop asymmetry, its directional content, and heterogeneity across narratives. Section 4 provides direct falsification evidence, interaction-IV evidence, and EPU-based external validation for the attention channel. Section 5 develops the attention network and quantifies its aggregate implications. Section 6 concludes.

2. Data and Measurement

For measuring each household's attention allocation and belief distortion, we draw on two data sources: linked MSC microdata recording household inflation expectations and open-ended news descriptions, and a real-time machine-learning benchmark for inflation that serves as a full-information reference.

2.1. Household Inflation Expectations and News Attention

We use microdata from the University of Michigan *Survey of Consumers* (MSC), a long-running, nationally representative survey that elicits inflation expectations alongside questions on news heard about business conditions. For inflation expectations, the survey asks: "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?" and then "By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?" The respondent's one-year inflation expectation is denoted $\pi^e_{i,t}$.

For reported "news heard" topics (up to two items), the survey asks: "During the last few months, have you heard of any favorable or unfavorable changes in business conditions?" If the respondent answers yes, the follow-up is "What did you hear? (Have you heard of any other favorable or unfavorable changes in business conditions?)." Respondents can report up to two items, which we refer to as a first response and a second response. This survey design is important for our analysis for two reasons. First, with at most two mentions, survey reports of news heard serve as a proxy for the news that occupied households' attention between interviews (Sahm and Sockin, 2016).[5]. Second, it permits precise measurement of which news narratives are added to or dropped from this attention set between interviews.

To construct the matched panel, we link respondents across consecutive MSC interviews and retain only those with inflation expectations in both rounds. Because the MSC uses a rotating-panel design, this yields a standard six-month interval between interviews for each household.[6] We therefore measure attention reallocations and forecast bias over a common six-month horizon.

By construction, each household reports news by sentiment (favorable or unfavorable) and by topic (e.g., elections or crime levels). The responses are then aggregated into fourteen topic families (broad groupings of related MSC codes), and the survey's sentiment classification, favorable (F) or unfavorable (UF), is retained. A response is coded as valid if the response maps to one topic family and one sentiment. Missing, uncodable, and ineligible responses are excluded in the sample. Appendix Table A1 reports the full code mapping and excluded residual categories.

2.2. Measuring Expectation Bias

We quantify households' expectation bias using a real-time machine-learning benchmark for inflation, denoted $\widehat{\pi}_t^{ML}$, following Bianchi et al. (2022). The benchmark is constructed to proxy the full-information rational-expectations forecast using only publicly observable data available at the interview date. We therefore define forecast bias as the deviation of

[5]When respondents attend to more than two narratives over the intervening period, the reported items could potentially capture the most salient or most recently recalled topics, though we have no direct evidence on the selection mechanism.

[6]Specifically, 99.6% of reinterviews occur exactly at the six-month mark, with a negligible remainder at five or seven months.

a respondent's survey expectation from $\widehat{\pi}_t^{ML}$. The bias is conceptually distinct from an expectation revision, which does not by itself indicate whether an expectation change moves a household toward or away from a rational expectation benchmark. Bianchi et al. (2022) provide the construction details and validation, and the discussion below summarizes only the benchmark features needed for measurement in this paper.

Formally, the benchmark is generated from

$$\pi_{t+h} \; = \; G(\mathcal{I}_t) \; + \; \varepsilon_{t+h}, \tag{1}$$

where $\mathcal{I}_t$ is a high-dimensional real-time information set and $G(\cdot)$ is a flexible machine-learning mapping estimated with rolling pseudo-out-of-sample validation and tuning. The Elastic Net implementation of Bianchi et al. (2022) is used, and a monthly series is constructed to match MSC interview timing. Internet Appendix Section B provides details on the inputs, timing, and real-time estimation protocols.

Three timing choices are important. First, each element of $\mathcal{I}_t$ is observed in real time, using vintages that respect release lags and revisions. Second, because predictors are mixed frequency, the benchmark is re-estimated and updated at monthly forecast dates using the most recent values observable at that date. Third, for interviews conducted during month t, the machine deadline is set conservatively to the first day of month t, so the machine uses information only through the end of month $t-1$. This convention is intentionally conservative and can handicap the machine relative to respondents interviewed later in month t.

Given $\widehat{\pi}_t^{ML}$, respondent i's forecast bias is $b_{i,t}^{ML} = \pi_{i,t}^{e} - \widehat{\pi}_t^{ML}$.[7] The main dependent variable is the six-month change in absolute forecast bias,

$$\Delta|b_{i,t}^{ML}| \; \equiv \; |b_{i,t}^{ML}| - |b_{i,t-1}^{ML}|.$$

A positive value indicates movement away from the machine benchmark between interviews, while a negative value indicates movement toward it. Signed changes, $\Delta b_{i,t}^{ML} = b_{i,t}^{ML} - b_{i,t-1}^{ML}$,

[7] We verify the rationality of $\widehat{\pi}_t^{ML}$ using the Mincer and Zarnowitz (1969) test, regressing realized inflation on the ML forecast and testing the joint null of zero intercept and unit slope. The benchmark passes this test; results available upon request.

are examined in robustness exercises.

Figure 1: Distribution of MSC Inflation Forecast Bias Across Subperiods

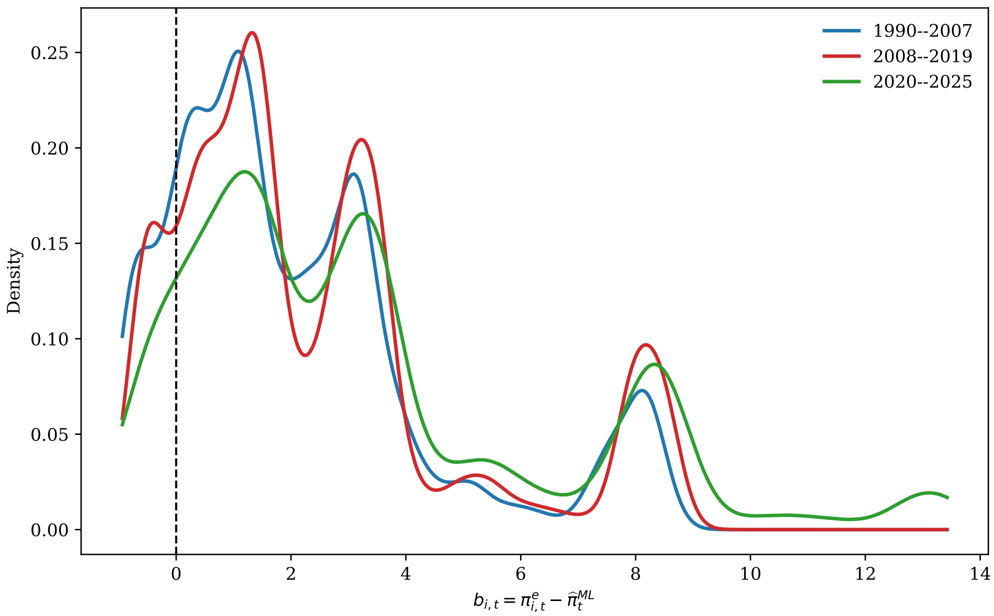

Notes. Kernel densities of estimated forecast bias $b_{i,t}^{ML} = \pi_{i,t}^{e} - \widehat{\pi}_t^{ML}$ are shown for 1990–2007, 2008–2019, and 2020–2025 across all households. The vertical dashed line marks zero bias.

Figure 1 plots the cross-sectional distribution of forecast biases $b_{i,t}^{ML}$ separately for three subperiods. Two features stand out. First, the distribution is right-skewed in every period: the mass concentrates near zero, but a persistent right tail reflects households that substantially over-predict inflation relative to the full-information machine benchmark (even in the 1990–2007 period, when realized inflation was low). Second, the cross-sectional distribution is remarkably stable across macroeconomic regimes. The Great Moderation (1990–2007), the post-GFC decade (2008–2019), and the recent inflation surge (2020–2025) produce nearly overlapping distributions. The 2020–2025 distribution shifts modestly rightward, consistent with the median forecast rising to around 5 percent during the inflation surge, but remains similarly right-skewed as earlier periods. This cross-regime stability suggests that upward forecast bias is a persistent feature of household inflation expectations rather than driven by a specific period.[8]

[8] Coibion and Gorodnichenko (2015) document systematic upward bias and information rigidities in MSC inflation expectations over a comparable sample period, confirming that the right tail in household forecasts is not attributable to any single macroeconomic episode.

Table 1 summarizes the matched samples. The full sample covers January 1990 through December 2025 and contains 40,362 matched interview pairs. The dual-response sample, which requires two valid coded mentions in both interviews, contains 10,035 pairs. Average movements are slightly toward the benchmark ($\overline{\Delta|b^{ML}|} = -0.193$ and -0.155 in the two samples). Because 34.8% of respondents report no valid mention and 25.3% report exactly one, the dual-response sample captures within-respondent changes in the composition of reported news more sharply. Unfavorable mentions account for roughly two-thirds of valid mentions in both samples, consistent with the negativity bias in news coverage (Soroka, 2012).

Table 1: Sample Summary Statistics

	Full sample	Dual-response sample
Observations	40,362	10,035
$\Delta\|b^{ML}\|$ (6m change)	-0.193	-0.155
Lag $\|b^{ML}\|$	2.950	2.798
Add UF (any)	0.352	0.534
Drop UF (any)	0.362	0.555
Add F (any)	0.231	0.370
Drop F (any)	0.225	0.350
Total moves	1.402	2.228
UF share in current reported mentions	0.636	0.629

Notes. "Full sample" refers to respondents matched across consecutive MSC interviews. "Dual-response sample" restricts the sample to respondents with two valid reported news mentions in both interviews. Rows report weighted means over January 1990–December 2025. Forecast bias $b^{ML}_{i,t}$ and its change $\Delta|b^{ML}_{i,t}|$ are defined in Section 2.2 and Internet Appendix Section B.

3. Attention Reallocation and Forecast Bias

To characterize the effects of attention reallocations on forecast bias, we link each respondent across adjacent interviews $(t-1, t)$ and ask how changes in the sentiment composition of reported news covary with the six-month change in absolute bias. For example, a respondent who mentioned labor market conditions at the first interview and war or geopolitical tensions at the second has dropped unfavorable labor-market news and added unfavorable geopolitical news over the six-month horizon. We formalize this with Add/Drop indicators below. The dependent variable is $\Delta|b^{ML}_{i,t}|$, where $b^{ML}_{i,t} \equiv \pi^e_{i,t} - \widehat{\pi}^{ML}_t$.

Let $\mathcal{R}_{i,t} \subseteq \{\text{UF, F}\}$ denote the set of sentiment labels present in respondent i's reported news at interview t: UF $\in \mathcal{R}_{i,t}$ if the respondent reported at least one unfavorable item, and F $\in \mathcal{R}_{i,t}$ if at least one favorable item was reported. Because respondents can report up to two items, $\mathcal{R}_{i,t}$ can equal $\{\text{UF}\}$, $\{\text{F}\}$, $\{\text{UF}, \text{F}\}$, or $\emptyset$. We define (attention) Add/Drop indicators as follows,

$$\text{Add}^s_{i,t} \equiv \mathbf{1}\{s \in \mathcal{R}_{i,t},\ s \notin \mathcal{R}_{i,t-1}\}, \qquad \text{Drop}^s_{i,t} \equiv \mathbf{1}\{s \notin \mathcal{R}_{i,t},\ s \in \mathcal{R}_{i,t-1}\}.$$

We refer to these as Add UF, Drop UF, Add F, and Drop F for the rest of the paper. Because respondents report at most two topics per interview, these indicators measure changes in reported attention composition rather than the arrival of objectively new information. We restrict $\mathcal{R}_{i,t}$ to topic families appearing in at least 0.1% of interviews to limit sensitivity to rarely reported categories.

The baseline regression specification is

$$\Delta|b^{ML}_{i,t}| = \beta_{UF}\,\text{Add}^{UF}_{i,t} + \delta_{UF}\,\text{Drop}^{UF}_{i,t} + \beta_F\,\text{Add}^F_{i,t} + \delta_F\,\text{Drop}^F_{i,t} + \Gamma' X_{i,t-1} + \tau_t + \psi_{g(i)} + \varepsilon_{i,t}, \quad (2)$$

where $X_{i,t-1}$ includes lagged absolute bias, lagged inflation expectations, the gap in months between interviews, and total topic moves. All specifications include interview-month fixed effects τ_t, demographic fixed effects $\psi_{g(i)}$, and standard errors are two-way clustered by respondent and month.

A natural concern is that the four attention indicators may be mechanically linked. For most transitions they are not: a respondent moving from $\{\text{UF}\}$ to $\{\text{UF}, \text{F}\}$ adds favorable news without dropping unfavorable news, while one moving from $\{\text{UF}, \text{F}\}$ to $\{\text{UF}\}$ drops favorable news without adding unfavorable news. The four indicators therefore vary independently across a large share of the sample.[9]

[9] One case does produce a mechanical link: a respondent who reports a single item in both interviews and switches sentiment entirely (moving from $\{\text{UF}\}$ to $\{\text{F}\}$) has Drop UF $= 1$ and Add F $= 1$ simultaneously by construction. This does not drive the results for two reasons. First, such pure-sentiment switches account for a minority of sample observations. Second, the dual-response subsample, restricted to households reporting two valid items in both interviews where a single-item sentiment switch cannot occur, yields estimates nearly identical to the full-sample baseline (Table 2, Panel A, columns 1–2), confirming that the mechanical link does not explain the asymmetry.

An alternative two-variable approach would use net changes in unfavorable and favorable mention counts, $\Delta n_{i,t}^{UF}$ and $\Delta n_{i,t}^{F}$, where $n_{i,t}^{UF}, n_{i,t}^{F} \in \{0, 1, 2\}$ aggregate first and second responses. This approach, however, imposes an implicit symmetry: increases and decreases in each count are constrained to have equal and opposite effects. As we show next, the data reject that symmetry.

3.1. Asymmetric Response to News Attention

Table 2 reports the first key results of the paper. Panel A shows that adding unfavorable news to the reported attention set increases forecast bias by 0.278 percentage points, while adding favorable news reduces it by 0.301 percentage points. Both estimates are highly significant and economically large. The mean of $\Delta|b^{ML}|$ is -0.193 (Table 1), so a 0.278 increase more than offsets the average reduction in absolute forecast bias between interviews and implies a net increase. This asymmetry is consistent with evidence in Kohlhas and Walther (2021) that agents allocate attention asymmetrically toward unfavorable signals relative to favorable ones. It also motivates the Add/Drop formulation directly: the symmetry imposed by net-count changes is not supported in the data.

The restricted dual-response sample yields the same result, with somewhat larger magnitudes: $\text{Add}_{i,t}^{UF} = 0.345$ and $\text{Add}_{i,t}^{F} = -0.278$. The larger magnitudes are consistent with attention reallocations being captured more directly when two explicit responses are observed in both interviews.

3.1.1. Persistence After Dropping Unfavorable Attention

Dropping unfavorable news, by contrast, does not reduce forecast bias. Panel A of Table 2 shows that the coefficient on Drop UF is small and statistically insignificant at 0.062 in the full sample. The dual-response sample shows the same result. This null result is informative because a topic can leave the reported news mentions without leaving households' belief formation. Because the MSC records at most two news mentions in each interview, a newly salient narrative can displace an earlier unfavorable topic in the survey response even if that earlier topic continues to affect expectations. This is consistent with evidence that adverse information leaves more persistent imprints on beliefs than favorable information (Kohlhas and Walther, 2021; Gennaioli et al., 2024). Internet Appendix Section C.3 further shows that interview pairs in which unfavorable mentions outnumber favorable mentions are substantially more persistent across consecutive interviews than the reverse, corroborating this interpretation in the MSC data.

Drop F behaves differently. In the full sample, the coefficient is negative and statistically significant at -0.126, the same sign and a similar order of magnitude as the average six-month decline in absolute bias reported in Table 1. Unlike the null Drop UF effect, this suggests that the bias-reducing effect of favorable news dissipates more quickly once it leaves the reported attention set. Internet Appendix Section C.3 supports this interpretation directly: favorable attention is substantially less persistent across consecutive interviews than unfavorable attention.

