The News that Moves Inflation Expectations*

Carola Binder Pascal Frank Jane Ryngaert

September 15, 2026

Abstract

This paper studies how news coverage shapes the transmission of inflation information to household expectations. We combine daily household inflation expectations with television news coverage of inflation and professional-forecast-based surprises around scheduled CPI and PCE releases. On days with high news coverage, expectations move more in the direction of revealed surprises and, when we include the high inflation period, increase above and beyond these surprises. The finds are consistent with a model of noisy updating and infrequent updating.

Keywords: Inflation expectations; expectations surveys; daily data; cable television; news coverage; local projection
JEL Codes: D83, D84, E31

*Binder: University of Texas at Austin, carola.binder@austin.utexas.edu, Frank: University of Notre Dame, pfrank@nd.edu, Ryngaert: University of Notre Dame, jryngaer@nd.edu. The authors did not receive financial support for this paper and have no conflict of interest to acknowledge. This paper subsumes the previously circulated working paper *The Causal Effect of News on Inflation Expectations.* We thank Paul Bousquet, Lukas Hack, Christian Matthes, Benjamin Pugsley, Eric Sims, and participants at the Expectations Conference in Asia and University College Dublin for useful comments.

1 Introduction

The news media play a large role in transmitting economic information to the public, thereby shaping their economic expectations.[1] In particular, many consumers encounter news about monetary policy and macroeconomic conditions from the media rather than from policy announcements or official data sources. Previous randomized controlled trial (RCT) evidence has examined the causal effect of particular media *messages* on household expectations, conditional on direct exposure in the experimental setting.[2]

Outside of the experimental setting, the effect of news on expectations is likely more complex. To claim that news coverage moves expectations could mean several different things. Suppose, for example, that a CPI release shows that inflation is higher than professional forecasters expected. Without news coverage, this might raise the inflation expectations only of households who are directly exposed to the release. Greater news coverage might simply amplify that response by exposing more households to the signal. But the coverage might also change expectations in other ways, for example by informing households about prior inflation developments they were not aware of. A household that had followed inflation closely would learn primarily about the latest surprise, while one with stale information would also revise to reflect information since its previous update. Thus, news coverage can affect expectations both by amplifying the response to newly available information and by prompting households to incorporate information they previously missed.

This paper examines the high frequency response of inflation expectations in high news and normal news settings, focusing on days when reliable information arrives. We allow for news to affect expectations in two ways; first, by amplifying the response to information received on that day, and second, by increasing the overall salience of inflation. We measure daily spikes in inflation coverage relative to recent levels using a daily series of cable news mentions of inflation. Following Binder, Campbell,

[1] See for example Carroll (2003), Krueger and Blinder (2004), Binder (2017b), Coibion et al. (2020), Larsen, Thorsrud, and Zhulanova (2021), Chahrour, Nimark, and Pitschner (2021), Binder (2017a), Ehrmann, Pfajfar, and Santoroc (2017).

[2] See, for example, Binder (2021a) and Coibion, Gorodnichenko, and Weber (2022).

and Ryngaert (2024), we then leverage daily responses from the Survey of Consumer Expectations to estimate a state-dependent local projection model on inflation data release days between 2013 and 2024.[3] This allows us to compare the effects of releases that receive elevated media coverage with those that do not. Restricting the sample to release days focuses the analysis on days on which information arrival is predetermined and not driven by audience demand, and further provides a proxy for the information innovation on a given day - the "surprise" the inflation print constitutes relative to the professional consensus forecast.

We find that high coverage amplifies the response of expectations to measured inflation surprises; on release days without unusually elevated inflation coverage, the response of the median expectation to core inflation surprises is small or short-lived, while on high-coverage release days the same surprises generate larger and more persistent movements in expectations. Furthermore, in the full sample, high-coverage inflation release dates are associated with higher household inflation expectations relative to standard-coverage releases. This means that after controlling for the new information contained in the release, expectations increased on highly covered release days. However, when we exclude the high inflation period 2021-2024 from the sample, the second effect is not present, suggesting that the estimated positive effect is specific to the high inflation period.

These findings are consistent with a model of expectations formation that combines noisy signals with infrequent updating. Exposure to the signal is exogenous and increasing in news intensity. When agents observe the signal, they access both the current information innovation or "surprise" relative to prior public information set as well as the history of inflation news since they were last able to update. High-coverage releases therefore prompt "catch-up" revisions above and beyond the response to information that is novel to the public. Elevated news coverage affects aggregate expectations by increasing the share of agents who update towards the signal *and* the share who make catch-up revisions. These two channels map to the estimated amplification and salience effects, respectively. Importantly, while the amplification effect should always be non-negative, the sign of the salience effect can be time-varying.

[3]The sample is determined by the availability of the cable news data.

If inflation is stable, we would expect positive and negative catch-up revisions to cancel out. If, on the other hand, inflation is increasing the average catch-up revision will be positive.

The model further predicts that agents who have been relatively less exposed to inflation-related news will accumulate possibilities for catch-up relative to those who are up-to-date on the news. Their expectations should therefore respond more to elevated coverage, even after controlling for the realized surprise. The recent inflation period and the accompanying partisan divide in news coverage allows us to test this prediction. Respondents exposed to more left-leaning news coverage encountered inflation related coverage less than respondents exposed to more right-leaning news during this period. We separate respondents into likely-Republican and likely-Democrat groups using their state of resident and data on cable news viewership by state. The capacity for catch-up revisions does differ across groups, with respondents from CNN-viewing states entering each period with lower priors than those in Fox-viewing states. Consistent with this, we see that likely-Democrat respondents revise their expectations up, even after controlling for the measured information innovation.

This paper contributes to three closely related literatures. First, it contributes to work on news media and expectation formation. Carroll (2003) develops an epidemiological model in which households gradually absorb macroeconomic expectations from professional forecasters through the news media, while Lamla and Lein (2014) provide evidence that media coverage affects consumers' inflation expectations. Related work studies how media coverage, central bank communication, and news selection shape beliefs and aggregate fluctuations (Binder, 2017b; Binder, 2017a; Larsen, Thorsrud, and Zhulanova, 2021; Chahrour, Nimark, and Pitschner, 2021; Nimark and Pitschner, 2019; Ehrmann, Pfajfar, and Santoroc, 2017). Most closely related is Chahrour, Shapiro, and Wilson, 2024, who examine how self-reported encounters with favorable and unfavorable price news affect inflation expectations in the Michigan Survey of Consumers. Our approach is complementary. Rather than relying on self-reported exposure or experimental information treatments, we study realized media coverage around scheduled inflation releases. This allows us to examine how realized media

coverage around scheduled inflation releases is related to household expectations, and whether coverage amplifies the pass-through of measured inflation surprises to household expectations in the field.

Second, the paper relates to the literature on information acquisition, attention, and heterogeneous expectation formation. Households' beliefs depend on the information they encounter and on the experiences they bring to the interpretation of aggregate news (Malmendier and Nagel, 2016; Kuchler and Zafar, 2019; D'Acunto et al., 2021; D'Acunto and Weber, 2024). Experimental studies show that households update inflation expectations when they are provided with macroeconomic information or central bank communication (Binder, 2021a; Coibion, Gorodnichenko, and Weber, 2022). Our contribution is to study a related mechanism in the field: whether naturally occurring variation in the intensity of news coverage changes the pass-through of public inflation signals to household expectations.

Third, the paper contributes to work on attention, salience, and state dependence in inflation expectations. When inflation is high, households pay more attention to inflation and media outlets devote more coverage to the topic (Pfäuti, 2025; Nimark, 2014). This makes the transmission of inflation news potentially state dependent. Our results show that high-coverage inflation releases are associated with larger expectation movements and that measured inflation surprises have larger effects when accompanied by unusually intense coverage. This suggests that the effects of public macroeconomic signals depend not only on their informational content, but also on the attention environment in which they arrive.

The paper is structured as follows. In Section 2, we motivate and develop our baseline econometric framework, a state-dependent local projection model. Section 3 describes the daily news coverage and expectations data as well as the construction of inflation surprises. Section 2.2 presents the baseline empirical results that are consistent with a model of expectation formation combining noisy signals with infrequent updating, as outlined in Section 5. In Section 6, we test the implications of the model by focusing on expectations from CNN- and Fox-leaning states. Section 7 concludes.

2 Econometric Framework

This section develops an econometric framework that separates two ways in which media coverage may shape the transmission of inflation news: by making inflation more *salient* and by *amplifying* the pass-through of a given inflation surprise.[4] First, we explain why regressing expectations on news coverage alone cannot separately identify these channels; such a regression omits the informational content of coverage and pools days with and without information arrival. This motivates our baseline model, a state-dependent local projection estimated on inflation data release days.

2.1 Challenges to Estimation

Suppose that inflation news may impact expectations independently by raising the salience of inflation to households or by increasing the likelihood that households view a given information innovation, λ_t, on day t. Assume that $\lambda_t \sim N(0, \sigma_\lambda^2)$ with $\lambda_t > 0$ meaning that the inflation signal is inflationary relative to the prior, and $\lambda_t < 0$ indicating that the signal is relatively disinflationary. We can express year-ahead inflation expectations as a daily local projection model of the following form:

$$E_{t+h}[\pi_{t+365,t}] = \beta_0^h + \beta_1^h news_t + \beta_2^h \lambda_t + \beta_3^h (news_t \times \lambda_t) + \Phi^h X_t + u_{t+h}. \quad (1)$$

The coefficient β_2^h measures the pass-through of the information innovation, λ_t, on a day without news coverage and β_3^h captures the additional pass-through due to news coverage, or the *amplification effect* of the news. We expect that both coefficients are greater than or equal to 0 as $\beta_2^h < 0$ would mean that expectations move in the opposite direction of information signals and $\beta_3^h < 0$ would mean that increased coverage would dampen the response to information. β_1^h captures the *salience effect* of the news or the effect on expectations of simply being more aware of inflation.

[4]Increaseed salience could drive expectations either up or down. It would increase expectations if households believe the media follows a news selection fuction which elevates stories about rising prices (Chahrour, Shapiro, and Wilson, 2024) or if consumers take a pessimistic view towards inflation news in general (Kamdar, 2018).

Importantly, while we will observe daily measures of $E_t[\pi_{t+365,t}]$ and $news_t$, we do not observe λ_t at a daily frequency. Instead, suppose that we estimate the effect of the news with the following naive model:

$$E_{t+h}[\pi_{t+365,t}] = \alpha^h + \delta^h news_t + \Phi^h X_t + u_{t+h}. \tag{2}$$

The coefficient δ^h captures the dynamic association between coverage and expectations, but it does not identify the effect of coverage for two distinct reasons.

The first is an *omitted variable bias*. The model in Equation (2) omits λ_t, so δ^h conflates the effect of coverage on exposure with the content of the news that receives coverage. If coverage is systematically related to the underlying signal — as it would be under a news-selection function prioritizing information about higher inflation (Nimark and Pitschner, 2019; Chahrour, Shapiro, and Wilson, 2024) — then δ^h will be biased by the omission of terms containing λ_t. Abstracting from controls, the bias is approximately:

$$\text{Bias}(\delta^h) = \beta_2^h \frac{\text{Cov}(news_t, \lambda_t)}{\text{Var}(news_t)} + \beta_3^h \frac{\text{Cov}(news_t, \lambda_t news_t)}{\text{Var}(news_t)}. \tag{3}$$

As $\beta_2^h \geq 0$ and $\beta_3^h \geq 0$, the sign of the bias depends on the correlation between coverage and the innovation. If more inflationary signals receive more coverage, then $\text{Cov}(news_t, \lambda_t) > 0$ and $\text{Cov}(news_t, \lambda_t news_t) > 0$, and the bias is positive.[5]

Any state dependence in δ^h may also interfere with the interpretation of this coefficient as the effect of the news on expectations. Suppose that on many days, new signals do not arrive and the media recirculates old information. Households may respond differently to coverage on these days and on days where truly novel information arrives (λ_t observed). Denote the response on days with new information as δ_I^h and the response on days without new information as $\delta_0^h \approx 0$, then the pooled estimate is a weighted average,

$$\delta^h \;\approx\; \alpha\, \delta_I^h + (1-\alpha)\, \delta_0^h \;\approx\; \alpha\, \delta_I^h, \tag{4}$$

[5]We show this in Appendix Section D.1.

where α is the share of the (residual) variation in coverage that occurs on days with an information arrival.[6] If such days are infrequent, α is small, and the full-sample estimate is strongly *attenuated* version of the response on days when coverage actually transmits new information.

Scheduled CPI and PCE release days address both problems. They provide observable, predetermined days on which a credible inflation signal arrives. This provides a set of days with information arrival in Equation (4) as well as a proxy for λ_t, the difference between the release and pre-release forecasts. Further, controlling for λ_t and $news_t \times \lambda_t$ mitigates omitted variable bias concerns. More importantly, including $news_t \times \lambda_t$ in the regression allows us to estimate the information amplification channel of the news, captured by β_3^h in Equation 5.

2.2 Baseline Model

Our baseline specification compares expectation dynamics across scheduled CPI and PCE release days with different media environments using the following state-dependent local projection:[7]

$$E_t[\pi_{t+365,t}] = \beta_0^h + \beta_1^h \mathbf{1}(news_t) + \beta_2^h \lambda_t + \beta_3^h (\lambda_t \times \mathbf{1}(news_t)) + \sum_{k=1}^{14} \Phi_k^{h\prime} X_{t-k} + \varepsilon_{t+h}. \quad (5)$$

Here $E_t[\pi_{t+365,t}]$ is the median 12-month-ahead inflation expectation of households. The variable λ_t is the inflation surprise associated with the CPI or PCE release on day t. The variable $\mathbf{1}(news_t)$ measures the inflation-related media environment on

[6] We show this decomposition using the Frisch–Waugh–Lovell theorem in Appendix Section D.2.

[7] Local projections estimate dynamic responses directly at each horizon (Jordà, 2005; Jordà, 2023), and state-dependence is usually implemented through interactions between a shock measure such as monetary policy surprises and indicators for the relevant state, allowing the impulse responses to differ across conditions such as expansions and recessions (Auerbach and Gorodnichenko, 2012; Ramey and Zubairy, 2018; Tenreyro and Thwaites, 2016; Ben Zeev, Ramey, and Zubairy, 2023). Recent work provides formal econometric treatment of state-dependent local projections (Gonçalves et al., 2024). We estimate Equation (5) on the subsample of CPI and PCE release days — equivalent to a local projection fully saturated in a release-day indicator. We are not interested in differentiating the response between release and non-release days; we condition on release days because the information innovation λ_t is observable only on those days.

the release day. In the baseline implementation below, $\mathbf{1}(news_t)$ is an indicator for unusually elevated inflation coverage; we also consider continuous measures of coverage in robustness exercises. The vector X_{t-k} includes calendar-day lags of the outcome and inflation-related news coverage. Lagged outcomes account for persistence in inflation expectations. Lags of news coverage control for pre-existing media attention, so that the news indicator captures unusually elevated coverage conditional on the recent coverage history, instead of a continuation of an ongoing news cycle.

