Local News and the Transmission of Monetary Policy*

Joao B. Duarte† Stephen Hansen‡ Afonso S. Moura§ Álvaro Ortiz¶
Tomasa Rodrigo‖

March 2026

Preliminary and Incomplete

Abstract

Experimental evidence shows that treating households with information about monetary policy changes what they buy. But does the organic local media environment—the coverage that arises from provincial newsrooms, not from a researcher's treatment—generate the same effect in practice? Using daily bank-transaction data across 52 provinces within a single currency union, paired with geocoded local news coverage, we provide the first observational evidence that it does: provinces with richer media coverage of interest-rate decisions respond to the same contractionary shock as if their households are better informed, with consumption declines roughly 18% smaller at the trough. The effect traces entirely to the province's persistent media environment—its structural coverage level in the week before the announcement. The shock-day coverage surge contributes nothing. This result holds when coverage is measured 14, 30, or even 60 days before the shock, a pattern harder to reconcile with time-varying confounders than with a structural information channel. Coverage volume, regardless of tone, is the stronger channel, consistent with an attention mechanism. Media density—outlets per capita—is the strongest time-invariant predictor of differential transmission. At

*We thank Ricardo Reis for detailed comments that substantially improved the paper, and seminar participants at Nova SBE, UCL, the European Central Bank, Banco de Portugal, and BBVA Research for helpful suggestions. The views expressed are those of the authors and do not necessarily reflect those of the European Central Bank, Banco de Portugal, or BBVA. All errors are our own.

†Nova School of Business and Economics. Email: joao.duarte@novasbe.pt

‡University College London and CEPR.

§European Central Bank and Banco de Portugal.

¶BBVA Research and CRIW.

‖BBVA Research.

the category level, all eleven non-food categories attenuate while food alone amplifies, consistent with information resolving uncertainty about the shock's persistence and dampening precautionary pullback on big-ticket items.

Keywords: Monetary policy transmission, media, household consumption, information frictions, local projections
JEL Codes: E52, E21, D83, R12

Contents

1 Introduction **5**

2 Related Literature **7**

3 Data **9**

3.1 Daily Consumption: Bank Transaction Data 9

3.2 Monetary Policy Shocks 10

3.3 Local News Coverage: GDELT 10

3.4 Province Characteristics 11

3.5 Province Contiguity Matrix 11

4 Empirical Strategy **11**

4.1 Baseline Specification 12

4.2 Smooth Local Projections 12

4.3 Interaction Specification 12

4.4 Inference 13

4.5 Identification 13

4.6 Spatial Specification 14

5 Results **15**

5.1 Baseline: Provincial Consumption Response 15

5.2 The News Channel 16

5.2.1 Composite News Index 16

5.2.2 Persistent Media Environment vs. Shock-Responsive Coverage 16

5.2.3 Decomposition: Coverage vs. Sentiment 17

5.3 Province-Level Moderators 18

5.3.1 News by Urban Status 20

5.4 Spatial Transmission 20

5.4.1 Spatial Lag Controls 20

5.4.2 News Diffusion Geography 21

5.4.3 Core-to-Periphery Timing 21

6 The Attenuation Puzzle **23**

6.1 What Theory Actually Predicts 24

6.2 Category-Level Evidence 27

7 Robustness and Extensions **29**
7.1 The Persistent Component Holds at Longer Lags 29
7.2 Discriminating Between Sticky Information and Rational Inattention 30
7.3 Food Amplification: Volume Channel and Endogeneity 30

8 Conclusion **30**

A Summary Statistics **36**

B Smoothing Parameter Selection **36**

C Province Contiguity Matrix **36**

D Inference Details **37**

E Robustness: Kitchen-Sink Specification **37**

F Additional Category-Level Results **38**

G Split-Sample Urban/Non-Urban Results **38**

H Robustness of the News Interaction **38**

I Persistent Media Environment vs. Shock-Responsive Coverage: Full Decomposition **40**

J Multiple Testing Correction for Category-Level Results **41**

K Formal Test of Urban vs. Non-Urban Difference **42**

1 Introduction

When households are experimentally treated with information about monetary policy, they change what they buy. Coibion et al. (2022) establish this in a randomised trial: inflation-expectations treatments move both beliefs and spending composition. But does the organic local media environment—the coverage that arises naturally from provincial newsrooms, not from a researcher's treatment—shape how households respond to interest rate changes in practice? Models of information frictions—whether sticky information (Mankiw and Reis, 2002; Reis, 2006), rational inattention (Sims, 2003; Maćkowiak and Wiederholt, 2009), or epidemiological updating (Carroll, 2003)—all predict that it should: heterogeneous information generates heterogeneous real responses to monetary policy. On the expectations side, these predictions find support (Carroll, 2003; Coibion and Gorodnichenko, 2015; Mankiw et al., 2004). On the spending side, only the experimental evidence of Coibion et al. exists. Whether naturally occurring variation in local media environments generates differential consumption responses to the same monetary shock has not, to our knowledge, been measured. Testing this with observational data demands both geographic granularity—to capture variation in information environments—and high frequency—to detect the *timing* of adjustment, the margin on which information-friction models make their sharpest predictions.

This paper provides the first observational evidence that provinces with richer media coverage of monetary policy respond to the same interest rate shock as if their households are better informed. We combine daily bank-transaction data across 52 Spanish provinces—the near-universe of account activity at one of the country's two largest retail banks—with a geocoded index of local interest-rate news coverage from GDELT. The daily frequency is not incidental: it is what allows us to measure whether higher-coverage provinces respond *earlier*, not just differently in magnitude—a prediction that quarterly or even monthly data cannot test. Spain's membership in the euro area provides a natural laboratory: the ECB sets a single interest rate, but the local information environment—how much provincial media reports on that decision—varies substantially across provinces. This within-monetary-union variation identifies the role of information in monetary transmission, free from the confound of differential national policy responses that plagues cross-country comparisons.

We estimate smooth local projections (Barnichon and Brownlees, 2019) with province fixed effects and find four main results.

(*i*) *Media coverage dampens monetary transmission.* A one-standard-deviation increase in the local news index reduces the consumption trough by approximately 18%—about 0.4 percentage points against a 2.1 p.p. baseline decline at roughly 55 days after the shock. The point estimate is stable when conditioning on province characteristics, and a decomposition

into the province's persistent media environment—its seven-day pre-shock average coverage level—and the shock-responsive spike attributes the entire effect to the persistent component ($t = 2.14$ for volume vs. $t = 0.28$ for the spike), a pattern that survives in a horse-race specification ($t = 2.14$ conditional on the spike). The structural local media ecosystem, not same-day coverage, drives the result.

(*ii*) *Volume matters more than tone.* Decomposing the news channel, coverage intensity—how much local media reports on rate changes, regardless of tone—is the stronger standalone moderator ($t = 2.85$ vs. $t = 1.81$ for the composite index). Both dimensions carry independent information in a horse-race specification ($t = 2.58$ for volume conditional on sentiment; $t = 2.75$ for sentiment conditional on volume), but the greater standalone power of pure volume is consistent with the probability-of-reception channel of Lamla and Vinogradov (2019): what matters most is whether households receive *any* signal about the rate change.

(*iii*) *The effect is not exclusively an urban phenomenon.* Among time-invariant province characteristics, media density—the number of outlets per capita—achieves the highest t-statistic as a standalone moderator ($t = 3.31$), followed by economic size; income and education are not significant at 90%. The dampening is present in both urban and non-urban provinces, with a larger point estimate in non-urban areas, though we cannot reject equality of the two coefficients ($p = 0.51$). That the effect appears in media-sparse provinces as well as media-rich ones is consistent with a genuine information channel, though media density may correlate with unobserved province characteristics.

(*iv*) *The category-level pattern points to reallocation, not uniform pullback.* Food—a non-durable, non-deferrable necessity—is the only category where news *amplifies* the monetary policy effect ($t = -2.23$); all eleven other categories show attenuation. The sign pattern rejects the null of random signs ($p = 0.003$), and six categories including food survive Benjamini–Hochberg FDR correction at $q < 0.10$. We interpret this through the lens of precautionary uncertainty: households in low-coverage provinces, unable to gauge the persistence of a rate change, cut deferrable spending broadly; households in high-coverage provinces, behaving as if they recognise the shock as transitory, maintain big-ticket purchases while adjusting food more actively. The aggregate dampening reflects this compositional shift—deferrable categories, which dominate total spending, drive the net result.

Information does not respect provincial borders. News coverage in neighbouring provinces is associated with local consumption responses of a similar magnitude, though less precisely estimated. Urban provinces reach their consumption trough five days before non-urban ones—a pattern consistent with geographic information diffusion, though we cannot formally distinguish it from other core-to-periphery transmission channels.

Our contribution is fourfold. First, we provide the first observational evidence that nat-

urally occurring local variation in media coverage of monetary policy is associated with differential household spending responses—extending the information-friction research programme, which has operated almost entirely through expectations data (Carroll, 2003; Coibion and Gorodnichenko, 2015; Mankiw et al., 2004) and experimental treatments (Coibion et al., 2022), to observed consumption. A persistent-versus-spike decomposition establishes that the province's structural media environment drives the entire attenuation, while the shock-responsive coverage spike contributes nothing; this result holds when media coverage is measured 14, 30, or even 60 days before the announcement, a pattern harder to reconcile with time-varying confounders than with a structural information channel. Second, the greater standalone power of volume over sentiment delivers the first consumption-side evidence for the attention channel of Lamla and Vinogradov (2019). Third, our finding that media density—a supply-side measure of local outlets per capita—is the strongest moderator complements Berger et al. (2011), who study editorial decisions about ECB coverage content; we show that media infrastructure is associated with differential real outcomes. Fourth, within the regional heterogeneity literature (Hauptmeier et al., 2020; Herreño and Pedemonte, 2022), we add the information dimension: the local media environment is the strongest standalone predictor of differential transmission across Spanish provinces, outperforming income, education, and economic size.

The remainder of the paper is organised as follows. Section 2 reviews the related literature. Sections 3 and 4 describe the data and empirical strategy. Section 5 presents the main findings. Section 6 discusses the category-level mechanism and its theoretical underpinnings. Section 7 presents robustness tests including longer-lag placebos, a test discriminating between sticky-information and rational-inattention models, and a food-amplification endogeneity check. Section 8 concludes.

2 Related Literature

This paper engages four literatures: media and monetary policy expectations, regional heterogeneity in monetary transmission, high-frequency consumption measurement, and spatial propagation of macroeconomic shocks.

Media, information, and monetary policy. Blinder et al. (2008) establish that most households learn about monetary policy through media, not central bank communications. Coibion et al. (2022) go further with an RCT embedded in the Nielsen Homescan panel: randomised information about Fed policy moves both expectations and spending, but untreated households update only when local media covers the announcement. Chahrour et al.

(2021) formalise this delegation of information acquisition in a model where editorial topic selection drives aggregate beliefs. Closest to our mechanism, Lamla and Vinogradov (2019) show that central bank announcements raise the probability that a household receives *any* signal about policy—an attention channel rather than a persuasion channel.

On the supply side, Berger et al. (2011) study what determines the extent and tone of print coverage of ECB decisions, finding that unexpected decisions and larger rate changes attract more reporting. We complement their analysis at a different margin: whereas they examine editorial decisions about content, we show that media *density*—the number of outlets per capita—shapes real consumption outcomes through the sheer volume of intermediation.

Bybee et al. (2024) apply a topic model to 800,000 Wall Street Journal articles and show that news attention tracks and forecasts business cycles at the aggregate level. We bring this logic to the sub-national scale, exploiting local variation in news attention to identify differential spending responses within a single country. Across this literature, spending evidence comes from experimental settings (Coibion et al., 2022) or focuses on expectations, financial prices, and aggregate data; we provide the first observational link from naturally occurring local media variation to differential real household consumption responses to monetary policy.

Regional heterogeneity in monetary transmission. Within currency areas, monetary policy affects regions unevenly. Hauptmeier et al. (2020) find larger consumption and employment responses in poorer NUTS3 regions of the euro area. Herreño and Pedemonte (2022) document the same gradient for US metropolitan areas, driven by housing, labour markets, and financial constraints. Pizzuto (2020) applies spatial local projections to US states and traces heterogeneity to industry composition and housing exposure.

These studies emphasise socioeconomic composition; we add the information dimension. Conditional on income, education, and economic size, the local media environment is the strongest standalone predictor of differential transmission across Spanish provinces. Income and education—the channels stressed in prior work—are individually weaker in our data. This does not contradict Herreño and Pedemonte (2022) or Hauptmeier et al. (2020), who work with different countries and lower-frequency data, but it suggests that, at least in the Spanish context, the information channel may partially account for variation previously attributed to socioeconomic factors.

