Bank of England

Central bank communications that reach the public

Staff Working Paper No. 1,194
July 2026

Eric Tong and Rennae Cherry

Bank of England

Staff Working Paper No. 1,194

Central bank communications that reach the public

Eric Tong[1] and Rennae Cherry[2]

Abstract

We identify central bank communication shocks designed to measure how policy messages reach households. The shocks are constructed from central bank text, where policy messages originate, and newspaper narratives, through which households encounter them. We compare central bank narratives with pre-announcement newspaper narratives and decompose the resulting narrative surprises into stance and information communication shocks. Applying the framework to the Bank of Canada, the Bank of England, and the Federal Reserve, we find that central bank narratives shape media coverage and move households' one-year-ahead inflation expectations. Tighter stance communication shocks lower inflation expectations, while expansionary information communication shocks raise them, especially when households' attention is high. Conventional shocks identified with high-frequency asset-price moves do not deliver these responses, underscoring the importance of measuring central bank communication as households experience it. Taken together, the results qualify the view that central bank communication rarely reaches the public, but also show that its effects depend on how central bank messages are received and perceived.

Key words: Central bank communication, event-study, textual analysis, households' inflation expectations.

JEL classification: E31, E52, E58.

---

(1) Bank of England. Email: eric.tong@bankofengland.co.uk
(2) Reserve Bank of New Zealand. Email: rennae.cherry@rbnz.govt.nz

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England, the Reserve Bank of New Zealand, or any of their committees. We are grateful to Marcus Buckmann, David Elliott, Andrew Gimber, Simon Lloyd, Ricardo Reis, and Leif Thorsrud for comments and suggestions.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1. Introduction

> 'After all, it is the everyday economic decisions of people and companies that we seek to influence with our policy and communication. If our language is not accessible, our policy will be less effective.'
>
> — Christine Lagarde

Can central bank communication reach the general public? A recent survey by Blinder et al. (2024) observes that almost no central bank communication reaches the public. Yet much of the evidence behind this view rests on monetary policy shocks identified from asset prices. Asset prices reflect the beliefs of financial market participants, but not necessarily those of households, who acquire information and form expectations differently (Weber et al., 2022; Coibion and Gorodnichenko, 2026). This paper therefore identifies communication shocks in a way that mirrors how households encounter central bank messages through media intermediation (Lamla and Maag, 2012; Ter Ellen et al., 2021).

Our identification framework builds on the asset-price event studies of Kuttner (2001); Gurkaynak et al. (2005); Gertler and Karadi (2015), but replaces prices with text as the building block of identification. Instead of using asset prices that reflect the underlying central bank messages, we use the central bank texts from which the messages originate. And instead of using pre-announcement asset prices to proxy for expected policy decisions, we use pre-announcement newspaper narratives to capture public policy expectations.[1] We convert both central bank and newspaper text into quantitative measures of policy tone using the dictionary approach of Apel et al. (2022). We then construct a narrative surprise as the gap between the central bank's narrative and the expected policy narrative in pre-announcement newspaper coverage, which is the textual analogue of the surprise in an asset-price event study.[2] We turn these numerical narrative surprises into structural shocks by applying the sign rotation of Jarociński and Karadi (2020) and Jarociński (2022). This decomposition yields two orthogonal components: a *stance communication shock*, which captures unexpected communication about the policy stance, and an *information communication shock*, which captures unexpected communication about the central bank's assessment of economic conditions. Both steps of identification—using and decomposing narrative surprises—are necessary to deliver impulse responses that are statistically significant and consistent with theory. Put differently, we show that neither traditional asset price shocks nor unrotated narrative surprises produce intuitive impulse responses.

We apply the identification framework to monetary policy announcements by the Bank of Canada, the Bank of England, and the Federal Reserve System between January 1995 and December 2024.[3] For each country, we estimate the effects of the communication shocks on households' one-year-ahead inflation expectations using smooth local projections (SLP), which improve estimation accuracy relative to standard local projections (Barnichon and Brownlees, 2019). To separate communication from policy actions—defined as changes in interest rates across the maturity spectrum induced by policy rate decisions—we control for the policy rate, including shadow rates around the effective lower bound, as well as Bu et al. (2021)'s measure of unexpected yield curve movements induced by policy announcements.[4] We also compute forecast error variance decompositions (FEVD) to gauge the contribution of communication shocks to variation in inflation expectations (Gorodnichenko and Lee, 2020).