3.2. Direction of Bias

Throughout the paper, the main focus is the change in **absolute** forecast bias, $\Delta|b_{i,t}^{ML}|$, because the central question is whether reallocating attention moves household expectations toward or away from the ML benchmark. Both upward and downward departures from the benchmark represent belief distortions, and households on opposite sides of the benchmark can both move farther away from it following different attention shifts. At the same time, the direction of bias is economically informative. To clarify the directional content of the main asymmetry, we examine signed expectation bias, $b_{i,t}^{ML} = \pi_{i,t}^{e} - \widehat{\pi}_{t}^{ML}$, where positive values indicate that households expect higher inflation than the benchmark (*overreaction*) and negative values indicate expectations below the benchmark (*underreaction*).[10]

Table 3 shows that the main asymmetry carries through directly to the sign of the distortion. In Columns (1) and (2), adding unfavorable news increases signed bias by 0.313 percentage points in the full sample, implying greater upward overreaction relative to the ML benchmark. By contrast, adding favorable news reduces signed bias by 0.346 percentage

[10]This definition of overreaction and underreaction is established by Bianchi et al. (2022) and adopted in subsequent work using the same ML benchmark framework (Bianchi et al., 2024, 2025).

Table 2: Attention Allocation and Expectation Bias

	Panel A: Baseline Asymmetry		Panel B: Split by Response Order	
	Full sample	Dual-response sample	Full sample	Dual-response sample
	(1)	(2)	(3)	(4)
Add UF	0.278***	0.345***		
	(0.048)	(0.086)		
Drop UF	0.062	0.150*	0.109**	0.181**
	(0.045)	(0.080)	(0.048)	(0.079)
Add F	-0.301***	-0.278***	-0.232***	-0.161
	(0.052)	(0.097)	(0.056)	(0.122)
Drop F	-0.126***	-0.136*	-0.086*	-0.113
	(0.046)	(0.079)	(0.048)	(0.080)
Add UF (Primary Response)			0.266***	0.323***
			(0.051)	(0.091)
Add UF (Secondary Response)			0.197***	0.275***
			(0.055)	(0.093)
Equality Test:				
Add UF (Primary − Secondary)			0.069	0.048
			(0.068)	(0.087)
Observations	40,362	10,035	40,362	10,035
R^2	0.371	0.410	0.371	0.410
Month FE	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes

Notes. The table reports weighted matched-pair regressions for news-attention reallocations and forecast bias. In both panels, the dependent variable is $\Delta|b_{i,t}^{ML}|$. All regression columns include baseline controls, month fixed effects, and demographic fixed effects, with two-way clustered standard errors by respondent and interview month reported in parentheses. Panel A uses the baseline Add/Drop indicators. Panel B replaces $\text{Add}_{i,t}^{UF}$ with first-response and second-response indicators, $\text{Add}_{i,t}^{UF,1}$ and $\text{Add}_{i,t}^{UF,2}$, while keeping the remaining regressors unchanged, and reports the test of $H_0 : \beta_{UF,1} - \beta_{UF,2} = 0$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

points, moving beliefs toward or below the benchmark. The squared-bias column shows that this asymmetry is not only directional but also reflects economically larger distortions in magnitude: Add UF raises squared bias by 3.099, whereas Add F lowers it by 3.271. The signed and squared results therefore both indicate that unfavorable attention amplifies upward overreaction, while favorable attention reduces it.

These findings connect to the overreaction and underreaction literature. Households that shift attention toward unfavorable news overreact upward relative to the ML benchmark, while those shifting toward favorable news underreact. Because unfavorable news accounts for a larger share of reported mentions in the sample (Table 1), overreaction is the more common direction at the individual level, consistent with Bordalo et al. (2020) and

Table 3: Direction of Bias and Forecast Errors

	(1) $\Delta b_{i,t}^{ML}$	(2) $\Delta (b_{i,t}^{ML})^2$	(3) $\Delta \lvert FE_{i,t} \rvert$	(4) $\Delta FE_{i,t}$
Add UF	0.313***	3.099***	0.317***	0.463***
	(0.050)	(0.627)	(0.059)	(0.076)
Add F	-0.346***	-3.271***	-0.226***	-0.446***
	(0.054)	(0.659)	(0.064)	(0.078)
Drop UF	0.074	1.119**	0.178***	0.088
	(0.047)	(0.553)	(0.051)	(0.057)
Drop F	-0.161***	-1.756***	-0.162***	-0.143**
	(0.050)	(0.570)	(0.057)	(0.067)
Observations	40,362	40,362	40,362	40,362
R^2	0.363	0.349	0.284	0.291
Month FE	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes

Notes. Columns (1) and (2) use the same baseline specification as Table 2 but replace $\Delta|b_{i,t}^{ML}|$ with the signed change in forecast bias, $\Delta b_{i,t}^{ML}$, and the change in squared forecast bias, $\Delta(b_{i,t}^{ML})^2$, respectively. Columns (3) and (4) replace the forecast-bias outcomes with realized inflation forecast error outcomes, $\Delta|FE_{i,t}|$ and $\Delta FE_{i,t}$. All columns use the same baseline controls, month fixed effects, demographic fixed effects, survey weights, and two-way clustered standard errors by respondent and month. The realized-forecast-error sample is truncated at the end of the panel because future realized inflation is required to construct $FE_{i,t}$. Dual-response specifications for columns (1) and (2) yield the same qualitative result and are omitted for compactness. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Kohlhas and Walther (2021).[11] Individual-level overreaction need not translate into aggregate overreaction: Bianchi et al. (2024) use the ML benchmark to show that when agents respond simultaneously to news events with offsetting implications, aggregate expectations can exhibit informational rigidity (Coibion and Gorodnichenko, 2015) even when individuals overreact.

The directional asymmetry also holds when predicting realized forecast errors rather than the bias: adding unfavorable news increases forecast errors, while adding favorable news reduces them (Columns (3) and (4)). Attention shifts toward unfavorable news thus drive beliefs toward overreaction and larger forecast errors regardless of the benchmark used. Internet Appendix Section C.1 further shows that response order does not drive these results.

[11] Section 5 shows that unfavorable narratives consistently occupy the focal hub position in the attention network, and the bias decomposition shows the attention contribution spikes during recessions and large negative news events.

3.3. Inflation versus Non-Inflation News

A natural next question is whether the bias asymmetry is driven mainly by inflation-related narratives, or whether non-price topic composition adds independent information. We evaluate this in two complementary exercises. First, we estimate topic-by-topic regressions that characterize heterogeneity in the bias effect across narratives. Second, we estimate a joint specification that ranks topic-level effects while holding other attention switches fixed.

Let $f \in \mathcal{F}$ index a topic family and $s \in \{UF, F\}$ index sentiment. Define the indicator that topic (f, s) is present in the reported attention set at interview t:

$$A_{i,t}^{f,s} \equiv \mathbf{1}\{\mathcal{R}_{i,t} \text{ contains at least one mention in family } f \text{ with sentiment } s\}.$$

The corresponding Add/Drop indicators are

$$\text{Add}_{i,t}^{f,s} \equiv \mathbf{1}\{A_{i,t}^{f,s} = 1,\ A_{i,t-1}^{f,s} = 0\}, \qquad \text{Drop}_{i,t}^{f,s} \equiv \mathbf{1}\{A_{i,t}^{f,s} = 0,\ A_{i,t-1}^{f,s} = 1\}.$$

These capture whether a given narrative topic f with sentiment s is added to or dropped from the respondent's reported attention set between interviews. For instance, a respondent moving from $\{\text{UF inflation}, \text{F market}\}$ to $\{\text{UF politics}, \text{F market}\}$ drops unfavorable inflation news and adds unfavorable political news, even though the overall sentiment composition of the reported attention set is unchanged. In that case, $\text{Add}_{i,t}^{\text{politics},UF} = 1$ and $\text{Drop}_{i,t}^{\text{inflation},UF} = 1$, while the broad sentiment-level indicators Add UF and Drop UF are both zero.

3.3.1. Topic-by-Topic Regressions

To characterize heterogeneity across topics, we estimate a separate regression for each (f, s):

$$\Delta|b_{i,t}^{ML}| = \beta_{f,s}\, \text{Add}_{i,t}^{f,s} + \delta_{f,s}\, \text{Drop}_{i,t}^{f,s} + \Gamma' X_{i,t-1} + \tau_t + \psi_{g(i)} + \varepsilon_{i,t}^{f,s}. \tag{3}$$

Figure 2 plots $\widehat{\beta}_{f,s}$ — the estimated association between adding narrative (f, s) and $\Delta|b^{ML}|$, conditional on baseline controls and fixed effects — restricting to categories with move frequency at least 0.1% and add-coefficient $p \leq 0.10$ to avoid overinterpreting infrequently reported narratives.

We highlight two key findings. First, the coefficient signs align with the aggregate asymmetry: all statistically significant Add-UF coefficients are positive, and most Add-F coefficients are negative, with the exception of favorable war news, which raises bias regardless of sentiment label.[12] This split shows that the baseline asymmetry does not merely hold on average, but operates consistently on a topic-by-topic basis.

Second, despite substantial within-sentiment dispersion across topic families, most coefficients are not statistically distinguishable from one another. This suggests that no single topic drives the aggregate effects documented in Table 2. Negative inflation news does not produce the largest bias. Adding unfavorable news about social issues or war generates larger point estimates for bias increases than adding unfavorable inflation news, although the differences are not statistically significant. A natural interpretation is that households map salient non-price bad news into broader beliefs about economic deterioration and future prices.[13]

3.3.2. Joint Specification: the Most Bias-Increasing Narratives

Figure 2 evaluates topics one at a time, but adding a new narrative often crowds out an existing one for respondents reporting two items. To isolate partial effects, we estimate a joint specification that conditions on all observed Add and Drop events simultaneously:

$$\Delta|b_{i,t}^{ML}| = \sum_{f\in\mathcal{F}}\sum_{s\in\{UF,F\}}\left(\beta_{f,s}\,\mathrm{Add}_{i,t}^{f,s} + \delta_{f,s}\,\mathrm{Drop}_{i,t}^{f,s}\right) + \Gamma' X_{i,t-1} + \tau_t + \psi_{g(i)} + \varepsilon_{i,t}. \tag{4}$$

Relative to Equation (3), $\beta_{f,s}$ is identified as a partial effect for topic family f with sentiment s: the incremental effect between adding narrative (f, s) to the attention set and

[12] Caldara et al. (2026) show across 44 countries that geopolitical risk raises realized inflation, with supply disruptions dominating deflationary demand effects. Coibion et al. (2025) provide survey-based evidence that households expecting longer geopolitical conflicts anticipate higher inflation, consistent with war-related news increasing forecast bias regardless of the sentiment label attached to it.

[13] This pattern is consistent with Shiller (2017)'s argument that non-economic stories regularly shape economic expectations. One interpretation is that unfavorable price news is bounded by the observed price level and therefore produces a more limited updating response, whereas unfavorable political and social narratives carry no direct correspondence to observable prices and may activate broader revisions to beliefs about economic conditions.

Figure 2: News Topic Attention and Forecast Bias

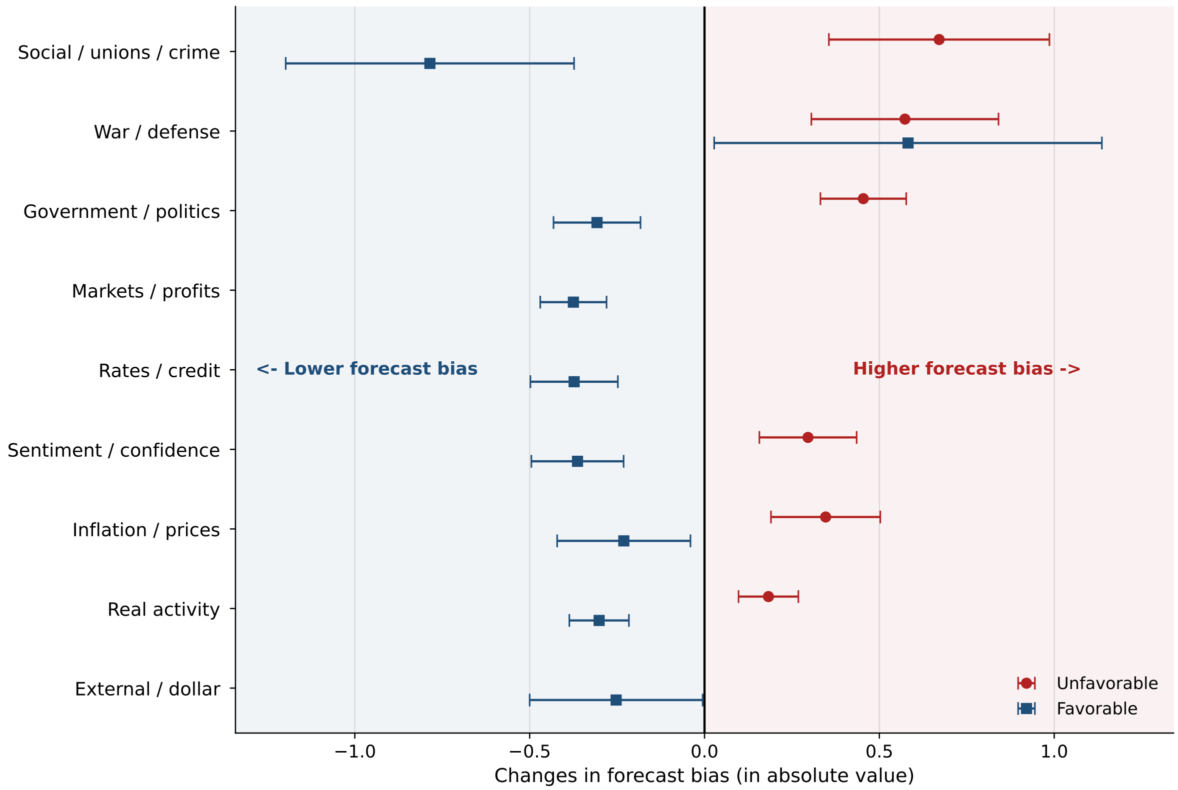

Notes. Each marker reports the estimated effect of adding attention to topic family f with sentiment s on the six-month change in absolute forecast bias, $\widehat{\beta}_{f,s}$, from Equation (3). A positive estimate implies that adding the corresponding topic is associated with a larger deviation from the benchmark, while a negative estimate implies movement toward the benchmark. Error bars are 90% confidence intervals based on two-way clustered standard errors by respondent and interview month. For readability, the figure displays topic-by-sentiment categories with move frequency $\mathbb{E}[\text{Add}^{f,s} + \text{Drop}^{f,s}] \geq 0.001$ and add-coefficient significance $p \leq 0.10$.

$\Delta|b^{ML}|$, holding fixed all other observed attention switches, controls, and fixed effects. This partial interpretation is important for the crowding mechanism: when one narrative is newly reported, another may be dropped, and the joint specification separates these concurrent effects.

Table 4 reports the largest Add and Drop coefficients from this joint specification after restricting the underlying attention set, $\mathcal{R}_{i,t}$, to topic families appearing in at least 1.0% of interviews in the full sample ($\mathbb{E}[\text{Add}^{f,s}] \geq 0.01$). We report the full-sample and dual-response estimates side by side.