High-frequency consumption data. Bank transaction data have only recently become available for macroeconomic research. In a companion paper (Buda et al., 2023), we construct daily consumption and investment aggregates from BBVA transaction records and show that

monetary policy transmits to real activity within days, not quarters. That paper operates at the national level; the present paper adds the regional dimension and asks how local information environments mediate the aggregate response. Ferrara et al. (2025) also use Spanish BBVA data but focus on aggregate nowcasting rather than monetary transmission. Our contribution to this strand is to show that high-frequency data can reveal not only faster transmission but also the information frictions that govern its cross-sectional distribution.

Spatial propagation. Ozdagli and Weber (2023) show that 55–85% of monetary policy effects on stock returns propagate through input-output production networks. In our setting, firm-level production data are unavailable, so geographic proximity serves as a proxy for network linkages. Following the methodological template of Pizzuto (2020), we augment local projections with spatial lags of both consumption and news coverage. The suggestive finding that news coverage in neighbouring provinces is associated with local consumption responses is consistent with information diffusion, though we cannot formally distinguish it from trade or production-network channels.

3 Data

We combine three data sources: daily provincial consumption from bank transactions, province-level monetary policy news coverage from GDELT, and a cross-section of time-invariant province characteristics.

3.1 Daily Consumption: Bank Transaction Data

The dependent variable $y_{i,t}$ is the year-on-year growth rate of real spending in province i on day t. The data originate from BBVA, one of the two largest retail banks in Spain, and cover the near-universe of account transactions across all 52 Spanish provinces (NUTS3 regions). Buda et al. (2023) validate these data for macroeconomic analysis, showing that the daily transaction index closely tracks official consumption statistics at the national level. Beyond the aggregate, we observe disaggregated spending across 12 COICOP categories: food, clothing, furnishings, health, transport, communications, recreation, education, restaurants, housing, alcohol/tobacco, and a residual "other" category.

The sample spans September 2018 to February 2026. We exclude the acute COVID-19 lockdown (14 March–30 October 2020), a non-monetary shock of unprecedented magnitude that would otherwise dominate the estimated impulse responses. The resulting panel contains 2,480 daily observations per province and 128,960 province-day observations.

3.2 Monetary Policy Shocks

The treatment variable shock_t is the median high-frequency monetary policy surprise on ECB announcement days, measured as the change in overnight index swap rates in a narrow window around the policy decision. This measure, now standard in the literature, isolates the unexpected component of monetary policy—the part not anticipated from prior communications, economic data, or forward guidance. On non-announcement days, $\text{shock}_t = 0$.

The standard deviation of shock_t on non-zero days is approximately 0.0397. We normalise all impulse responses to a one-standard-deviation shock, so that estimated coefficients represent the consumption response to a typical monetary surprise.

3.3 Local News Coverage: GDELT

Local media coverage of monetary policy comes from the GDELT (Global Database of Events, Language, and Tone) project, which monitors news media worldwide in real time. We extract all articles from Spanish outlets that mention interest rates, monetary policy, or ECB decisions, geocoded to the province level.

The raw GDELT variable combines coverage volume and sentiment tone: positive values indicate favourable or explanatory coverage, negative values indicate unfavourable or alarmist coverage, and the magnitude reflects volume. We standardise this index to mean zero and unit variance at the province level.

For the decomposition analysis, we construct four derived variables: (*i*) *coverage intensity*, $|\text{news}_{i,t}|$, the absolute value of the composite index, capturing the volume of interest-rate reporting regardless of tone; (*ii*) *sentiment direction*, $\text{sign}(\text{news}_{i,t}) \in \{-1, 0, +1\}$, the direction of coverage tone; (*iii*) *positive coverage*, $\max(0, \text{news}_{i,t})$, favourable coverage only; and (*iv*) *negative coverage*, $|\min(0, \text{news}_{i,t})|$, unfavourable coverage only.

Coverage intensity and sentiment direction are near-orthogonal: $\text{Cor}(|\text{news}|, \text{sign}(\text{news})) \approx -0.04$. This near-zero correlation ensures that the horse-race specification between volume and sentiment is well-identified.

Validation. GDELT offers three advantages over hand-collected press archives or proprietary media databases. First, automated real-time monitoring captures a broad cross-section of Spanish online media—regional newspapers, local broadcast transcripts, and online-only outlets—without the selection bias of curated databases focused on national outlets, though coverage is more complete for outlets with an online presence. Second, the geocoding algorithm assigns articles to provinces based on named entities in the text, generating province-day variation without relying on outlet location (which would conflate media supply with

demand). Third, the tone score provides a continuous sentiment measure validated against human coding in multiple languages, including Spanish (Leetaru and Schrodt, 2013). The main limitation is that GDELT captures online media presence rather than household media consumption; we address this by using media density as a complementary, supply-side measure.

3.4 Province Characteristics

We merge a cross-section of province-level characteristics from the Spanish National Statistics Institute (INE) and related sources: (*i*) *urban status*, a binary indicator for 11 urban provinces (those containing a city with population above 500,000 or classified as a major metropolitan area); (*ii*) *media density*, the number of media outlets per capita, capturing the density of the local media landscape; (*iii*) *income*, the log of per capita income; (*iv*) *education*, the share of population with tertiary education; and (*v*) *economic size*, the log of gross value added (GVA).

All continuous moderators are standardised to z-scores (mean zero, unit variance) for comparability. Summary statistics appear in Appendix Table 6.

3.5 Province Contiguity Matrix

For the spatial analysis, we construct a 52×52 queen-contiguity matrix $\mathbf{W}$ from the official INE NUTS3 boundary shapefiles, using the `mapSpain` and `spdep` R packages. Two provinces are neighbours if they share a common border or vertex. The matrix is row-normalised so that each row sums to one. Five provinces have no land neighbours (Baleares, Las Palmas, Tenerife, Ceuta, and Melilla) and are excluded from the spatial analysis, leaving 47 mainland provinces. Details of the construction appear in Appendix C.

4 Empirical Strategy

We estimate smooth panel local projections following Barnichon and Brownlees (2019), augmented with province fixed effects and moderator interactions. We present the baseline specification, the interaction design, and the inference procedure below. Appendix B details the cross-validation procedure for the smoothing parameter.

4.1 Baseline Specification

For each horizon $h = 0, 1, \ldots, 90$ days, we estimate:

$$y_{i,t+h} = \alpha_i + \beta^h \cdot \text{shock}_t + \boldsymbol{\gamma}^{h\prime}\mathbf{X}_{i,t} + u_{i,t+h} \tag{1}$$

where $y_{i,t+h}$ is year-on-year real spending growth in province i at horizon h; α_i are province fixed effects, implemented via within-transformation; shock_t is the ECB monetary policy surprise on day t; and $\mathbf{X}_{i,t} = (y_{i,t-1}, y_{i,t-7}, y_{i,t-14})'$ are lagged dependent variables that absorb serial correlation.

The sequence $\{\hat{\beta}^h\}_{h=0}^{90}$ traces out the impulse response function (IRF) of provincial consumption to a monetary policy shock. All IRFs are scaled to a one-standard-deviation shock (≈ 3.97 basis points).

4.2 Smooth Local Projections

Following Barnichon and Brownlees (2019), we impose smoothness on the IRF by expanding the horizon-varying coefficients in a B-spline basis and penalising their second differences. Instead of 91 separate regressions, we solve a single penalised least-squares problem:

$$\hat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta},\boldsymbol{\gamma}} \left\{ \sum_{h=0}^{H} \sum_{i,t} (y_{i,t+h} - \alpha_i - \beta^h \text{shock}_t - \boldsymbol{\gamma}^{h\prime}\mathbf{X}_{i,t})^2 + \lambda \sum_{h=2}^{H} (\Delta^2 \beta^h)^2 \right\} \tag{2}$$

where $\Delta^2\beta^h = \beta^h - 2\beta^{h-1} + \beta^{h-2}$ is the second difference and $\lambda > 0$ governs the smoothness–fit trade-off. The penalty discourages sharp kinks without imposing a parametric functional form.

We select λ by five-fold cross-validation on the aggregate (province-averaged) consumption series, yielding $\lambda^* = 0.2$. Appendix B reports the cross-validation curve and sensitivity of the main results to alternative values.

4.3 Interaction Specification

To test whether a province characteristic z_i moderates monetary policy transmission, we augment the baseline:

$$y_{i,t+h} = \alpha_i + \beta^h \cdot \text{shock}_t + \delta^h \cdot \text{shock}_t \times \tilde{z}_i + \boldsymbol{\gamma}^{h\prime}\mathbf{X}_{i,t} + u_{i,t+h} \tag{3}$$

where $\tilde{z}_i = z_i - \bar{z}$ is the de-meaned moderator. De-meaning ensures that β^h retains its interpretation as the response at the sample mean of z, while δ^h captures how a one-unit

increase in z shifts the impulse response.

When z varies at daily frequency (e.g., the GDELT news index), the interaction shock$_t \times z_{i,t}$ is time-varying; when z is a time-invariant province characteristic (e.g., media density), identification comes purely from cross-sectional variation.

The IRF at a specific moderator value U is:

$$\text{IRF}^h(U) = \hat{\beta}^h + \hat{\delta}^h \cdot (U - \bar{z}) \tag{4}$$

We evaluate IRFs at one standard deviation below, at, and one standard deviation above the moderator mean, labelled "low," "average," and "high."

4.4 Inference

Because all provinces share the same monetary policy surprise on any given day, residuals are cross-sectionally correlated. We cluster standard errors by date. The cluster-robust variance estimator is:

$$\hat{V} = (\mathbf{Z}'\mathbf{Z})^{-1} \left(\sum_{t=1}^{T} \mathbf{Z}_t' \hat{\mathbf{u}}_t \hat{\mathbf{u}}_t' \mathbf{Z}_t \right) (\mathbf{Z}'\mathbf{Z})^{-1} \tag{5}$$

where $\mathbf{Z}_t$ stacks all regressors (shock and interactions) and $\hat{\mathbf{u}}_t$ collects the residuals across provinces on date t. Standard errors at arbitrary moderator values U follow from a Jacobian mapping from the stacked coefficient vector to the IRF evaluated at U; details appear in Appendix D.

All figures display two-tier confidence bands: 68% ($\hat{\beta}^h \pm 1.00 \times \hat{\sigma}^h$) and 90% ($\hat{\beta}^h \pm 1.645 \times \hat{\sigma}^h$), following convention in the local projections literature (see, e.g., Jordà, 2005; Barnichon and Brownlees, 2019). We note that our headline composite news interaction ($t = 1.81$) is significant at the 10% but not the 5% level; the persistent-media-environment and coverage-intensity specifications clear conventional thresholds.

4.5 Identification

Identification rests on two pillars: high-frequency monetary policy surprises and the local projection estimator. The shocks of Buda et al. (2023), constructed from changes in overnight index swap rates in narrow windows around ECB announcements, isolate the unexpected component of policy by differencing out all pre-announcement information. This high-frequency identification strategy, following Kuttner (2001) and Gürkaynak et al. (2005), ensures orthogonality between the shocks and macroeconomic conditions at the time of announcement—resolving the simultaneity problem that confounds direct estimation of mon-

etary policy effects (see the discussion in Stock and Watson, 2018; Ramey, 2016).

The local projection estimator (Jordà, 2005) delivers a direct, horizon-by-horizon estimate of the impulse response without specifying a full dynamic system. As Stock and Watson (2018) show, local projections with external instruments yield consistent causal estimates under the standard relevance and exogeneity conditions. In our panel, province fixed effects absorb time-invariant unobserved heterogeneity and clustering by date accounts for the common-shock structure, together delivering valid inference for the causal effect of monetary surprises on provincial consumption.

Two features of the design warrant emphasis. First, the interaction specification studies *moderation*, not treatment assignment: the monetary policy shock is exogenous by construction, and we ask whether cross-provincial variation in news coverage moderates its pass-through. The news variable need not satisfy the same exogeneity as the shock, but the interaction must be uncorrelated with unobserved province characteristics that independently moderate transmission. Temporal precedence of news relative to consumption at horizon h is necessary but not sufficient; the key assumption is that cross-provincial variation in news environments does not proxy for omitted moderators. This assumption is plausible because GDELT coverage reflects editorial decisions by local outlets about which national stories to relay—decisions shaped by newsroom capacity and competing local events—though we cannot fully exclude the possibility that local economic conditions influence both media coverage and consumption responses.

Second, we probe this assumption in three ways: the kitchen-sink specification controls for observable province×shock interactions (Appendix E); media density—a time-invariant, supply-side measure—performs well as an independent moderator; and a decomposition of the news index into a persistent component (seven-day pre-shock average) and a shock-responsive spike yields a decisive attribution: the persistent component alone achieves $t = 2.14$ for volume, while the spike is essentially zero ($t = 0.28$; Section 5.2.2), confirming that the structural local media environment, rather than same-day coverage, drives the result. None of these tests definitively rules out confounding from unobserved province characteristics, but collectively they narrow the scope of plausible alternatives.