The coefficient β_2^h measures the pass-through of the inflation surprise on normal-coverage release days ($\mathbf{1}(news_t) = 0$). The interaction coefficient β_3^h measures the additional pass-through on high-coverage release days. Hence, the total response to a given surprise when $\mathbf{1}(news_t) = 1$ is $\beta_2^h + \beta_3^h$. We refer to β_3^h as the *amplification effect*. Finally, β_1^h captures the *salience effect*: the conditional difference in expectation dynamics between high- and normal-coverage release days when the measured inflation surprise is zero, holding lagged observables fixed.

Conditioning on CPI and PCE release days ensures that every observation corresponds to the scheduled arrival of public inflation information and focuses the analysis on variation in the media environment across such events. Because release days are scheduled in advance, the timing of the signal is predetermined. Media coverage is not randomly assigned, but comparing release days with different coverage intensity allows us to test whether measured inflation surprises pass through more strongly to household expectations when the releases are more intensely covered.

The next section describes the data used to construct each component of Equation (5).

2.3 Smooth local projections

We estimate dynamic responses using smooth local projections. Standard local projections estimate a separate regression for each horizon, which provides flexibility but can produce noisy impulse responses, especially at longer horizons and with high-frequency outcomes (Ramey, 2016; Jordà, 2023). This concern is particularly relevant in our setting because daily household inflation expectations are volatile and

the estimation sample is restricted to CPI and PCE release days (Binder, Campbell, and Ryngaert, 2024).

Introducing mild restrictions across horizons can substantially improve efficiency while largely preserving the flexibility of the local projection approach (Jordà, 2023). We therefore adopt the smooth local projections methodology of Barnichon and Brownlees (2019), which regularizes impulse response estimates using penalized B-splines following Eilers and Marx (1996). Specifically, we approximate the sequence of horizon-specific coefficients $\{\beta^h\}_{h=0}^H$ using a B-spline basis expansion:

$$\beta^h \approx \sum_{k=1}^{K} b_k B_k(h), \tag{6}$$

where $B_k(h)$ denote B-spline basis functions spanning the forecast horizon. The spline coefficients b_k are estimated via a generalized ridge estimator. We implement this building on the `SmoothLP` package provided by Bousquet (2024), modified to allow the state-dependent interactions in Equation (5).

Following Barnichon and Brownlees (2019), we select the penalization parameter using a data-driven cross-validation procedure for a fine grid from 0.00001 to 1. In our application, the data chooses the upper bound of about 1, which is then scaled by a factor of 0.01. This "undersmoothing" takes into account that inference for smoothed estimates is still an open issue and that smoothing, by introducing bias, mechanically narrows the confidence intervals (Barnichon and Brownlees, 2019; Bousquet, 2024).[8]

3 Data

Our empirical analysis combines three daily data components over the period June 2013 to September 2024: surprises around scheduled CPI and PCE releases, cable-television coverage of inflation, and household inflation expectations. This section describes the construction of each measure and then outlines the smooth local

[8]While removing noise, our results are robust when using unrestricted local projections; see Appendix section A.1. In Appendix section C, we provide implementation details for the smooth local projection method.

projection procedure used to estimate the dynamic responses.

3.1 Inflation Surprises

To measure the informational content of inflation releases, we construct surprise measures as the difference between realized core inflation and the corresponding professional forecast:

$$\lambda_t^k = \pi_t^{k,\text{realized}} - \pi_t^{k,\text{forecast}}, \qquad k \in \{\text{CPI}, \text{PCE}\}.$$

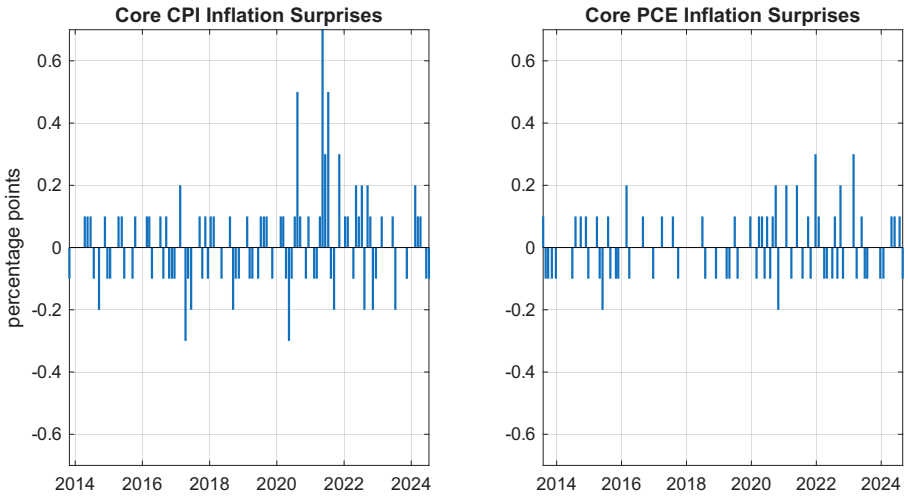
A positive value of λ_t^k indicates that inflation was higher than expected, while a negative value indicates that inflation was lower than expected. The forecast we use is the median forecast from Bloomberg's survey of economists, using the latest survey value available before the release. The resulting forecast error isolates the component of the release that was unexpected relative to the professional consensus, rather than the level of inflation itself.

We use these professional forecast errors as an empirical proxy for the news content of the release. Release-day media coverage frequently frames inflation data relative to market or professional expectations, by using language such as "higher than expected" or "cooler than expected". The surprise measure therefore aligns with the benchmark used by the media to characterize the informational content of a release, even if households do not directly observe the professional consensus itself or share this forecast as their own prior.

Following Chahrour, Shapiro, and Wilson (2024), we focus on surprises to *core* inflation rather than headline inflation. Core inflation developments are more likely to be emphasized, interpreted, and contextualized in release-day news coverage, making higher-than-expected or lower-than-expected core inflation a salient signal about underlying inflation pressure. By contrast, households may directly observe some non-core price changes, such as gasoline or food prices, and may therefore rely less on news media to learn about those components.[9]

[9]In Appendix section **??**, we show that our main results are robust to using alternative surprise measures, including headline inflation surprises, the change in year-over-year headline inflation

Figure 1: Inflation Surprise Series



Notes: Daily inflation surprise series constructed as the difference between realized core CPI (PCE) inflation and Bloomberg Survey forecasts. Sample period: June 1, 2013 to September 30, 2024.

Figure 1 plots the surprise series for both core CPI and core PCE inflation. Most surprises are modest, typically between 0.1 and 0.2 percentage points, but forecast errors increased substantially during the post-pandemic inflation surge. The largest positive core CPI surprise occurred on May 12, 2021, when April core CPI exceeded expectations by 0.7 percentage points. The largest core PCE surprises, around 0.3 percentage points, occurred in late 2021.

A natural concern is that PCE surprises may contain little new information because PCE releases follow CPI releases within the same month. If CPI releases fully revealed underlying inflation developments, subsequent PCE surprises would be largely predictable. To assess this, we regress core PCE surprises on core CPI surprises. The estimated coefficient is small and statistically insignificant (0.09, s.e. 0.06), indicating that CPI surprises do not systematically predict subsequent PCE surprises. PCE releases therefore retain an unexpected component not mechanically explained by the earlier CPI release.

relative to the preceding release, and the gap between the reported inflation rate and the trailing average of household expectations.

3.2 News Coverage

The Stanford Cable TV News Analyzer allows researchers to compute the length of time that words are spoken on the three major cable news stations—CNN, Fox, and MSNBC (Hong et al. 2021). The underlying dataset, provided by the Internet Archive's TV News Archive, includes nearly 24-7 video coverage of these stations, with accompanying transcripts, beginning on January 1, 2010. Our primary measure of interest is the number of seconds per day that the word "inflation" is heard on cable television, summed across the three stations. For text searches the recorded time is only the time in which the word is said, not the length of the surrounding discussion; since "inflation" takes roughly one second to say, each second of coverage corresponds to about one mention.

Panel (a) of Figure 2 plots this daily series over our sample. Inflation coverage is minimal for most of the pre-pandemic period and rises sharply during the 2021–22 inflation surge. In 2013–2020, 30% of days have zero mentions and 79% have at most four seconds, with average coverage of 3 seconds per day. From 2021 to 2024 coverage is far higher: only 2% of days have zero mentions, and average coverage is 99 seconds per day with a standard deviation of 109 seconds.

News indicator. We use this continuous series, denoted $news_t$, to construct a binary news indicator, $\mathbf{1}(news_t)$, that flags days on which inflation coverage is unusually intense relative to its recent baseline. We interpret this indicator as a measure of media amplification: it captures days on which the information environment surrounding inflation is unusually intense.

Specifically, a day t is classified as a high-coverage day if

$$news_t \;>\; \overline{news}_{t-10:t-1} \;+\; \sigma_{t-180:t-1},$$

where $\overline{news}_{t-10:t-1}$ is the average coverage over the previous 10 days and $\sigma_{t-180:t-1}$ is the standard deviation over the previous 180 days. This time-varying threshold adapts to changing reporting intensity, accommodating both the low-coverage period prior to 2021 and the post-pandemic surge, during which overall coverage increased

Table 1: Summary Statistics for Inflation Release Days

	All releases	CPI releases	PCE releases
Number of release days	255	134	121
Share with high news coverage (%)	21.18	32.84	8.26
Mean news coverage (seconds)	59.36	79.94	36.57
Median news coverage	4	5	4
Standard deviation of news coverage	125.26	154.74	75.35
Mean inflation surprise	0.007	0.012	0.001
Mean absolute inflation surprise	0.075	0.088	0.06
Standard deviation of inflation surprise	0.116	0.136	0.09
Corr (news coverage, inflation surprise)	0.11	0.146	-0.039

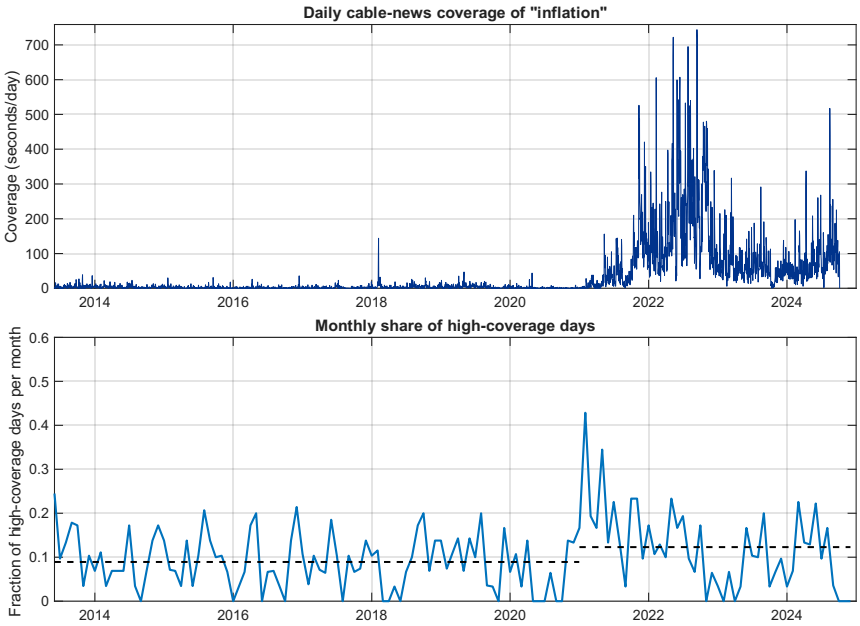
Notes: The table reports summary statistics for scheduled CPI and PCE release days. High news coverage is defined according to the news indicator, $seconds_t > \overline{seconds}_{t-10:t-1} + \sigma_{t-180:t-1}$. Inflation news coverage is measured as seconds of television coverage of "inflation." Inflation surprises are realized core inflation minus the corresponding pre-release professional forecast, in percentage points.

substantially. Panel (b) of Figure 2 reports the monthly share of high-coverage days. Such days occur throughout the sample; the average monthly share is 8.9% during 2013–2020 and 12.3% during 2021–2024, and high-coverage days account for 10.1% of all days in the sample.[10]

High-coverage days are not a relabeling of release days: they capture variation in coverage intensity both on and off scheduled release dates. Only about 13% of high-coverage days fall on a CPI or PCE release—even though release days make up just 6.3% of the sample—and, conversely, only a subset of releases are amplified: 32.8% of CPI release days qualify as high-coverage days, compared with 8.3% of PCE release days. In the baseline analysis, we use scheduled release days together with the news indicator to study whether inflation surprises have larger effects when a release occurs in a high-coverage media environment.

[10]Appendix A.4 demonstrates robustness of our main results when using alternative definitions of the news indicator, or a continuous measure.

Figure 2: Cable-TV Inflation Coverage and the News Indicator



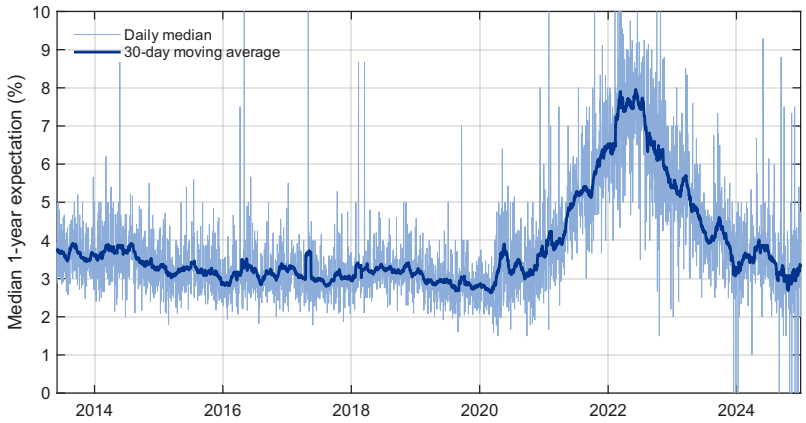
Notes: Panel (a) shows the seconds per day that the word "inflation" is said on CNN, Fox News, and MSNBC combined. Panel (b) shows the share of days per month classified as high-coverage days under the baseline news indicator ($seconds_t > \overline{seconds}_{t-10:t-1} + \sigma_{t-180:t-1}$); dashed lines mark the average monthly shares during 2013–2020 (8.9%) and 2021–2024 (12.3%). Sample period: June 1, 2013 to September 30, 2024.

3.3 Inflation Expectations Data

Our data on inflation expectations come from the Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE). Binder, Campbell, and Ryngaert (2024) document that, although this is a monthly survey, respondents take it throughout the month and their characteristics are consistent across the month. Respondents also do not appear to select into taking the survey around particular events, such as FOMC announcements or data releases. We therefore construct a daily frequency.