The joint specification confirms and sharpens the findings from individual topic estimates. As shown in Table 4, nearly all topics associated with an increase in forecast bias

Table 4: Topic Attention Allocation and Forecast Bias (Joint Specification)

			Full sample (Baseline)		Dual-response sample	
Rank	Topic family	Sent.	Freq.	Est.	Freq.	Est.
Panel A: Largest increases in $\Delta\lvert b^{ML}\rvert$ (ADD only)						
1	Social / unions / crime	UF	0.012	0.622***	0.023	0.510**
				(0.182)		(0.239)
2	War / defense	UF	0.011	0.440***	0.020	0.376
				(0.160)		(0.229)
3	Government / politics	UF	0.065	0.347***	0.109	0.328***
				(0.068)		(0.103)
4	Inflation / prices	UF	0.048	0.255***	0.083	0.476***
				(0.089)		(0.170)
5	Sentiment / confidence	UF	0.040	0.251***	0.054	0.389***
				(0.080)		(0.145)
6	Housing	UF	0.014	0.171	0.026	0.300
				(0.136)		(0.239)
7	Real activity	UF	0.150	0.104**	0.175	0.044
				(0.045)		(0.096)
8	External / dollar	UF	0.024	0.005	0.048	0.089
				(0.088)		(0.123)
Panel B: Largest decreases in $\Delta\lvert b^{ML}\rvert$ (ADD only)						
1	Sentiment / confidence	F	0.037	-0.320***	0.055	-0.366***
				(0.075)		(0.126)
2	Rates / credit	F	0.024	-0.316***	0.043	-0.229*
				(0.074)		(0.120)
3	Markets / profits	F	0.042	-0.305***	0.080	-0.261***
				(0.052)		(0.095)
4	Government / politics	F	0.037	-0.286***	0.069	-0.329***
				(0.070)		(0.124)
5	Real activity	F	0.118	-0.243***	0.167	-0.138
				(0.047)		(0.101)
6	Inflation / prices	F	0.012	-0.190*	0.025	-0.085
				(0.110)		(0.155)
7	Rates / credit	UF	0.022	-0.028	0.041	-0.095
				(0.096)		(0.154)
8	Markets / profits	UF	0.044	-0.027	0.068	-0.100
				(0.076)		(0.140)

Notes. The table reports the largest estimated coefficients from a joint specification that includes Add/Drop indicators for each topic-family×sentiment category between matched interviews. The row is sorted by the magnitude of the coefficients (in absolute value). Panel A lists the Add categories associated with the largest increases in the change in absolute forecast bias, $\Delta|b^{ML}|$. Panel B lists the largest decreases. "Freq." is the sample frequency of the corresponding Add indicator. "Full sample (Baseline)" uses the full sample. "Dual-response sample" restricts to pairs where both mentions are observed in both interviews. All specifications include the baseline controls and fixed effects from Table 2. Standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

are unfavorable mentions, while nearly all topics associated with a decrease in bias are favorable mentions. In Panel A, the largest bias increases stem from adding unfavorable news about social issues, unions, and crime (0.622), war and defense (0.440), and government and politics (0.347).[14] Unfavorable inflation and price news also shows a positive coefficient of 0.255. Because these effects span multiple topic families, the bias-increasing role of unfavorable attention is not confined to a single category. The largest coefficients are not inflation-specific. Instead, social and political narratives dominate in magnitude, pointing to a broad unfavorable-sentiment channel.

Panel B shows symmetric bias-reducing effects for favorable topics: −0.320 for sentiment and confidence, −0.316 for rates and credit, and −0.305 for markets and profits. As shown in Table 1, these magnitudes are large relative to the average six-month change in absolute bias. Together, Panels A and B identify the specific bias-reducing narratives that are crowded out of households' limited attention sets when unfavorable attention dominates. Bias rises most when households shift attention toward unfavorable topics rather than drawing favorable narratives concerning sentiment, interest rates, or market performance.

The dual-response columns align closely with these baseline findings. The Add-UF coefficient for inflation and prices increases to 0.476 and remains statistically significant, while the Add-F coefficient for sentiment and confidence is −0.366 and significant. More broadly, the conditional asymmetry documented in Section 3.1 extends beyond inflation-specific narratives. Both sentiment and topic matter: the same topic family can produce substantially different coefficients depending on whether the attention shift is favorable or unfavorable. Because these coefficients are jointly estimated, they capture the incremental role of each news topic. As shown in Internet Appendix Table A3, the same ranking is robust in sim-

[14] A related concern is that the unfavorable government and politics coefficient reflects partisan differences in inflation expectations. Existing work shows that inflation expectations differ systematically by political affiliation and by whether respondents are politically aligned with the party in power (Gillitzer et al., 2021; Binder et al., 2024). If those partisan differences also affect the propensity to report unfavorable political news, they could mechanically inflate the coefficient on government and politics. We address this concern in two ways. First, the demographic fixed effects $\psi_{g(i)}$ include political affiliation, absorbing time-invariant partisan differences in forecast bias. Second, we estimate a specification that adds an interaction between political affiliation and out-of-power status. The interaction term is not robust, while the main government and politics coefficient remains stable, suggesting that the baseline result is not driven by partisan differences.

pler family-by-family matched-pair regressions using the exact same baseline pipeline: broad non-price narratives such as social conditions, war, and politics generate larger bias increases than inflation news itself, while favorable sentiment, rates, and market narratives are more bias-reducing than favorable inflation news.

3.3.3. Response Order and Additional Robustness

Chahrour et al. (2025) find that first and second MSC responses can carry different information for expectation formation. We find that the attention effect on expectation bias is not driven by response order. As shown in Panel B of Table 2, the Add UF coefficients are similar whether the unfavorable topic appears first or second in the interview, and the hypothesis of equal first- and second-response coefficients cannot be rejected. As a result, what matters is the content of the reported attention set, not where the topic appears in the response sequence. Internet Appendix Section C.2 further shows the full response-order estimates and confirms that the effects on forecast bias do not hinge on response order.

The results are also robust to a battery of specification checks reported in the Internet Appendix. Table A4 shows that the Add-UF and Add-F coefficients are nearly unchanged when we control for current mention counts or changes in mention counts between interviews. Generic adding or dropping of any mention, regardless of sentiment, has little explanatory power. This indicates that forecast bias responds to attention allocation across narratives rather than to reporting more news. Table A7 shows stability across sample restrictions, weights, and control sets, and Table A8 shows that the asymmetry remains highly robust across inflation regimes and demographic groups.

The asymmetry also survives direct controls for within-household shifts in broad sentiment. Internet Appendix Table A5 adds changes in the Index of Consumer Sentiment and in personal financial outlook to the baseline specification. The Add-UF and Add-F coefficients remain economically large and precisely estimated, which is difficult to reconcile with a pure recall story in which households simply report more unfavorable news when they become generally pessimistic.

The baseline results cannot be reproduced when the information content of attention allocation is removed. Internet Appendix Table A6 further shows that randomly reassigning sentiment labels and replacing current switches with future-switch placebos produce no significant effects, ruling out mechanical or spurious explanations for the Add-UF and Add-F estimates.

4. The Attention Channel

Section 3 shows that attention composition and forecast bias move together within households. A natural concern is that both respond to common underlying conditions: a household experiencing a sharp rise in gasoline or grocery prices may simultaneously revise inflation expectations upward and start to pay attention toward unfavorable economic news, generating a spurious correlation between Add UF and $\Delta|b^{ML}|$ that is unrelated to the attention channel. We address this using an external news-shock exposure $Q_{i,t}$, which combines a common aggregate news-sentiment innovation with predetermined heterogeneity in household exposure to that innovation.

More specifically, for each household i, news exposure $Q_{i,t} = \xi_i \times \text{shock}_t$ compares households interviewed in the same month who differ in their predetermined sensitivity to the news environment. Because shock_t is common to all households within a month, variation in $Q_{i,t}$ identifies differences in ξ_i that were fixed before the current shock realized. The aggregate component shock_t is a one-year AR residual from the Daily News Sentiment Index (DNSI) of Shapiro et al. (2022), a daily lexical measure of U.S. economic news tone constructed from Factiva newspaper coverage, where higher values indicate more positive sentiment.[15] The household component ξ_i is a predetermined exposure weight for unfavorable non-price news, estimated from fixed demographics over the first ten years of the sample and held fixed thereafter.[16] Figure 3 summarizes this identification strategy and motivates the reduced-form pass-through tests in the next section.

[15] We aggregate the daily DNSI to the monthly level by averaging over days in the month preceding the MSC interview deadline.

[16] We fit a logit for an indicator equal to one if the respondent reports any unfavorable non-price mention at the first interview in the linked pair, using age group, education group, income tercile, sex, Census region, and political affiliation. The exposure weight is clipped to $[0.01, 0.99]$.

Figure 3: Schematic illustration of the attention channel

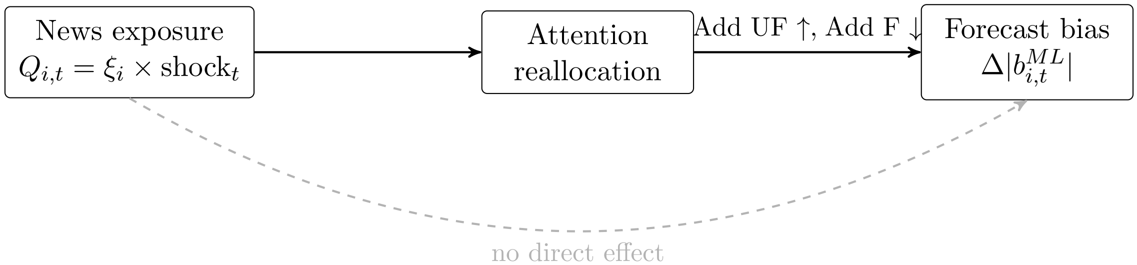

Notes. The figure summarizes the empirical design in Section 4. News exposure $Q_{i,t} = \xi_i \times \text{shock}_t$ shifts attention composition (solid arrows), and those shifts move forecast bias in the direction of the Add asymmetry. The dashed arc summarizes the null result in Table 5: we find no significant direct effect of $Q_{i,t}$ on bias in the subsample where reported attention is unchanged between interviews.

4.1. Preliminary Reduced-Form Evidence

We start with reduced-form evidence. Table 5 reports estimates on the pass-through from news-shock exposure to attention and forecast bias. Columns (1) and (2) show that households with higher $Q_{i,t}$, that is, greater exposure to positive news-sentiment innovations, are less likely to add an unfavorable narrative and more likely to add a favorable one. Column (3) shows that the same exposure measure is associated with a decline in absolute forecast bias over the six-month interview window.

The key comparison comes in Columns (4) and (5), which restrict attention to households whose reported attention composition is *unchanged* between interviews. Among these interview pairs, news exposure has no significant effect on forecast bias, even in months with large DNSI innovations. This comparison matters because a substantial literature shows that aggregate news sentiment shapes inflation expectations (Carroll, 2003; Larsen et al., 2021; Binder et al., 2025; Chahrour et al., 2025). As shown in Table 5, the pass-through to forecast bias appears to operate only when news enters households' reported attention sets.

4.2. Interaction-IV: Ambient News versus Attention Reallocation

The reduced-form evidence suggests that the same news-sentiment exposure measure shifts both attention composition and forecast bias. We next use this variation to estimate the effect of attention reallocations induced by these news-sentiment shocks through a Bartik-style 2SLS specification (Bartik, 1991; Goldsmith-Pinkham et al., 2020). More specifically,

Table 5: Preliminary Reduced-Form Evidence on News-Sentiment Pass-Through

Sample / outcome	(1) Add UF Full sample	(2) Add F Full sample	(3) $\Delta\lvert b^{ML}\rvert$ Full sample	(4) $\Delta\lvert b^{ML}\rvert$ No realloc.	(5) $\Delta\lvert b^{ML}\rvert$ No realloc., top 10%
$Q_{i,t}$ (DNSI × exposure)	-0.125***	0.116***	-0.455*	-0.779	-0.570
	(0.031)	(0.029)	(0.238)	(0.508)	(0.841)
Observations	40,362	40,362	40,362	10,225	946
R^2	0.314	0.226	0.368	0.381	0.484
Month FE	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes	Yes
WT Weights	Yes	Yes	Yes	Yes	Yes

Notes. All columns use the matched-pair MSC sample with the same baseline controls, month fixed effects, demographic fixed effects, and survey weights as Table 2. News exposure is $Q_{i,t} = \xi_i \times \text{shock}_t$, where shock_t is the monthly DNSI innovation and ξ_i is the predetermined household exposure weight. Columns (1)–(2) report reduced-form regressions for the probability of adding any unfavorable or favorable narrative. Column (3) uses $\Delta|b^{ML}_{i,t}|$ in the full sample. Columns (4)–(5) use $\Delta|b^{ML}_{i,t}|$ in the subsample with no Add or Drop switch of any kind, with Column (5) further restricting to months in the top decile of absolute shock realizations. Standard errors are two-way clustered by respondent and interview month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

for each endogenous variable $D_{i,t} \in \{\text{Add UF}, \text{Add F}\}$, we estimate:

$$D_{i,t} = \pi\, Q_{i,t} + \Gamma' X_{i,t-1} + \tau_t + \psi_{g(i)} + \nu_{i,t}, \tag{5}$$

$$\Delta|b^{ML}_{i,t}| = \beta\, \widehat{D}_{i,t} + \Gamma' X_{i,t-1} + \tau_t + \psi_{g(i)} + \varepsilon_{i,t}, \tag{6}$$

where $X_{i,t-1}$ contains the baseline controls from the main specification, τ_t are interview-month fixed effects, and $\psi_{g(i)}$ are demographic fixed effects. Standard errors are two-way clustered by respondent and interview month.

Table 6 reports the 2SLS estimates. Panel A uses DNSI as the news-sentiment shock and compares two definitions of the endogenous attention reallocation: Column (1) restricts Add UF and Add F to non-price topics, while Column (2) uses all topic families. The first stage is strong in both columns, with F-statistics of 16 to 18 for both Add UF and Add F. The second-stage estimates are nearly identical across the two columns: instrumented Add UF raises $\Delta|b^{ML}|$ and instrumented Add F lowers it, with coefficient magnitudes close across the non-price and all-topic specifications. The similarity across columns suggests that attention reallocations induced by news-sentiment shocks have sizable effects on forecast bias, and that the baseline asymmetry does not depend only on inflation or price narratives.

A concern is that shock_t proxies for macro conditions rather than the narrative environment. Panel B of Table 6 addresses this by replacing shock$_t$ with AR residuals from four FRED-MD series (McCracken and Ng, 2016): CPI inflation, unemployment, the federal funds rate, and the oil price. These macro placebos do not reproduce the DNSI result. Some deliver a moderate first stage for Add UF, but none delivers strong first-stage relevance for both Add UF and Add F. The second-stage coefficients are also unstable in sign and magnitude across the macro series.[17] This contrast suggests that the DNSI instrument is picking up variation in the narrative environment rather than generic macro sensitivity. Internet Appendix Table A10 further shows that the 2SLS estimates are nearly unchanged when demographic groups are allowed to follow separate linear time trends.

Panel C of Table 6 reports reduced-form exclusion checks in the subsample where reported attention is unchanged between interviews. In those columns, exposure to DNSI innovations has no significant effect on forecast bias, including in months with the largest news-sentiment shocks. This result supports the interpretation that the aggregate news shock matters for forecast bias when it changes reported attention composition, as illustrated in Figure 3.

4.3. News-Level Evidence: EPU Spikes and Attention Composition

The IV evidence uses aggregate sentiment shocks and this section provides complementary narrative-level evidence using category-specific economic policy uncertainty (EPU) spikes that are externally identified rather than constructed from MSC or DNSI data. We match the six categorical EPU indices of Baker et al. (2016) to the closest MSC topic family,[18] identify spike months as AR residual exceedances above the 90th percentile with a minimum nine-month separation, and estimate a local-linear RD in calendar time (Imbens and Lemieux, 2008). The running variable r_{it} is the distance in months between the MSC reinterview and the nearest spike, so restricting to $|r_{it}| \leq \ell$ compares pairs reinterviewed just before versus

[17] This result holds across a grid of 132 macro-placebo specifications using eleven FRED-MD series spanning prices, real activity, labor markets, interest rates, and financial conditions.

[18] Monetary policy uncertainty to *rates/credit*, taxes and government spending uncertainty to *government/politics* (treated as separate categories with distinct spike episodes), national security uncertainty to *war/defense*, regulation uncertainty to *markets/profits*, trade policy uncertainty to *external/dollar*.