4.6 Spatial Specification

To examine geographic spillovers, we augment the baseline with spatial lags of the dependent variable:

$$y_{i,t+h} = \alpha_i + \beta^h \text{shock}_t + \phi_1^h (\mathbf{W}y)_{i,t-1} + \phi_7^h (\mathbf{W}y)_{i,t-7} + \boldsymbol{\gamma}^{h\prime} \mathbf{X}_{i,t} + u_{i,t+h} \tag{6}$$

where $(\mathbf{W}y)_{i,t} = \sum_j W_{ij} y_{j,t}$ is the row-normalised average consumption of province i's neighbours.

A news-diffusion extension adds both own and neighbour news interactions:

$$y_{i,t+h} = \alpha_i + \beta^h \text{shock}_t + \delta_1^h \text{shock}_t \times \text{news}_{i,t} + \delta_2^h \text{shock}_t \times (\mathbf{W}\text{news})_{i,t} + \boldsymbol{\gamma}^{h\prime}\mathbf{X}_{i,t} + u_{i,t+h} \quad (7)$$

where δ_1^h captures the own-province news effect and δ_2^h the neighbour-province effect. A positive δ_2 indicates that monetary policy information spills across provincial borders.

5 Results

We present results in four stages: the baseline impulse response (Section 5.1), the news channel and its decomposition (Section 5.2), province-level moderators (Section 5.3), and spatial transmission (Section 5.4).

5.1 Baseline: Provincial Consumption Response

Before asking whether local news moderates monetary transmission, we establish that there is a well-identified baseline to moderate. Figure 1 plots the impulse response of provincial consumption to a one-standard-deviation contractionary ECB monetary policy shock. The response is negative and hump-shaped: consumption growth begins to decline within the first week, reaches a trough of approximately -2.1 percentage points around day 55, and partially recovers by day 90 but remains below zero.

The shape and timing accord with the "long and variable lags" view of monetary transmission, but the daily frequency reveals a faster onset than quarterly data would suggest. The response is statistically significant at the 90% level from approximately day 10 through day 80, with the tightest confidence bands near the trough. For context, this provincial-level response is comparable in magnitude and timing to the aggregate estimates of Buda et al. (2023) using the same BBVA data at the national level—confirming that the panel structure does not introduce artefacts relative to the aggregate.

The baseline is stable across subsamples and robust to alternative smoothing parameters (Appendix B).

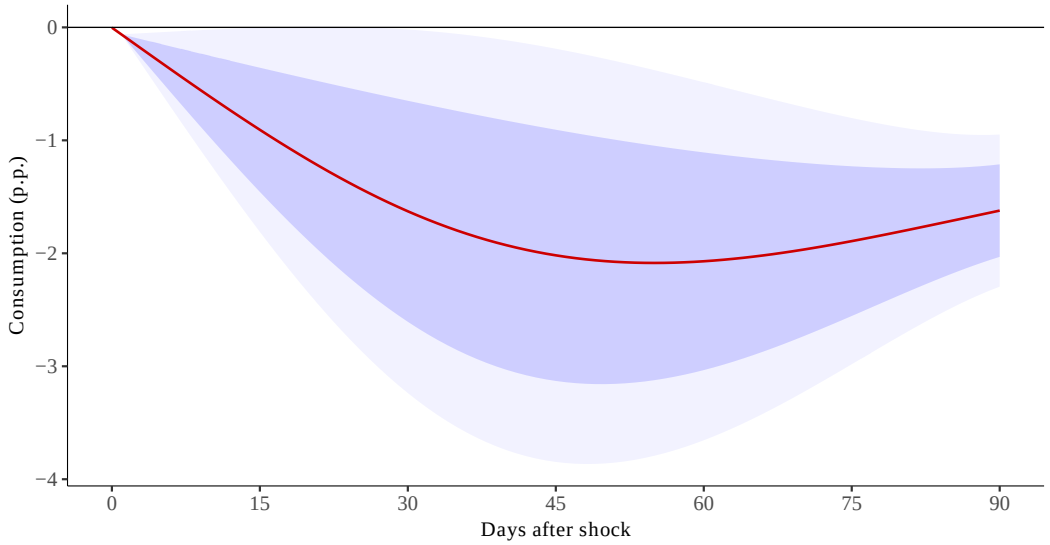


Figure 1: Baseline: Provincial Consumption Response to a 1-s.d. ECB Monetary Policy Shock

Notes: Smooth panel local projection with 52 province fixed effects. Y-axis in percentage points. Shaded areas: 68% and 90% confidence intervals (cluster-by-date SEs). Horizon: 90 days. COVID-19 lockdown excluded. IRF scaled to a one-standard-deviation shock (≈ 3.97 bp).

5.2 The News Channel

5.2.1 Composite News Index

Figure 3 presents the central result. At the peak of the interaction (around day 43), the coefficient is +0.0038 ($t = 1.81$, significant at 10% but not 5%): provinces with one-standard-deviation higher news coverage experience a consumption decline approximately 18% smaller than the baseline trough. The effect builds gradually over the first six weeks and persists through day 90.

The point estimate is stable when conditioning on province characteristics. In a kitchen-sink specification that simultaneously controls for interactions of the shock with urban status, income, education, GVA, and media density, the news interaction remains positive and economically meaningful, though precision declines modestly ($t = 1.76$; Appendix E).

5.2.2 Persistent Media Environment vs. Shock-Responsive Coverage

A natural concern is that the contemporaneous news index captures media coverage *of* the ECB decision itself—articles published within hours of the announcement—rather than the pre-existing local information environment. If so, the interaction would partly reflect a mechanical correlation between shock size and media response.

We address this through a decomposition that isolates two distinct components of the news index. The *persistent* component is the province's seven-day pre-shock moving average, $\overline{\text{news}}_{i,t-7:t-1}$, capturing the structural level of local interest-rate coverage—a feature of the provincial media ecosystem that is predetermined relative to the ECB announcement. The *spike* component is the within-day deviation from that baseline, $\text{news}_{i,t} - \overline{\text{news}}_{i,t-7:t-1}$, capturing the surge in coverage triggered by the announcement itself. By construction, the two components sum to the contemporaneous index and the persistent component is exogenous to the current shock.

Figure 2 reports the results using coverage volume. The contrast is decisive. In standalone specifications, the persistent component produces a strong and significant interaction ($t = 2.16$), while the spike contributes essentially nothing ($t = 0.28$). The horse-race specification—including both components simultaneously—confirms the attribution: the persistent component retains its significance ($t = 2.14$) and the spike remains negligible ($t = 0.54$).

This decomposition transforms the identification argument. Rather than showing merely that predetermined news "also works," it establishes that the province's structural media environment does *all* the work. The shock-responsive coverage spike—the component most vulnerable to reverse-causality concerns—plays no role. The result is consistent with media density as a supply-side driver: what matters is whether a province has the media infrastructure to sustain regular interest-rate coverage, not whether outlets happen to run a story on announcement day. It does not, however, rule out that the persistent media environment correlates with other unobserved province characteristics that independently moderate transmission. We report the full decomposition, including signed-index results, in Appendix I.

5.2.3 Decomposition: Coverage vs. Sentiment

The composite news index confounds two potentially distinct channels: coverage *volume*—how much local media reports on interest rates—and coverage *sentiment*—whether the tone is reassuring or alarming. We separate them.

Figures 4 and 5 report the decomposition results across six specifications.

(*i*) Coverage intensity $|\text{news}_{i,t}|$: peak interaction +0.0055 at day 50, $t = 2.85$. (*ii*) Composite $\text{news}_{i,t}$: peak +0.0038 at day 43, $t = 1.81$ (reference). (*iii*) Positive coverage $\max(0, \text{news}_{i,t})$: peak +0.0077 at day 46, $t = 1.93$. (*iv*) Negative coverage $|\min(0, \text{news}_{i,t})|$: peak +0.0016 at day 90, $t = 1.00$. (*v*) Coverage conditional on sentiment: +0.0059, $t = 2.58$. (*vi*) Sentiment conditional on coverage: +0.0058, $t = 2.75$.

Three findings emerge. First, coverage intensity is a stronger standalone moderator than the composite ($t = 2.85$ vs. $t = 1.81$). The composite, by netting positive against negative

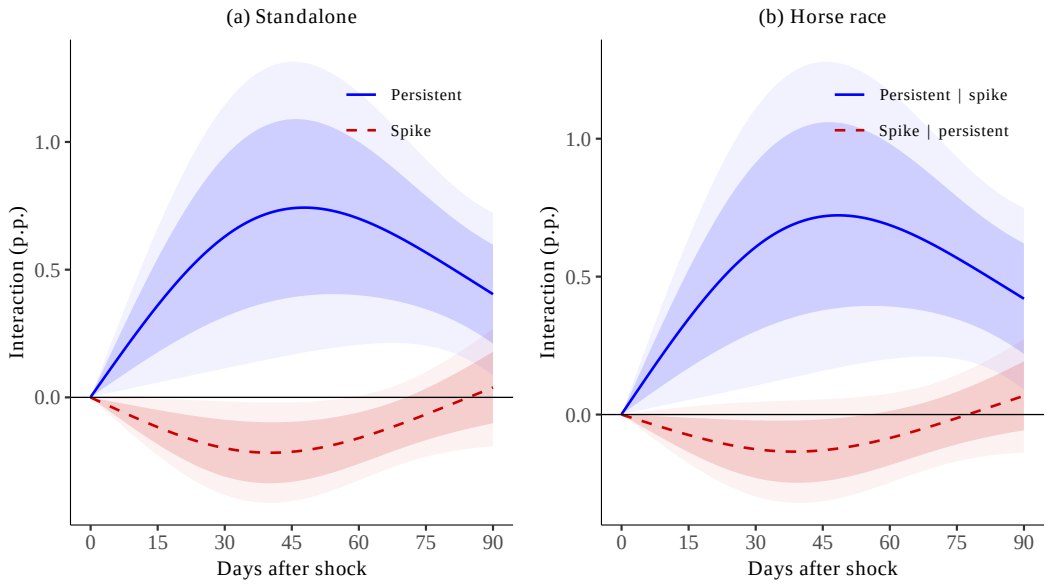


Figure 2: Persistent Media Environment vs. Shock-Responsive Coverage Spike (Volume)

Notes: Panel (a): interaction coefficients from separate specifications using the persistent component (seven-day pre-shock average, solid blue) and the shock-responsive spike (within-day deviation, dashed red) of coverage volume. Panel (b): horse-race specification including both components simultaneously. Persistent drives the entire attenuation ($t = 2.14$); spike contributes nothing ($t = 0.54$). Shaded areas: 68% and 90% CIs. Cluster-by-date SEs.

coverage, introduces noise that weakens the signal.

Second, volume and sentiment carry independent information. In the horse race (Specifications 5 and 6), coverage conditional on sentiment ($t = 2.58$) and sentiment conditional on coverage ($t = 2.75$) both remain significant. Neither channel subsumes the other.

Third, there is an asymmetry by sentiment direction. Favourable or explanatory coverage ($+0.0077$, $t = 1.93$) is substantially larger than alarmist coverage ($+0.0016$, $t = 1.00$), though we cannot formally reject equality of the two coefficients and the negative-coverage estimate may partly reflect low power.

These patterns are consistent with the attention mechanism of Lamla and Vinogradov (2019): the greater standalone power of volume suggests that *receiving any signal* about the rate change matters, though the horse race confirms that tone also carries independent information.

5.3 Province-Level Moderators

One might worry that the media channel proxies for observable province characteristics. If richer or more educated provinces happen to have both more media coverage and better consumption-smoothing for unrelated reasons, the news effect would be spurious. We address

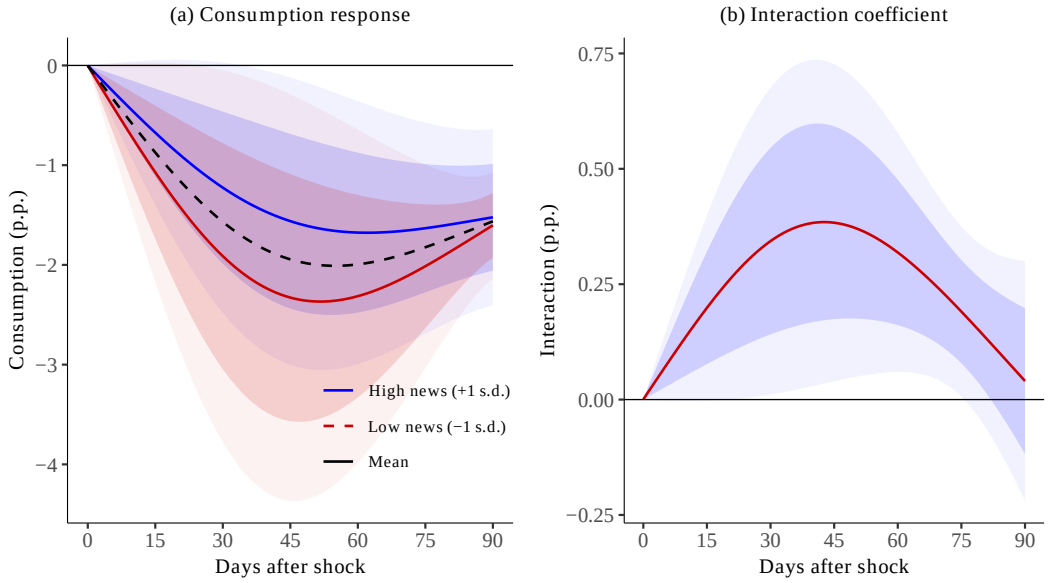


Figure 3: News Channel: Consumption Response at Low vs. High News Coverage

Notes: Panel (a): consumption IRFs from the interaction specification (Equation 3) with the GDELT news index as moderator; "Low" and "High" correspond to ±1 s.d. from the mean; shaded areas: 68% and 90% CIs. Panel (b): interaction coefficient $\hat{\delta}^h$ in percentage points. Cluster-by-date SEs.

this concern directly.