[1] Our event-study window is longer by design. We measure newspaper narratives over the three days before each central bank announcement, rather than over the few minutes typically used in high-frequency financial market studies. This choice reflects the slower and more intermediated way in which monetary policy information reaches households (Coibion and Gorodnichenko, 2026).

[2] For each country, we draw on a broad set of national newspapers as a measure of the information set available to households (Koop and di Vettimo, 2023).

[3] The sample starts in January 1997 for Canada and January 2000 for the United Kingdom, reflecting the availability of inflation expectations and control variables.

[4] Brennan et al. (2024) evaluate frontier monetary policy shock measures and find that Bu et al. (2021) holds up well relative to alternatives. Other aspects of the empirical specification follow Jordà and Taylor (2025) and Ramey (2016).

A linchpin of effective communication is households' attention to central bank messages, but the theoretical effect of attention is ambiguous (Coibion and Gorodnichenko, 2026). On the one hand, greater attention may make households more likely to notice central bank messages and incorporate them into their expectations. On the other hand, attention may rise just when inflation is high and monetary policy appears to be falling short, weakening the credibility of central bank communication. To examine this ambiguity, we augment the baseline SLP by allowing households' attention to interact with the communication shocks. Consistent with our use of newspapers, we follow Fisher et al. (2022) and measure attention from the share of newspaper articles devoted to monetary policy and inflation.[5]

Our empirical results are threefold. First, central bank narratives shape media narratives. Regression analysis shows that roughly 30–60 per cent of the central bank narrative surprise passes through to subsequent media coverage. Central bank messages enter public discourse as mediated by the press.

Second, when measured properly, central bank communication reaches households. A tightening stance communication shock, normalised to a one per cent fall in the national stock price index, lowers households' annualised one-year-ahead inflation expectations by around 0.1 percentage point—about half the average absolute monthly change in household expectations across the three countries. Rather than displaying the inertia documented in Blinder et al. (2024), households respond almost immediately, and the response persists for 5–12 months before fading out. We also find significant effects of information communication shocks in the United Kingdom. A unitary information communication shock, normalised to a one per cent rise in the FTSE stock price index, raises households' inflation expectations persistently. The effect peaks at 0.1 percentage point 14 months after the shock and fades out only after 19 months. Communication shocks not only affect inflation expectations, but also account for a significant share of their variation. In the United Kingdom and the United States, FEVD results show that communication shocks explain about ten per cent of the unexplained variation in households' inflation expectations at the 12-month horizon. This magnitude is in line with estimates commonly associated with conventional monetary policy shocks.

Third, attention matters. In the United Kingdom and the United States, high attention amplifies the effect of stance communication shocks and brings forward their impact. Information communication shocks also increase inflation expectations when attention is high in the United Kingdom and Canada—the latter effect is absent in the baseline specification that does not allow for state dependence. That said, not all responses line up with theory. In Canada, stance communication shocks are effective when attention is low, while in the United States, attention does not appear to shape the transmission of information communication shocks.

This paper makes two contributions. First, building on Ter Ellen et al. (2021), we show that narrative surprises must be separated before they can be interpreted. A surprise in central bank narrative may reflect news about policy stance or about the central bank's view of the economy, and treating the two as one obscures the effect of communication. Once orthogonalised, the communication shocks are distinct from conventional monetary policy shocks, both in their correlations and in the impulse responses they produce. This distinction matters for households: neither conventional shocks nor unrotated narrative surprises generate responses of inflation expectations that are consistent with theory. The decomposition therefore allows us to trace the dynamic effects of central bank communication on household inflation expectations in a macroeconometric setting. In doing so, it qualifies the view that central bank communication has limited direct reach among households, and complements experimental evidence based on controlled information treatments (D'Acunto et al., 2024; Coibion et al., 2022).

Second, we show that communication can serve as an additional policy tool alongside policy rates, but only conditionally. Its effects depend on how the message is perceived and on the information environment in which it is received. The contrast between stance and information communication shocks shows that a

[5] We estimate both a continuous-interaction SLP and a binary regime-switching SLP to assess whether the results hold up across alternative ways of modelling attention.

message read as news about policy stance may move expectations in different directions from one read as news about the central bank's economic outlook. Transmission also depends on whether households are attentive to monetary policy and inflation news when the message arrives. For policymakers seeking to manage inflation expectations, the implication is that communication strategy must account for not only what central banks say, but also how their messages are perceived and received.