Table 6: News Narrative versus Economic Fundamentals: 2SLS Estimates

	Panel A: News Sentiment		Panel B: Macro Placebos			Panel C: Exclusion	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Estimator	2SLS	2SLS	2SLS	2SLS	2SLS	OLS	OLS
Instrument $Q_{i,t}$	DNSI	DNSI	CPI	FFR	OIL	DNSI	DNSI
Endog.	Non-price	All	Non-price	Non-price	Non-price	No Add	Top 10%
AddUF	2.245*	2.065*	-2.662	1.473	-9.006	–	–
	(1.337)	(1.137)	(4.743)	(1.088)	(26.247)		
AddF	-2.397*	-2.374*	4.568	-2.243	8.512	–	–
	(1.416)	(1.394)	(9.535)	(1.823)	(20.606)		
$\Delta\lvert b^{ML}\rvert$ on Q						-0.375	-0.187
						(0.373)	(0.482)
Observations	40,362	40,362	40,362	40,362	40,362	19,720	1,928
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
WT Weights	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FS F (AddUF)	18.43	16.40	9.62	14.01	5.89	–	–
FS F (AddF)	16.30	15.85	3.71	1.74	1.27	–	–

Notes. Columns (1)–(5) report 2SLS estimates using $Q_{i,t} = \xi_i \times \text{shock}_t$, where ξ_i is the demographic-predicted propensity to report unfavorable non-price news (held fixed after the first interview in the linked pair) and shock_t is a one-year AR innovation from the indicated external series. The row "Endog." indicates how the endogenous attention-reallocation variables (AddUF and AddF) are defined in each 2SLS column: "Non-price" excludes inflation/prices/gas topic families, while "All" uses all topic families. Columns (6)–(7) report reduced-form (OLS) placebos that regress $\Delta|b^{ML}|$ on $Q_{i,t}$ in the subsample with no AddUF/AddF (non-price) reallocations, and column (7) further restricts to months in the top decile of $|\text{shock}_t|$. In these placebo columns, the "Endog." row therefore refers to the sample restriction rather than an endogenous regressor. All columns include baseline controls, month FE, demo FE, WT weights, and two-way clustering by (id, month). Standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

just after the shock. The estimating equation is

$$Y_{it} = \underbrace{\beta\,\mathbb{1}\{r_{it} \geq 0\}}_{\text{jump at EPU spike}} + \rho\, r_{it} + \varrho\,\mathbb{1}\{r_{it} \geq 0\} \cdot r_{it} + X_{it}'\xi + \psi_e + \psi_{\text{demo}} + \varepsilon_{it}, \tag{7}$$

where ψ_e is an event fixed effect, ψ_{demo} is a demographic fixed effect, and X_{it} contains the baseline controls. The coefficient β captures the jump at the spike, estimated separately for the matched MSC topic-family add indicator (Panel A) and for $\Delta|b^{ML}|$ (Panel B) at bandwidths $\ell \in \{3, 6\}$.

Table 7 reports the RD estimates. Panel A shows that categorical EPU innovations coincide with statistically significant shifts in MSC composition toward matched unfavorable topics. At bandwidth $\ell = 6$, jumps are positive and significant for monetary policy, taxes, government spending, national security, and regulation, with magnitudes from roughly one to six percentage points. Trade policy is the consistent exception: estimated jumps are small and insignificant at both bandwidths. This exception is informative because the trade-matched MSC family is the least frequently reported category. If a category rarely appears in reported attention, it cannot be crowded in.

Panel B reports analogous jumps for $\Delta|b^{ML}|$. At $\ell = 6$, spike months in monetary policy (0.313^{***}), government spending (0.393^{***}), national security (0.646^{***}), and regulation (0.375^{***}) are associated with positive and significant increases in forecast bias. Taxes is positive but marginal, and trade remains insignificant.

The two panels align closely: categories that shift MSC composition in Panel A are precisely those that show bias increases in Panel B, while trade, which fails to shift attention, also fails to move bias.[19] Magnitude rankings are broadly consistent across the two panels. This comparison matters because it shows that EPU spikes must reach reported attention sets to raise forecast bias. The categories with the largest effects, politics, national security, and regulation, are also the categories households most frequently report, which makes them more likely to enter attention sets when external uncertainty spikes.

5. Aggregate Implications

What are the aggregate implications of household-level attention allocations? Answering this requires first measuring the aggregate attention environment. Section 4 shows that the ambient news environment does not necessarily move bias on its own: what matters is what actually enters households' reported attention sets. Aggregating the household-level results therefore requires tracking the distribution of reported attention across households rather than the news environment itself. To do so, we first construct a monthly *attention network* from MSC reported mentions, and then ask whether concentration around unfavorable narratives amplifies the sensitivity of bias to attention reallocations, and quantify how much of aggregate forecast bias is accounted for by past and present attention allocations.

[19] Because each household is reinterviewed approximately six months after their initial interview, the measured bias change $\Delta|b^{ML}|$ integrates information over the full inter-survey window rather than responding sharply at the spike month. A "donut" specification for $\ell = 6$ that excludes reinterview months within one month of the spike produces quantitatively similar results.

Table 7: Regression discontinuity around categorical EPU spikes

Domain	Panel A: $1\{\text{AddUF}_{\text{matched}}\}$ (1) $h = 3$	(2) $h = 6$	Domain	Panel B: $\Delta\lvert b^{ML}\rvert$ (3) $h = 3$	(4) $h = 6$
GovSpend (16 spikes)	0.020	0.024^{***}	GovSpend (16 spikes)	0.602^{***}	0.393^{***}
	(0.014)	(0.009)		(0.173)	(0.111)
Monetary (19 spikes)	0.008	0.012^{**}	Monetary (19 spikes)	0.149	0.313^{***}
	(0.007)	(0.005)		(0.158)	(0.105)
NatSec (12 spikes)	0.025^{***}	0.033^{***}	NatSec (12 spikes)	0.571^{***}	0.646^{***}
	(0.007)	(0.005)		(0.189)	(0.140)
Regulation (16 spikes)	0.037^{***}	0.055^{***}	Regulation (16 spikes)	0.458^{**}	0.375^{***}
	(0.010)	(0.007)		(0.187)	(0.119)
Taxes (14 spikes)	−0.002	0.019^{**}	Taxes (14 spikes)	0.098	0.224^{*}
	(0.014)	(0.009)		(0.168)	(0.118)
Trade (11 spikes)	0.021	−0.011	Trade (11 spikes)	−0.034	−0.118
	(0.017)	(0.011)		(0.170)	(0.120)
Observations	7,326	13,427	Observations	6,810	12,800
R^2	0.043	0.037	R^2	0.355	0.342
Event FE	Yes	Yes	Event FE	Yes	Yes
Demographic FE	Yes	Yes	Demographic FE	Yes	Yes
Lagged controls	Yes	Yes	Lagged controls	Yes	Yes
Survey weights	Yes	Yes	Survey weights	Yes	Yes
Respondent clusters	Yes	Yes	Respondent clusters	Yes	Yes

Notes: Local-linear regression discontinuity around categorical EPU spike months, estimated separately by category and stacked across spike episodes. The running variable is the MSC reinterview month minus the spike month, in months, and the reported coefficient is the estimated jump at the cutoff ($r = 0$). Spikes are months in which the AR(12) residual of each categorical EPU series exceeds the 90th percentile, with a minimum nine-month separation between consecutive events. The number of spike episodes per category is in parentheses next to each row label. The sample is restricted to MSC linked pairs with a six-month reinterview gap. Columns (1) and (3) use bandwidth $\ell = 3$, and columns (2) and (4) use $\ell = 6$. Panel A reports results for the indicator for adding a matched-topic unfavorable news item. Panel B reports results for $\Delta\lvert b^{ML}\rvert$. Event fixed effects are category times episode fixed effects, and demographic fixed effects are age × sex × education cells. Lagged controls include lagged absolute bias, lagged inflation expectation, interview gap, and total topic moves. Standard errors clustered by respondent in parentheses. The Observations and R^2 rows report means across the six category-specific regressions. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.1. The Attention Network

To start, we build a monthly attention network from the MSC interviews. Each node is a topic-family × sentiment pair. Two nodes are linked when respondents report both topics in the same interview. Node size is measured by $p_{u,t}$, the survey-weighted mention share in month t, and link strength is measured by $w_{uv,t}$, the survey-weighted share of linked mentions. Recomputing these objects month by month allows the network to track how the aggregate attention structure shifts over time.

Figure 4 illustrates the attention network in two contrasting months (the highest- and lowest-concentration observations in the sample). The left panel shows January 1992, the highest-concentration month, during the sluggish recovery from the 1990 to 1991 recession. Attention concentrates on a single dominant unfavorable "Real Activity" node, creating an *unfavorable focal hub* with a pronounced hub-and-spoke configuration. When one negative narrative accounts for most reported mentions, households are less likely to co-mention the favorable topics in rates, markets, and sentiment that Section 3.1 associates with bias reduction. This is the same crowding mechanism emphasized by Nimark (2014) and Chahrour et al. (2021), where concentration around a dominant topic leaves less room for alternative narratives.

The right panel shows October 2019, the lowest-concentration month in the sample. Trade negotiations, a Federal Reserve rate cut, and the House impeachment inquiry were all active simultaneously, with no single narrative accounting for a dominant share of mentions. Because attention divides across many nodes, respondents are more likely to report favorable and unfavorable topics together, which offsets their net impact on forecast bias.

5.2. Attention Network and Forecast Bias

The household-level evidence implies a testable aggregate prediction: unfavorable narratives should move forecast bias more strongly in months when aggregate attention concentrates on one dominant unfavorable topic. We summarize the monthly attention environment with two variables. $H_t = \text{Hub}_t$ indicates that the most-mentioned node in month t is unfavorable, and Z_t is the standardized share of respondents who report two valid news mentions. Together, these variables distinguish months dominated by one unfavorable narrative from months in which attention is more diffuse.

Using H_t and Z_t, we write forecast bias, $b^{ML}_{i,t} \equiv \pi^e_{i,t} - \widehat{\pi}^{ML}_t$, as the sum of a persistence term and a contemporaneous attention term:

$$b^{ML}_{i,t} \;=\; \gamma\, b^{ML}_{i,t-1} \;+\; \kappa_t \sum_{u \in \mathcal{A}_{i,t}} \mu_{s_u,t}\, \lambda_u \;+\; \varepsilon_{i,t}, \qquad \gamma \in [0,1), \tag{8}$$

where $\mathcal{A}_{i,t}$ is the set of topics household i reports at interview t, $\lambda_u = 1/|\mathcal{A}_{i,t}|$ is the equal

Figure 4: Attention Network: High vs. Low Concentration

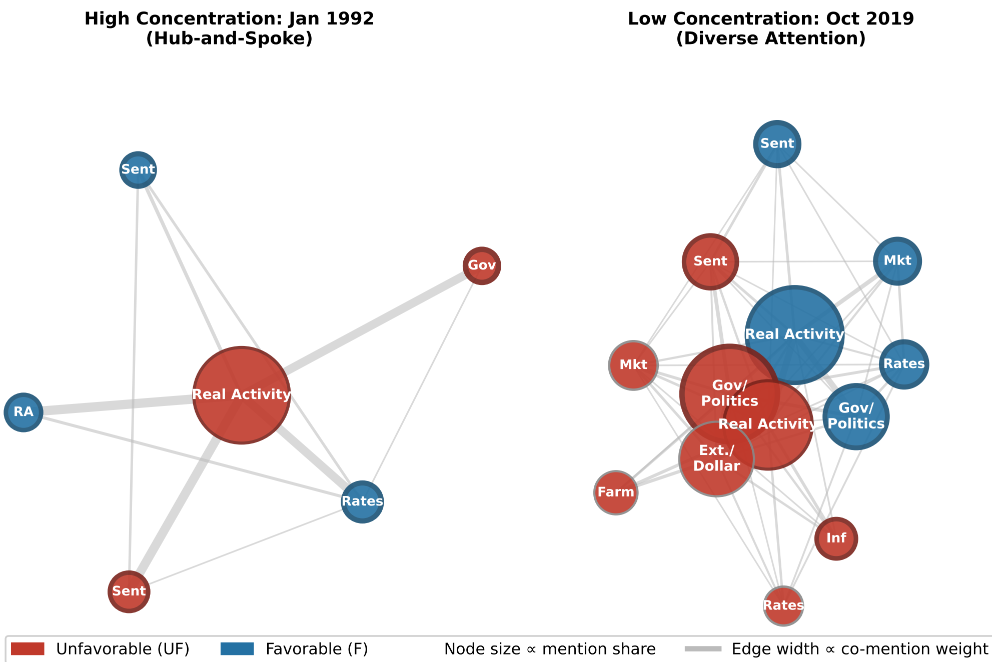

Notes. Each node is a topic-family × sentiment pair, red nodes are unfavorable (UF) and blue nodes are favorable (F). Node size is proportional to weighted mention share. Node position reflects the strength of links between topics, with stronger links drawing nodes closer. Line thickness is proportional to the aggregated survey weight on those links. Nodes with weighted mention share below 2% of total monthly mentions are omitted. Left panel: January 1992 (highest share of the most-mentioned topic). Right panel: October 2019 (lowest concentration among months with at least 500 weighted mentions).

attention weight on each reported topic, and $\varepsilon_{i,t}$ is mean-zero noise. The distortion weights $\mu_{s,t}$ capture how the monthly network state modulates the effect of each sentiment:

$$\mu_{UF,t} = \mu_{UF} + \omega_{UF,H} H_t + \omega_{UF,Z} Z_t, \qquad \mu_{F,t} = \mu_F + \omega_{F,H} H_t + \omega_{F,Z} Z_t, \qquad \mu_{UF} > 0, \quad \mu_F \leq 0. \tag{9}$$

When aggregate attention concentrates around an unfavorable hub ($H_t = 1$), unfavorable narratives receive more weight and favorable narratives are less likely to offset them. The parameter κ_t scales up the attention term to account for the at-most-two-mention ceiling in the MSC, so that the decomposition reflects the full attention effect rather than only the reported share.

The key prediction is that unfavorable attention raises forecast bias more in months with an unfavorable focal hub. Table 8 tests this prediction by interacting the baseline Add UF and Add F indicators with H_t. The interaction between the hub indicator and Add UF is positive and remains significant after controlling for aggregate uncertainty and benchmark volatility. These estimates suggest that forecast bias is more sensitive to unfavorable attention when the attention network is organized around a dominant unfavorable narrative. By contrast, the corresponding interaction for favorable attention is not robust once those controls are included.

This evidence also disciplines the decomposition that follows. The baseline Add UF and Add F coefficients, together with their interactions with H_t and Z_t, pin down the sentiment weights and their state dependence. The persistence parameter γ is estimated from a separate fixed-effects regression of current forecast bias on lagged forecast bias. The month-specific correction κ_t is pinned down from the observed shares of interviews with zero, one, and two valid news mentions. These inputs allow the decomposition to separate inherited bias from the effect of current attention allocation in each month. Internet Appendix Section E.1 reports the calibration details.

Table 8: Attention Network Concentration and Forecast Bias

	(1) Base	(2) +Aggregate EPU	(3) $+\lvert\Delta\widehat{\pi}_t^{ML}\rvert$	(4) Both
Add UF × Hub	0.440***	0.367***	0.364***	0.279**
	(0.108)	(0.125)	(0.106)	(0.124)
Add F × Hub	0.115*	0.073	0.094	0.049
	(0.068)	(0.068)	(0.073)	(0.072)
Observations	40,362	40,362	40,362	40,362
R^2	0.376	0.376	0.376	0.376
Baseline controls	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes
Two-way clusters (id, month)	Yes	Yes	Yes	Yes

Notes. The dependent variable is $\Delta|b_{i,t}^{ML}|$ between matched interviews. The key coefficient is Add $\text{UF}_{i,t} \times \text{Hub}_t$, where Hub_t equals one when the top mentioned node in month t is unfavorable. Column (1) reports the base specification. Column (2) adds interactions with aggregate EPU (standardized). Column (3) adds interactions with $|\Delta\widehat{\pi}_t^{ML}|$. Column (4) includes both. All columns include baseline respondent controls, month fixed effects, and demographic fixed effects. Standard errors are two-way clustered by respondent and month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

We next examine whether position in the monthly attention network helps explain which households experience larger increases in forecast bias within the same month. Internet Appendix Table A11 reports these estimates using $\text{DistUF}_{i,t}$, the minimum shortest-path distance from any reported topic to the nearest unfavorable topic active in month t, where distance is measured by the number of links connecting two topics in the monthly attention network. Larger values therefore indicate that the reported topics lie farther from the unfavorable hub in the attention network. Column (1) shows that each one-step reduction in distance to the unfavorable hub is associated with a 0.381 percentage-point increase in absolute forecast bias. This result continues to hold after conditioning on the specific topics households report. It therefore shows that the same attention network that amplifies aggregate bias also helps explain which households experience larger increases within a given month.