Table 1 presents four province characteristics as standalone interaction moderators—not as controls for the news channel, but as independent candidates for explaining heterogeneity in monetary transmission. Each is estimated in a separate specification.

Media density achieves the highest t-statistic among province characteristics ($t = 3.31$), followed by economic size ($t = 1.87$); income and education are significant only at the 68% level. Media density is a time-invariant, supply-side measure—the number of media outlets per capita in the province—distinct from the daily GDELT news content index. That a structural measure of media infrastructure outperforms socioeconomic characteristics in standalone specifications is consistent with a genuine information channel, though we note that media density may correlate with unobserved province characteristics (population density, service-sector share) that independently moderate transmission.

The income result merits comment. Contrary to the cross-country findings of Herreño and Pedemonte (2022) and Hauptmeier et al. (2020), who report that poorer regions respond more to monetary policy, our point estimate suggests the opposite direction (higher income → slightly stronger response), though it is insignificant at the 90% level. One possible explanation is Spain's mortgage structure: wealthier provinces tend to have greater exposure to variable-rate mortgages, making them more interest-rate sensitive through the cash-flow

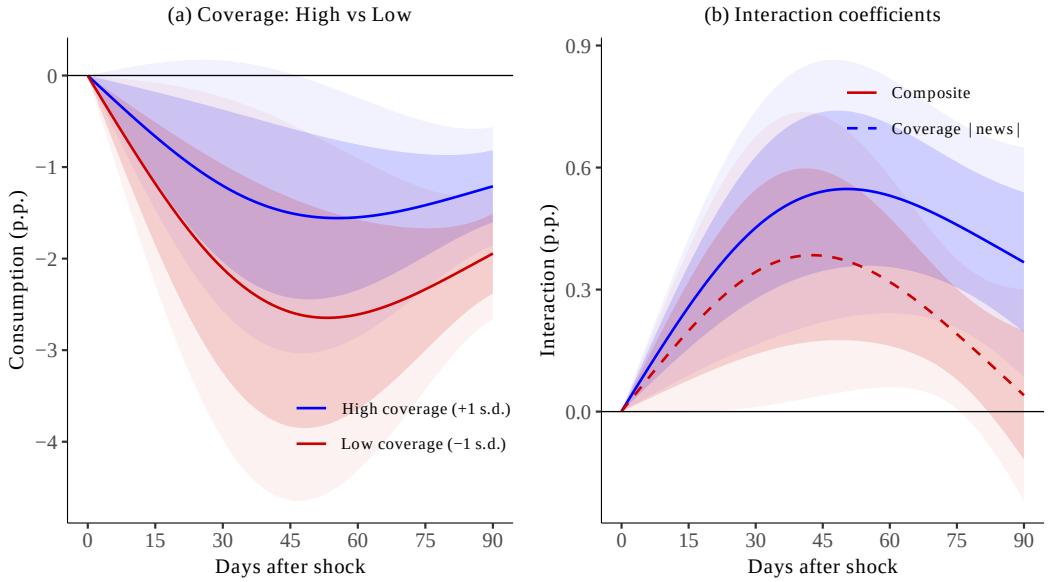


Figure 4: Decomposition: Coverage Volume vs. Composite News Index

Notes: Panel (a): consumption IRFs at high and low coverage intensity (±1 s.d.); shaded areas: 68% and 90% CIs. Panel (b): interaction coefficients $\hat{\delta}^h$ from separate specifications using coverage intensity |news| and the composite index. Coverage volume ($t = 2.85$) is the stronger moderator.

channel, a mechanism distinct from the financial-constraints channel emphasised in the cross-country literature.

5.3.1 News by Urban Status

As a direct test of the urbanisation confound, we estimate the news-interaction LP separately on the 11 urban and 41 non-urban provinces. The point estimate is positive in both subsamples: +0.0045 ($t = 1.98$) in non-urban provinces and +0.0026 ($t = 1.44$) in urban ones. The difference is not statistically significant ($t_{\text{diff}} = 0.66$, $p = 0.51$; Appendix K), and the urban subsample lacks power with only 11 provinces. We conclude that the news effect is present in both subsamples and is not driven exclusively by urban provinces, but we cannot claim a significant difference between the two groups.

5.4 Spatial Transmission

5.4.1 Spatial Lag Controls

A common monetary policy shock—transmitted through national financial markets—should hit all provinces approximately simultaneously. Spatial lags of consumption, if they matter, would indicate trade or production-network propagation rather than the direct financial

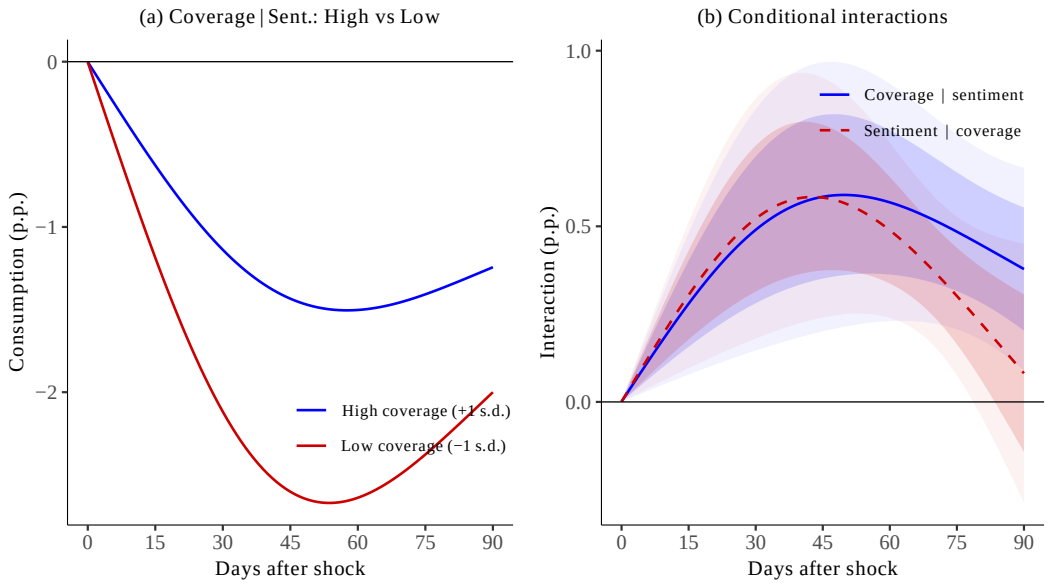


Figure 5: Horse Race: Coverage Controlling for Sentiment vs. Sentiment Controlling for Coverage

Notes: Panel (a): approximate consumption IRFs at high vs. low coverage, conditional on sentiment. Panel (b): conditional interaction coefficients with 68% and 90% CIs. Both survive the horse race (coverage|sentiment: $t = 2.58$; sentiment|coverage: $t = 2.75$).

channel. Adding lagged neighbour consumption ($(\mathbf{W}y)_{i,t-1}$ and $(\mathbf{W}y)_{i,t-7}$) as controls barely changes the estimated IRF: the trough remains virtually unchanged at -2.10 p.p.. Spatial structure does not absorb the direct effect.

5.4.2 News Diffusion Geography

While production linkages do not mediate the direct monetary policy effect, *information* linkages may mediate the news channel. The dual-news specification (Equation 7) tests this directly. Table 2 reports:

The neighbour-news point estimate is larger in magnitude than the own-news coefficient, though less precisely estimated ($t \approx 1.49$, significant only at the 68% level) due to the averaging inherent in the spatial lag. One interpretation is geographic information diffusion: shared media markets, commuting patterns, and social networks create cross-province information flows. We cannot, however, formally distinguish information diffusion from correlated economic shocks or trade linkages between neighbouring provinces.

5.4.3 Core-to-Periphery Timing

Estimating separate IRFs for urban and non-urban mainland provinces reveals a timing differential suggestive of core-to-periphery propagation. Urban provinces reach their con-

Table 1: Province-Level Moderators of Monetary Policy Transmission

Moderator ($\tilde{z}_i$)	Peak $\hat{\delta}^h$ ($\times 10^4$)	Peak day
Media density (outlets per capita)	6.14^{**}	90
	(1.85)	
Economic size (log GVA)	5.44^{**}	90
	(2.90)	
Income (log per capita)	-7.71^{*}	50
	(4.89)	
Education (% tertiary)	-4.42^{*}	90
	(3.00)	

Notes: Each row reports a separate interaction specification $y_{i,t+h} = \alpha_i + \beta^h \varepsilon_t^{mp} + \delta^h \varepsilon_t^{mp} \times \tilde{z}_i + \gamma' X_{i,t} + u_{i,t+h}$, where $\tilde{z}_i$ is the de-meaned, standardised province characteristic. $N = 128{,}960$ province-day observations in all specifications. Positive coefficients indicate attenuation (weaker MP effect). Standard errors (in parentheses) clustered by date. *, **: 68% and 90% confidence intervals exclude zero.

Table 2: News Diffusion Geography: Own vs. Neighbour-Province News

News channel	Peak $\hat{\delta}^h$ ($\times 10^4$)	Peak day
Own-province news ($\hat{\delta}_1$)	42.25^{**}	42
	(23.31)	
Neighbour-province news ($\hat{\delta}_2$)	52.09^{*}	45
	(34.96)	

Notes: Both interactions estimated simultaneously. Neighbour news is the row-normalised weighted average of adjacent provinces' GDELT news indices, using queen contiguity. Sample: 47 mainland provinces. Standard errors (in parentheses) clustered by date. *, **: 68% and 90% confidence intervals exclude zero.

sumption trough at day 48 (-0.0238), five days earlier than non-urban provinces (day 53, -0.0212). The urban trough is 12% deeper, but urban provinces recover faster: by day 90, the urban consumption effect is -0.0117 versus -0.0149 for non-urban provinces.

This pattern is consistent with urban centres—which have denser media environments and more diverse information sources—absorbing the monetary shock earlier and more sharply, though we note that faster urban responses could also reflect denser financial markets, greater variable-rate mortgage exposure, or other structural differences.

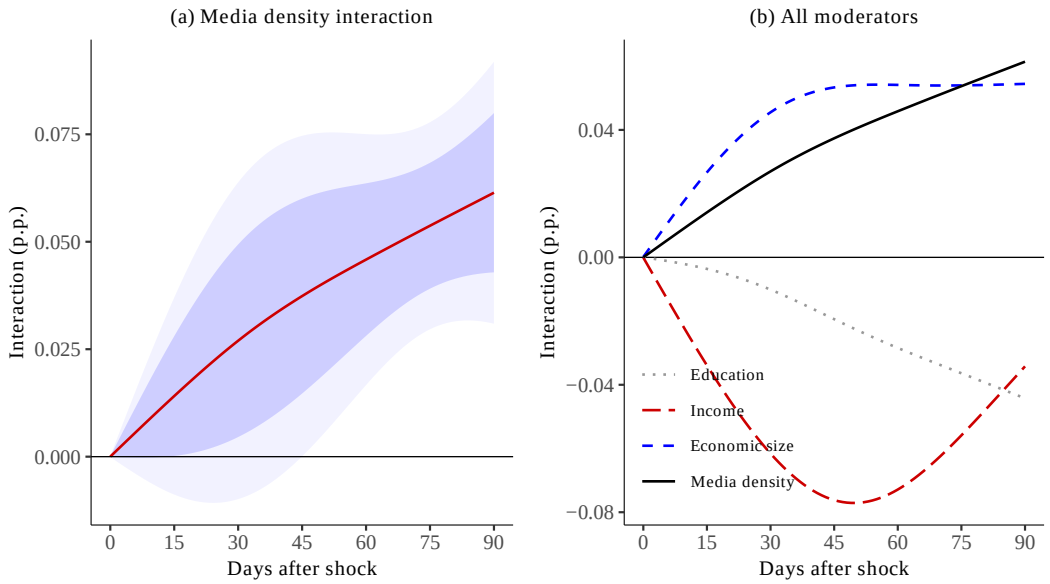


Figure 6: Province-Level Moderators: Interaction Coefficients

Notes: Panel (a): media density interaction coefficient $\hat{\delta}^h$ in percentage points; shaded areas: 68% and 90% CIs (cluster-by-date SEs). Panel (b): interaction coefficients from four separate moderator specifications. Media density ($t = 3.31$) achieves the highest t-statistic. Positive values indicate dampening.