We subject our baseline results to a battery of robustness and validation exercises, which fall into three categories. First, we examine alternative constructions of the communication shocks. These exercises include varying the media window; using alternative dictionaries; and benchmarking our measures against existing US monetary policy shocks.[6] We also use ChatGPT-5.4 to classify narrative tone, drawing on the contextual understanding of large language models. The resulting classifications line up closely with those from our dictionary approach, which is relatively transparent and replicable. Second, we assess specification robustness by varying the policy-action controls, lag structure, and expectation measures. Third, we test whether the state-dependent results hold up under alternative measures of attention, including a residualised measure of monetary policy attention purged of inflation attention. The main results remain intact.

The remainder of the paper is structured as follows. Section 2 outlines the identification framework. Section 3 describes the corpus and data. Section 4 presents the main findings. Section 5 reports the robustness and validation exercises. Section 6 concludes.

2. Empirical framework

2.1. Identification strategy

We construct communication shocks using two building blocks.

First, we construct a text-based measure of policy stance. Our baseline dictionary draws on Apel et al. (2022), which is referenced in Shapiro and Wilson (2022), Malmendier et al. (2021), and Hubert and Labondance (2021). The method counts instances in which hawkish or dovish modifiers appear near key economic terms related to inflation, economic activity, and employment.[7] We expand on Apel et al. (2022) by incorporating additional key terms identified in Aruoba and Drechsel (2025), who develop a selection procedure that yields 296 economic concepts comprising unigrams, bigrams, and trigrams. From their work, we extract synonyms related to inflation, economic activity, and employment and add them to our set of key terms.

Second, we require a proxy for households' beliefs about monetary policy stance before the central bank announcement. For this purpose, we apply the same dictionary to newspaper narratives prior to the policy events. Newspapers remain relevant in the modern media landscape: printed and online press remain the most important sources of information about the ECB for euro area households (Blinder et al., 2024), and newspapers have long served as watchdogs on central banks (Koop and di Vettimo, 2023). Furthermore, newspaper coverage reflects both the newsworthiness of central bank messages and the editorial team's efforts to meet households' demand for monetary policy news—hence shaping their expectations (Ter Ellen et al., 2021; Fisher et al., 2022).

From the two building blocks, we construct narrative surprises as the difference between central bank tone and newspaper tone. Specifically, let

$$hawkishness_t^{CB} = \frac{hawk_t^{CB} - dove_t^{CB}}{hawk_t^{CB} + dove_t^{CB}} \quad \text{and} \quad hawkishness_t^{News} = \frac{hawk_t^{News} - dove_t^{News}}{hawk_t^{News} + dove_t^{News}}$$

[6] We take the union of predictor sets across multiple Swanson studies to build in a broad set of predictors.

[7] Modifiers are adjectives that signal policy stance. We use a seven-word window to link modifiers to nearby terms, with negation terms within the same window reversing the classification.

be the net hawkishness extracted from central bank communication and newspaper coverage, respectively. We define the narrative surprises as

$$narrative\ surprises_t = \left(hawkishness_t^{CB} - hawkishness_{t:w^-}^{News}\right),$$

where the subscript of $hawkishness^{News}$ denotes hawkishness compiled over a period of w^- working days prior to each announcement day t. Smaller values of w^- are likelier to capture media's short-run focus immediately before the announcement, while larger values reflect a broader focus over that period. However, larger windows also incorporate information further away from the event day t, potentially undermining the event study identification. For these reasons, and in light of evidence that general public's expectations adjust gradually, we adopt $w^- = 3$ as our baseline and conduct robustness checks around this window (Coibion et al., 2018). Although a three-day window is longer than that typically used in the high-frequency identification (HFI) literature, we consider it appropriate given our focus on households, whose expectations may respond only after central bank information has diffused through the media and been incorporated into beliefs (Coibion and Gorodnichenko, 2026). For reference, Ter Ellen et al. (2021) use a ten-day window in constructing their shocks.