Survey respondents provide forecasts in two formats: numerical point forecasts and density forecasts. Our baseline measure of inflation expectations is the daily cross-sectional median of the one-year-ahead numerical point forecast. Figure 3 plots this daily series together with a 30-day moving average. The median expectation is stable near 3% through 2020, rises sharply during the 2021–22 inflation surge to a peak of roughly 7.5% in mid-2022, and declines thereafter.

Figure 3: Daily Consumer Inflation Expectations



Notes: Daily cross-sectional median of the one-year-ahead numerical point forecast from the Survey of Consumer Expectations (light line), with a 30-day moving average (dark line). Sample period: June 1, 2013 to September 30, 2024.

4 Empirical Results

In this section, we present our baseline evidence on two distinct roles of news coverage: the salience effect, captured by β_1^h, and the amplification effect, captured by β_3^h. We also show that

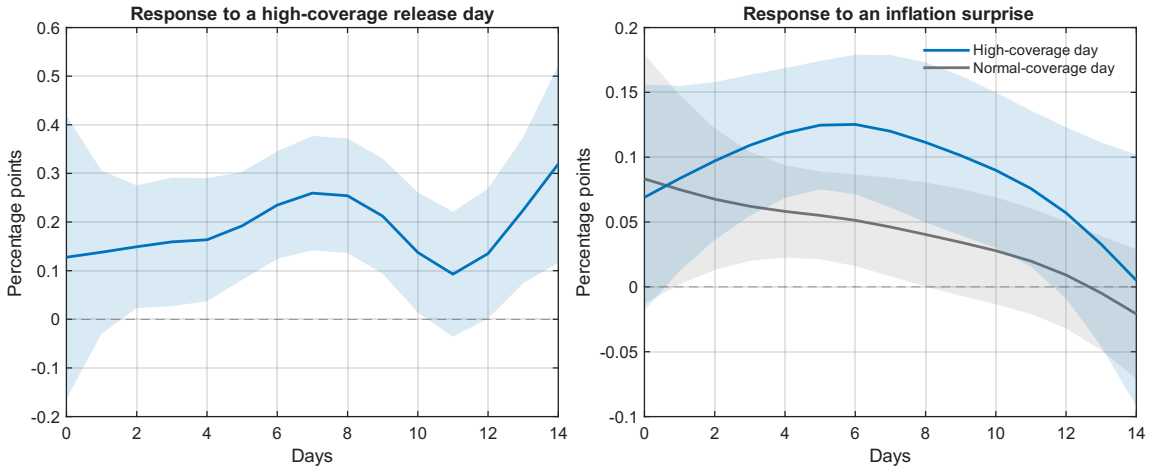
4.1 The Two Effects of the News

We estimate the baseline local projection from Equation (5) using the daily median 12-month-ahead point forecast as our measure of household inflation expectations. The model pools CPI and PCE releases and includes the news indicator, the inflation surprise, and their interaction. Because the model is estimated conditional on release days, the omitted category is a normal-coverage release day. As discussed in Section 2.2, $\{\beta_1^h\}_{h=0}^{14}$ traces the conditional response associated with a high-coverage release day when the measured surprise is zero, while $\{\beta_3^h\}_{h=0}^{14}$ captures the additional pass-through of a given surprise in a high-coverage environment.

Figure 4 reports the estimated dynamic responses. The left panel plots the

response associated with a high-coverage release day, $\{\beta_1^h\}_{h=0}^{14}$, when the measured inflation surprise is zero. The response is positive, builds over the first week to a peak of approximately 0.25 percentage points after seven days, and is statistically significant over most of the two-week horizon. Thus, even in the absence of a measured inflation surprise, a high-coverage release day is associated with a persistent increase in household inflation expectations. [11]

Figure 4: Response of Household Inflation Expectations to a High-Coverage Release Day and an Inflation Surprise



Notes: Dynamic response of the daily median 12-month-ahead inflation expectation to a high-coverage release day (left panel) and to a 0.1 percentage point core inflation surprise (right panel) on normal-coverage (dashed grey) and high-coverage (solid blue) release days. Responses are in percentage points. Shaded areas denote 90% confidence intervals based on Newey–West HAC standard errors. Impulse responses are estimated using smooth local projections. Five-fold cross-validation selects a penalty parameter of 1 from the candidate grid. Applying an undersmoothing factor of 0.01 yields a penalty parameter of 0.01 for the reported estimates. Sample period: June 1, 2013 to September 30, 2024.

The right panel examines the amplification effect, β_3^h. The dashed gray line shows the response to a 0.1 percentage point core inflation surprise on a normal-coverage

[11] Appendix section A shows that this result is robust to alternative definitions of inflation surprises, alternative definitions of the news indicator, and to an alternative measure of inflation expectations.

release day, $\{0.1\beta_2^h\}_{h=0}^{14}$, while the solid blue line shows the corresponding response on a high-coverage release day, $\{0.1(\beta_2^h + \beta_3^h)\}_{h=0}^{14}$. The vertical difference between the two paths at each horizon is therefore $0.1\beta_3^h$, the estimated amplification effect.

If coverage did not affect the pass-through of the information innovation, the two response paths would coincide. Instead, they differ in both magnitude and persistence. On a normal-coverage release day, a 0.1 percentage point surprise raises expectations by approximately 0.08 percentage points on impact, and the response fades to zero within roughly eight days. On a high-coverage release day, the response builds to a peak of approximately 0.13 percentage points after six days and remains statistically significant for roughly eleven days. The estimated amplification effect is therefore small on impact but becomes more pronounced over the first week, as the response under high coverage becomes larger and more persistent.

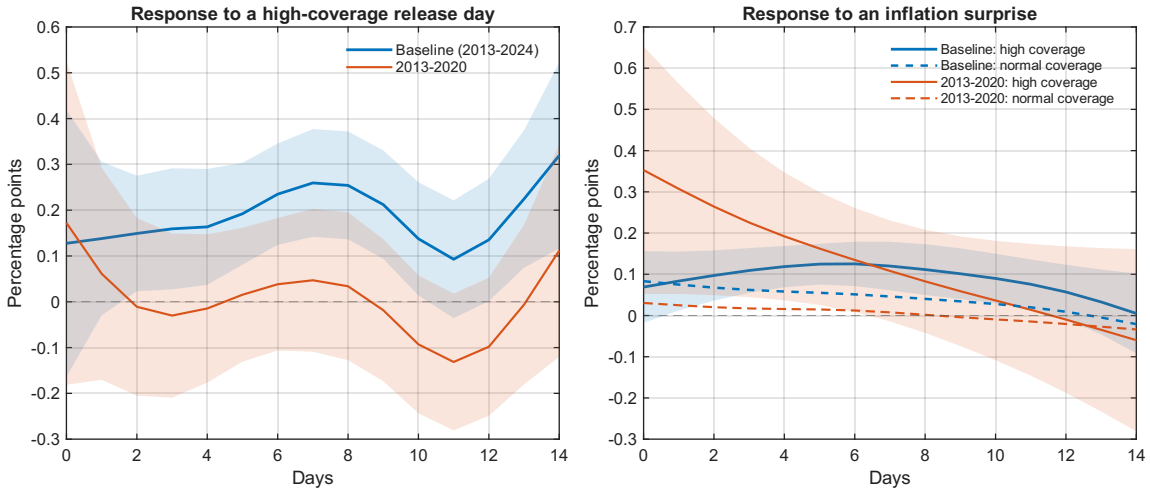
Taken together, the two panels document distinct roles for news coverage. The positive β_1^h response means that, in our sample, high coverage induces households to increase their expectations even when the current release contains no measured surprise. The positive β_3^h indicates an amplification effect: a given inflation surprise passes through more strongly when the release receives intense coverage. The transmission of macroeconomic information to households therefore depends not only on the informational content of the release, but also on the media environment in which it arrives.

4.2 The Role of the High Inflation Period

Our sample period includes the post-pandemic inflation surge in 2021–2022. To assess whether the results depend on the inclusion of this period, we re-estimate Equation (5) restricting the estimation sample to 2013–2020. Figure 5 compares these estimates with those obtained over the full baseline sample, 2013–2024.

The left panel shows that the persistent positive response of high-coverage days relative to normal coverage days is not apparent in the 2013–2020 sample. The estimates of β_1^h fluctuate around zero, and the confidence bands include zero at every horizon. This suggests that the estimated positive effect is time-varying and specific

Figure 5: Response of Inflation Expectations: Baseline and 2013–2020 Sample



Notes: The figure compares estimates for the baseline sample, 2013–2024 (blue), with estimates for 2013–2020 (orange). The left panel reports the response of the daily median one-year-ahead inflation expectation to a high-coverage release day when the measured inflation surprise is zero, β_1^h. The right panel reports responses to a 0.1 percentage-point core inflation surprise on high-coverage release days, $0.1(\beta_2^h + \beta_3^h)$ (solid), and normal-coverage release days, $0.1\beta_2^h$ (dashed). The solid-minus-dashed difference within each sample is the amplification effect, $0.1\beta_3^h$. Shaded areas show 90% confidence intervals.

to the high inflation period.

The amplification effect of an information effect, however, remains visible in the earlier sample, particularly at short horizons. In the right panel, the estimated response to a 0.1 percentage-point core inflation surprise on a high-coverage release day is approximately 0.35 percentage points on impact, compared with approximately 0.03 percentage points on a normal-coverage release day.

These results distinguish two empirical patterns. Amplification remains visible when the high inflation years are excluded, whereas the persistent positive response associated with high coverage when the measured surprise is zero does not. In the next section, we develop a model that provides an interpretation of these two effects of news.

5 A Model of Noisy Information with Expectations "Catch-up"

We find that data releases with increased media coverage influence household expectations in two ways. They amplify the effect of measured inflation surprises and they induce upward revision even after controlling for the measured surprise. A model of expectation formation as a mixture of noisy information and infrequent updating can explain our empirical findings. In this model, news increases the probability that households observe information and allows households with stale information sets a chance to "catch-up," leading to additional revisions when expectations have fallen behind.

First, consider the expectations of an agent who observes an inflation signal every day and rationally updates her prior:

$$E_t^{Obs.}[\pi_{t+365,t}] = (1-\kappa)E_{t-1}^{Obs.}[\pi_{t+365,t}] + \kappa\rho s_t, \tag{7}$$

$E_t^{Obs.}[\pi_{t+365,t}]$ is the daily expectation of inflation over the next 12 months on day t and $E_{i,t-1}^{Obs.}[\pi_{t+365,t}]$ is the expectation over the same 12-month period one day prior.[12]

[12] $E_{t-1}^{Obs.}[\pi_{t+365,t}]$ will depend on the household's expectation of year-ahead inflation on day $t-1$,

The noisy signal $s_t = \pi_{t,t-365} + u_t$, with $u_t \sim N(0, \sigma_u^2)$ communicates information about the *current* rate of inflation and $\kappa > 0$ captures how strongly the agent updates her expectatons in the direction of the signal. As is standard in noisy information models, households form expectations about the current level of inflation and project this belief forward according to the data-generating process for inflation. Accordingly, ρ denotes the *annual* persistence of the inflation process.

Let λ_t denote the approximate signal innovation, that is the "surprise" that the signal constitutes to the previous prior expectation:

$$\lambda_t \approx \rho s_t - E_{t-1}^{Obs.}[\pi_{t+365,t}]$$

positive values of λ_t should lead to upward revisions in expectations, while negative values prompt downward revisions. We can therefore rewrite Equation 7 as

$$E_t^{Obs.}[\pi_{t+365,t}] = E_{t-1}^{Obs.}[\pi_{t+365,t}] + \kappa\lambda_t, \tag{8}$$

Suppose now that not all households observe the signal every day. The prior will then depend on the number of days since the household observed the signal: $E_{t-j}^{Obs.}[\pi_{t+365,t}]$. Respondents who do not observe the signal continue with this prior. The effect of the signal on those who observe it is now two-fold. First, they can update their information set to the previous day's prior - that is they can "catch-up" on all inflation related news. Second, given the $t-1$ prior, they proceed with rational signal processing as described in Equation 7. The resulting expectation from the $t-j$ prior are:

$E_{t-1}^{Obs.}[\pi_{t+364,t-1}]$, and the household's belief about how this expectation translates to inflation over the period $t+365, t$. Due to the high frequency of our data, we assume that $E_{t-1}^{Obs.}[\pi_{t+364,t-1}] \approx E_{t-1}^{Obs.}[\pi_{t+365,t}]$ and the lagged dependent variable can proxy for the prior expectation.These horizons overlap with the exception of days $t-1$ and $t+365$ and therefore differ only in the daily innovations on day $t+365$ and the daily innovation on $t-1$.

$$E^{Obs.}_{t,t-j}[\pi_{t+365,t}] = E^{Obs.}_{t-1}[\pi_{t+365,t}] + \kappa\lambda_t$$
$$= \underbrace{E^{Obs.}_{t-1}[\pi_{t+365,t}] - E^{Obs.}_{t-j}[\pi_{t+365,t}]}_{Revision\ of\ Prior} + \underbrace{E^{Obs.}_{t-j}[\pi_{t+365,t}]}_{Pre-Signal\ Prior} + \underbrace{\kappa\lambda_t}_{Update} \quad (9)$$

Note that while all individuals who observe the signal in time t will have the same expectation (a simplifying assumption that could be relaxed by allowing idiosyncratic noise in observed signals), their forecast revisions will differ and depend on j. Denote the revision of prior, or "catch-up" term as $C_{t-1,t-j}$

$$E^{Obs.}_{t,t-j}[\pi_{t+365,t}] - E^{Obs.}_{t-j}[\pi_{t+365,t}] = C_{t-1,t-j} + \kappa\lambda_t \quad (10)$$

For those who update every day, $j = 1$ and $C_{t-1,t-j} = 0$. If $j > 1$, $C_{t-1,t-j}$ describes the update from a stale information set. If inflation has changed significantly since the last time the household viewed the signal, the magnitude of this term may be large.