5.3. Bias Decompositions

We next quantify how much observed aggregate forecast bias is associated with persistence and contemporaneous attention, and how those components vary over time. More specifically, we implement the following decomposition:

$$P_{i,t} \;\equiv\; \hat{\gamma}\, b^{ML}_{i,t-1}, \qquad A_{i,t} \;\equiv\; \hat{\kappa}_t \sum_{u \in \mathcal{A}_{i,t}} \hat{\mu}_{s_u,t}\, \lambda_u, \qquad R_{i,t} \;\equiv\; b^{ML}_{i,t} - P_{i,t} - A_{i,t},$$

where P is persistence, A is contemporaneous attention, and R is a residual. We allocate contributions to $|b^{ML}_{i,t}|$ by averaging over all orderings of $\{P, A, R\}$ following the Shapley-Owen-Shorrocks decomposition as detailed in Audoly et al. (2025).

Figure 5 reports the decomposition results. Panel (a) shows the full sample and Panel (b) shows the dual-response sample. In the full sample, persistence accounts for 42 percent of absolute forecast bias on average, contemporaneous attention for 16 percent, and the residual for 42 percent. In the dual-response sample, the attention share rises to 30 percent and the residual falls to 31 percent, while persistence remains close to 38 percent.

The higher attention share among dual-response households is consistent with more complete measurement of attention composition when both reported mentions are observed.

Together, persistence and contemporaneous attention account for 58 percent of observed absolute forecast bias in the full sample and about 68 percent in the dual-response sample. The persistence component is itself largely inherited from earlier attention draws: iterating equation (8) backward shows that $P_{i,t}$ is a geometrically weighted sum of past attention components.

These averages also mask substantial time variation. The attention component is small for much of the pre-2005 period, rises in recession months, and surges again at the onset of the COVID-19 pandemic. Table A12 in the Internet Appendix shows that, in the full sample, the attention share reaches 31 percent in NBER recession months, averages 21 percent over 2005 to 2020, and rises to 48 percent post-2020, compared with about 1 percent before 2005.

Part of the low pre-2005 attention share likely reflects measurement. In those years a larger fraction of interviews contained no valid news response, which directly lowers the measured attention component. A small decomposition share in that period therefore does not imply that the attention channel was absent. As reported in Table A8, the Add-UF and Add-F coefficients in the pre-2005 subsample retain their expected signs and remain economically similar in magnitude. The reduced-form evidence therefore points to the same asymmetry before 2005, even though the decomposition assigns a smaller role to attention in those years.

The rise in the attention share post-2020 may appear to conflict with the substantial expansion of media outlets since the 1990s. If more sources are available, households' limited attention sets might be expected to contain more diverse news rather than more concentrated coverage. More outlets need not produce more diverse reported attention, however: when a sufficiently large event occurs, coverage concentrates simultaneously across outlets, making the focal hub more pronounced rather than less. Each of the major events underlying the post-2020 surge in the attention component (the pandemic, the subsequent inflation episode, and escalating geopolitical tensions) generated a dominant unfavorable narrative that crowded out alternative topics regardless of the number of outlets available.

The time variation is strongest in months when the attention network has an unfavorable hub and reporting intensity is high. In those months, the attention component becomes much larger, as in (9). In lower-concentration periods, such as the mid-2000s and post-GFC expansions, the attention term is much smaller. Months in which attention concentrates around an unfavorable hub therefore account for a disproportionate share of aggregate bias fluctuations, especially during recessions and other large negative-news events.

Figure 5: Decomposition of Forecast Bias: Full Sample and Dual-Response Sample

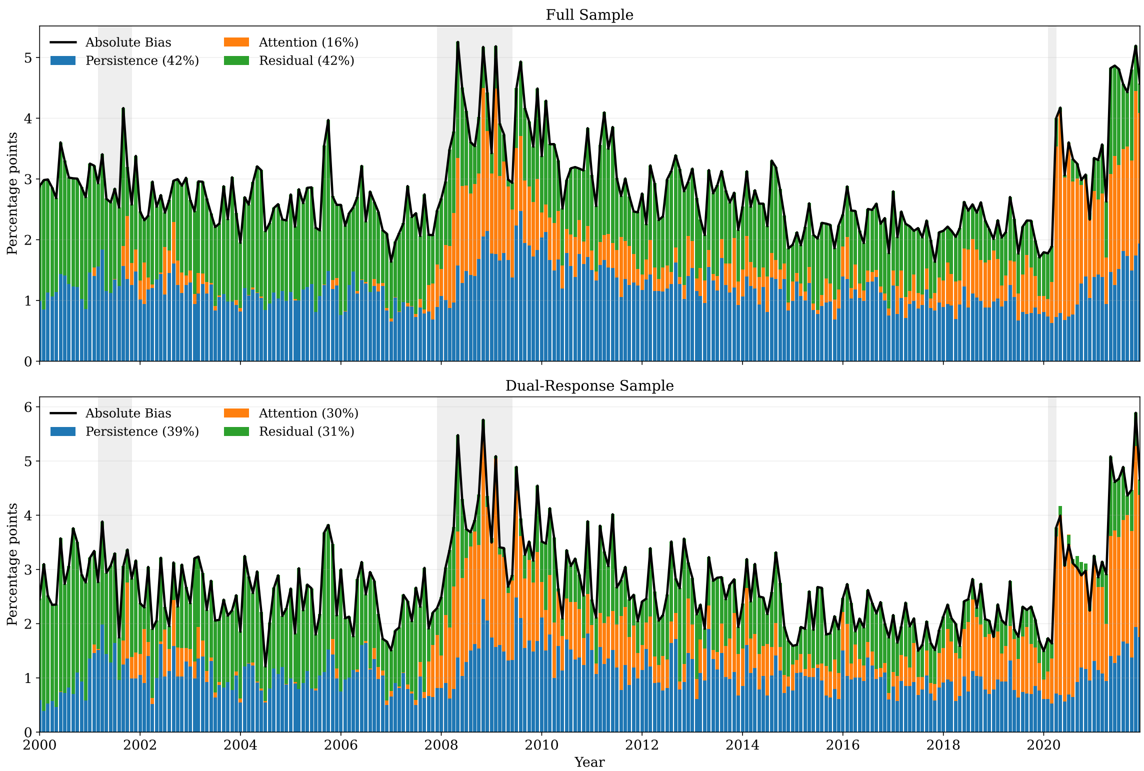

Notes. Each panel shows the monthly decomposition of aggregate absolute bias into persistence (P_t, blue), attention (A_t, orange), and residual (R_t, gray) via the Shapley-Owen-Shorrocks decomposition. The three components sum to aggregate absolute bias in each month by construction. Panel (a): full matched-pair sample (N = 40,362). Panel (b): respondents who reported two valid news mentions at the current interview (N = 16,821). Baseline: negative-binomial latent-K with $\hat{\alpha} = 0.75$, two-state attention loads $H_t = \text{Hub}_t$ and normalized intensity measure Z_t, and month-specific κ_t. NBER recessions are shaded. Average shares over the full sample are reported in the figure legend.

6. Conclusion

This paper studies how within-household attention reallocations affect inflation forecast bias. Using linked MSC interviews and a real-time machine-learning benchmark as a disciplined full-information anchor, we show that adding unfavorable news raises forecast bias, while adding favorable news lowers it by a similar amount. Unfavorable attention generates overreaction while attention to favorable news produces underreaction. The largest bias increases come not from inflation news itself, but from unfavorable attention to social, political, and geopolitical narratives.

Among households whose reported attention composition is unchanged between interviews, aggregate news sentiment has no effect on forecast bias, even in months with large news-sentiment innovations. Similarly, EPU spikes that shift household attention toward matched unfavorable topics also raise bias, while categories that rarely enter reported attention sets move neither attention nor bias. Ambient news that does not alter reported attention composition leaves forecast bias unchanged.

The aggregate evidence points in the same direction. Months with an unfavorable focal hub amplify the effect of unfavorable attention on bias, and the decomposition indicates that attention and its persistence account for a substantial share of aggregate forecast bias. In the full sample, these two components account for 58 percent of average absolute bias, and in the dual-response sample they account for up to 70 percent. These magnitudes suggest that household forecast bias captures economically meaningful information-processing distortions rather than arbitrary survey noise.

One interpretation of the Add/Drop asymmetry is that it reflects optimal attention allocation under an asymmetric loss function: risk-averse households whose disutility from unexpectedly high inflation exceeds the cost of caution will rationally weight unfavorable signals more heavily, producing upward bias relative to the full-information benchmark. Our 2SLS estimates, however, show that attention reallocations induced by exogenous news-sentiment shocks produce bias increases of similar magnitude to the OLS estimates, a finding difficult to reconcile with a fully rational account: if the asymmetry were explained entirely by precautionary risk adjustment, exogenously induced attention shifts would be expected to produce smaller pass-through to forecast bias. We leave formal welfare analysis of attention allocation under asymmetric loss to future work.

Finally, our results add a timing dimension to the literature on central bank communication (Coibion et al., 2022; Weber et al., 2025; Kim and Binder, 2023): pass-through to inflation expectations may depend not only on the content of communication, but also on *when* it is delivered. In months dominated by an unfavorable focal hub, favorable bias-reducing narratives are crowded out of households' limited attention sets. Central bank communication that must compete with dominant unfavorable narratives for household attention may therefore pass through to inflation expectations less effectively.

References

Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar, "The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment," *The Review of Economics and Statistics*, 2016, *98* (3), 503–523.

Audoly, Richard, Rory McGee, Sergio Ocampo-Diaz, and Gonzalo Paz-Pardo, "A Practitioner's Note on the Shapley-Owen-Shorrocks Decomposition," *FRB of New York Staff Report*, 2025, (1163).

Baker, Scott R, Nicholas Bloom, and Steven J Davis, "Measuring economic policy uncertainty," *The quarterly journal of economics*, 2016, *131* (4), 1593–1636.

Bartik, Timothy J., *Who Benefits from State and Local Economic Development Policies?*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 1991.

Bianchi, Francesco, Do Q. Lee, Sydney C. Ludvigson, and Sai Ma, "The Prestakes of Stock Market Investing," Working Paper 34420, National Bureau of Economic Research 2025.

_ , Sydney C. Ludvigson, and Sai Ma, "Belief Distortions and Macroeconomic Fluctuations," *American Economic Review*, 2022, *112* (7), 2269–2315.

_ , _ , and _ , "What Hundreds of Economic News Events Say About Belief Overreaction in the Stock Market," Working Paper 32301, National Bureau of Economic Research 2024.

Binder, Carola, Pascal Frank, and Jane M. Ryngaert, "The Causal Effect of News on Inflation Expectations," Working Paper 34088, National Bureau of Economic Research August 2025.

__ , Rupal Kamdar, and Jane M. Ryngaert, "Partisan Expectations and COVID-Era Inflation," *Journal of Monetary Economics*, 2024, *148.*

Bordalo, Pedro, John J. Conlon, Nicola Gennaioli, Spencer Yongwook Kwon, and Andrei Shleifer, "Memory and Probability," *Quarterly Journal of Economics*, 2023, *138* (1), 265–311.

__ , Nicola Gennaioli, and Andrei Shleifer, "Diagnostic Expectations and Credit Cycles," *Journal of Finance*, 2018, *73* (1), 199–227.

__ , __ , and __ , "Overreaction and Diagnostic Expectations in Macroeconomics," *Journal of Economic Perspectives*, 2022, *36* (3), 223–244.

__ , __ , Yueran Ma, and Andrei Shleifer, "Overreaction in Macroeconomic Expectations," *American Economic Review*, 2020, *110* (9), 2748–2782.

Bracha, Anat and Jenny Tang, "Inflation Levels and (In)Attention," *The Review of Economic Studies*, 2025, *92* (3), 1564–1594.

Caldara, Dario, Sarah Conlisk, Matteo Iacoviello, and Maddie Penn, "Do Geopolitical Risks Raise or Lower Inflation?," *Journal of International Economics*, 2026, *159*, 104188.

Carroll, Christopher D., "Macroeconomic Expectations of Households and Professional Forecasters," *The Quarterly Journal of Economics*, 2003, *118* (1), 269–298.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia, "Learning from Potentially Biased Statistics," *Brookings Papers on Economic Activity*, 2016, *52* (1), 59–108. Spring.

Chahrour, Ryan, Adam Hale Shapiro, and Daniel J. Wilson, "News Selection and Household Inflation Expectations," Working Paper 33837, National Bureau of Economic Research May 2025.

_ , Kristoffer Nimark, and Stefan Pitschner, "Sectoral Media Focus and Aggregate Fluctuations," *American Economic Review*, 2021, *111* (12), 3872–3922.

Coibion, Olivier and Yuriy Gorodnichenko, "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 2015, *105* (8), 2644–2678.

_ , _ , and Michael Weber, "Monetary Policy Communications and Their Effects on Household Inflation Expectations," *Journal of Political Economy*, 2022, *130* (6), 1537–1584.

_ , _ , and Rupal Kamdar, "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature*, 2018, *56* (4), 1447–1491.

_ , _ , Dimitris Georgarakos, and Geoff Kenny, "The Impact of Geopolitical Risk on Consumer Expectations and Spending," Working Paper 34195, National Bureau of Economic Research 2025.

Correa, Ricardo, Ai He, Christoph Herpfer, and Ugur Lel, "The Rising Tide Lifts Some Interest Rates: Climate Change, Natural Disasters and Loan Pricing," *Journal of Finance*, 2025. forthcoming.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber, "Exposure to Grocery Prices and Inflation Expectations," *Journal of Political Economy*, 2021, *129* (5), 1615–1639.

Gennaioli, Nicola and Andrei Shleifer, "What Comes to Mind," *The Quarterly Journal of Economics*, 2010, *125* (4), 1399–1433.

— , **Marta Leva, Raphael Schoenle, and Andrei Shleifer**, "How Inflation Expectations De-Anchor: The Role of Selective Memory Cues," Working Paper 32633, National Bureau of Economic Research 2024.

Gentzkow, Matthew and Jesse M. Shapiro, "Media Bias and Reputation," *Journal of Political Economy*, 2006, *114* (2), 280–316.

Gillitzer, Christian, Nalini Prasad, and Tim Robinson, "Political Attitudes and Inflation Expectations: Evidence and Implications," *Journal of Money, Credit and Banking*, 2021, *53* (4), 605–634.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Hamilton, James D., "Why You Should Never Use the Hodrick-Prescott Filter," *Review of Economics and Statistics*, 2018, *100* (5), 831–843.

Imbens, Guido W. and Thomas Lemieux, "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 2008, *142* (2), 615–635.

Kim, Gwangmin and Carola Binder, "Learning-through-Survey in Inflation Expectations," *American Economic Journal: Macroeconomics*, 2023, *15* (2), 254–278.

Kohlhas, Alexandre N. and Ansgar Walther, "Asymmetric Attention," *American Economic Review*, 2021, *111* (9), 2879–2925.

Larsen, Vegard H., Leif Anders Thorsrud, and Julia Zhulanova, "News-driven Inflation Expectations and Information Rigidities," *Journal of Monetary Economics*, 2021, *117*, 507–520.

Link, Sebastian, Andreas Peichl, Christopher Roth, and Johannes Wohlfart, "Attention to the Macroeconomy," Discussion Paper 18699, CEPR December 2023.

Ludvigson, Sydney C. and Serena Ng, "The Empirical Risk-Return Relation: A Factor Analysis Approach," *Journal of Financial Economics*, 2007, *83* (1), 171–222.

Maćkowiak, Bartosz, Filip Matejka, and Mirko Wiederholt, "Rational Inattention: A Review," *Journal of Economic Literature*, 2023, *61* (1), 226–273.

Malmendier, Ulrike and Stefan Nagel, "Learning from Inflation Experiences," *The Quarterly Journal of Economics*, 2016, *131* (1), 53–87.