Because our aim is to isolate unexpected communication, we further purge the narrative surprises of variation that is predictable from information available before the event. Following Bauer and Swanson (2023a), Bauer and Swanson (2023b), and Swanson (2024), we estimate

$$narrative\ surprise_t = \alpha + \beta' X_{t-} + u_t,$$

where t indexes monetary policy events, X_{t-} denotes pre-event macroeconomic and financial information, and u_t is the residual. To be conservative, X_{t-} is defined as the union of the predictor sets used in these studies, including macroeconomic surprises, expected inflation, recent inflation dynamics, labour-market conditions, stock-market returns, yield-curve movements, commodity prices, government bond yields, and financial conditions. Appendix F lists the predictors in detail. We use the residual u_t as the orthogonalised narrative surprise. This step ensures that the subsequent decomposition is applied to the unexpected component of central bank communication rather than to narrative variation that could have been anticipated from the pre-announcement information set.

Narrative surprises may combine two distinct types of central bank communication. A hawkish narrative may signal a tighter policy stance, but it may also reveal that the central bank has a stronger assessment of economic conditions (Nakamura and Steinsson, 2018). For example, a phrase such as 'expanding economic growth' may point to stronger demand and a more hawkish policy outlook, while also conveying favourable news about the economy.[8]

Asset pricing theory gives these two components opposite implications for equity prices (Jarociński and Karadi, 2020). A tighter policy stance signal should lower equity prices by raising discount rates and reducing expected future dividends. By contrast, favourable central bank information about economic conditions should raise equity prices, as market participants revise up their expectations of future activity and cash flows. We therefore use the sign of the comovement between narrative surprises and stock price surprises to separate the two components. The component of narrative surprise that is negatively correlated with stock price surprises is classified as a *stance communication shock*; the component that is positively correlated with stock price surprises is classified as an *information communication shock*. We refer to the second component as an information *communication* shock, rather than simply an information shock, to distinguish it from the information shocks in Nakamura and Steinsson (2018) and Jarociński and Karadi (2020), which are identified from interest rate movements in response to central bank communication.

[8]The same intuition follows from Bauer and Swanson (2023a) and Bauer and Swanson (2023b): if central banks tend to tighten in response to strong economic data, agents may read hawkish communication as a signal of central bank optimism.

Formally, we implement this decomposition using the sign-rotation approach of Jarociński (2022). Let s_t denote the stock price surprise on the day of the central bank announcement. We decompose narrative surprises into two orthogonal components:

$$\textit{narrative surprises}_t = \textit{stance communication shocks}_t + \textit{information communication shocks}_t,$$

where the decomposition satisfies

$$M = UC, \quad \text{with} \quad U'U = \text{diagonal matrix} \quad \text{and} \quad C = \begin{pmatrix} 1 & c_{\text{Stance}} < 0 \\ 1 & c_{\text{Info}} > 0 \end{pmatrix}, \tag{1}$$

where $M = (\textit{narrative surprises}, s)$ and

$$U = (\textit{stance communication shocks}, \quad \textit{information communication shocks})$$

are both $T \times 2$ matrices, with T denoting the number of central bank announcements. The two columns of U are mutually orthogonal. The first column of C ensures that the two shocks add up to the observed narrative surprise, while the second column captures how each component maps into the stock-price surprise.

The decomposition in Eq. (1) is not unique, because multiple rotation angles can satisfy the sign restrictions. To obtain a point estimate, we follow Jarociński (2022) and use the Median Target rotation, which takes the median admissible rotation angle. At this rotation, the coefficients c_{Stance} and c_{Info} have the same absolute value.

To recap, our identification strategy is as follows:

1. **Identify monetary policy events**
   We include monetary policy announcements for each of the three central banks that we covered—Bank of Canada, Bank of England, and the Federal Reserve.[9] To minimise contamination, we exclude dates on which other major economic news—unrelated to monetary policy—was released.[10]
2. **Extract narrative surprises**
   Using our combined dictionary, we construct $\textit{hawkishness}_t^{CB}$, $\textit{hawkishness}_{t:w^-}^{News}$, and $\textit{narrative surprises}_t$. Following Apel et al. (2022), we count instances in which a modifier appears within seven words of the target term in the same sentence. We then purge the narrative surprises of variation predictable from pre-event macroeconomic and financial information, following the predictability exercises of Bauer and Swanson (2023a), Bauer and Swanson (2023b), and Swanson (2024).
3. **Decompose orthogonised narrative surprises**
   We combine the narrative surprises with stock price responses on the same day, and apply the sign-rotation method to separate the narrative surprises into *stance communication shocks* and *information communication shocks*.[11] We account for the timing of announcements as follows when calculating stock price responses:

[9] Recent evidence shows that speeches and testimony can be important sources of monetary policy news for financial markets (Swanson and Jayawickrema, 2024). We focus on scheduled monetary policy announcements to maintain a cleaner event-study design across countries. Speeches and testimony often bundle monetary policy with financial stability, regulatory, institutional, fiscal, and broader economic issues, making it harder to isolate monetary policy communication from other central bank communication.

[10] For Canada, we remove days with releases of quarterly annualised GDP, year-on-year CPI, or the unemployment rate. For the United Kingdom, we exclude days with quarter-on-quarter GDP, year-on-year CPI, or jobless-claims releases. For the United States, we exclude days with annualised quarter-on-quarter GDP, month-on-month CPI, or nonfarm-payroll releases.

[11] The stock-price indices are the S&P/TSX Composite for Canada, the FTSE 100 for the United Kingdom, and the S&P 500 for the United States.

- If the central bank event takes place during or before the stock market open, we measure surprises as the change in the closing price of stock price compared to the closing price the day before the announcement.
- For events that take place after the market close, the surprise is measured as the difference between the closing price on the day after the event compared to the day of the event.

4. **Rescale shocks and aggregate to monthly frequency**
   We rescale the daily shocks such that a unit of the stance (information) communication shock corresponds to a one per cent decrease (increase) in the stock price level. We then aggregate daily shocks into monthly series to align with the frequency of our downstream local projections. Months without central bank events are assigned a value of zero.

2.2. Estimating impulse responses

We estimate dynamic responses using the smooth local projections (SLP) approach of Barnichon and Brownlees (2019). Relative to unrestricted local projections, SLP imposes smoothness across horizons and improves precision without requiring a fully specified dynamic model. Let θ_h denote the response at horizon h. The SLP estimator approximates the horizon profile of the impulse response by a linear combination of B-spline basis functions:

$$\theta_h \approx \sum_{k=1}^{K} b_k B_k(h), \tag{2}$$

where $B_k(h)$ denotes the k-th basis function and b_k is the associated coefficient. The corresponding local projection is

$$Y_{t+h}^{(i)} = \sum_{k=1}^{K} b_k^{(i)} B_k(h) Shock_t^{(i)} + \Gamma_h^{(i)\prime} X_t^{(i)} + u_{t+h}^{(i)}, \qquad i = 1, \ldots, N. \tag{3}$$

where $Y_{t+h}^{(i)}$ denotes the one-year-ahead household inflation expectation in country i, and

$$Shock_t \in \{stance\ communication\ shock_t, information\ communication\ shock_t\}.$$

Following Barnichon and Brownlees (2019), we use cubic B-splines and allow for a relatively rich basis to reduce approximation bias. The degree of smoothing is governed by a shrinkage parameter, which we select by five-fold cross-validation to balance bias and variance. Eq. (3) is estimated separately for each country. The superscript i indicates that the variables and coefficients belong to the country-i time-series regression. To avoid notational clutter, we omit this superscript in the remainder of the paper.

The vector X_t includes a constant, lags of the outcome variable, annualised inflation, log industrial production, log commodity prices, and the contemporaneous Oxford Covid-19 Stringency Index, which accounts for the exceptional conditions during the pandemic (Hale et al., 2021). Following Jordà and Taylor (2025), we choose the lag length using information criteria evaluated at the first horizon, $h = 1$, and add one additional lag for lag augmentation.

A central concern is that the estimated effects of communication shocks may reflect the policy actions that accompany central bank communication, rather than the informational content of the communication itself. To disentangle words from actions, we implement two safeguards. First, all specifications control for the short-term policy rate, augmented with a shadow rate during the zero lower bound period. Second, we control for monetary policy actions using the country-specific shocks of Bu et al. (2021), who summarise announcement-day yield-curve movements with a single-factor measure of the overall policy-action stance.

Conditional on these controls, the estimated responses can be interpreted as the effects of central bank communication holding contemporaneous policy actions fixed.