6 The Attenuation Puzzle

Our main result—dampening rather than amplification—poses a theoretical puzzle, because two forces pull in opposite directions. *Intertemporal substitution* should make informed households react *more*: in the New Keynesian Euler equation (Woodford, 2003), better knowledge of the rate path strengthens the motive to defer purchases, and in rational inattention models (Sims, 2003; Maćkowiak and Wiederholt, 2009), more attention to a signal mechanically increases the response. A second force may work against this: *precautionary uncertainty*. If uninformed households observe the downstream consequences of tightening without understanding their cause, they may pull back broadly on deferrable items because they cannot distinguish a transitory rate change from a permanent deterioration. Households in high-coverage provinces, behaving as if they recognise the shock as temporary, would face no such motive. Which force dominates depends on the perceived persistence of the shock. We conjecture—and find category-level evidence consistent with the conjecture—that for the high-frequency surprises in our sample, precaution exceeds substitution. We review three frameworks that support this interpretation, then turn to the evidence.

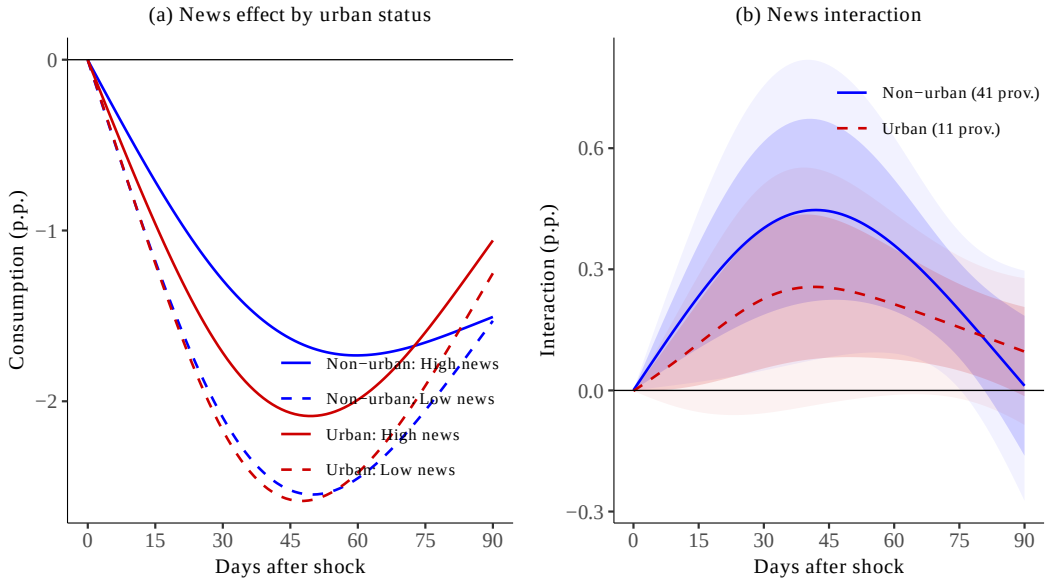


Figure 7: News Attenuation by Urban Status

Notes: Panel (a): approximate consumption IRFs at high vs. low news for urban and non-urban subsamples. Panel (b): news interaction coefficients estimated separately on the 11 urban and 41 non-urban subsamples; 68% and 90% CIs. Attenuation is larger in non-urban provinces ($t = 1.98$) than urban ($t = 1.44$).

6.1 What Theory Actually Predicts

Anticipated policy has smaller real effects. In the sticky information model of Mankiw and Reis (2002), *announced disinflations are less contractionary than surprise disinflations*: when agents have already incorporated a policy change into their plans, it has smaller real consequences. The model concerns firm pricing rather than household consumption, but a loose analogy can be drawn.

In our setting, provinces with more media coverage may effectively receive the rate change as partially anticipated policy—households have read about interest-rate developments and absorbed their implications. Provinces with less coverage experience it as more of a surprise. The direction of the prediction matches our finding: more information → more anticipation → smaller real effects. We stress, however, that the mapping is informal—Mankiw and Reis (2002) model pre-announcement of future policy, whereas our mechanism concerns media coverage of a shock that has already occurred.

Inattentive consumers and excess smoothness. Why does time-invariant media *density*—not just daily coverage variation—emerge as a powerful moderator? Reis (2006) offers a micro-foundation. Consumers optimally choose how frequently to update their consumption plans, trading off the cost of acquiring information against the benefit of accuracy.

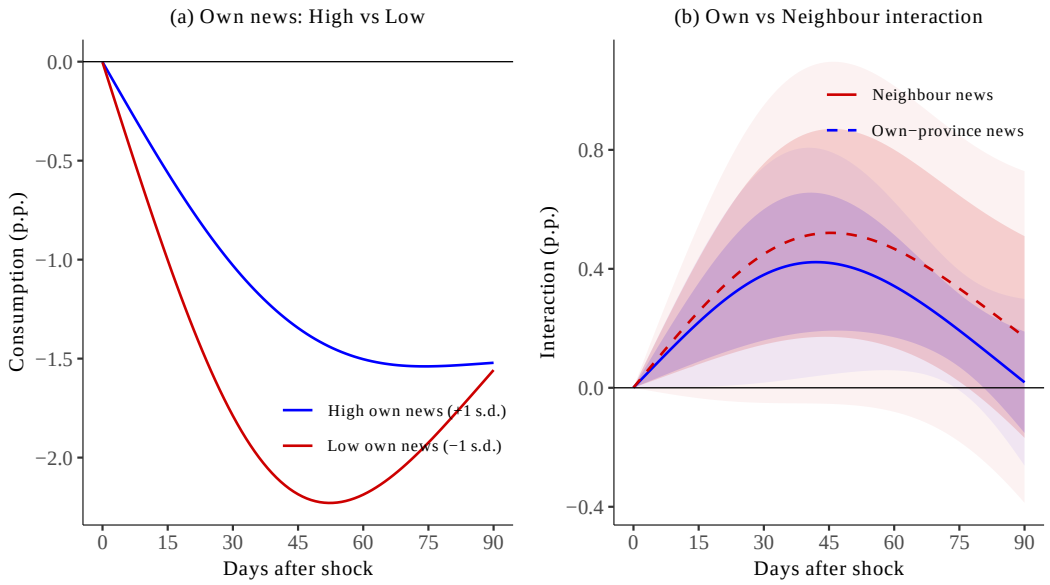


Figure 8: News Diffusion Geography: Own vs. Neighbour News Effects

Notes: Panel (a): approximate consumption IRFs at high vs. low own-province news from the dual-news specification. Panel (b): interaction coefficients for own-province and neighbour-province news (Equation 7); 68% and 90% CIs. 47 mainland provinces. Cluster-by-date SEs.

Between updates they are inattentive, following their previous plan regardless of intervening news.

Media density may proxy for the cost side of this trade-off: denser media environments lower the effective cost of information acquisition, potentially inducing more frequent updating. Updating consumers adjust to permanent income and therefore recognise transitory shocks as transitory. The result is not that informed consumers react more; it is that they react *optimally*—and for a transitory monetary shock, the optimal response is smaller than the uninformed precautionary pullback. That media density achieves the highest t-statistic among province characteristics ($t = 3.31$) is consistent with this interpretation, though we cannot directly test whether consumers in high-media-density provinces actually update more frequently.

Comparison with CGW spending results requires caution. Coibion et al. (2022) also find spending effects: information treatments that raise inflation expectations lead to higher non-durable spending (intertemporal substitution) but *lower* durable purchases—an *amplification* of the durable decline. Our setting yields the opposite sign on durables: information about rate changes leads to a *smaller* decline in durable consumption. The comparison, however, is not apples to apples. CGW's treatment is inflation-expectations information; ours is media coverage of interest-rate decisions—fundamentally different infor-

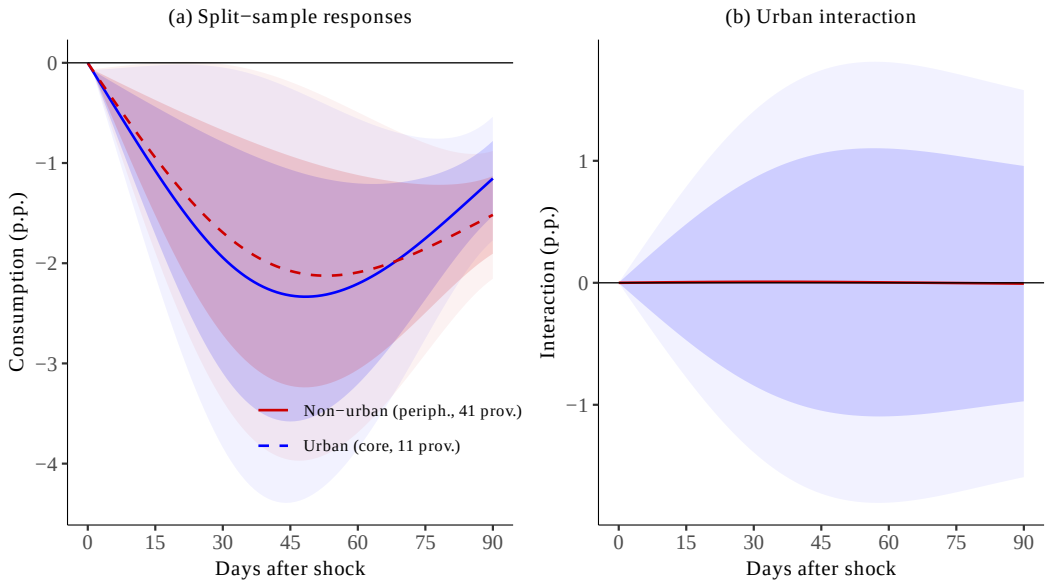


Figure 9: Core-to-Periphery: Urban vs. Non-Urban Consumption Responses

Notes: Panel (a): separate baseline IRFs for urban and non-urban subsamples; shaded areas: 68% and 90% CIs. Panel (b): urban interaction coefficient $\hat{\delta}^h$ in percentage points. Urban provinces trough 5 days earlier ($h = 48$ vs. $h = 53$). Cluster-by-date SEs.

mation content. CGW's inflation information triggers pessimism about economic conditions, reinforcing the pullback in durables. Our interest-rate news helps households assess a rate change as transitory, dampening the precautionary motive to defer purchases. What both results share is that information changes the *composition* of spending, not just its level. But the direction of that compositional shift depends on what households learn.

Precision reduces precautionary behaviour. Why does coverage *volume* ($t = 2.85$) outperform the composite index ($t = 1.81$) as a standalone moderator? If framing mattered most, sentiment should dominate. Instead, what matters is whether households receive *any signal at all*—exactly the prediction of Lamla and Vinogradov (2019), who show that news exposure "improves precision of perceptions and expectations, increases confidence and lowers dispersion of beliefs." The consumption implication is direct: less uncertainty reduces precautionary behaviour. Without coverage, a household facing tighter conditions cannot gauge how severe or persistent the shock is and cuts spending broadly. With coverage, it forms precise beliefs about the rate change's magnitude and duration, eliminating the uncertainty that drives precaution.

The horse-race specification—where volume conditional on sentiment ($t = 2.58$) and sentiment conditional on volume ($t = 2.75$) both survive—confirms that attention and precision operate as complementary channels rather than substitutes.

Synthesis. All three frameworks point in the same direction: information may *eliminate* a large precautionary response rather than *activate* a large substitution response. Table 3 summarises the two channels. The textbook substitution channel is present but modest for shocks perceived as transitory. The precautionary channel is potentially large: uninformed households, unable to judge whether tighter conditions reflect a temporary rate move or a permanent deterioration, may pull back broadly on deferrable spending. If the precautionary margin dominates, the net effect of information is dampening. The standard prediction—more information amplifies the response—would hold when the shock is permanent and substitution is the binding margin. For shocks perceived as transitory, information allows households to identify the shock as temporary and adjust *less*, not more.