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

McCracken, Michael W. and Serena Ng, "FRED-MD: A Monthly Database for Macroeconomic Research," *Journal of Business & Economic Statistics*, 2016, *34* (4), 574–589.

Mincer, Jacob and Victor Zarnowitz, "The Evaluation of Economic Forecasts," 1969, pp. 3–46.

Nimark, Kristoffer P., "Man-Bites-Dog Business Cycles," *American Economic Review*, 2014, *104* (8), 2320–2367.

_ **and Stefan Pitschner**, "News Media and Delegated Information Choice," *Journal of Economic Theory*, 2019, *181*, 160–196.

Pfajfar, Damjan and Emiliano Santoro, "News on Inflation and the Epidemiology of Inflation Expectations," *Journal of Money, Credit and Banking*, 2013, *45* (6), 1045–1067.

Pfäuti, Oliver, "The Inflation Attention Threshold and Inflation Surges," Technical Report 2025-06, University of Texas at Austin, EMPCT Working Paper Series August 2025. Revised working paper version.

Sahm, Claudia and Michael Sockin, "“Limited Attention” and Inflation Expectations of Households," FEDS Notes, Board of Governors of the Federal Reserve System 2016.

Schmidt, Torsten, Henrik Müller, Jonas Rieger, Tobias Schmidt, and Carsten Jentsch, "Inflation Perception and the Formation of Inflation Expectations," Ruhr Economic Papers 1025, RWI – Leibniz-Institut für Wirtschaftsforschung 2023.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson, "Measuring news sentiment," *Journal of econometrics*, 2022, *228* (2), 221–243.

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

Sims, Christopher A., "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Soroka, Stuart N., "The Gatekeeping Function: Distributions of Information in Media and the Real World," *The Journal of Politics*, 2012, *74* (2), 514–528.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent H. Meyer, Saten Kumar, and Yuriy Gorodnichenko, "Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings," *Econometrica*, 2025, *93* (1), 229–264.

Internet Appendix

A. Michigan "News Heard" Codes and Topic-Family Mapping

Table A1 maps MSC news-heard codes into the topic-family $\times$ sentiment classification used in the main text. The survey's coding architecture separates sentiment by construction: codes below 50 correspond to favorable mentions, while codes of 50 and above correspond to unfavorable mentions. The raw codes are aggregated into fourteen topic families, and each valid code is assigned to exactly one family. Following standard practice, generic "no changes heard" responses (code 0) and residual categories that do not correspond to substantive news, missing, invalid, or indeterminate codes, are excluded. This coding rule isolates changes in economically meaningful news heard while respecting the survey design that allows respondents to report up to two news topics.

B. Details of the Machine-Learning Benchmark Construction

This appendix documents the inputs, timing rules, and rolling-estimation protocol used to construct the monthly machine benchmark $\widehat{\pi}_t^{ML}$ used in Section 2.2. The benchmark series is treated as a pre-computed data input and is not re-estimated within the respondent-level regressions.

At each forecast month t, the information set $\mathcal{I}_t$ includes macroeconomic and financial predictors as they would have been observed in real time, accounting for release lags and revisions. Predictors are transformed to stationary units prior to estimation, following Bianchi et al. (2022).

Implementation follows three timing rules. For interviews conducted in month t, the machine information cutoff is set at the first day of month t, so the machine uses information only through the end of month $t-1$. At each month t, the Elastic Net model is estimated on the historical sample available at that date with tuning parameters chosen on a hold-out validation window. The resulting true out-of-sample forecast $\widehat{\pi}_t^{ML}$ is then stored, and the procedure rolls forward one month. This protocol ensures that deviations between survey

Table A1: Classification of Survey News Codes into Topic Families

Topic Family	**Favorable Codes**	**Unfavorable Codes**	**Representative Survey Responses**
Government & Politics	10, 12–19	50, 52–59	Elections, administration changes, taxes, fiscal policy, deficits, specific government spending programs, and reduced military spending.
War & Defense	11	51	Increases in military spending, escalations in war, and heightened geopolitical tensions.
Real Activity & Employment	20–25, 27–29, 45	60–65, 67–68, 85	Employment levels, consumer demand, production (GNP), plant openings/closings, wages, and specific industry performance.
Housing	–	69	Real estate slumps, falling home prices, and housing market declines.
Inflation & Prices	31, 32	71, 72	Changes in aggregate price levels, including both inflation and deflation.
Rates & Credit	30, 33, 37	73, 77	Interest rates, monetary policy tightness, and general credit availability.
Markets & Corporate Profits	35, 36	74–76	Corporate profit levels, earnings, and stock market fluctuations.
External & Dollar	38	78	Balance of payments and the devaluation or strength of the dollar.
Price & Wage Controls	39	79	Imposition or removal of wage and price controls.
Social, Unions & Crime	40, 41	80, 81	Race relations, crime levels, labor unrest, excessive wage demands, and union disputes.
Sentiment & Confidence	42–44, 47–48	82–84, 87–88	General economic optimism or pessimism, lack of confidence, and broad references to approaching "good times" or "bad times."
Agriculture	46	86	Agricultural conditions, crop yields, droughts, and farm prices.
Energy & Environment	49	89	Energy crises, pollution, and natural resource constraints.
Business Scandals	–	90	Corporate, business, and accounting scandals.

Notes. The table outlines the mapping of the raw Michigan Survey of Consumers news-heard codes into fourteen aggregated topic families. The survey architecture assigns codes below 50 to favorable mentions and codes of 50 or above to unfavorable mentions. All remaining codes are assigned to exactly one topic family. Generic "no changes heard" responses (code 0) and missing, invalid, or indeterminate codes are excluded from the reported news.

expectations and the benchmark reflect only predictive content available in real time.

The predictor panel is mixed frequency: some inputs are quarterly, while MSC expectations and many macro-finance indicators are monthly. $\mathcal{I}_t$ is formed at monthly forecast dates using the most recent values observable under the release calendar at that date.

The vector W'_{jt} contains the following survey and macroeconomic variables: lagged forecasts $F^{(i)}_{jt-k}[y_{jt+h-k}]$ for $k = 1, 2$, lagged forecasts of other types $F^{(s \neq i)}_{jt-1}[y_{jt+h-1}]$, the cross-sectional variance and skewness of lagged survey expectations, and trend inflation π_{t-1}, measured by the median SPF ten-year inflation forecast from 1991:Q4 onward and by an exponentially weighted moving average of past inflation prior to that date, detrended GDP and employment following Hamilton (2018), and a nowcast of the i-th percentile of inflation spanning the period $t-h$ to t.

The model also controls for one-month lagged inflation, $y_{t-1,t-h-1}$. The vector $\hat{G}'_{jt}$ incorporates factors extracted from three high-dimensional datasets: real-time macroeconomic factors $G_{M,t-k}$ from a panel of 92 monthly series, financial factors $G_{F,t-k}$ from a panel of 147 monthly financial indicators as in Ludvigson and Ng (2007), and monthly factors $G^{M}_{D,t-k}$ aggregated from daily financial data covering 87 daily indicators including exchange rates, volatility indices, credit and yield spreads, equity indices, and commodity prices.

B.1. Benchmark implementation checks and robustness

Bianchi et al. (2022) provides the benchmark construction and validation. Appendix Table A2 reports implementation checks tied to this paper's estimation sample. Panel A summarizes forecast diagnostics for the ML benchmark and three transparent alternatives, AR(12), trailing-mean, and random-walk inflation forecasts, evaluated on months in the MSC-linked sample with realized 12-month CPI inflation available. Panel B recomputes forecast bias using each benchmark definition and re-estimates the baseline Add UF and Add F specification with the same controls, fixed effects, weighting, and two-way clustering. This appendix evidence is presented as a sanity check on implementation and a robustness check on the headline estimates, not as a re-validation of the BLM benchmark.

C. Additional Evidence and Checks for the Baseline Results

C.1. Response Order and Sentiment Persistence

This section presents supplementary evidence discussed in Section 3. Panel B of Table 2 shows that the Add UF effect is robust to response order: it is indistinguishable across first and second responses. We also report descriptive evidence that interview pairs in which unfavorable mentions outnumber favorable mentions are substantially more persistent across

Table A2: Benchmark Implementation Checks and Robustness

Panel A: Forecast diagnostics (12-month CPI inflation)				
	(1)	(2)	(3)	(4)
Benchmark	RMSE	MAE	Mean	SD
ML	1.799	1.240	1.803	0.392
AR(12)	1.990	1.474	2.891	1.379
Trail mean	1.804	1.241	2.193	0.588
RW	1.881	1.376	2.221	1.242
Panel B: Baseline coefficients by benchmark choice				
	(5)	(6)	(7)	(8)
	ML	AR(12)	Trail mean	RW
Add UF	0.280***	0.183***	0.272***	0.233***
	(0.047)	(0.045)	(0.046)	(0.045)
Add F	-0.304***	-0.207***	-0.273***	-0.238***
	(0.053)	(0.049)	(0.051)	(0.049)
Observations	39,588	39,588	39,588	39,588
R^2	0.373	0.373	0.367	0.384
Month FE	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes
Two-way clusters (id, month)	Yes	Yes	Yes	Yes

Notes. Panel A reports RMSE, MAE, mean forecast, and forecast standard deviation for each benchmark. Evaluation uses months in the MSC-linked sample with realized 12-month inflation observed. Panel B re-estimates the baseline regression with $\Delta|b_{i,t}^{ML}|$ recomputed relative to each benchmark. Add UF and Add F coefficients are reported with two-way clustered standard errors in parentheses. Panel B uses the intersection sample across benchmark-specific outcomes to keep the estimation sample identical across rows. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

consecutive interviews than interview pairs in which favorable mentions outnumber unfavorable mentions, consistent with the weak Drop-UF effect in Section 3.1.1.

C.2. First vs. Second Response

One potential concern is that the order of the news household reported, rather than the actual informational content, drives the baseline coefficients. For example, respondents who mention unfavorable news first may differ systematically from those who report it second, so order itself could mechanically drive the results. This concern parallels the identification discussion in Chahrour et al. (2025), where response order in MSC helps separate different sources of variation in reported news and expectation updating. Panel B of Table 2 addresses this by splitting $\text{Add}_{i,t}^{UF}$ into mutually exclusive first- versus second-response indicators. Let

$UF^{(k)}_{i,t}$ denote an indicator that mention $k \in \{1, 2\}$ at interview t contains an unfavorable topic when valid. Then

$$\begin{aligned}\mathrm{Add}^{UF,1}_{i,t} &\equiv \mathbf{1}\{\mathrm{Add}^{UF}_{i,t} = 1,\ UF^{(1)}_{i,t} = 1\},\\ \mathrm{Add}^{UF,2}_{i,t} &\equiv \mathbf{1}\{\mathrm{Add}^{UF}_{i,t} = 1,\ UF^{(1)}_{i,t} = 0,\ UF^{(2)}_{i,t} = 1\},\end{aligned}$$

so that $\mathrm{Add}^{UF}_{i,t} = \mathrm{Add}^{UF,1}_{i,t} + \mathrm{Add}^{UF,2}_{i,t}$ by construction. The specification becomes

$$\begin{aligned}\Delta|b^{ML}_{i,t}| = \beta_{UF,1}\,\mathrm{Add}^{UF,1}_{i,t} + \beta_{UF,2}\,\mathrm{Add}^{UF,2}_{i,t} + \delta_{UF}\,\mathrm{Drop}^{UF}_{i,t}\\ + \beta_F\,\mathrm{Add}^{F}_{i,t} + \delta_F\,\mathrm{Drop}^{F}_{i,t} + \Gamma' X_{i,t-1} + \tau_t + \psi_{g(i)} + \varepsilon_{i,t}.\end{aligned} \tag{10}$$

Estimated coefficients on adding attention to unfavorable news are positive in both response orders. In the full sample, the first- and second-response coefficients are very similar at 0.266 and 0.197. In the dual-response sample they are 0.323 and 0.275, respectively. The hypothesis $H_0 : \beta_{UF,1} - \beta_{UF,2} = 0$ cannot be rejected in either sample. Thus, forecast bias increases whenever an unfavorable topic is added to the reported attention set, regardless of whether it is reported first or second. This contrasts with Chahrour et al. (2025), who find that response order matters for expectation *revisions*. In our setting, response order carries little additional information for forecast *bias*: what matters is whether households reallocate attention toward unfavorable versus favorable news.

C.3. Persistence in Sentiment Composition

We report descriptive evidence on persistence in sentiment composition for the dual-response sample. We define a respondent to be UF-dominant when unfavorable news mentions exceed favorable mentions, and F-dominant when the reverse holds.

Unfavorable dominance is substantially more persistent than favorable dominance within household attention sets. The probability of remaining unfavorable-dominant across consecutive interviews is 64.6%, compared to 45.6% for remaining favorable-dominant, yielding a gap of 19 percentage points. Mixed reporting news sentiment is also common in the sample: the probability of reporting a favorable topic conditional on reporting an unfavorable one, Pr(any F | any UF), is 34.3%, and in 26.6% of all interviews, households report a mixed news sentiment. This difference is consistent with broader evidence that unfavorable information can have more persistent effects on expectations than favorable information.[20]

These summary statistics help rationalize the baseline regression results. In particular, the stronger persistence of unfavorable dominance is consistent with the weak Drop-UF response: removing unfavorable topics from the attention set does not immediately reverse the associated bias. At the same time, the prevalence of mixed sentiments within attention sets implies that the Add-UF/Add-F asymmetry is not a corner case driven by a single dominant topic family.

C.4. Narrative Heterogeneity

Table A3 reports separate Add and Drop effects for each of the fourteen MSC topic families, complementing the joint specification in Section 3. Panel A reports the unfavorable narratives and Panel B reports the favorable narratives. The appendix estimates confirm the main-text ranking: broad non-price unfavorable narratives generate the largest increases in forecast bias, while favorable sentiment, rates, and market narratives are the most bias-reducing.

C.5. Additional Falsification and Robustness Checks

Attention Intensity versus Allocation. One concern is that the baseline coefficients capture attention intensity rather than attention allocation over reported "news heard" topics. Households that report more mentions, or become more engaged between interviews, may be more likely to report unfavorable news mechanically. If so, the Add UF coefficient could reflect how much households attend to news rather than what type of news they attend to. Bracha and Tang (2025) document that consumers pay more attention to inflation when inflation is high, and that this variation in attention intensity affects how strongly

[20] Bordalo et al. (2020) show that individual forecasters overreact through diagnostic expectations, generating persistence in forecast errors that is larger for negative surprises. Kohlhas and Walther (2021) document that agents allocate attention asymmetrically, placing excess weight on unfavorable signals, which mechanically amplifies the persistent influence of unfavorable news on beliefs. Gennaioli et al. (2024) provide a complementary memory-based account using the Michigan Survey of Consumers together with the New York Fed Survey of Consumer Expectations: selective recall of past high-inflation episodes, triggered by recency cues, generates persistent upward bias that is difficult to reverse once activated.

Table A3: Grouped Narrative Effects on Forecast Bias

	(1) Full sample	(2) Dual-response sample
Panel A: Unfavorable narratives		
Social / unions / crime	0.671***	0.542**
	(0.188)	(0.251)
War / defense	0.573***	0.465**
	(0.159)	(0.219)
Government / politics	0.454***	0.417***
	(0.073)	(0.110)
Inflation / prices	0.346***	0.533***
	(0.093)	(0.173)
Panel B: Favorable narratives		
Sentiment / confidence	-0.363***	-0.447***
	(0.078)	(0.116)
Rates / credit	-0.373***	-0.336***
	(0.074)	(0.112)
Markets / profits	-0.375***	-0.367***
	(0.056)	(0.088)
Inflation / prices	-0.231**	-0.163
	(0.113)	(0.148)
Observations	40,362	10,035
R^2	0.369	0.407
Month FE	Yes	Yes
Demo FE	Yes	Yes
WT Weights	Yes	Yes

Notes. Each coefficient comes from a separate weighted matched-pair regression of $\Delta|b_{i,t}^{ML}|$ on the indicated Add-family indicator, the corresponding Drop-family indicator, the baseline controls, month fixed effects, and demographic fixed effects. The sample construction, matching rules, weights, and fixed effects are identical to the baseline specifications in Table 2. "Dual-response sample" restricts to matched pairs with two valid news mentions in both interviews. Standard errors are two-way clustered by respondent and interview month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

expectations respond to inflation information.