To incorporate the role of news in expectations, let p_t denote the probability that a household is exposed to the inflation signal on date t, either directly through the news or indirectly through conversations, social media, or other forms of interpersonal transmission. The aggregate expectation for respondents who last viewed the signal j periods ago is then:

$$E^{Obs.}_{t,t-j}[\pi_{t+365,t}] = p_t\left[C_{t-1,t-j} + \kappa\lambda_t\right] + E^{Obs.}_{t-j}[\pi_{t+365,t}] \quad (11)$$

To find the aggregate expectation for the economy, first define the share of respondents who last updated at j as ω_j such that $\sum_{j=1}^{\infty}\omega_j = 1$. The average expectation is then the weighted sum of expectations:

$$
\begin{aligned}
\bar{E}_t[\pi_{t+365,t}] &= \sum_{j=1}^{\infty} \omega_j E^{Obs.}_{t,t-j}[\pi_{t+365,t}] \\
&= \sum_{j=1}^{\infty} \omega_j p_t \left[C_{t-1,t-j} + \kappa\lambda_t\right] + \sum_{j=1}^{\infty} \omega_j E^{Obs.}_{t-j}[\pi_{t+365,t}] \\
&= p_t \sum_{j=1}^{\infty} \omega_j C_{t-1,t-j} + p_t\kappa\lambda_t \sum_{j=1}^{\infty} \omega_j + \bar{E}_{t-1}[\pi_{t+365,t}] \\
&= p_t\bar{C}_{t-1} + p_t\kappa\lambda_t + \bar{E}_{t-1}[\pi_{t+365,t}] \qquad (12)
\end{aligned}
$$

We assume that p_t is increasing in news coverage of inflation such that more intense coverage increases the probability that households encounter the signal. For simplicity and to match the binary news variable from Section 3.2, allow p_t to take one of two possible values:

$$
p_t = \begin{cases} p_L \text{ if } \mathbf{1}(news_t) = 0 \\ p_H \text{ if } \mathbf{1}(news_t) = 1 \end{cases}
$$

where $p_H - p_L > 0$. We can then rewrite Equation 12 as:

$$
\begin{aligned}
\bar{E}_t[\pi_{t+365,t}] &= p_L\bar{C}_{t-1} + (p_H - p_L)\bar{C}_{t-1}\mathbf{1}(news_t) \\
&\quad + p_L\kappa\lambda_t + (p_H - p_L)\kappa\lambda_t\mathbf{1}(news_t) + \bar{E}_{t-1}[\pi_{t+365,t}] \qquad (13)
\end{aligned}
$$

providing approximate mappings to the coefficients discussed in Section 2.2. Specifically, $(p_H - p_L)\bar{C}_{t-1}$ is the additional average catch-up that occurs on a high news day relative to a normal day, β_1, and $(p_H - p_L)\kappa$ is the additional amplification of the signal, β_3. As $\kappa \geq 0$ and $(p_H - p_L) > 0$, $\beta_3 \geq 0$. The sign of the salience effect, β_1, however, depends on the sign of $\bar{C}_{t-1}$ and is potentially time-varying. More specifically, it depends on the information households have previously missed. During periods of rising inflation, households that have not followed the news may accumulate a larger "backlog" of missed information pointing to higher inflation than during periods of low and stable inflation rates. Elevated coverage can then prompt upward

catch-up revisions. This is consistent with our finding in Section 4.2 of a positive salience effect in the full sample, but not for the 2013–2020 sample. In the next section, we exploit partisan differences in prior inflation expectations to shed further light on this mechanism.

6 Prior News Exposure and the Catch-Up Response

Our model predicts that households' response to elevated news coverage depends on the information they have already incorporated into their expectations, which may differ across households. In particular, in Equation (13), $\bar{C}_{t-1}$ is the "catch up" averaged across consumers who last updated on all days $t-j$. During a period of rising inflation, as experienced from 2021–2023, that "catch-up" term should be positive, and should be larger for respondents whose last update was longer ago. That is, respondents whose priors are further below an informed benchmark should have more room for upward catch-up, and therefore should exhibit a larger salience response.

Here we consider whether differential catch-up in response to elevated coverage can help explain the striking partisan dynamics in inflation expectations during the 2021–2023 inflation surge. Binder, Kamdar, and Ryngaert (2024) document that Republican expectations increased more than Democratic expectations, and moreover, that responses on release days often differed significantly along partisan lines. Respondents do not report political affiliation in the SCE. We use instead the partisan tilt of different news stations and differences in viewership by state to construct groupings of likely Republican-leaning and likely Democrat-leaning households.[13]

[13]Using data from Nationscape, a large public opinion survey conducted in 2020 (Tausanovitch and Vavreck, 2021), on news viewership in each state, we categorize states as either predominantly Fox-viewing or predominantly CNN-viewing. The relative partisan viewership of these stations mean that our Fox-viewing and CNN-viewing samples provide our Republican-leaning and Democrat-leaning samples, respectively.

The capacity for revision in this period differs along partisan lines. Figure 6 plots the difference in inflation coverage and in prior expectations over 2021–2024. The top panel shows that Fox News devoted more time to cover news about inflation compared to CNN, and the difference is particularly pronounced in the period of the post-pandemic inflation surge. The bottom panel shows the difference in median prior inflation expectations, in Fox and CNN states, $\bar{E}_{Fox,t-1} - \bar{E}_{CNN,t-1}$.[14] In general, prior expectations were higher in Fox-leaning states during the surge, correlating with the more intense coverage in these states.[15]

These differences in prior expectations are a proxy for the differences in the average catch-up gap implied by the model:[16]

$$\bar{C}_{CNN,t-1} - \bar{C}_{Fox,t-1} = \bar{E}_{Fox,t-1} - \bar{E}_{CNN,t-1}.$$

The higher prior expectations in Fox-leaning states shown in Figure 6 suggests that respondents in CNN-leaning states entered high-coverage release days with greater scope for upward catch-up revisions. A testable prediction is therefore a larger salience response in CNN-leaning states: $\beta_1^{CNN} > \beta_1^{Fox}$.

We test this prediction directly by re-estimating Equation (5), replacing the aggregate outcome with the median one-year-ahead point forecast computed separately for respondents in Fox-leaning and CNN-leaning states.[17] We construct network-specific high-coverage indicators for Fox News and CNN using the same definition as in the baseline analysis and define a joint high-coverage indicator equal to one only when inflation coverage is elevated on both networks. The response of both groups is

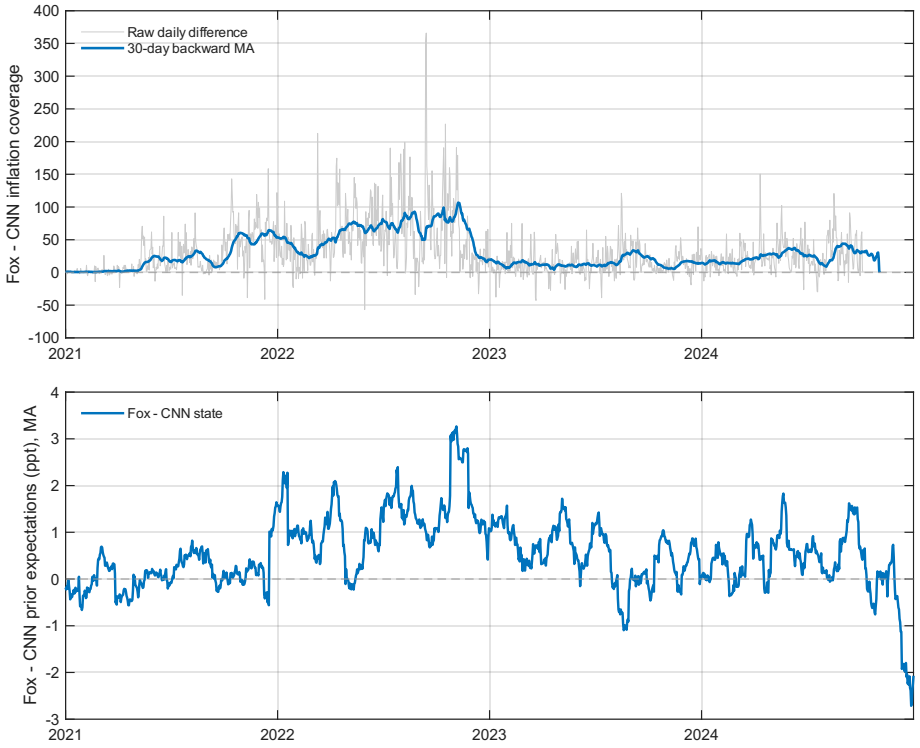
[14]For respondents interviewed in t, we use the one-year-ahead point forecasts from the previous monthly interview and compute the median for the CNN and the Fox News group. This is our proxy for the prior inflation expectations.

[15]The correlation coefficient of the coverage difference and the prior expectations difference over the 2021–2024 window is 0.53.

[16]To see this, consider the definition of the catch-up term in Equation (9), where $C_{t-1,t-j} = E_{t-1}^{\text{Obs}} - E_{t-j}^{\text{Obs}}$. Averaging within each group as in Equation (12) gives $\bar{C}_{g,t-1} = E_{t-1}^{\text{Obs}} - \bar{E}_{g,t-1}$. With a common benchmark E_{t-1}^{Obs} across groups, the difference between CNN and Fox states in average catch-up gaps therefore equals the difference of Fox and CNN states in average prior expectations.

[17]The estimation uses the full baseline sample, from June 1, 2013 to September 30, 2024.

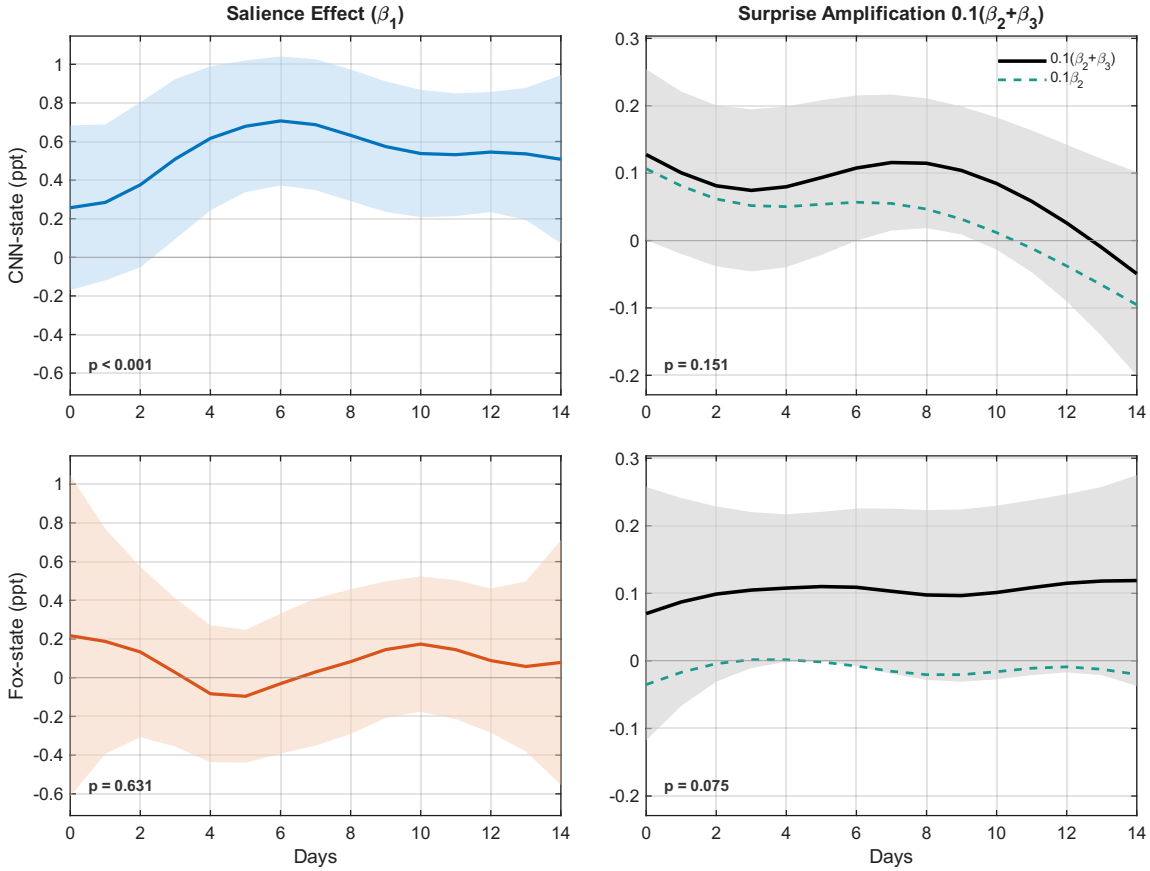
Figure 6: Differences in Inflation News Coverage and Prior Expectations: Fox Minus CNN



Notes: Both panels report Fox-CNN differences over January 2021 to December 2024. *Top:* daily difference in inflation news coverage together with the 30-day backward moving average; *Bottom:* 30-day backward moving average of the difference in prior one-year ahead inflation expectations across Fox and CNN states. The prior is measured as the median over the respondent's previous monthly interview.

therefore evaluated around the same high-coverage release days.[18] Figure 7 reports the resulting estimates of the salience and amplification effects.

Figure 7: Salience and Inflation-Surprise Responses in CNN- and Fox-Leaning States



Notes: The left column reports the salience response, β_1^h, for respondents in CNN-leaning states (top) and Fox-leaning states (bottom). The right column reports the response to a 0.1 percentage-point core inflation surprise on joint high-coverage release days, $0.1(\beta_2^h + \beta_3^h)$ (solid), and on other release days, $0.1\beta_2^h$ (dashed). The vertical difference between these paths is the amplification effect, $0.1\beta_3^h$. Shaded areas denote 90% confidence intervals based on Newey–West HAC standard errors. Impulse responses are estimated using smooth local projections. Five-fold cross-validation selects a penalty parameter of 1 from the candidate grid. Applying an undersmoothing factor of 0.01 yields a penalty parameter of 0.01 for the reported estimates. Sample period: June 1, 2013 to September 30, 2024. Panel p-values test whether the average solid-line response over horizons $h = 0, \ldots, 14$ equals zero.

[18]The set of controls includes 14 calendar day lags of the group-specific outcome variable as well as 14 lags of inflation coverage in Fox News and CNN (in seconds).

Table 2: Average Responses and Differences by News Viewership

Response	CNN-leaning states	Fox-leaning states	CNN minus Fox difference	p-value of difference
Salience, β_1	0.530^{***}	0.080	0.450	0.145
	(0.157)	(0.161)	(0.312)	
Amplification, $0.1\beta_3$	0.047	0.116^{*}	-0.068	0.717
	(0.058)	(0.069)	(0.188)	
Total under high coverage,	0.070	0.100^{*}	-0.030	0.853
$0.1(\beta_2 + \beta_3)$	(0.052)	(0.058)	(0.158)	

Notes: Entries report average estimated responses over horizons $h = 0, \ldots, 14$, in percentage points, with standard errors in parentheses. The second and third rows normalize the inflation surprise to 0.1 percentage points. A joint high-coverage indicator equals one when inflation coverage is elevated on both CNN and Fox News. Estimates use the baseline specification with a smoothing penalty of 1. Within-group inference uses the full covariance matrix across horizons. Inference for CNN-minus-Fox differences uses a moving-block bootstrap with 1,000 replications and a block length of 90 days. Differences are computed before rounding. Stars refer to two-sided tests against zero: $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$. The final column reports two-sided p-values for differences between groups.