Table 3: How Information Changes the Household Response

	Without information	With information
Perception	Shock arrives as surprise Uncertain about persistence Uncertain about magnitude	Shock arrives as anticipated policy Recognises shock as transitory Has precise beliefs about magnitude
IS channel	Weak (don't know rate path)	Present but small (transitory shock)
Precautionary channel	Strong (can't rule out permanent)	Absent (persistence resolved)
Net effect	Large pullback (precaution $\gg$ IS)	Small adjustment (IS only)
Durables	Broad precautionary cut	Maintain (no precautionary motive)
Food	Barely adjust	Actively optimise (reallocation)

6.2 Category-Level Evidence

The theory predicts that better-informed households attenuate deferrable spending (where the option value of waiting depends on perceived persistence) but may amplify non-deferrable spending (where informed reallocation can be actively optimised). Table 4 presents the evidence.

The pattern across categories discriminates among mechanisms. All eleven non-food categories show positive interaction coefficients—attenuation—with seven significant at the 90% level (clothing, transport, education, restaurants, health, housing, furnishings) and four more at the 68% level (recreation, miscellaneous, communication, alcohol). The consistency of signs across categories with very different baseline responses, purchase frequencies, and durability classifications is hard to explain as noise: a one-sided binomial test of the sign pattern (11 positive, 1 negative) rejects the null at $p = 0.003$, and six categories including food survive Benjamini–Hochberg FDR correction at $q < 0.10$ (Appendix J). Food stands alone: news *amplifies* the effect ($t = -2.23$). The key distinction is not durability per se

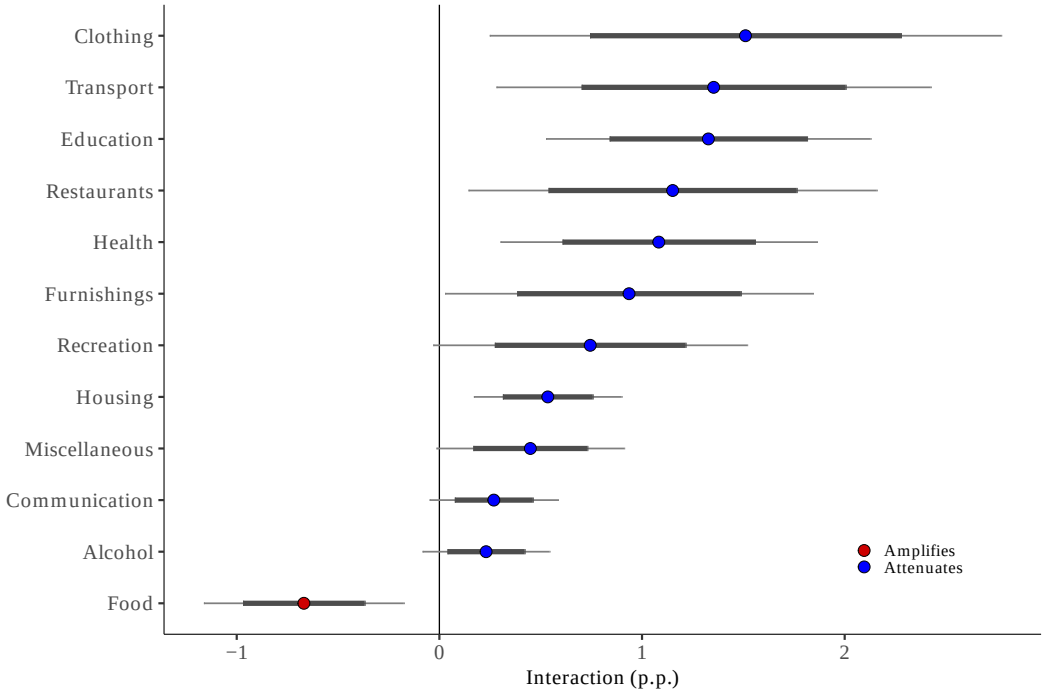


Figure 10: News Interaction by Consumption Category

Notes: Peak news interaction coefficients by COICOP category, in percentage points. Positive values (blue) indicate attenuation; negative values (red) indicate amplification. Thick bars: 68% CIs; thin bars: 90% CIs. Cluster-by-date SEs. Food is the only category where news amplifies the monetary policy effect.

but *deferrability*: categories where households can adjust the timing of purchases attenuate; food—the quintessential non-deferrable necessity—amplifies.

Durables and discretionary goods: attenuation. Deferrable purchases are precisely where the two channels collide. An uninformed household cannot tell whether rates will stay high permanently or revert quickly, and responds with a broad precautionary pullback on big-ticket items. An informed household faces a modest substitution motive but *no* precautionary uncertainty. The net result is attenuation: informed households cut deferrable spending less, not because they ignore the rate change, but because they correctly assess its persistence.

Food: amplification. Food is non-durable and non-deferrable, so the precautionary channel has little purchase: uninformed households barely adjust food spending (the baseline trough is near zero). The negative interaction coefficient ($t = -2.23$) indicates that informed households make a *larger* food adjustment. One interpretation is that households maintaining deferrable spending face a tighter effective budget and optimise food more ac-

tively; another is that informed households, understanding the rate path, make deliberate adjustments to non-deferrable necessities that uninformed households do not. We observe total food spending, not within-category substitution, so we cannot identify the specific behavioural margin.

Resolution. Under this interpretation, the category-level evidence resolves the apparent tension between theory and our main finding. For food—non-deferrable, with no precautionary margin—the substitution channel operates cleanly and informed households respond *more*. For deferrable categories, the precautionary channel dominates and informed households respond *less*. The aggregate dampening reflects the fact that deferrable categories dominate total spending.

The bottom line: households in high-coverage provinces do not react less—they react *differently*. The local information environment appears to replace a broad precautionary pullback with a targeted reallocation. We stress that this interpretation, while consistent with the category-level pattern, rests on informal reasoning rather than a formal model; direct evidence on household uncertainty and within-category substitution would sharpen the test.

7 Robustness and Extensions

7.1 The Persistent Component Holds at Longer Lags

If the persistent media environment captures a structural information advantage, the result should hold when measuring coverage further before the shock. We construct four persistent-component moderators using coverage volume averaged over 7, 14, 30, and 60 days before each ECB announcement.

Table 5: Longer-Lag Placebos: Persistent Coverage Volume at Different Pre-Shock Windows

Pre-shock window	Peak $\hat{\delta}^h$	Peak t-stat	Peak horizon
7-day (reference)	0.0075	2.04	$h = 47$
14-day	0.0072	1.85	$h = 45$
30-day	0.0066	1.73	$h = 43$
60-day	0.0066	1.83	$h = 43$

Notes: Each row is a separate panel LP using the indicated pre-shock moving average of coverage volume as the moderator. Cluster-by-date SEs with CR1 correction.

The coefficient barely decays from the 7-day to the 60-day window (0.0075 to 0.0066), and

the t-statistic remains above 1.7 throughout. Media coverage measured two months before an ECB decision predicts differential consumption responses almost as well as coverage measured one week before. This is harder to reconcile with time-varying local economic conditions driving both coverage and the consumption response: a province's economic situation 60 days before a rate change is unlikely to independently moderate the response to that specific shock.

7.2 Discriminating Between Sticky Information and Rational Inattention

The results above are consistent with a broad class of information-friction models. We report one test that can begin to discriminate between the two leading families.

Under sticky information, the updating probability λ is a structural parameter independent of the shock. The attenuation ratio should be constant across shock magnitudes. Under rational inattention, agents allocate more attention to larger shocks, so the attenuation should be stronger for large shocks. We test this by adding a triple interaction: $\text{shock}_t \times \text{news}_{i,t} \times \mathbf{1}[|\text{shock}_t| > \text{median}]$. The coefficient γ on this term is insignificant (average $t = 1.18$ over $h = 20$–70), indicating that the attenuation does not vary with shock magnitude. This is consistent with a structural λ parameter as in Mankiw and Reis (2002), though we note that the test has limited power given the small number of large shocks.

A back-of-envelope calculation maps the interaction coefficient to the model's structural parameter: a one-standard-deviation increase in news coverage is associated with a $\Delta\lambda \approx 0.19$ increase in the effective updating probability, implying that the median province has roughly half its households operating on recently updated plans.

7.3 Food Amplification: Volume Channel and Endogeneity

The food amplification result is stronger when driven by coverage volume ($t = -2.74$) than by the composite index ($t = -2.32$), mirroring the aggregate pattern: it is the *amount* of coverage, not its tone, that shapes the food response. Moreover, the food amplification survives when using 7-day lagged news as the moderator ($t = -1.56$, same sign and peak timing as the contemporaneous specification), reducing concern about same-day endogeneity.

8 Conclusion

Coibion et al. (2022) showed that when you tell people about monetary policy, they change what they buy. The open question was whether naturally occurring media environments

generate the same effect. Using daily consumption data from 52 provinces within a single currency union, paired with geocoded news coverage of monetary policy, we provide the first observational evidence that provinces with richer media coverage respond to the same interest rate shock as if their households are better informed.

The persistent-versus-spike decomposition is the paper's sharpest result. The province's structural media environment—its coverage level in the week, or even the month, before the announcement—drives the entire attenuation. The announcement-day coverage surge contributes nothing. Coverage volume, not sentiment, is the stronger channel. Media density—outlets per capita—outperforms income and education as a predictor of differential transmission. At the category level, all eleven non-food categories attenuate while food alone amplifies, consistent with information resolving uncertainty about shock persistence.

Whether these results generalise beyond the Spanish setting—and whether they survive direct evidence on media *consumption* rather than media *coverage*—remains to be established. What the evidence strongly suggests is that the same monetary shock, hitting the same economy on the same day, produces detectably different consumption responses depending on the local media environment. It is the media environment provinces already inhabit, not the coverage that arrives on announcement day, that appears to shape how they respond.

We do not have an instrument for media coverage. The identification instead rests on a mosaic of evidence, each piece cutting a different confound. The persistent-versus-spike decomposition shows that the pre-existing media environment—not the announcement-day reaction—drives the result, addressing reverse causality. The 60-day lag placebo shows the effect holds when coverage is measured two months before the shock, making time-varying economic conditions an unlikely driver. The kitchen-sink specification shows the news interaction survives conditioning on province income, education, urban status, economic size, and media density simultaneously. Media density—a time-invariant, supply-side measure—is the single strongest standalone moderator, ruling out the hypothesis that day-to-day coverage variation proxies for transient economic conditions. The effect is present in non-urban provinces, ruling out the hypothesis that it reflects urban financial sophistication. Coverage volume dominates sentiment, which is harder to explain through a financial-literacy channel (where the tone of news should matter) and more naturally consistent with an attention channel (where any signal suffices). And the attenuation does not vary with shock magnitude, consistent with a structural rather than endogenous information channel. No single test is decisive, but collectively they narrow the space of plausible confounders substantially.

One gap remains: we observe media *coverage* (supply) but not media *consumption* (demand). We are currently assembling province-level data on media readership, internet pen-

etration, and survey-based news habits to close this link directly. Additional limitations include the COICOP classification, which is coarse—finer disaggregation within durables would sharpen the transitory-shock mechanism test—and the restriction to card spending, which misses cash purchases and financial decisions such as mortgage refinancing.

These limitations point toward a natural research agenda. Provincial-level expectations data—from the CIS barometers or the ECB Consumer Expectations Survey—would test directly whether media coverage helps households form more accurate beliefs about the persistence of rate changes. Finer category data within durables would reveal whether attenuation concentrates in the most deferrable subcategories. Data on provincial mortgage structure (variable vs. fixed rate) would illuminate the interaction between cash-flow and information channels.

References

Barnichon, R. and Brownlees, C. (2019). Impulse response estimation by smooth local projections. *Review of Economics and Statistics*, 101(3):522–530.

Berger, H., Ehrmann, M., and Fratzscher, M. (2011). Monetary policy in the media. *Journal of Money, Credit and Banking*, 43(4):689–709.

Buda, G., Carvalho, V. M., Corsetti, G., Duarte, J. B., Hansen, S., Moura, A. S., Ortiz, Á., Rodrigo, T., Rodríguez Mora, J. V., and Alves da Silva, G. (2023). The short lags of monetary policy. *Working Paper, R&R Journal of Political Economy*.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*, 46(4):910–945.

Bybee, L., Kelly, B., Manela, A., and Xiu, D. (2024). Business news and business cycles. *Journal of Finance*, 79(5):3419–3472.

Chahrour, R., Nimark, K., and Pitschner, S. (2021). Sectoral media focus and aggregate fluctuations. *American Economic Review*, 111(12):3872–3922.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, 130(6):1537–1584.

Ferrara, L., Mogliani, M., Sahuc, J.-G., and Paredes, J. (2025). High-frequency monitoring of the euro area economy. *Working paper, Banque de France/BBVA*.

Hauptmeier, S., Holm-Hadulla, F., and Nikalexi, K. (2020). Monetary policy and regional inequality. *ECB Working Paper Series*, No. 2440.

Herreño, J. and Pedemonte, M. (2022). The geographic effects of monetary policy. *Federal Reserve Bank of Cleveland Working Paper*, No. 22-15.

Lamla, M. J. and Vinogradov, D. V. (2019). Central bank announcements: Big news for little people? *Journal of Monetary Economics*, 108:21–38.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, 117(4):1295–1328.