Table A4 tests this directly. Panel A reports the composition coefficients and the sample splits, while Panel B reports the intensity controls and placebo indicators. Columns (2) and (3) in Panel A augment the baseline with respondent-level intensity controls. Column (2) controls for the current number of valid mentions (0–2). Column (3) controls for within-pair changes in mention count. Both intensity controls are small and statistically insignificant, while Add UF and Add F remain close to baseline magnitudes and highly significant.

Columns (4)–(7) test whether composition effects are confined to specific inflation regimes. Columns (4) and (5), following Bracha and Tang (2025), split the sample by below- versus

Table A4: Attention Intensity versus Attention Composition

	Control for Intensity		By Inflation		By Inf Mentions		Placebo
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)

	(1) Base	(2) $+n$	(3) $+\Delta n$	(4) Low	(5) High	(6) Low	(7) High	(8) Placebo
Panel A: Composition indicators								
Add UF	0.278***	0.285***	0.262***	0.254***	0.315***	0.265***	0.292***	
	(0.048)	(0.050)	(0.056)	(0.063)	(0.076)	(0.067)	(0.067)	
Add F	-0.301***	-0.295***	-0.316***	-0.254***	-0.348***	-0.164**	-0.441***	
	(0.052)	(0.058)	(0.064)	(0.070)	(0.075)	(0.066)	(0.078)	
Drop UF	0.062	0.060	0.080	0.121*	-0.001	0.101*	0.040	
	(0.045)	(0.046)	(0.052)	(0.063)	(0.061)	(0.058)	(0.060)	
Drop F	-0.126***	-0.128***	-0.111**	-0.113*	-0.146**	-0.029	-0.229***	
	(0.046)	(0.046)	(0.051)	(0.064)	(0.061)	(0.066)	(0.061)	
Panel B: Intensity controls / placebo								
Mention count		-0.008						
(current)		(0.021)						
Δ Mention count			0.015					
			(0.027)					
Add any mention								0.044
(placebo)								(0.035)
Drop any mention								-0.033
(placebo)								(0.035)
Observations	40,362	40,362	40,362	20,254	20,108	20,250	20,112	40,362
R^2	0.371	0.371	0.371	0.415	0.341	0.378	0.374	0.368
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes. The table reports weighted matched-pair regressions of $\Delta|b_{i,t}^{ML}|$ on attention-composition measures and intensity controls. Column (1) reports the baseline specification. Columns (2) and (3) add the current number of valid news mentions and the within-pair change in that count, respectively. Columns (4) and (5) re-estimate the baseline specification separately in below- and above-median realized 12-month CPI inflation months, following Bracha and Tang (2025). Columns (6) and (7) instead split the sample by whether the monthly share of respondents mentioning any inflation-related news is below or above its sample median. Column (8) replaces the sentiment-composition regressors with placebo indicators for adding or dropping any valid news mention regardless of sentiment. All columns include baseline controls, month fixed effects, and demographic fixed effects. Standard errors are two-way clustered by respondent and month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

above-median realized 12-month CPI inflation. Columns (6) and (7) instead split by the monthly share of respondents mentioning inflation-related news. In all four subsamples, Add UF and Add F remain qualitatively unchanged. Add F is somewhat larger in high-regime subsamples under both splits,[21] but the central composition result is not driven by

[21]The larger Add F coefficient in high-inflation months is consistent with favorable news carrying more weight when households are already paying closer attention to economic conditions, though we do not

high-inflation or high-salience episodes.

Column (8) provides an intensity placebo by replacing sentiment-specific composition indicators with two generic indicators: whether the respondent adds any valid mention and whether the respondent drops any mention, regardless of sentiment. Neither is significant. Changes in survey engagement, absent sentiment content, are not associated with changes in forecast bias. Table A4 separates this paper's composition channel from a pure intensity channel in Bracha and Tang (2025): the key variable is what is added to and dropped from the attention set, not simply how much news is reported.

Latent sentiment and recall. An additional concern is that households may report unfavorable news not because attention has shifted, but because a more pessimistic underlying outlook makes negative news easier to recall. Table A5 addresses this by adding within-household changes in the Index of Consumer Sentiment and in a personal-financial-outlook measure to the baseline regression. Both controls are predictive of changes in absolute bias, but the Add UF and Add F coefficients remain large, statistically significant, and close to their baseline magnitudes.

Placebos and Permutation Tests. Table A6 reports placebo and permutation tests addressing spurious correlation, reverse causality, and aggregate shock concerns.

Panel A evaluates whether the estimates are driven by the mechanical construction of the variables or by aggregate shocks. Column 2 shows that randomly swapping unfavorable and favorable sentiment labels while holding the underlying topic-family switches fixed yields coefficients for Add UF and Add F close to zero. This rules out spurious effects from sentiment assignment. To address aggregate shocks, we perform permutations within months and demographic groups by shuffling either the dependent variable (Columns 3 and 4) or the news indicators (Columns 5 and 6). Across these alternative specifications, the placebo mean coefficients are substantially smaller than the baseline magnitudes and fail to reproduce the baseline signs. This confirms that the results reflect respondent-level attention reallocations

emphasize this difference given sample-split noise.

Table A5: Baseline Asymmetry with Sentiment and Personal-Outlook Controls

	(1) Baseline	(2) +ΔICS	(3) +ΔPEXP	(4) +ΔICS, +ΔPEXP
Add UF	0.278***	0.238***	0.277***	0.237***
	(0.048)	(0.047)	(0.048)	(0.047)
Add F	-0.301***	-0.266***	-0.300***	-0.265***
	(0.052)	(0.051)	(0.052)	(0.051)
Drop UF	0.062	0.086*	0.064	0.086*
	(0.045)	(0.045)	(0.045)	(0.045)
Drop F	-0.126***	-0.161***	-0.128***	-0.162***
	(0.046)	(0.045)	(0.046)	(0.045)
ΔICS	–	-0.005***	–	-0.006***
		(0.001)		(0.001)
ΔPEXP	–	–	0.018*	-0.018*
			(0.010)	(0.011)
Observations	40,362	40,362	40,362	40,362
R^2	0.371	0.374	0.371	0.374
Month FE	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes
WT Weights	Yes	Yes	Yes	Yes

Notes. The dependent variable is $\Delta|b^{ML}_{i,t}|$ between matched interviews. All columns use the same matched-pair sample, baseline controls, month fixed effects, demographic fixed effects, survey weights, and two-way clustered standard errors as Table 2. Column (2) adds the within-household change in the Index of Consumer Sentiment (ICS). Column (3) adds the within-household change in $PEXP$, a survey measure of personal financial outlook. Column (4) includes both controls. The table asks whether the Add UF / Add F asymmetry survives after conditioning on shifts in broad sentiment or personal financial outlook that could otherwise rationalize selective recall of unfavorable news. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

rather than generic month-level comovement.

Panel B implements additional falsification tests. Regressing current bias changes on future attention switches yields economically small and statistically insignificant coefficients (Column 7), which alleviates concerns about reverse causality (e.g., changes in bias driving household news attention switches). Column 8 shows that replacing the sentiment-specific Add and Drop indicators with a generic topic-family-change indicator produces a small and insignificant coefficient, verifying that the baseline estimates are tied strictly to sentiment rather than random topic switching. Finally, Column 9 shows we cannot reject the symmetry restriction $\beta_{UF}+\beta_F = 0$ at any conventional significance level, consistent with approximately offsetting magnitudes for Add UF and Add F despite the asymmetric Add and Drop dynamics.

Table A6: Placebos, Falsification, and Permutation Tests

Panel A: Permutation Tests (Empirical Null Distributions)						
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline Estimate	**Label Perm.** (UF ↔ F)	**Y-Shuffles (Dep. Var.)** w/in Month	w/in Mo×Demo	**X-Shuffles (Regressors)** w/in Month	w/in Mo×Demo
Add UF	0.278***	-0.012	0.005	0.171	-0.001	0.159
	(0.048)	(0.019)	(0.067)	(0.045)	(0.034)	(0.025)
Add F	-0.301***	-0.008	0.003	-0.151	0.005	-0.177
	(0.052)	(0.018)	(0.054)	(0.039)	(0.033)	(0.029)
Observations	40,362	40,362	40,362	40,362	40,362	40,362
R^2	0.371	0.368	0.041	0.147	0.368	0.369
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Falsification Tests (Alternative Specifications)				
	(7) **Lead Test (Future Shifts)**		(8) **Non-Sentiment Turnover**	(9) **Symmetry Test**
	Future Add UF	Future Add F	Any family change	$H_0 : \beta_{UF} + \beta_F = 0$
Estimate	-0.047	0.061	-0.021	-0.023
	(0.042)	(0.049)	(0.032)	(0.081)
Observations	22,821		40,362	40,362
R^2	0.391		0.368	0.371
Month FE	Yes		Yes	Yes
Demo FE	Yes		Yes	Yes

Notes. The dependent variable is the six-month change in absolute forecast bias. Column (1) reports baseline estimates from the main specification with month and demographic fixed effects and two-way clustered standard errors by respondent and month. In Panel A, columns (2)–(6) report placebo means across 50 replications, with cross-replication standard deviations in parentheses. "Label Perm." randomly swaps UF and F labels within observation. "Y-Shuffles" permute the dependent variable, and "X-Shuffles" permute attention-shift regressors. Panel B reports falsification regressions (standard errors in parentheses): lead test using future shifts, non-sentiment turnover placebo, and the symmetry restriction. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$ (stars shown only for standard regression estimates).

Sample restrictions, weighting, and controls. Table A7 varies sample restrictions, survey weighting, and control sets. The Add UF coefficient ranges from 0.26 to 0.29 across all alternatives and Add F remains comparably stable with the opposite sign.

Subgroup analysis. Table A8 re-estimates the baseline separately for college-educated respondents, older respondents, households with below-median initial bias, and the early half of the sample. The asymmetry holds across all subgroups.

Household-level state dependence. Table A9 shows that the Add UF effect does not depend on whether a household's past bias is above or below the median, ruling out a mechanical mean-reversion explanation.

Table A7: Robustness of Baseline Estimates

	(1) Baseline	(2) Movers Only	(3) Unweighted	(4) Trim Lag Bias	(5) Drop Rare Families
Add UF	0.278***	0.262***	0.289***	0.286***	0.237***
	(0.048)	(0.050)	(0.043)	(0.047)	(0.044)
Add F	-0.301***	-0.311***	-0.308***	-0.324***	-0.274***
	(0.052)	(0.054)	(0.047)	(0.049)	(0.049)
Drop UF	0.062	0.067	0.079**	0.066	0.055
	(0.045)	(0.046)	(0.039)	(0.048)	(0.043)
Drop F	-0.126***	-0.143***	-0.135***	-0.162***	-0.079*
	(0.046)	(0.048)	(0.044)	(0.048)	(0.044)
Observations	40,362	30,139	40,362	36,359	40,362
R^2	0.371	0.383	0.370	0.241	0.370
Month FE	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes	Yes

Notes. The dependent variable is the six-month change in absolute forecast bias. Column (1) reports the full baseline specification. Column (2) restricts the sample to pairs with at least one attention change ("movers"). Column (3) removes sampling weights. Column (4) trims the 5th and 95th percentiles of lagged absolute bias to limit the influence of outliers. Column (5) excludes topic families appearing in less than 10% of total interviews. Standard errors are two-way clustered by respondent and month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A8: Robustness Across Household Subsamples

	(1) Baseline	(2) College Degree	(3) Age $\geq$ 55	(4) Low Lag Bias	(5) Early Period	(6) Late Period
Add UF	0.278***	0.280***	0.244***	0.181***	0.237***	0.371***
	(0.048)	(0.045)	(0.075)	(0.054)	(0.075)	(0.068)
Add F	-0.301***	-0.290***	-0.464***	-0.285***	-0.185***	-0.383***
	(0.052)	(0.051)	(0.076)	(0.054)	(0.068)	(0.073)
Drop UF	0.062	0.107**	0.052	-0.044	0.054	0.096
	(0.045)	(0.050)	(0.065)	(0.057)	(0.070)	(0.063)
Drop F	-0.126***	-0.136***	-0.153**	-0.172***	-0.084	-0.155***
	(0.046)	(0.046)	(0.077)	(0.051)	(0.075)	(0.060)
Observations	40,362	30,121	18,491	20,278	20,246	19,656
R^2	0.371	0.377	0.369	0.114	0.370	0.382
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. The dependent variable is the six-month change in absolute forecast bias. Column (1) reports the full sample baseline. Column (2) restricts to respondents with at least some college education. Column (3) restricts to older respondents ($\geq$ 55 years). Column (4) restricts to respondents whose baseline absolute bias is below the sample median. Column (5) restricts to the earlier half of the survey sample period. Column (6) restricts to matched pairs from 2010 through 2025, covering the post-GFC recovery, the low-inflation decade, and the 2021–22 inflation surge. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A9: Household-Level State Dependence in Add-UF/Add-F Effects

	(1) Baseline	(2) +AboveMed interactions	(3) Dual-response
Add UF	0.374***	0.375***	0.465***
	(0.050)	(0.049)	(0.097)
Add F	-0.370***	-0.283***	-0.225**
	(0.054)	(0.054)	(0.103)
Drop UF	0.086*	0.089**	0.201**
	(0.044)	(0.044)	(0.082)
Drop F	-0.179***	-0.181***	-0.117
	(0.048)	(0.047)	(0.082)
AboveMed		-0.049	0.368**
		(0.089)	(0.170)
Add UF × AboveMed		-0.000	-0.095
		(0.072)	(0.150)
Add F × AboveMed		-0.236***	-0.343**
		(0.074)	(0.153)
Observations	40,362	40,362	10,035
R^2	0.336	0.336	0.341
Month FE	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes

Notes. AboveMed equals one when lagged absolute bias exceeds the within-month median. Standard errors are two-way clustered by respondent and month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

D. Additional Evidence and Checks for the Attention Channel

D.1. Additional IV Diagnostics

Table A10 reports the additional IV diagnostics. The DNSI interaction-IV is also stable when we allow demographic groups to follow separate linear time trends. This addresses the concern that the instrument might proxy for slow-moving, group-specific changes in inflation beliefs rather than shifts in the narrative environment itself.

E. Additional Evidence and Checks for the Aggregate Implications

E.1. Calibration Details for the Decomposition Framework

This appendix documents the calibration procedure for the decomposition framework used in Section 5. Step 1 uses weighted reduced-form regressions to identify the attention parameters $(\mu_{UF}, \mu_F, \omega_{UF,H}, \omega_{UF,Z}, \omega_{F,H}, \omega_{F,Z})$. Step 2 estimates the persistence parameter γ

Table A10: Interaction-IV Robustness to Demographic-Specific Time Trends

	Endog. = Add UF (non-price)		Endog. = Add F (non-price)	
	(1)	(2)	(3)	(4)
	Baseline	+ Demo trends	Baseline	+ Demo trends
2SLS coefficient	2.245*	2.234*	-2.397*	-2.391*
	(1.337)	(1.327)	(1.416)	(1.399)
First-stage F	18.43	18.38	16.30	15.09
Observations	40,362	40,362	40,362	40,362
R^2	-0.063	-0.062	-0.059	-0.059
Month FE	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes
Demo trends	No	Yes	No	Yes
WT Weights	Yes	Yes	Yes	Yes

Notes. The table reports DNSI-based 2SLS estimates using the same matched-pair sample and predetermined exposure weight as Table 6. Columns (1) and (3) reproduce the non-price Add UF and Add F specifications. Columns (2) and (4) add demographic-specific linear time trends, implemented as interactions between demographic fixed-effect cells and a linear month trend. All columns include the baseline controls, month fixed effects, demographic fixed effects, survey weights, and two-way clustered standard errors. The first-stage F statistic is for the excluded instrument in the corresponding first stage. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

in a separate fixed-effects regression. Step 3 fits a month-level negative-binomial model to recover the latent-count correction (α, κ_t). Step 4 selects the sentiment-bias parameter ϕ via simulation-based grid search.