The estimates of β_1^h differ markedly across the two groups in a manner consistent with news exposure prompting the revision of a stale prior. For respondents in CNN-leaning states β_1^h builds to a peak of approximately 0.7 percentage points after six days and remains positive throughout. For respondents in Fox-leaning states, this estimated effect is considerably smaller and statistically indistinguishable from zero. Together with Figure 6, this pattern is consistent with respondents in Republican-leaning states having less scope for catch-up. Table 2 summarizes the average responses over the 14-day horizon.

An alternative explanation for the salience pattern is that elevated coverage generates a larger increase in exposure among respondents in CNN-leaning states. Equation 13 shows that $\beta_1^h \approx (p_H - p_L)\bar{C}_{g,t-1}$. It may be that $p_H - p_L$ is smaller for Fox viewers because Fox was already so elevated during this period that high coverage days did not substantially change the probability of exposure. In the context of our model, this would imply that the amplification response, β_3, is then also larger for

CNN-states. The right column of Figure 7 plots the responses to an inflation surprise on both high- and normal- coverage days. Their difference is the additional response attributable to amplification, $0.1\beta_3^h$. However, we *do not observe* that $\beta_3^{CNN} > \beta_3^{Fox}$. Rather, we reject the null that $\beta_3^{Fox} = 0$ at the 10% significance level and fail to reject the null that $\beta_3^{CNN} = 0$.

Taken together, these results are consistent with news exposure shaping the scope for catch-up. During the inflation surge, respondents in Fox-leaning states had higher prior expectations in a media environment with more inflation coverage. In the group-specific estimates, their salience response to high-coverage releases is smaller, while their amplification response is not. We conclude that these results support the catch-up interpretation.

7 Conclusion

In this paper, we examine how news coverage shapes the transmission of inflation information to household expectations. Combining daily household expectations with television news coverage and surprises around scheduled inflation releases, we find that expectations respond more strongly to inflation surprises when coverage is unusually elevated. High-coverage releases are also associated with higher expectations conditional on the measured surprise. These findings distinguish the amplification effect of new information from an additional salience response, due to households updating a stale information set.

A model combining noisy information with infrequent updating provides an interpretation of these patterns. Greater coverage increases the share of households who incorporate the current surprise and catch up on information they previously missed. The salience response therefore depends on prior information: during rising inflation, those whose expectations have fallen behind may make larger upward revisions when coverage becomes elevated. Differences in prior expectations and estimated responses across CNN-leaning and Fox-leaning states provide suggestive evidence consistent with this mechanism.

Our findings suggest that understanding the transmission of public inflation

information requires considering both the content of a release and the information households have already incorporated. Information that is familiar to attentive observers can still move aggregate expectations when wider coverage brings it to a broader audience. The media environment thus helps explain why a given inflation surprise can be associated with different expectation responses across episodes and groups.

References

Auerbach, Alan J. and Yuriy Gorodnichenko (2012). "Measuring the Output Responses to Fiscal Policy". In: *American Economic Journal: Economic Policy* 4.2, pp. 1–27.

Barnichon, Regis and Christian Brownlees (2019). "Impulse Response Estimation by Smooth Local Projections". In: *The Review of Economics and Statistics* 101.3, pp. 522–530.

Ben Zeev, Nadav, Valerie A. Ramey, and Sarah Zubairy (2023). "Do Government Spending Multipliers Depend on the Sign of the Shock?" In: *AEA Papers and Proceedings* 113, pp. 382–387.

Binder, Carola (2017a). "Fed Speak on Main Street: Central bank communication and household expectations". In: *Journal of Macroeconomics* 52, pp. 238–251.

— (2017b). "Federal Reserve Communication and the Media". In: *Journal of Media Economics* 30.4, pp. 191–214.

— (2021a). "Presidential Antagonism and Central Bank Credibility". In: *Economics and Politics* 33.2, pp. 244–263.

Binder, Carola, Jeffrey Campbell, and Jane Ryngaert (2024). "Consumer Inflation Expectations: Daily Dynamics". In: *Journal of Monetary Economics* 145.

Binder, Carola Conces, Rupal Kamdar, and Jane M. Ryngaert (Nov. 2024). "Partisan expectations and COVID-era inflation". In: *Journal of Monetary Economics*. Inflation in the COVID Era and Beyond 148, p. 103649. ISSN: 0304-3932. DOI: 10.1016/j.jmoneco.2024.103649. URL: https://www.sciencedirect.com/science/article/pii/S0304393224001028 (visited on 07/12/2025).

Bousquet, Paul (2024). *SmoothLP: A Package for Smooth Local Projections*. GitHub repository. Accessed: 2025-02-28. URL: https://github.com/paulbousquet/SmoothLP.

Carroll, Christopher (2003). "Macroeconomic Expectations of Households and Professional Forecasters". In: *Quarterly Journal of Economics* 118.1, pp. 269–298.

Chahrour, Ryan, Kristoffer Nimark, and Stefan Pitschner (2021). "Sectoral media focus and aggregate fluctuations". In: *American Economic Review* 111 (12), pp. 3872–3922.

Chahrour, Ryan, Adam Shapiro, and Daniel Wilson (2024). "News Selection and Household Inflation Expectations". In: *Federal Reserve Bank of San Francisco Working Paper*.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber (2022). "Monetary policy communications and their effects on household inflation expectations". In: *Journal of Political Economy* 130.6.

Coibion, Olivier et al. (2020). "Inflation Expectations as a Policy Tool?" In: *Journal of International Economics* 24788.

D'Acunto, Francesco and Michael Weber (Aug. 2024). "Why Survey-Based Subjective Expectations Are Meaningful and Important". en. In: *Annual Review of Economics* 16.Volume 16, 2024. Publisher: Annual Reviews, pp. 329–357. ISSN: 1941-1383, 1941-1391. DOI: 10.1146/annurev-economics-091523-043659. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-economics-091523-043659 (visited on 06/16/2026).

D'Acunto, Francesco et al. (2021). "Exposure to Grocery Prices and Inflation Expectations". In: *Journal of Political Economy* 129.5, pp. 1615–1639.

Ehrmann, Michael, Damjan Pfajfar, and Emiliano Santoroc (2017). "Consumers' Attitudes and Their Inflation Expectations". In: *International Journal of Central Banking* 13 (1), pp. 225–259.

Eilers, Paul and Brian Marx (1996). "Flexible Smoothing with B-Splines and Penalties". In: *Statistical Science* 11, pp. 89–102.

Gonçalves, Silvia et al. (2024). "State-Dependent Local Projections". In: *Journal of Econometrics*.

Hong, James et al. (2021). "Analysis of Faces in a Decade of US Cable TV News". In: *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery Data Mining*. Association for Computing Machinery.

Jordà, Òscar (2005). "Estimation and Inference of Impulse Responses by Local Projections". In: *American Economic Review* 95.1, pp. 161–182.

— (2023). "Local Projections for Applied Economics". In: *Annual Review of Economics* 15. Publisher: Annual Reviews, pp. 607–631.

Kamdar, Rupal (2018). "The Inattentive Consumer: Sentiment and Expectations". In: *Manuscript*.

Krueger, Alan B. and Alan Blinder (2004). "What Does the Public Know about Economic Policy, and How Does It Know It?" In: *Brookings Papers on Economic Activity* 35, pp. 327–397.

Kuchler, Theresa and Basit Zafar (2019). "Personal Experiences and Expectations about Aggregate Outcomes". In: *The Journal of Finance* 74.5, pp. 2491–2542.

Lamla, Michael J. and Sarah M. Lein (2014). "The Role of Media for Consumers' Inflation Expectation Formation". In: *Journal of Economic Behavior & Organization* 106, pp. 62–77.

Larsen, Vegard H., Leif Anders Thorsrud, and Julia Zhulanova (2021). "News-driven inflation expectations and information rigidities". In: *Journal of Monetary Economics* 117, 507–520.

Malmendier, Ulrike and Stefan Nagel (2016). "Learning from Inflation Experiences". In: *The Quarterly Journal of Economics* 131.1, pp. 53–87.

Nimark, Kristoffer P. (2014). "Man-Bites-Dog Business Cycles". In: *American Economic Review* 104.8, 2320–67.

Nimark, Kristoffer P. and Stefan Pitschner (2019). "News media and delegated information choice". In: *Journal of Economic Theory* 181, pp. 160–196.

Pfäuti, Oliver (2025). *The Inflation Attention Threshold and Inflation Surges*.

Plagborg-Møller, Mikkel and Christian K. Wolf (2021). "Local Projections and VARs Estimate the Same Impulse Responses". en. In: *Econometrica* 89.2, pp. 955–980. ISSN: 0012-9682. (Visited on 03/17/2024).

Ramey, Valerie (2016). "Handbook of Macroeconomics". In: vol. 2. Cambridge: NBER. Chap. Macroeconomic Shocks and their Propagation, pp. 71–162.

Ramey, Valerie A. and Sarah Zubairy (2018). "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data". In: *Journal of Political Economy* 126.2, pp. 850–901.

Tausanovitch, Chris and Lynn Vavreck (2021). *Democracy Fund + UCLA Nationscape Project*.

Tenreyro, Silvana and Gregory Thwaites (2016). "Pushing on a String: U.S. Monetary Policy Is Less Powerful in Recessions". In: *American Economic Journal: Macroeconomics* 8.4, pp. 43–74.

Yotzov, Ivan et al. (2026). "The Speed of Firm Response to Inflation". In: *Journal of the European Economic Association* 24.2, pp. 736–768. DOI: 10.1093/jeea/jvaf044.

Online Appendix

The News that Moves Inflation Expectations

Carola Binder, Pascal Frank, and Jane Ryngaert

September 14, 2026

This online appendix presents robustness checks, additional results, and methodological details. First, we demonstrate robustness along a number of dimensions: unrestricted estimation, different measures of inflation surprises, different measures of news coverage and household expectations. Additional results present the responses of market-based inflation expectations, household uncertainty and disagreement, as well as comparisons between full-sample and release-day estimates. We also document the smooth local projection implementation and provide derivations underlying the empirical framework.

A Robustness Exercises

This appendix section examines the robustness of our baseline results along four dimensions. First, Appendix A.1 compares the smooth local projection estimates with unrestricted, horizon-by-horizon estimates. Second, Appendix A.2 replaces the baseline core inflation surprise with three alternative measures of the inflation signal. Third, Appendix A.3 uses the median expectation implied by households' density forecasts as an alternative outcome and excludes days with particularly few respondents. Finally, Appendix A.4 considers both a continuous measure of elevated coverage and a binary indicator constructed using regime-specific coverage volatility. Taken together, these exercises show that the main findings are not specific to the smoothing procedure, the measurement of inflation expectations, the inclusion of low-response days, or the particular definition of elevated news coverage.

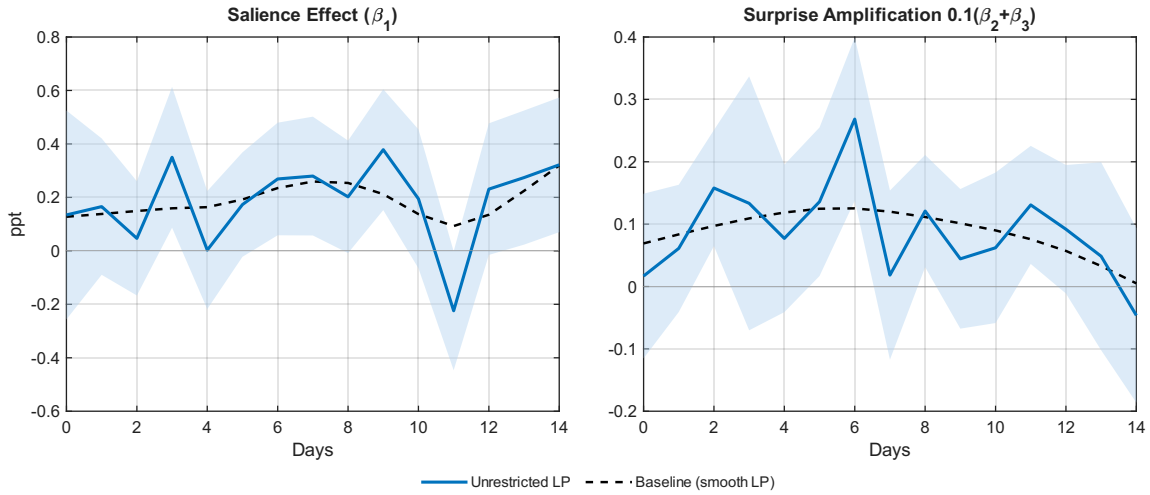
A.1 Unrestricted Local Projections

Our baseline specification uses smooth local projections to reduce the sampling noise arising from the relatively small release-day sample and the volatility of daily inflation expectations. To assess whether smoothing drives the results, we re-estimate the baseline specification using unrestricted local projections, estimating each horizon separately without imposing cross-horizon smoothness.

Figure A.1 compares the unrestricted estimates with the baseline smooth responses. The left panel shows that the unrestricted estimate of the high-coverage effect is positive at nearly all horizons, although it displays considerably more horizon-to-horizon variation. The response is statistically different from zero at several horizons, including multiple horizons between days 3 and 10. The smooth estimate captures the persistent positive tendency in the unrestricted estimates while attenuating sharp one-day fluctuations, such as the reversal at horizon 11.

The right panel reports the response to a 0.1 percentage point inflation surprise on a high-coverage release day. The unrestricted point estimates are positive over most of the first 13 days and reach their largest value around day 6, but are again substantially more jagged and less precisely estimated than the smooth response.

Figure A.1: Robustness Exercise: Unrestricted Local Projections



Notes: The solid blue lines report unrestricted local projection estimates, and the dashed black lines report the corresponding baseline smooth local projection estimates. The left panel shows the response of the daily median 12-month-ahead inflation expectation to a high-coverage release day, β_1^h. The right panel shows the response to a 0.1 percentage point core inflation surprise on a high-coverage release day, $0.1(\beta_2^h + \beta_3^h)$. Responses are expressed in percentage points. Shaded areas denote 90% confidence intervals for the unrestricted estimates based on heteroskedasticity-robust standard errors. Sample period: June 1, 2013 to September 30, 2024.

The baseline estimate captures the central pattern of these horizon-specific estimates rather than reproducing their short-run fluctuations.

Overall, the positive response to a high-coverage release day and the positive surprise response on such days are also present when the responses are estimated without smoothing. Smoothing primarily regularizes the noisy horizon-specific estimates and makes their underlying dynamic pattern easier to interpret; it does not generate the broad qualitative results.

A.2 Alternative Measures of the Inflation Signal

Our baseline measure of the inflation signal, λ_t, is the difference between realized core inflation and the corresponding pre-release consensus forecast of professional

forecasters. This measure isolates the component of an inflation release that is unexpected relative to the information set of professional forecasters. Households, however, may respond to the inflation number reported in the news rather than evaluate the release relative to a professional forecast. Consistent with this distinction, Yotzov et al. (2026) find that firms in the United Kingdom adjust their inflation perceptions and expectations in response to changes in reported headline inflation, but not to inflation surprises relative to professional forecasts. Although their analysis concerns firms rather than households, it motivates examining measures that more closely reflect the inflation information presented in news coverage.