Gürkaynak, R. S., Sack, B., and Swanson, E. T. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1):55–93.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544.

Leetaru, K. and Schrodt, P. A. (2013). GDELT: Global data on events, location, and tone, 1979–2012. *ISA Annual Convention*, 3(4):1–49.

Ozdagli, A. and Weber, M. (2023). Monetary policy through production networks: Evidence from the stock market. *Review of Financial Studies*, 36(3):1105–1155.

Pizzuto, P. (2020). The heterogeneous effects of US monetary policy on regional employment. *Working Paper*.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier.

Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics*, 53(8):1761–1800.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *Economic Journal*, 128(610):917–948.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

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

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

Mankiw, N. G., Reis, R., and Wolfers, J. (2004). Disagreement about Inflation Expectations. *NBER Macroeconomics Annual*, 18:209–248.

A Summary Statistics

B Smoothing Parameter Selection

The smoothing parameter λ in the penalised local projection controls the trade-off between fit and smoothness. We select λ by five-fold cross-validation on the aggregate (province-averaged) consumption series, which avoids over-fitting to province-level noise.

The procedure is as follows. We divide the sample of non-COVID dates into five temporal folds. For each candidate $\lambda \in \{0.001, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, \ldots, 1.0, 1.5, 2.0\}$, we estimate the aggregate smooth LP on four folds and compute the mean squared prediction error on the held-out fold. The cross-validation curve is U-shaped, with a minimum at $\lambda^* = 0.2$.

At $\lambda = 0.2$, the IRF is smooth enough to suppress day-to-day noise but flexible enough to capture the hump-shaped dynamics. The main results are robust to $\lambda \in [0.1, 0.5]$; at $\lambda < 0.05$, the IRF becomes noisy, and at $\lambda > 1.0$, the hump shape is over-smoothed toward a monotone decline.

C Province Contiguity Matrix

The queen-contiguity matrix $\mathbf{W}$ is constructed from official INE NUTS3 boundary shapefiles using the following procedure:

We first download province polygons via the `mapSpain` R package (`esp_get_prov()`), which provides official INE boundaries for all 52 Spanish provinces. We then compute queen contiguity using `spdep::poly2nb()` with `queen = TRUE`: two provinces are neighbours if their boundaries share at least one point (vertex or edge). Finally, we row-normalise using `spdep::nb2mat()` with `style = "W"`, so that each row sums to one and W_{ij} represents province j's weight in province i's neighbourhood average.

Five provinces have no land neighbours and receive zero rows in $\mathbf{W}$: Baleares (ES-PM), Las Palmas (ES-GC), Tenerife (ES-TF), Ceuta (ES-CE), and Melilla (ES-ML). These provinces are excluded from all spatial specifications. The resulting 47×47 submatrix has zero diagonal and each row sums to one. Note that while the binary queen contiguity relation is symmetric, row-normalisation makes $\mathbf{W}$ asymmetric: $W_{ij} = 1/d_i \neq 1/d_j = W_{ji}$ whenever provinces i and j have different numbers of neighbours.

As a robustness check, we have verified that: (i) the matrix contains 222 non-zero entries (111 unique neighbour pairs); (ii) the average province has 4.7 neighbours; and (iii) the results are qualitatively similar under rook contiguity (shared edges only, excluding shared vertices).

D Inference Details

Cluster-by-date standard errors. Because all provinces are exposed to the same monetary policy shock on any given date, residuals are cross-sectionally correlated. We cluster standard errors by date, which is equivalent to treating each calendar day as a single cluster containing all 52 provinces.

In the smooth LP framework, the coefficients are projections of the stacked coefficient vector onto a B-spline basis. The cluster-robust variance is computed on the stacked system and then mapped to horizon-specific standard errors using the spline basis matrix. Specifically, if $\hat{\boldsymbol{\theta}}$ is the vector of spline coefficients and $\mathbf{B}$ is the $91 \times K$ spline basis matrix, then $\hat{\beta}^h = \mathbf{B}[h, \cdot]\hat{\boldsymbol{\theta}}$ and $\text{Var}(\hat{\beta}^h) = \mathbf{B}[h, \cdot]\hat{V}_\theta \mathbf{B}[h, \cdot]'$.

Standard errors at arbitrary moderator values. For the interaction specification, the IRF at moderator value U is $\text{IRF}^h(U) = \hat{\beta}^h + \hat{\delta}^h(U - \bar{z})$. The variance at U is:

$$\text{Var}(\text{IRF}^h(U)) = \text{Var}(\hat{\beta}^h) + (U - \bar{z})^2 \text{Var}(\hat{\delta}^h) + 2(U - \bar{z})\text{Cov}(\hat{\beta}^h, \hat{\delta}^h)$$

This is implemented via a Jacobian matrix $\mathbf{J}(U)$ that maps from the full coefficient vector $(\hat{\boldsymbol{\theta}}_\beta, \hat{\boldsymbol{\theta}}_\delta)$ to the IRF at U:

$$\mathbf{J}(U) = [\mathbf{B} \;\; (U - \bar{z})\mathbf{B}]$$

and the variance is $\text{Var}(\text{IRF}(U)) = \mathbf{J}(U)\hat{V}_\theta \mathbf{J}(U)'$.

E Robustness: Kitchen-Sink Specification

Table 7 reports the news interaction coefficient from specifications that progressively add controls for other province-by-shock interactions:

The news interaction is stable across specifications, ranging from 38.5 to 41.7 ($\times 10^4$). Adding province-by-shock controls for urbanisation, media density, and economic size does not absorb the news channel; if anything, the coefficient increases slightly, confirming that the news effect is not a proxy for observable province characteristics. When coverage intensity |news| is used instead of the composite, the coefficient is larger and more significant (not shown), consistent with the decomposition results in Section 5.2.

F Additional Category-Level Results

Table 4 in the main text reports the full set of category-level results for all 12 COICOP categories. The key patterns are: (i) food is the only amplifying category ($t = -2.23$); (ii) durable and discretionary categories (clothing, transport, restaurants, furnishings) show the largest attenuation; (iii) categories with small baseline responses (housing, communication) show positive but imprecisely estimated interactions. The ordering is consistent with the deferrability mechanism: categories where purchase timing is most discretionary show the strongest attenuation.

G Split-Sample Urban/Non-Urban Results

We estimate separate impulse responses for the 11 urban and 41 non-urban provinces. The urban subsample reaches its trough at -0.0238 on day 48 and recovers to -0.0117 by day 90. The non-urban subsample troughs at -0.0212 on day 53 and recovers to -0.0149 by day 90. The urban response is slightly deeper but less persistent; urban confidence bands are wider due to the smaller subsample (11 provinces).

When the news interaction is estimated separately, the attenuation is larger in non-urban provinces (interaction $+0.0045$, $t = 1.98$) than in urban ones ($+0.0026$, $t = 1.44$).

The point estimate of the news interaction is approximately 74% larger in the non-urban subsample, though as the formal test in Appendix K shows, this difference is not statistically significant.

H Robustness of the News Interaction

We conduct three robustness exercises for the news interaction coefficient, each designed to probe a different dimension of fragility.

Leave-one-out by shock date. We re-estimate the baseline news interaction 49 times, each time zeroing out the monetary policy shock on one of the 49 non-zero shock dates in the non-COVID estimation sample.[1] Figure 11 plots the resulting peak interaction t-statistic for each iteration. Of the 49 iterations, 48 maintain significance at the 90% level ($t > 1.645$), with t-statistics ranging from 1.72 to 2.08. The single exception is the 12 March 2020 ECB emergency meeting, which occurred at the onset of the COVID-19 pandemic. Dropping this

[1] The full sample contains 54 non-zero shock dates, of which 5 fall within the COVID exclusion window. The LOO exercise iterates over the 49 dates used in estimation.

date reduces the t-statistic to 0.65, reflecting the unusually large shock (+19 basis points) hitting the economy just as the lockdown period begins. Excluding this outlier, the news interaction is stable across all remaining shock dates, confirming that the result is not driven by any single ECB decision.

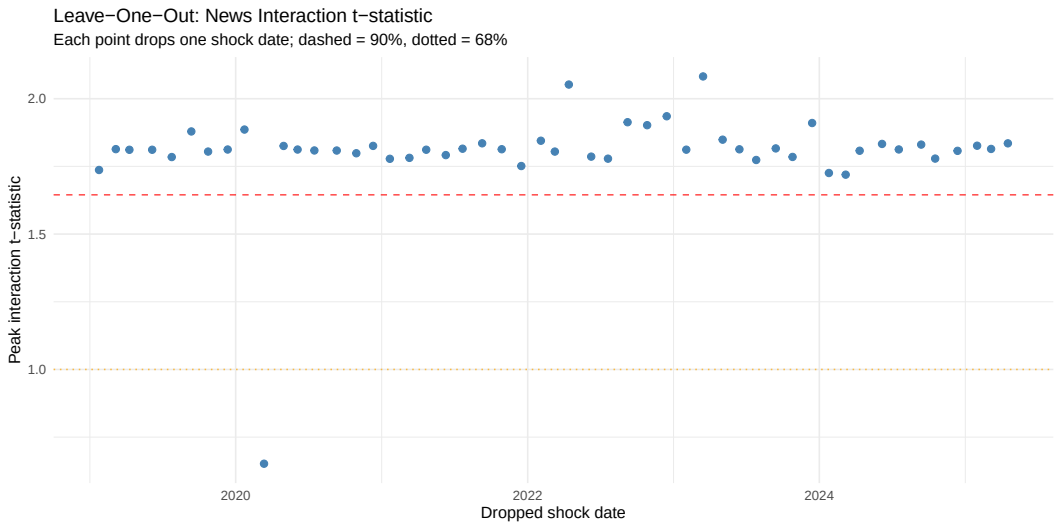


Figure 11: Leave-One-Out: News Interaction t-Statistic by Dropped Shock Date

Notes: Each point shows the peak news interaction t-statistic when one shock date is dropped from the estimation. Dashed red line: 90% significance threshold ($t = 1.645$); dotted orange line: 68% threshold ($t = 1.00$). The outlier near $t = 0.65$ corresponds to the 12 March 2020 ECB emergency meeting.

Smoothing parameter sensitivity. Table 8 reports the peak news interaction t-statistic across a grid of 12 smoothing parameter values $\lambda \in [0.01, 2.0]$. The result is significant at the 90% level for every value of λ, with t-statistics ranging from 1.74 ($\lambda = 0.5$) to 2.08 ($\lambda = 0.01$). The cross-validated optimum $\lambda^* = 0.2$ yields $t = 1.81$, sitting comfortably within this range. The peak interaction day shifts from $h = 39$ at low λ (noisier IRFs) to $h = 71$ at high λ (over-smoothed), while the coefficient magnitude declines monotonically from 4.52×10^{-3} to 2.71×10^{-3}. These patterns are mechanical: lower smoothing allows the IRF to spike early, while higher smoothing spreads the effect over more horizons.

Placebo timing test. We shift all monetary policy shocks forward by 30, 60, 90, and 120 days and re-estimate the news interaction using these placebo shocks. If the news interaction captures a genuine moderating effect tied to the actual timing of ECB decisions, it should vanish when shocks are misaligned with announcement dates. Table 9 reports the results: none of the four placebo shifts produces a significant interaction. The t-statistics are 1.49 (+30 days), 1.56 (+60 days), −0.58 (+90 days), and −0.61 (+120 days). The non-negligible

estimates at shorter shifts (30–60 days) likely reflect autocorrelation in the news index and slow-moving province-level trends; the sharp attenuation at 90 and 120 days, where the estimates are both small and of the opposite sign, provides more reassuring evidence that the effect is tied to actual shock timing rather than spurious province-level trends.

I Persistent Media Environment vs. Shock-Responsive Coverage: Full Decomposition

Section 5.2.2 reports the main decomposition using coverage volume. Here we present the full set of results, including both the signed composite index and the coverage volume decomposition, and retain the lagged-news specifications as supporting evidence.

Decomposition design. We split the news index into a *persistent* component (the seven-day pre-shock moving average $\overline{\text{news}}_{i,t-7:t-1}$) and a *spike* component (the within-day deviation $\text{news}_{i,t} - \overline{\text{news}}_{i,t-7:t-1}$). The persistent component captures the structural level of local interest-rate coverage—a feature of the provincial media ecosystem that is predetermined relative to the ECB announcement. The spike captures the surge in coverage triggered by the announcement itself. We estimate both standalone specifications (one moderator at a time) and horse-race specifications (both moderators simultaneously).

Table 10 reports the results. Panel A (coverage volume) delivers the sharpest contrast: the persistent component achieves $t = 2.16$ standalone and $t = 2.14$ in the horse race, while the spike is essentially zero ($t = 0.28$ standalone, $t = 0.54$ conditional). Panel B (signed composite) tells a consistent but less dramatic story: the persistent component dominates ($t = 2.00$ standalone, $t = 2.27$ in the horse race), while the spike is weaker and insignificant ($t = 1.42$ standalone, $t = 1.58$ conditional).