Step 1: Attention parameters. Let $y_{i,t} \equiv \Delta|b_{i,t}^{ML}|$ and let (H_t, Z_t) denote the two monthly states used in the baseline, with $H_t = \text{Hub}_t$ and Z_t equal to the standardized monthly two-slot share (Intensity). The reduced-form regression for attention coefficients is

$$\begin{aligned} y_{i,t} = \beta_{UF}\,\text{AddUF}_{i,t} + \beta_F\,\text{AddF}_{i,t} + \beta_{UF\times H}\,(\text{AddUF}_{i,t}H_t) \\ + \beta_{UF\times Z}\,(\text{AddUF}_{i,t}Z_t) + \beta_{F\times H}\,(\text{AddF}_{i,t}H_t) + \beta_{F\times Z}\,(\text{AddF}_{i,t}Z_t) \\ + X_{i,t}'\delta + \alpha_{m(t)} + \alpha_{g(i)} + \varepsilon_{i,t}, \end{aligned}$$

estimated by weighted `AbsorbingLS` with month and demographic fixed effects. The control vector $X_{i,t}$ contains lagged absolute bias, lagged expectation, interview gap, and total topic moves. Define $\lambda_{i,t} = 1/|\mathcal{A}_{i,t}|$ for non-empty reported sets and

$$\bar{\lambda}_{UF} = \mathbb{E}[\lambda_{i,t} \mid \text{AddUF}_{i,t} = 1], \quad \bar{\lambda}_F = \mathbb{E}[\lambda_{i,t} \mid \text{AddF}_{i,t} = 1],$$

$$\bar{\lambda}_{UF\times H} = \mathbb{E}[\lambda_{i,t}H_t \mid \text{AddUF}_{i,t} = 1], \quad \bar{\lambda}_{UF\times Z} = \mathbb{E}[\lambda_{i,t}Z_t \mid \text{AddUF}_{i,t} = 1],$$

$$\bar{\lambda}_{F\times H} = \mathbb{E}[\lambda_{i,t} H_t \mid \text{AddF}_{i,t} = 1],$$

$$\bar{\lambda}_{F\times Z} = \mathbb{E}[\lambda_{i,t} Z_t \mid \text{AddF}_{i,t} = 1].$$

The mapping implemented in code is

$$\hat{\mu}_{UF} = \frac{\hat{\beta}_{UF}}{\bar{\lambda}_{UF}}, \qquad \hat{\mu}_F = \frac{\hat{\beta}_F}{\bar{\lambda}_F},$$

$$\hat{\omega}_{UF,H} = \frac{\hat{\beta}_{UF\times H}}{\bar{\lambda}_{UF\times H}}, \qquad \hat{\omega}_{UF,Z} = \frac{\hat{\beta}_{UF\times Z}}{\bar{\lambda}_{UF\times Z}},$$

$$\hat{\omega}_{F,H} = \frac{\hat{\beta}_{F\times H}}{\bar{\lambda}_{F\times H}},$$

$$\hat{\omega}_{F,Z} = \frac{\hat{\beta}_{F\times Z}}{\bar{\lambda}_{F\times Z}}.$$

Reported standard errors for these mapped parameters are the regression standard errors scaled by the same denominators. Using the baseline linked-pair reduced-form regression, this mapping produces $\hat{\mu}_{UF} = 0.267$ and $\hat{\mu}_F = -0.552$. The two estimates are close in magnitude and opposite in sign, consistent with offsetting effects of unfavorable and favorable attention on forecast bias.

Step 2: Persistence parameter. The persistence parameter is estimated in a separate weighted fixed-effects regression,

$$b^{ML}_{i,t} = \gamma\, b^{ML}_{i,t-1} + X'_{i,t}\eta$$
$$+ \alpha_{m(t)} + \alpha_{g(i)} + u_{i,t},$$

again with `AbsorbingLS` and robust covariance. For numerical stability in decomposition, the code clips $\hat{\gamma}$ to $[0, 0.999]$. This regression produces $\hat{\gamma} = 0.672$, implying that about two-thirds of any month's forecast bias carries forward six months to the next interview. The interaction loadings estimated in Step 1 yield $(\hat{\omega}_{UF,H}, \hat{\omega}_{UF,Z}, \hat{\omega}_{F,H}, \hat{\omega}_{F,Z}) = (0.264, 1.346, 0.163, 0.354)$.

Step 3: Latent-count correction (κ_t). For latent topic intensity, the current-interview slot count is $n^{slot}_{i,t} = 1\{\text{NEWS1 valid}\} + 1\{\text{NEWS2 valid}\} \in \{0, 1, 2\}$. Monthly weighted shares are $\hat{p}_{j,t} = \mathbb{E}_w[1\{n^{slot}_{i,t} = j\}]$ for $j \in \{0, 1, 2\}$. The latent count K_t is modeled as a draw from a negative binomial distribution with mean m_t and variance parameter α, so

$$\begin{aligned}
\Pr(K_t = 0) &= (1 + \alpha m_t)^{-1/\alpha}, \\
\Pr(K_t = 1) &= \frac{m_t}{1 + \alpha m_t} \Pr(K_t = 0), \\
\Pr(K_t \geq 2) &= 1 - \Pr(K_t = 0) - \Pr(K_t = 1).
\end{aligned}$$

Given α, we invert $\hat{p}_{0,t}$ to obtain $m_t(\alpha) = \big(\hat{p}_{0,t}^{-\alpha} - 1\big)/\alpha$, then compute $p_{1,t}^{mod}(\alpha)$ and $p_{2,t}^{mod}(\alpha)$ from the formulas above. The selected $\hat{\alpha}$ minimizes

$$\sum_t W_t \Big[(p_{1,t}^{mod}(\alpha) - \hat{p}_{1,t})^2 + (p_{2,t}^{mod}(\alpha) - \hat{p}_{2,t})^2 \Big],$$

where W_t is the monthly total survey weight. Month-level latent intensity is then $\hat{\kappa}_t = \mathbb{E}[K_t]/\mathbb{E}[\min(K_t, 2)]$ under the fitted NB2 model, and the reported average latent correction is $\bar{\kappa} = \sum_t W_t \hat{\kappa}_t / \sum_t W_t$.

Step 4: Selection parameter (ϕ). Finally, ϕ is calibrated by simulation on each observed monthly network $\mathcal{G}_t$. Node sentiment is $s_u = +1$ for unfavorable and $s_u = -1$ for favorable nodes. For each month, the simulation uses empirical $\Pr(N_t = 0) = \hat{p}_{0,t}$, $\Pr(N_t = 2) = \hat{p}_{2,t}$, and $\Pr(N_t = 1) = 1 - \hat{p}_{0,t} - \hat{p}_{2,t}$. Conditional on drawing at least one topic, the first node is drawn from

$$\Pr(a = u \mid t, \phi) \propto p_{u,t} \exp(\phi s_u).$$

For two-topic draws, the second node is drawn from

$$\Pr(b = v \mid a = u, t, \phi) \propto w_{uv,t} \exp(\phi s_v),$$

with fallback to the first-draw distribution if u has no active neighbors. For each ϕ on the fixed grid $\{0, 0.01, 0.02, 0.05, 0.08, 0.1, 0.15, 0.2, 0.3, 0.5\}$, the code simulates monthly moments and aggregates them with monthly survey weights. The selected $\hat{\phi}$ minimizes

$$\begin{aligned}
\mathcal{L}(\phi) = &\Big(\overline{\text{shareF}}^{\,mod}(\phi) - \overline{\text{shareF}}^{\,data}\Big)^2 \\
&+ \Big(\overline{\text{shareUF}}^{\,mod}(\phi) - \overline{\text{shareUF}}^{\,data}\Big)^2 \\
&+ \Big(\overline{\text{top1}}^{\,mod}(\phi) - \overline{\text{top1}}^{\,data}\Big)^2.
\end{aligned}$$

E.2. Network Proximity and Forecast Bias

Network proximity to the attention-network hub predicts household-level bias changes in the expected direction, complementing the hub-state dependence result in Table 8.

Conditional on the first reported topic, the second draw is heavily weighted toward its local neighborhood in the attention network. When an unfavorable hub forms, households reporting nearby topics are thus more likely to draw additional bias-increasing unfavorable news as their second mention. We test this proximity effect using the minimum shortest-path distance from any of a household's reported topics to the nearest unfavorable topic active in month t, denoted $\text{DistUF}_{i,t}$.[22] Table A11 confirms the proximity prediction: greater $\text{DistUF}_{i,t}$ predicts smaller changes in absolute forecast bias, and this holds under progressively finer demographic-by-time fixed effects and leave-one-out distance measures.

Table A11: Network Position Predictors of Changes in Absolute Bias

	(1) Baseline	(2) Age×Educ FE	(3) +Age×Educ×YQ FE	(4) +LOO Dist
Distance to UF block	-0.381***	-0.383***	-0.390***	-0.390***
	(0.066)	(0.066)	(0.067)	(0.067)
Add UF	0.201***	0.199***	0.179***	0.179***
	(0.066)	(0.066)	(0.066)	(0.066)
Add F	-0.295***	-0.295***	-0.303***	-0.303***
	(0.069)	(0.069)	(0.069)	(0.069)
Observations	26,921	26,921	26,921	26,921
R^2	0.354	0.354	0.355	0.355
Month FE	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes
Two-way clusters (id, month)	Yes	Yes	Yes	Yes

Notes. The dependent variable is $\Delta|b_{i,t}^{ML}|$ between matched interviews. $\text{DistUF}_{i,t}$ is the minimum shortest-path distance from the respondent's reported topics to any unfavorable topic active in month t in the unweighted attention network. Column (1) is the baseline specification. Columns (2) to (3) add progressively finer demographic-by-time fixed effects. Column (4) replaces distance with a leave-one-out measure. All columns control for Add $\text{UF}_{i,t}$, Add $\text{F}_{i,t}$, lagged absolute bias, lagged expectation, interview gap, and total topic moves, and absorb month and demographic fixed effects. Standard errors are two-way clustered by respondent and month. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

[22] Distances are undefined when a household's reported topics share no path with any unfavorable topic in month t, which accounts for the smaller sample size in Table A11.

E.3. Additional Decomposition Evidence

Table A12 reports robustness of the Shapley-Owen-Shorrocks decomposition to alternative sample definitions, including dropping the post-2020 period, restricting to the dual-response sample, and varying the κ_t correction. The attention and persistence shares are stable across all alternatives.

Table A12: Latent-K Decomposition Robustness Across Samples

Sample	Observations	Share(P)	Share(A)	Share(R)	$\bar{\kappa}$
Full sample	40,362	0.417	0.163	0.420	1.489
Dual sample	16,821	0.388	0.301	0.310	1.489
Pre-2005	13,096	0.419	0.014	0.567	1.489
2005-2020	25,015	0.424	0.206	0.369	1.489
Post-2020	4,330	0.333	0.482	0.186	1.489
NBER recessions only	3,785	0.390	0.306	0.304	1.489

Notes. All rows use the same baseline two-state parameter estimates $(\mu_{UF}, \mu_F, \omega_{UF,H}, \omega_{UF,Z}, \omega_{F,H}, \omega_{F,Z}, \gamma, \phi, \alpha_{NB2})$ with $H_t = \text{Hub}_t$ and $Z_t = \text{Intensity}$. Month-specific κ_t from the latent-count NB2 mapping is applied in every row, and $\bar{\kappa}$ is reported only as a descriptive weighted mean within each sample.

E.4. Attention Network Robustness Checks

Table A13 reports robustness checks for both mechanism channels to verify that results are not driven by coding or sample restrictions.

For the results with an unfavorable focal hub, Panel A varies inference, fixed effects, support, and state-variable construction. Across columns (1)–(5), Add $\text{UF}_{i,t} \times Z_t$ remains positive and tightly clustered between 0.176 and 0.197. Column (6), which orthogonalizes unfavorable share relative to concentration, still yields a positive and significant interaction of 0.247, indicating that composition is not a by-product of concentration. Column (7), restricted to denser-network months (months with at least 20 active topic nodes and 30 network links), rises to 0.306, consistent with stronger crowding when network links between unfavorable and favorable topics are richer.

For the micro-level evidence, Panel B performs analogous checks. The baseline distance-to-UF coefficient remains negative under leave-one-out distances, edge-threshold filtering (keeping only links with at least five linked mentions), giant-component restrictions (keep-

Table A13: Robustness Across Concentration and Topology Channels

Panel A: Focal Hub Amplification							
	(1) Baseline	(2) Cluster (by month)	(3) Month FE only	(4) Winsorize Z (1/99)	(5) Trim tails of Z	(6) UF-share residual	(7) Dense months $n \geq 20, e \geq 30$
Add UF × Z_t	0.176***	0.176***	0.176***	0.179***	0.197***	0.247***	0.306***
	(0.040)	(0.041)	(0.040)	(0.040)	(0.042)	(0.065)	(0.061)
Add F × Z_t	0.108***	0.108***	0.108***	0.105***	0.095**	0.108***	0.027
	(0.040)	(0.041)	(0.040)	(0.039)	(0.038)	(0.042)	(0.080)
Observations	40,362	40,362	40,362	40,362	39,652	40,362	10,803
R^2	0.375	0.375	0.374	0.376	0.374	0.375	0.397
Baseline controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	No	Yes	Yes	Yes	Yes
Two-way clusters (id, month)	Yes	No	Yes	Yes	Yes	Yes	Yes
Cluster by month	No	Yes	No	No	No	No	No
Panel B: Distance to Hub							
	(1) Baseline	(2) LOO	(3) Backbone edges ≥ 5	(4) Giant comp. only	(5) + Content controls	(6) Fill unreach. + indicator	(7) Sparse screen $n \geq 20, e \geq 30$
Add UF$_{i,t}$	0.201***	0.201***	0.226***	0.200***	−0.014	0.317***	0.103
	(0.066)	(0.066)	(0.067)	(0.066)	(0.094)	(0.062)	(0.117)
Add F$_{i,t}$	−0.295***	−0.295***	−0.335***	−0.296***	−0.229**	−0.292***	−0.211**
	(0.069)	(0.069)	(0.070)	(0.069)	(0.094)	(0.065)	(0.098)
DistUF$_{i,t}$	−0.381***	−0.381***	−0.324***	−0.381***	−0.417***	−0.251***	−0.480***
	(0.066)	(0.066)	(0.056)	(0.066)	(0.071)	(0.055)	(0.125)
Observations	26,921	26,921	26,119	26,883	26,921	40,362	7,597
R^2	0.354	0.354	0.354	0.355	0.359	0.375	0.366
Baseline controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Two-way clusters (id, month)	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes. Panel A reports interaction coefficients from the Table 8 specification. "Winsorize Z (1/99)" caps Z_t at the 1st and 99th percentiles, and "Trim tails of Z" drops months in the tails of the Z_t distribution. "UF-share residual" uses the residual from regressing the month-level unfavorable-share index on the concentration measure and interacts Add UF / Add F with this residual. "Dense months" restricts the sample to months with at least n nodes and e edges in the month-level network. Panel B reports the Table A11 coefficients on Add UF, Add F, and distance to the unfavorable block. "LOO" recomputes distances on a leave-one-out monthly network that removes respondent i's unique edge contributions where applicable. "Backbone edges ≥ 5" constructs the monthly network using only edges with weight at least 5. "Giant comp. only" restricts to respondents whose mentioned nodes lie in the largest connected component. "Fill unreach. + indicator" assigns unreachable observations distance equal to the month-specific maximum distance plus one and includes an indicator for unreachable cases. "Sparse screen" restricts to months with at least n nodes and e edges. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

ing only respondents whose attention topics lie in the largest connected component), and distance-missingness adjustments (filling in distances for unreachable topic pairs). Across columns (1)–(5), coefficients range from -0.417 to -0.324 around the baseline -0.381, with consistent sign and magnitude. The filled-distance specification in column (6) is -0.251 and significant, and the corresponding estimate in column (7) is -0.480. This result is unlikely to be driven by months with fewer network links. It appears to be a stable structural feature of the attention network, consistent with narrative-network amplification (Shiller, 2017).