We consider three alternatives. First, we replace the core inflation surprise with the corresponding headline inflation surprise relative to the pre-release professional consensus forecast. Second, we use the change in the reported year-over-year inflation rate relative to the preceding release of the same series, for both CPI and PCE separately:

$$\Delta\pi_t^{\text{YoY}} = \pi_t^{\text{YoY}} - \pi_{\text{t}-1}^{\text{YoY}}.$$

Third, we construct a proxy that is close to households by computing an inflation gap equal to the year-over-year inflation rate reported on the release date minus the trailing 30-day average of household one-year-ahead inflation expectations:

$$gap_t^{\text{SCE}} = \pi_t^{\text{YoY}} - \overline{E}_{t-1}^{30\text{d}}[\pi_{t+365,t}].$$

Unlike the first two measures, it captures the distance between the current inflation signal and households' recent beliefs and therefore provides an alternative proxy for how 'surprising' an information release may have been for households.

Figure A.2 reports the results. For comparability, the amplification responses (right column) are normalized to a one standard deviation increase in the corresponding signal measure. One standard deviation is approximately 0.1 percentage points for both the baseline core surprise and the headline surprise, and approximately 0.3 percentage points for the change in year-over-year inflation as well as 1 percentage points for the household-relative inflation gap.

The left column shows that the estimated response to a high-coverage release day

remains essentially unchanged when using all three alternative measures. Using the household-relative inflation gap produces an even larger salience response.

The right column reports the response to each alternative inflation signal on a high-coverage release day. The response remains positive and close to the baseline under all three alternative proxies. The amplification effect is even somewhat larger when the signal is measured using the headline surprise or the change in year-over-year inflation, and remains persistently positive when the signal is measured relative to households' recent expectations. Thus, the positive response to high coverage and the transmission of inflation information on high-coverage release days do not depend on measuring the release exclusively through core inflation surprises relative to professional forecasts.

A.3 Robustness to Outcome Measurement and Daily Sample Size

Our baseline outcome is the daily median of households' numerical point forecasts for inflation over the next 12 months. We consider two robustness exercise for this measure. First, numerical point forecasts may not fully reflect the information contained in respondents' subjective probability distributions. We therefore replace each respondent's point forecast with the mean implied by the respondent's density forecast and calculate the daily median of these density-implied means. Second, because the number of survey responses varies across days, the daily median may be measured less precisely on days with particularly few respondents. We address this concern by retaining the baseline point-forecast outcome but excluding days with fewer than ten valid responses.

Figure A.3 presents the results, with the baseline smooth local projection estimates shown as dashed lines for comparison. The top row uses the median density-implied mean as the outcome. The estimated high-coverage response remains positive and follows a similar dynamic pattern to the baseline, although it is somewhat larger at longer horizons. The response to an inflation surprise on a high-coverage day also closely tracks the baseline through its peak around days 6–7, but declines more

quickly after approximately day 10. The divergence at the end of the horizon is imprecisely estimated.

The bottom row excludes days with fewer than ten respondents. Both the high-coverage response and the response to an inflation surprise are nearly unchanged relative to the baseline, although the confidence intervals widen somewhat because of the smaller sample. These exercises indicate that the central results are not driven by the use of numerical point forecasts or by a small number of days on which the daily median is constructed from very few respondents.

A.4 Alternative Measures of Elevated News Coverage

Our baseline high-coverage indicator equals one when inflation coverage exceeds its 10-day moving average by more than one trailing 180-day standard deviation. We examine whether the results depend on either the binary nature of this measure or the rolling standard deviation used to define unusually elevated coverage.

First, we replace the binary indicator with the continuous standardized coverage measure

$$z_t = \frac{coverage_t - \overline{coverage}_{t-1}^{10}}{s_{t-1}^{180}},$$

where $\overline{coverage}_{t-1}^{10}$ is the 10-day moving average and s_{t-1}^{180} is the 180-day moving standard deviation. The baseline indicator can therefore be written as $\mathbf{1}\{z_t > 1\}$. In the continuous specification, we replace the indicator and its interaction with the inflation surprise by z_t and $z_t \times \lambda_t$. We also include 14 lags of z_t rather than 14 lags of the raw coverage measure. To facilitate comparison with the binary specification, we evaluate the responses at $z_t = 2.57$, the average of z_t across all baseline high-coverage days.

Second, we keep the binary high-coverage indicator but replace the rolling 180-day standard deviation with a fixed, regime-specific standard deviation. Specifically, we define

$$news_t^R = \mathbf{1}\left\{coverage_t > \overline{coverage}_{t-1}^{10} + s_{R(t)}\right\},$$

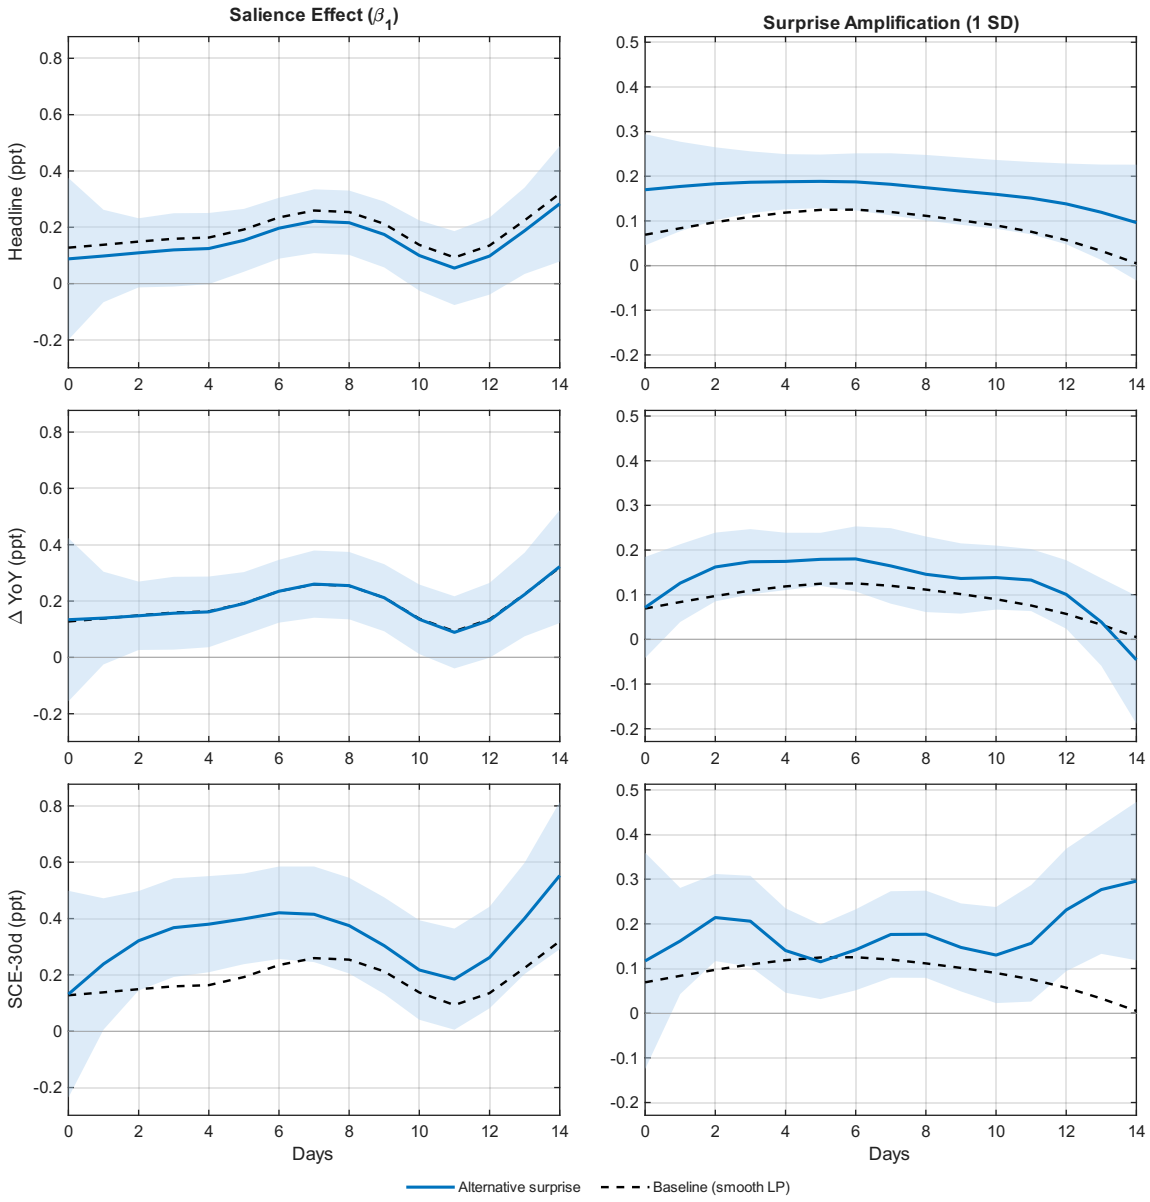
where $s_{R(t)}$ is the standard deviation of inflation coverage within the relevant inflation

regime. It equals approximately 5.46 seconds during the relatively calm 2013–2020 period and 108.9 seconds during the elevated-inflation period of 2021–2024. Figure A.4 compares the resulting estimates with the baseline responses.

In the continuous specification, shown in the top row, the response to elevated coverage is initially less pronounced but becomes positive and statistically significant between approximately days 5 and 10. Its peak magnitude is similar to the baseline estimate. The response to an inflation surprise in a high-coverage environment also remains positive over most of the horizon, but is smaller than the baseline response after the first few days.

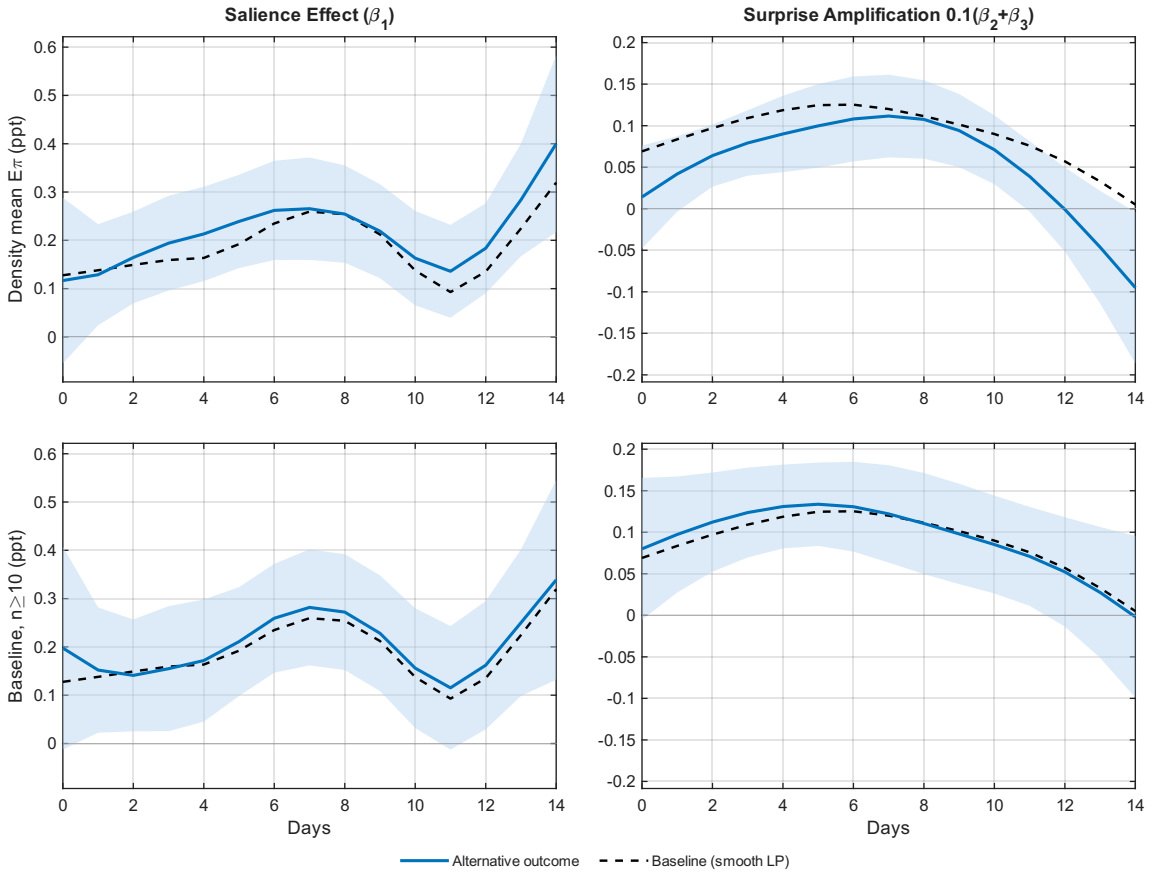
The regime-specific binary indicator, shown in the bottom row, produces results that are close to the baseline. The high-coverage response is almost identical, while the response to an inflation surprise is even somewhat larger and more persistent than in the baseline specification. Overall, these robustness exercises confirm that the results do not depend on using a binary indicator based on a rolling 180-day standard deviation.