The volume decomposition is the more informative test because it isolates the *quantity* of coverage—how many articles appear—from directional content. A province that sustains high coverage volume in the week before the ECB acts has a structurally active media ecosystem for interest-rate news. The fact that this persistent volume does all the work, while the announcement-day surge contributes nothing, rules out the reverse-causality channel that would plague the contemporaneous specification.

Supporting evidence: lagged specifications. Table 11 reports a complementary test using the one-day lag and seven-day moving average of the composite index as moderators in place of the contemporaneous measure.

Both predetermined measures produce larger t-statistics than the contemporaneous specification ($t = 1.86$ and $t = 2.00$ vs. $t = 1.79$). The coefficient magnitudes are somewhat smaller—the contemporaneous measure captures both the pre-existing media environment and the same-day coverage boost—but the standard errors fall faster, consistent with the predetermined measures removing noise from the contemporaneous coverage response. These lagged results are a special case of the broader decomposition: the seven-day moving average is precisely the persistent component used in Table 10, and its success foreshadowed the decisive attribution that the full decomposition delivers.

J Multiple Testing Correction for Category-Level Results

Table 4 in the main text reports news interaction coefficients for 12 COICOP categories. With multiple hypotheses tested simultaneously, there is a concern that some individually significant results may reflect chance variation rather than genuine heterogeneity. We address this in two ways.

First, Table 12 reports Benjamini–Hochberg (BH) adjusted q-values, which control the false discovery rate (FDR) at the stated level. Six of twelve categories remain significant at $q < 0.10$ after adjustment: education ($q = 0.077$), housing ($q = 0.077$), health ($q = 0.077$), food ($q = 0.077$), transport ($q = 0.091$), and clothing ($q = 0.098$). Critically, food—the sole amplifying category and the centrepiece of the reallocation interpretation—survives the BH correction with $q = 0.077$. At a more permissive threshold of $q < 0.20$, eleven of twelve categories are significant.

Second, we test whether the *pattern* of signs across categories is consistent with chance. Under the null hypothesis that news has no systematic effect on category-level monetary transmission, each category's interaction coefficient would be equally likely to be positive or negative. We observe 11 positive (attenuating) and 1 negative (amplifying) out of 12 categories. A one-sided binomial test against $H_0 : \Pr(\text{positive}) = 0.5$ yields $p = 0.003$. The probability of observing 11 or more categories with the same sign under random noise is vanishingly small. This confirms that the category-level pattern—near-universal attenuation with food as the sole exception—reflects a systematic phenomenon rather than a statistical artefact.

K Formal Test of Urban vs. Non-Urban Difference

Section 5.3 reports that the news attenuation effect is larger in non-urban provinces (+0.0045, $t = 1.98$) than in urban provinces (+0.0026, $t = 1.44$). However, a difference in significance is not the same as a significant difference. We test $H_0 : \delta_{\text{non-urban}} = \delta_{\text{urban}}$ directly.

Using the split-sample estimates and their standard errors (which are independent across subsamples), we construct:

$$t_{\text{diff}} = \frac{\hat{\delta}_{\text{non-urban}} - \hat{\delta}_{\text{urban}}}{\sqrt{\text{SE}^2_{\text{non-urban}} + \text{SE}^2_{\text{urban}}}} = \frac{0.0045 - 0.0026}{\sqrt{0.0023^2 + 0.0018^2}} = \frac{0.0019}{0.0029} = 0.66$$

The resulting p-value is 0.51—we cannot reject the null that the news effect is equal across urban and non-urban provinces. The point estimate of the difference is positive (news matters more in non-urban provinces), but the samples are too small to distinguish this from zero at conventional levels.

We therefore adopt a more cautious interpretation in the main text: the news attenuation effect is present in both subsamples and is not driven exclusively by urban provinces. The point estimate is larger in the non-urban subsample—consistent with the marginal value of information being higher where it is scarcer—but we do not claim a statistically significant difference between the two groups.

Table 4: News Interaction by Consumption Category (COICOP)

Category	Type	Trough	Peak $\hat{\delta}^h$ ($\times 10^4$)	Effect
Panel A: Attenuating categories				
Clothing	Semi-durable	-0.068	151.1^{**}	Attenuates
			(76.7)	
Transport	Semi-durable	-0.065	135.4^{**}	Attenuates
			(65.1)	
Education	Non-durable	-0.042	132.7^{**}	Attenuates
			(48.7)	
Restaurants	Discretionary	-0.050	115.1^{**}	Attenuates
			(61.2)	
Health	Non-durable	-0.025	108.3^{**}	Attenuates
			(47.5)	
Furnishings	Durable	-0.034	93.6^{**}	Attenuates
			(55.2)	
Recreation	Discretionary	-0.039	74.4^{*}	Attenuates
			(47.1)	
Housing	Non-durable	≈ 0	53.5^{**}	Attenuates
			(22.2)	
Miscellaneous	Semi-durable	-0.027	44.9^{*}	Attenuates
			(28.1)	
Communication	Semi-durable	-0.005	26.9^{*}	Attenuates
			(19.2)	
Alcohol	Non-durable	-0.018	23.0^{*}	Attenuates
			(19.0)	
Panel B: Amplifying categories				
Food	Non-durable	≈ 0	-66.9^{**}	Amplifies
			(30.0)	

Notes: Each row reports a separate category-level local projection with the GDELT news index as interaction moderator. "Trough" is the peak negative IRF response. Positive interaction coefficients indicate attenuation; negative indicate amplification. Standard errors (in parentheses) clustered by date. *, **: 68% and 90% confidence intervals exclude zero.

Table 6: Summary Statistics

Variable	N	Mean	Std. dev.	Min	Median	Max
Panel A: Province-day variables						
$y_{i,t}$ (YoY consumption growth)	128,960	0.020	0.052	−0.129	0.011	0.432
MP shock ε_t^{mp} (all days)	128,960	0.000	0.006			
MP shock ε_t^{mp} (non-zero)	2,548					
s.d. of non-zero shocks			0.039			
News index (composite)†	140,036	−0.014	1.021	−11.671	0.020	10.492
|News| (coverage intensity)†	140,036	0.669	0.772	0.000		
Panel B: Province characteristics (cross-section)						
Urban (binary)	52	0.21	0.41	0	0	1
Media density (outlets/cap.)	52	4.79	1.57	0.40		9.00
Net income per capita (€)	52	13,465	1,719	10,605		18,325
Tertiary education (%)	52	35.7	5.0	25.2		48.1
GVA (€ millions, 2022)	52	24,106	39,830	1,608		243,266
Provinces	52					
Urban	11					
Non-urban	41					

Notes: Panel spans September 2018–February 2026, excluding the COVID-19 lockdown period (14 March–30 October 2020). Non-COVID sample: 128,960 province-day observations. †News variables are reported for the full sample period (140,036 observations); all other Panel A variables use the non-COVID estimation sample. Monetary policy shocks are high-frequency surprises around ECB announcements; 49 non-zero shock dates. Certain min/max values omitted for confidentiality.

Table 7: News Interaction: Robustness to Province-Level Controls

	Peak news $\hat{\delta}^h$	
Specification	$\times 10^4$	t-stat
(1) News only	38.5**	1.81
	(21.2)	
(2) + shock × urban	37.7**	1.82
	(20.8)	
(3) + shock × urban × media	39.1**	1.77
	(22.1)	
(4) + shock × urban × media × GVA	41.7**	1.76
	(23.7)	

Notes: Each row reports the peak news interaction coefficient ($\hat{\delta}^h$ at $h = 43$) from the smooth panel LP with progressively more province×shock controls. All specifications include province fixed effects, lagged dependent variables (y_{t-1}, y_{t-7}, y_{t-14}), and cluster-by-date standard errors (in parentheses). *, **: 68% and 90% confidence intervals exclude zero.

Table 8: Smoothing Parameter Sensitivity: News Interaction

λ	Peak $\hat{\delta}^h \times 10^3$	SE $\times 10^3$	Peak day	t-stat
0.01	4.52	2.17	39	2.08
0.02	4.38	2.14	40	2.05
0.05	4.24	2.13	41	1.99
0.10	4.09	2.14	42	1.91
0.15	3.97	2.15	42	1.85
0.20†	**3.85**	**2.12**	**43**	**1.81**
0.30	3.63	2.06	44	1.76
0.50	3.30	1.90	46	1.74
0.75	3.04	1.71	49	1.78
1.00	2.88	1.57	52	1.83
1.50	2.73	1.39	60	1.96
2.00	2.71	1.33	71	2.03

Notes: Each row reports the peak news interaction coefficient, its cluster-by-date standard error, the horizon at which the peak occurs, and the associated t-statistic, estimated using the indicated value of λ. †Cross-validated optimum. All t-statistics exceed 1.645 (90% significance). Coefficients are scaled by the shock standard deviation (3.97 basis points).

Table 9: Placebo Timing Test: Shifted Monetary Policy Shocks

Shift (days)	Peak $\hat{\delta}^h \times 10^3$	SE $\times 10^3$	Peak day	t-stat
+30	1.23	0.82	62	1.49
+60	3.86	2.48	90	1.56
+90	−0.61	1.04	49	−0.58
+120	−1.12	1.83	90	−0.61

Notes: Each row shifts all ECB monetary policy shocks forward by the indicated number of days and re-estimates the news interaction specification. The peak interaction coefficient (in absolute value), its standard error, the horizon, and the t-statistic are reported. None of the placebo shifts achieves significance at the 90% level. Coefficients are scaled by the shock standard deviation (3.97 basis points).

Table 10: Persistent vs. Spike Decomposition: Full Results

Specification	Peak $\hat{\delta}^h \times 10^3$	SE $\times 10^3$	Peak day	t-stat
Panel A: Coverage volume				
Persistent (standalone)	7.43	3.45	48	**2.16**
Spike (standalone)	0.39	1.40	90	0.28
Persistent | spike	7.22	3.37	48	**2.14**
Spike | persistent	0.68	1.25	90	0.54
Panel B: Signed composite				
Persistent (standalone)	2.83	1.41	46	**2.00**
Spike (standalone)	3.21	2.27	39	1.42
Persistent | spike	3.04	1.34	46	**2.27**
Spike | persistent	3.38	2.14	39	1.58

Notes: Each row reports the peak interaction coefficient from a smooth panel LP. "Persistent" = seven-day pre-shock average; "Spike" = within-day deviation from that average. Panel A uses coverage volume |news|; Panel B uses the signed composite index. "| spike" and "| persistent" denote conditional (horse-race) specifications including both components simultaneously. All moderators standardised to z-scores. Cluster-by-date SEs. Coefficients scaled by shock s.d. (3.97 bp).

Table 11: Lagged Moderator Specifications

News measure	Peak $\hat{\delta}^h \times 10^3$	SE $\times 10^3$	Peak day	t-stat
Contemporaneous $\text{news}_{i,t}$	3.84	2.14	43	1.79
Lagged $\text{news}_{i,t-1}$	3.26	1.75	44	1.86
7-day pre-shock MA	2.84	1.41	46	**2.00**

Notes: Each row reports the peak news interaction coefficient from a separate smooth panel LP where the moderator is the indicated news measure. Lagged news $(t-1)$ captures the local media environment on the day before the ECB announcement. The 7-day pre-shock moving average captures the average local interest-rate coverage in the week preceding the shock. All moderators are standardised to z-scores. Cluster-by-date SEs. Coefficients scaled by shock s.d. (3.97 bp).

Table 12: Multiple Testing Correction: COICOP Category Interactions

Category	$\hat{\delta}^h \times 10^3$	t-stat	p-value	BH q-value	Bonf. p
Education	13.27	2.73	0.006	0.077	0.077
Housing	5.35	2.41	0.016	0.077	0.190
Health	10.83	2.28	0.023	0.077	0.271
Food	−6.69	−2.23	0.026	0.077	0.310
Transport	13.54	2.08	0.038	0.091	0.453
Clothing	15.11	1.97	0.049	0.098	0.587
Restaurants	11.51	1.88	0.060	0.103	0.721
Furnishings	9.36	1.70	0.090	0.135	1.000
Miscellaneous	4.49	1.59	0.111	0.137	1.000
Recreation	7.44	1.58	0.114	0.137	1.000
Communication	2.69	1.40	0.163	0.178	1.000
Alcohol	2.30	1.21	0.226	0.226	1.000

Notes: Each row reports the peak news interaction coefficient from a separate category-level smooth panel LP. p-values are two-sided, computed from the t-statistic. BH q-values control the false discovery rate using the Benjamini–Hochberg procedure. Bonferroni p-values are conservative (multiply raw p by 12, cap at 1). Coefficients scaled by shock s.d.