Figure A.2: Robustness to Alternative Measures of the Inflation Signal



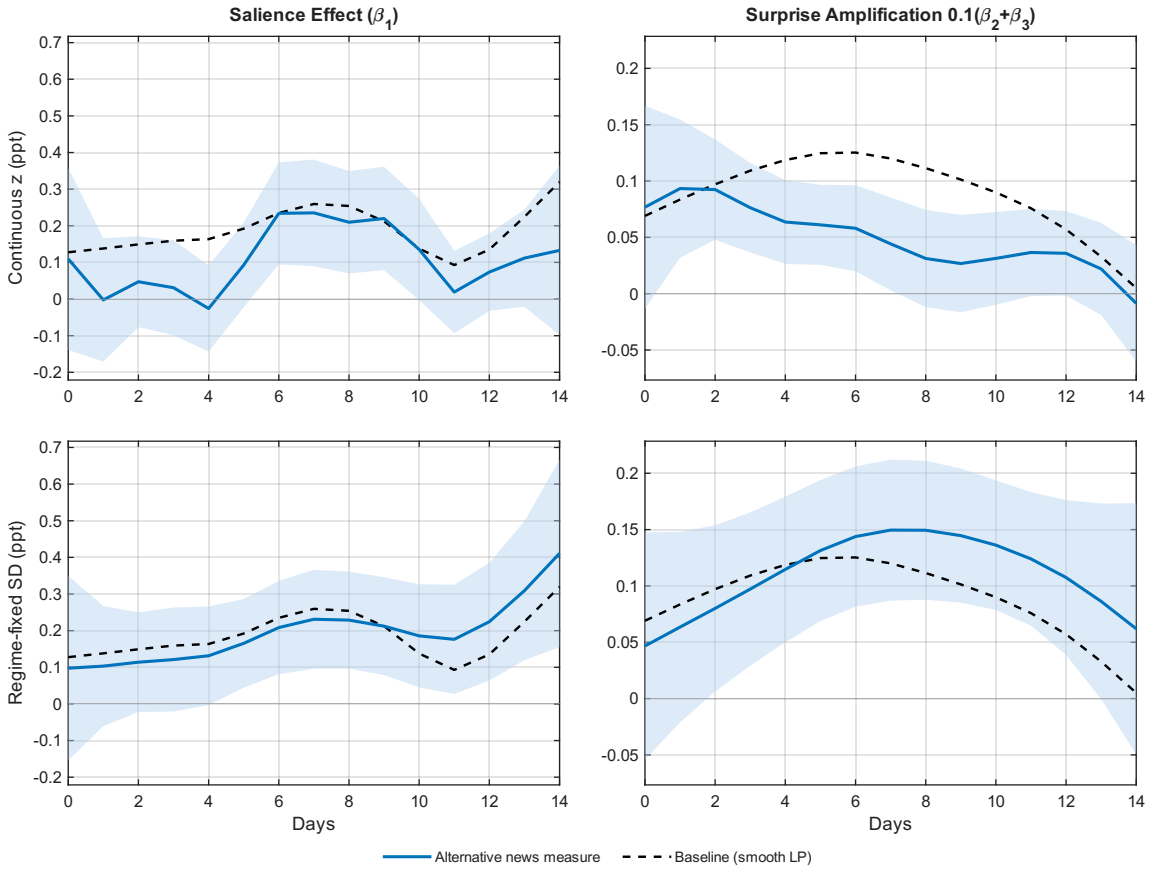
Notes: The solid blue lines report smooth local projection estimates using alternative measures of the inflation signal; the dashed black lines reproduce the corresponding baseline estimates using core inflation surprises relative to professional forecasts. From top to bottom, the alternative measures are the headline inflation surprise relative to the professional consensus, the change in the reported year-over-year inflation rate relative to the preceding release, and the gap between reported year-over-year inflation and the trailing 30-day average of household one-year-ahead inflation expectations. The left column reports the response to a high-coverage release day, β_1^h. The right column reports the response to a one-standard-deviation increase in the corresponding inflation signal on a high-coverage release day. Shaded areas denote 90% confidence intervals for the alternative specifications based on heteroskedasticity-robust standard errors. Sample period: June 1, 2013 to September 30, 2024.

Figure A.3: Robustness to Outcome Measurement and Daily Sample Size



Notes: The solid blue lines report smooth local projection estimates from the alternative specifications, and the dashed black lines reproduce the corresponding baseline estimates. The top row replaces the daily median numerical point forecast with the daily median of the respondent-level means implied by households' density forecasts. The bottom row retains excludes days with fewer than ten valid responses. The left column reports the response to a high-coverage release day, β_1^h. The right column reports the response to a 0.1 percentage point core inflation surprise on a high-coverage release day, $0.1(\beta_2^h + \beta_3^h)$. Shaded areas denote 90% confidence intervals for the alternative specifications based on heteroskedasticity-robust standard errors. Sample period: June 1, 2013 to September 30, 2024.

Figure A.4: Alternative Measures of Elevated News Coverage



Notes: The solid blue lines report smooth local projection estimates using alternative measures of elevated inflation coverage, and the dashed black lines reproduce the corresponding baseline estimates. The top row replaces the binary high-coverage indicator with the standardized continuous measure. The responses are evaluated at $z_t = 2.57$. The bottom row uses a binary indicator constructed with fixed regime-specific standard deviations for 2013–2020 and 2021–2024. Shaded areas denote 90% confidence intervals for the alternative specifications based on heteroskedasticity-robust standard errors. Sample period: June 1, 2013 to September 30, 2024.

B Additional Results

B.1 Validation with market-based expectations.

If the core inflation surprises measure economically meaningful information shocks, they should lead to a reaction of financial markets. To validate this, we estimate the dynamic response of market-based inflation expectations, measured by the 5-Year Breakeven Inflation Rate.[19]

Specifically, we estimate the following local projection model conditional on release days ($S_t = 1$):

$$\pi_{t+h}^{\text{exp,5Y}} = \beta_0^h + \beta_{CPI}^h \lambda_t^{CPI} + \beta_{PCE}^h \lambda_t^{PCE} + \sum_{j=1}^{7} \phi_j^h \pi_{t-j}^{\text{exp,5Y}} + \varepsilon_{t+h}. \tag{14}$$

Here, λ_t^k denotes the core inflation surprise associated with release $k \in \{\text{CPI}, \text{PCE}\}$. We include seven lags of breakeven inflation to account for persistence. The sequence $\{\beta_k^h\}_{h=0}^{14}$ traces out the dynamic response of market expectations to inflation surprises. All responses are scaled to correspond to a 0.1 percentage point surprise. Inference is based on heteroskedasticity-robust standard errors.[20]

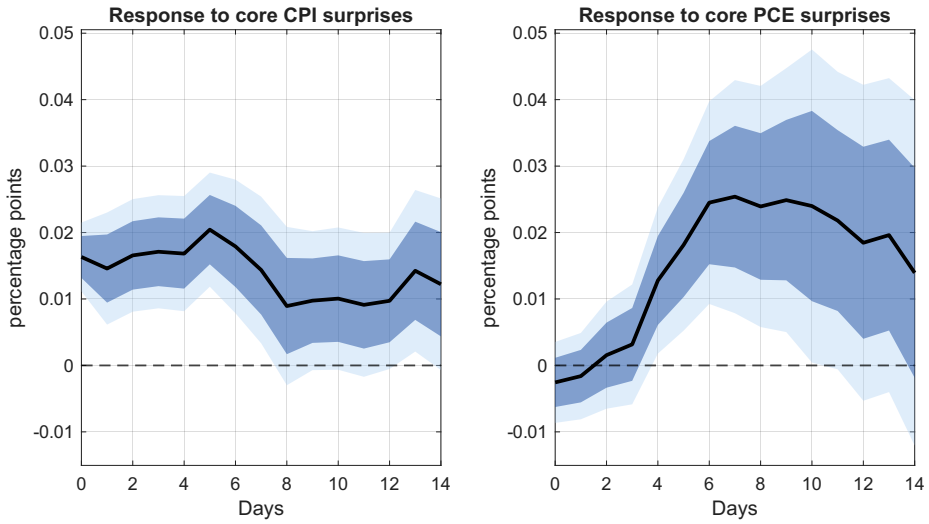
Figure B.1 reports the results. Market-based inflation expectations respond significantly to both CPI and PCE surprises. Breakeven inflation reacts immediately to core CPI surprises and exhibits a delayed but significant response to core PCE surprises. Despite differences in timing, the magnitudes are similar: a 0.1 percentage point surprise raises five-year inflation expectations by approximately 0.02 percentage points for several days.

Taken together, these findings indicate that both core CPI and core PCE surprises constitute distinct and economically meaningful information shocks. Differences in media coverage documented in the previous subsection therefore reflect variation in amplification rather than differences in the underlying informational content of the

[19]The breakeven rate is constructed from 5-Year Treasury Constant Maturity Securities (FRED ID: DGS5) and 5-Year Treasury Inflation-Indexed Constant Maturity Securities (FRED ID: DFII5).

[20]Plagborg-Møller and Wolf (2021) show that heteroskedasticity-robust standard errors are valid in lag-augmented local projection models.

Figure B.1: Response of Market Inflation Expectations



Notes: Dynamic response of the 5-Year Breakeven Inflation Rate to core CPI and core PCE surprises. Shaded areas denote 68% and 90% confidence intervals based on heteroskedasticity-robust standard errors. Sample period: June 1, 2013 to September 30, 2024.

releases.

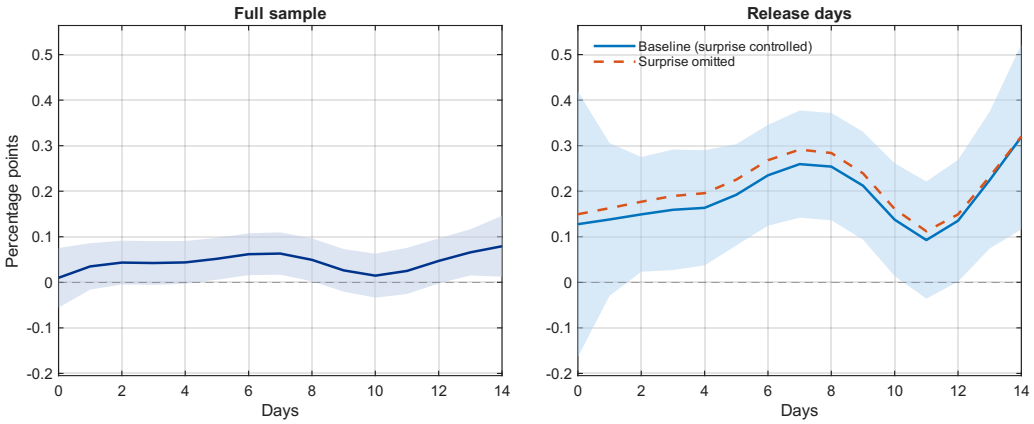
B.2 Reconciling the Full-Sample and Release-Day Estimates

The response to a high-coverage release day documented in Section 4.1 is substantially larger than the association obtained from a naive full-sample regression of inflation expectations on news coverage. This subsection uses the comparison to quantify the two forces discussed in Section 2: composition and omitted variable bias.

The left panel of Figure B.2 reports the full-sample estimate of Equation (2). It measures the average difference in the path of inflation expectations following high- and normal-coverage days across the full daily sample. The estimated response is positive but modest, peaking at approximately 0.06 percentage points, roughly one-fifth of the corresponding release-day estimate.

Equation (4) helps explain this difference. The full-sample coefficient is a variance-weighted average of the coverage differential on days when new inflation information arrives, δ_I^h, and the differential on other days, δ_0^h. If coverage has little effect on

Figure B.2: Reconciling the Full-Sample and Release-Day Responses



Notes: Left panel: full-sample response of the daily median 12-month-ahead inflation expectation to a high-coverage day, estimated using Equation (2). Right panel: response to a high-coverage release day when controlling for the inflation surprise and its interaction with high coverage (solid line; baseline β_1^h) and when omitting both terms (dashed line). Both panels use the same vertical scale, and responses are expressed in percentage points. Shaded areas denote 90% confidence intervals based on Newey–West HAC standard errors. Impulse responses are estimated using smooth local projections with the penalty parameter set to its baseline value of 1. Sample period: June 1, 2013 to September 30, 2024.

expectations when it does not coincide with a new inflation signal, so that $\delta_0^h \approx 0$, the full-sample coefficient is attenuated relative to δ_I^h. The degree of attenuation depends on the share of residual variation in coverage that occurs on information-arrival days.

Scheduled CPI and PCE release days constitute a directly observable subset of such days. This does not imply that inflation-relevant information arrives only on scheduled inflation-release days. Other macroeconomic releases, policy developments, or energy price shocks may also convey information about inflation that is not captured by our release-surprise measure. The comparison nevertheless illustrates why an estimate pooling all days can be substantially smaller than one that conditions on observable inflation-information arrivals.

The right panel isolates omitted variable bias by holding the sample fixed at release days. The dashed line reports the response estimated without the inflation surprise and its interaction with high coverage, while the solid line reproduces the

baseline estimate of β_1^h, which includes both terms. Omitting the surprise terms raises the peak response from approximately 0.26 to 0.29 percentage points.

The direction of this difference is consistent with the omitted variable bias derived in Equation (3). High-coverage releases tend to be associated with somewhat larger inflation surprises. When the surprise and its interaction with coverage are omitted, part of the expectations response to those surprises is therefore attributed to the high-coverage indicator. Quantitatively, however, the resulting difference is small: it amounts to approximately 0.03 percentage points at the peak.

Two conclusions follow. First, the full-sample and release-day coefficients answer different questions. The former averages the coverage differential across heterogeneous days, many of which do not coincide with an observable arrival of new inflation information. Its smaller magnitude is therefore consistent with the composition mechanism and does not contradict the larger release-day response. Second, within the release-day sample, omitting the inflation surprise shifts the estimated high-coverage coefficient in the predicted direction, but the resulting bias is quantitatively modest. The two forces work in opposite directions: composition attenuates the full-sample coefficient, whereas the positive association between coverage and inflation surprises biases it upward. In this application, the composition effect clearly dominates.

B.3 Household Uncertainty and Disagreement

The baseline analysis focuses on the response of median inflation expectations. As an extension, we examine whether high coverage also affects households' subjective uncertainty about future inflation and the degree of disagreement across households. We re-estimate the baseline local projection using two alternative outcomes. Household uncertainty is measured by the daily median of the log standard deviation of respondents' subjective 12-month-ahead inflation distributions. Household disagreement is measured by the log interquartile range of respondents' point forecasts. Because both dependent variables are expressed in logs, the responses are reported as percentage changes.

Figure B.3 presents the results. The top row reports the responses of household

uncertainty and the bottom row those of disagreement. The left column shows the response to a high-coverage release day, holding the inflation surprise fixed. The right column compares the response to a 0.1 percentage point inflation surprise on normal- and high-coverage release days.

The top-left panel shows that a high-coverage release day is followed by a gradual increase in household uncertainty. The response is muted on impact but builds over the following week, reaching approximately 5 percent. It is statistically significant from roughly day 6 through day 13. Thus, high coverage is not followed by an immediate compression of subjective uncertainty. Instead, households become more uncertain about future inflation at medium horizons.

The top-right panel shows that a positive inflation surprise raises household uncertainty under both normal and high coverage. The point estimates are somewhat larger under high coverage at short and medium horizons, but the difference between the two response paths is modest. The estimates therefore provide clearer evidence that inflation surprises raise uncertainty than that high coverage amplifies this response.

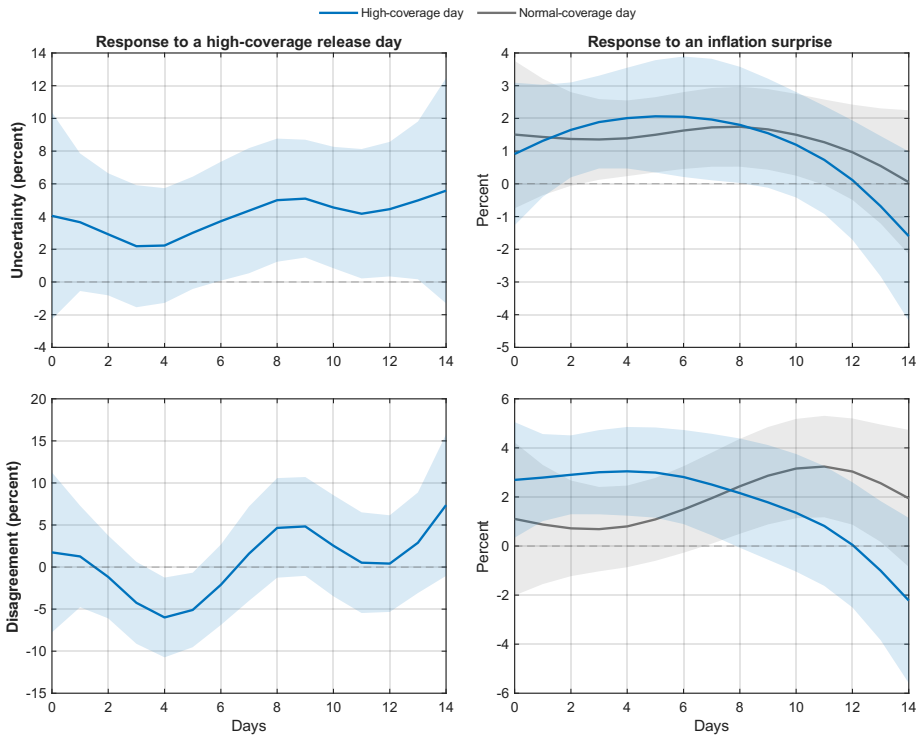
The results for disagreement are more muted. The response to a high-coverage release day in the bottom-left panel is statistically indistinguishable from zero on impact, turns temporarily negative around days 4–5, and is again indistinguishable from zero thereafter. We therefore find no evidence that high-coverage release days produce a persistent widening of disagreement across households.

The bottom-right panel indicates that positive inflation surprises are followed by greater disagreement under both media environments. The point estimates suggest that disagreement responds earlier on high-coverage release days, whereas the response under normal coverage emerges with a delay. The two paths, however, do not provide clear evidence of a persistent difference in the magnitude of the response across coverage states.

Overall, these additional results suggest that high-coverage release days are followed by greater household uncertainty at medium horizons, but not by a persistent increase in disagreement. Inflation surprises raise both uncertainty and disagreement, although the evidence that high coverage amplifies these distributional responses is

weaker than the corresponding evidence for median inflation expectations.

Figure B.3: Responses of Household Uncertainty and Disagreement



Notes: Dynamic responses of household uncertainty (top row; daily median of the log standard deviation of respondents' subjective one-year-ahead inflation distributions) and disagreement (bottom row; log interquartile range of respondents' point forecasts). The left column reports the response to a high-coverage release day. The right column reports the response to a 0.1 percentage point core inflation surprise on normal-coverage (dashed grey) and high-coverage (solid blue) release days. Responses are expressed in percent. Shaded areas denote 90% confidence intervals based on Newey–West HAC standard errors. Impulse responses are estimated using smooth local projections with the penalty parameter set to its baseline value of 1. Sample period: June 1, 2013 to September 30, 2024.

C Details on the Smooth Local Projection Method

We estimate the baseline dynamic responses using the smooth local projection method of Barnichon and Brownlees (2019). Our implementation extends the Stata routine of Bousquet (2024) to accommodate the release-day sample, three jointly smoothed response paths, and inference for their linear combinations. Appendix **??** describes the corresponding release-date decomposition.

C.1 Penalized estimation and selection of the smoothing parameter

Smooth local projections represent each response path using a B-spline basis. For the three regressors of interest—the high-coverage indicator, the inflation surprise, and their interaction—we write

$$\beta_j(h) = B(h)'b_j, \qquad j = 1, 2, 3, \quad h = 0, \ldots, H, \tag{15}$$

where $B(h)$ contains the basis functions evaluated at horizon h. We set $H = 14$ and use a cubic B-spline basis with knots at the integer horizons, yielding $K_b = 17$ basis functions for each path.

Let θ collect the spline coefficients and the horizon-specific intercepts and control coefficients. Stacking observations across release dates and horizons gives the estimator

$$\widehat{\theta}(\rho) = \arg\min_{\theta} \left\{ \frac{1}{n} \|\mathbf{y} - X\theta\|^2 + \rho\, \theta' P \theta \right\} = (X'X + n\rho P)^{-1} X'\mathbf{y}, \tag{16}$$

where n is the number of stacked observations. We use ρ for the smoothing parameter to distinguish it from the inflation surprise λ_t.

The penalty matrix P is block diagonal, with a block $D_2'D_2$ for each of the three response paths and zeros for the intercepts and controls. Here, D_2 takes second differences of the spline coefficients. The three paths share a common smoothing parameter but have separate coefficients. Larger penalties shrink each path toward a linear function of the horizon; penalties approaching zero yield the unrestricted

B-spline fit. The second-difference penalty and the cubic degree of the basis are separate choices.

We select the smoothing parameter jointly for the three paths by five-fold cross-validation. Release dates are divided into five contiguous blocks of approximately equal size, with all horizons from a release date assigned to the same block. Each candidate is evaluated by averaging the mean squared prediction errors across the five omitted blocks.[21] The candidate grid is

$$\rho \in \{0.00001, 0.0001, 0.001, 0.002, \ldots, 0.01, 0.1, 1\}.$$

For the baseline estimates, the criterion is minimized at $\rho_{\mathrm{CV}} = 1$, the upper endpoint of this grid. Following the routine's undersmoothing rule, we use

$$\rho_{\mathrm{used}} = 0.01\,\rho_{\mathrm{CV}} = 0.01 \tag{17}$$

for the reported estimates. Undersmoothing is intended to reduce penalization bias; it does not by itself guarantee nominal coverage of confidence intervals.

C.2 Extensions to the Implementation of Bousquet (2024)

Release-day estimation. We restrict regression origins to scheduled CPI and PCE release days. Outcome leads and lagged controls are constructed using the full daily data before imposing this restriction. Thus, y_{t+h} refers to the outcome h days after release t, and lagged controls refer to preceding days rather than preceding releases. The specification includes 14 lags of the outcome and 14 lags of inflation news coverage in seconds. Their coefficients remain unpenalized and unrestricted across horizons, while common across coverage states at each horizon.

Jointly smoothed response paths. For non-IV specifications, the original routine smooths the response to one regressor and treats the remaining positional variables

[21] We retain the original routine's cross-validation normalization: the penalty in the training-fold normal equations is $n_{\mathrm{train}}\rho(4/5)P$, where n_{train} counts the stacked training observations.

as controls. Our extension assigns separate spline blocks to $news_t$, λ_t, and $news_t\lambda_t$, and estimates their paths jointly. We retain the regressors' original units: the news indicator represents a change from zero to one, and surprise responses are reported for a 0.1 percentage point inflation surprise.

Linear combinations and inference. The response to a 0.1 percentage point inflation surprise on a high-coverage release day is

$$q_h = 0.1(\beta_2^h + \beta_3^h), \tag{18}$$

where β_2^h is the surprise response on a normal-coverage release day and β_3^h is the additional response under high coverage. We extend the routine to compute this combination and its uncertainty from the joint coefficient covariance matrix.

The current implementation uses a heteroskedasticity-robust (HC0) covariance estimate for the stacked penalized regression. The variance of the combined response is

$$\widehat{\mathrm{Var}}(\widehat{q}_h) = 0.1^2 \left[\widehat{\mathrm{Var}}(\widehat{\beta}_2^h) + \widehat{\mathrm{Var}}(\widehat{\beta}_3^h) + 2\widehat{\mathrm{Cov}}(\widehat{\beta}_2^h, \widehat{\beta}_3^h)\right]. \tag{19}$$

The covariance term accounts for the joint estimation of the two coefficient paths. Pointwise 90% confidence intervals, conditional on the selected smoothing parameter, are constructed as

$$\widehat{q}_h \pm z_{0.95}\sqrt{\widehat{\mathrm{Var}}(\widehat{q}_h)}.$$

D Decomposition of the Naive Coverage Coefficient

D.1 Omitted Variable Bias in the Naive Coverage Model

In this subsection, we derive the OVB that arises when the underlying inflationary signal and its interaction with news coverage are excluded from the regression. First, we derive the bias given that $news_t$ is a binary indicator for elevated coverage. We

then extend the derivation to a continuous news coverage measure.

As shown in Section 2, suppose that we can express inflation expectations as a local projection model such that:

$$E_{t+h}[\pi_{t+365,t}] = \beta_0^h + \beta_1^h news_t + \beta_2^h \lambda_t + \beta_3^h (news_t \times \lambda_t) + \Phi^h X_t + u_{t+h}, \tag{20}$$

where $news_t \in \{0, 1\}$ indicates elevated news coverage and λ_t denotes the underlying innovation. Note that $E[\lambda_t] = 0$ and assume that $\text{Cov}(news_t, u_{t+h}) = 0$.

Consider also the naive model:

$$E_{t+h}[\pi_{t+365,t}] = \alpha^h + \delta^h news_t + \Phi^h X_t + u_{t+h}, \tag{21}$$

which omits both λ_t and its interaction with news coverage.

The population coefficient on $news_t$ is

$$\begin{aligned} \delta^h &= \frac{\text{Cov}(news_t, y_{t+h})}{\text{Var}(news_t)} \\ &= \beta_1^h + \beta_2^h \frac{\text{Cov}(news_t, \lambda_t)}{\text{Var}(news_t)} + \beta_3^h \frac{\text{Cov}(news_t, news_t \lambda_t)}{\text{Var}(news_t)}. \end{aligned} \tag{22}$$

Consequently, the omitted-variable bias relative to β_1^h is

$$\text{Bias}(\delta^h) \equiv \delta^h - \beta_1^h = \beta_2^h \frac{\text{Cov}(news_t, \lambda_t)}{\text{Var}(news_t)} + \beta_3^h \frac{\text{Cov}(news_t, news_t \lambda_t)}{\text{Var}(news_t)}. \tag{23}$$

For the binary coverage indicator, the two covariance terms in Equation (23) are directly related. Let

$$p = E[news_t] = \Pr(news_t = 1), \qquad 0 < p < 1.$$

Because $news_t^2 = news_t$,

$$\begin{aligned}\text{Cov}(news_t, news_t\lambda_t) &= E[news_t^2\lambda_t] - E[news_t]E[news_t\lambda_t] \\ &= (1-p)E[news_t\lambda_t]. \end{aligned} \tag{24}$$

Moreover, because $E[\lambda_t] = 0$, it follows immediately that

$$E[news_t\lambda_t] = \text{Cov}(news_t, \lambda_t).$$

It follows that

$$\text{Cov}(news_t, news_t\lambda_t) = (1-p)\,\text{Cov}(news_t, \lambda_t). \tag{25}$$

Substituting Equation (25) into Equation (23) yields

$$\text{Bias}(\delta^h) = \left[\beta_2^h + (1-p)\beta_3^h\right] \frac{\text{Cov}(news_t, \lambda_t)}{\text{Var}(news_t)}. \tag{26}$$

The sign of the bias is therefore transparent. If more inflationary innovations are more likely to generate elevated coverage, then

$$\text{Cov}(news_t, \lambda_t) > 0.$$

For $\beta_2^h, \beta_3^h \geq 0$, with at least one inequality strict, Equation (26) implies

$$\text{Bias}(\delta^h) > 0.$$

The naive coefficient consequently overstates the pure effect of elevated coverage, β_1^h, because it also incorporates the more inflationary content of the news selected into elevated-coverage days and the amplified response to that content.

D.2 Composition: a variance-weighted average

The composition identity is exact and purely algebraic; it requires no distributional assumptions. Let $\tilde{n}_t$ and $\tilde{E}_t$ denote the residuals of $news_t$ and $E_{t+h}[\pi_{t+365,t}]$ from an OLS projection on the controls X_t (which include a constant, so $\sum_t \tilde{n}_t = 0$).

Step 1 (FWL). The naive coefficient equals the simple-regression slope of $\tilde{E}$ on $\tilde{n}$,

$$\hat{\delta}^h = \frac{\sum_t \tilde{n}_t \tilde{E}_t}{\sum_t \tilde{n}_t^2}.$$

Step 2 (partition). Split the sample into days with an information arrival, I, and the remaining days, I^c:

$$\sum_t \tilde{n}_t \tilde{E}_t = \sum_{t \in I} \tilde{n}_t \tilde{E}_t + \sum_{t \in I^c} \tilde{n}_t \tilde{E}_t, \qquad \sum_t \tilde{n}_t^2 = \sum_{t \in I} \tilde{n}_t^2 + \sum_{t \in I^c} \tilde{n}_t^2.$$

Step 3 (within-group slopes). Define

$$\delta_I^h = \frac{\sum_{t \in I} \tilde{n}_t \tilde{E}_t}{\sum_{t \in I} \tilde{n}_t^2}, \qquad \delta_0^h = \frac{\sum_{t \in I^c} \tilde{n}_t \tilde{E}_t}{\sum_{t \in I^c} \tilde{n}_t^2},$$

so that $\sum_{t \in I} \tilde{n}_t \tilde{E}_t = \delta_I^h \sum_{t \in I} \tilde{n}_t^2$ and $\sum_{t \in I^c} \tilde{n}_t \tilde{E}_t = \delta_0^h \sum_{t \in I^c} \tilde{n}_t^2$.
Step 4 (substitute). Combining Steps 1–3,

$$\hat{\delta}^h = \frac{\delta_I^h \sum_{t \in I} \tilde{n}_t^2 + \delta_0^h \sum_{t \in I^c} \tilde{n}_t^2}{\sum_t \tilde{n}_t^2} = \pi\, \delta_I^h + (1 - \pi)\, \delta_0^h, \qquad \pi \equiv \frac{\sum_{t \in I} \tilde{n}_t^2}{\sum_t \tilde{n}_t^2} \in [0, 1].$$

This is an exact finite-sample identity: the pooled coefficient is the variance-weighted average of the within-group slopes, with weight π equal to the share of the residual sum of squares of coverage contributed by information-arrival days. Writing $s_I^2 = \frac{1}{n_I} \sum_{t \in I} \tilde{n}_t^2$ and $s^2 = \frac{1}{n} \sum_t \tilde{n}_t^2$,

$$\pi = \frac{n_I}{n} \cdot \frac{s_I^2}{s^2},$$

the frequency of information-arrival days times their relative residual variance.
Step 5 (approximation). If coverage moves expectations little when no signal arrives,

$\delta_0^h \approx 0$, then $\hat{\delta}^h \approx \pi\, \delta_I^h$, which is Equation (4). Because information-arrival days are infrequent, π is small and the full-sample estimate is a strongly attenuated version of δ_I^h.

Remark (binary indicator). When $news_t$ is a 0/1 coverage indicator and controls are omitted, $\tilde{n}_t^2 = news_t$, and π reduces to the share of high-coverage days that fall on information-arrival days.