We then ask whether the transmission of communication shocks depends on the information environment in which the announcement is received. In particular, central bank communication can affect expectations only if households notice and process the message. We capture this dimension by introducing monetary policy and inflation attention, Att^{MP} as predetermined state variables.

The effect of Att^{MP} is a priori ambiguous. On the one hand, communication can affect expectations only if agents notice and process the message. Higher attention may therefore strengthen transmission by increasing the likelihood that households incorporate central bank communication into their beliefs. On the other hand, attention to monetary policy and inflation tends to rise in periods of heightened and volatile inflation. In such environments, central bank communication may be less effective if agents place greater weight on realised inflation or private signals that conflict with the central bank's message. This distinction is closely related to the selective inattention mechanism emphasised by Coibion and Gorodnichenko (2026): higher attention may strengthen the transmission of central bank communication if households listen more attentively to central bank messages, but it may also amplify adverse interpretations of central bank communication if households become attentive just when inflation is high or policy appears less effective.

To explore this question, we estimate non-linear smooth local projections à la Ascari and Haber (2022):

$$\begin{aligned} Y_{t+h} = &\sum_{k=1}^{K} b_{0,k} B_k(h) Shock_t + \sum_{k=1}^{K} b_{1,k} B_k(h) \left(Shock_t \times Att_{t-1}^{MP} \right) \\ &+ \delta_h Att_{t-1}^{MP} + \Gamma_h' X_t + u_{t+h}. \end{aligned} \tag{4}$$

We use lagged state variables so that the conditioning information is predetermined with respect to the communication shock. We use a continuous state-dependent specification rather than splitting observations into discrete high- and low-state regimes. This approach preserves information in the state variables, avoids imposing arbitrary cutoff rules, and is consistent with the idea that attention affects the transmission of central bank communication continuously.

2.3. Estimating forecast error variance decomposition

In addition to estimating dynamic responses, we assess the quantitative significance of shocks within the local projections framework. We follow Gorodnichenko and Lee (2020), who propose an estimator based on the coefficient of determination (R^2). This approach evaluates the extent to which a given shock accounts for the variation in the dependent variable. The R^2 measure has several desirable traits. Simulation evidence suggests that it performs better than alternatives based on the sum of squared impulse responses. Moreover, when shocks are measured with error, the estimator is downward-biased, implying that the resulting estimates provide a conservative lower bound.[12] Appendix C elaborates on the FEVD algorithm in detail.

3. Corpus and data

This section outlines the construction of our corpus and the scope of our data. We examine the monetary policy announcements of the Bank of Canada, the Bank of England, and the Federal Reserve—from the mid-1990s to the end of 2024. The start date varies across institutions, reflecting differences in communication practices, document availability, and institutional reforms over time.

[12] This complements the upper-bound FEVD estimates in Plagborg-Møller and Wolf (2017, 2022).

3.1. Central bank corpus

Coverage for the **Bank of Canada** begins in 1997. Before 2000, policy announcements were made on an ad hoc basis. Since then, the Bank has adopted a schedule of eight fixed announcement dates each year, reducing uncertainty around the timing of decisions (Desroches et al., 2024).

The **Bank of England** corpus begins in 2000. Until 2015, MPC policy statements were typically brief, often limited to a sentence unless the Bank Rate was changed. A more comprehensive communication framework was introduced with the launch of Super Thursday in 2015, which bundles the policy statement (Monetary Policy Summary), minutes, and voting record. We therefore apply a text-richness screen to the pre-2015 UK announcements, retaining only events with sufficient narrative content to construct a meaningful dictionary-based measure while preserving informative pre-Super-Thursday statements.

The **Federal Reserve** corpus covers the period 1995–2024. It includes post-meeting FOMC statements, which began in February 1994 and became standard practice from 2000 onwards, irrespective of whether the policy rate changed.

3.2. Newspapers

To gauge public sentiment and expectations surrounding monetary policy, we construct a newspaper-based dataset using Factiva's global archive. To achieve broad coverage of both elite and general readership, we include major daily newspapers with diverse editorial profiles (Figure 1). Articles are included if the full name of the central bank appears in either the headline or lead paragraph.[13] The list of newspapers used is broadly consistent with those employed by Caldara and Iacoviello (2022) and Baker et al. (2016), and offers wider coverage than other studies such as Husted et al. (2020), Koop and di Vettimo (2023), Kalamara et al. (2022), and Acosta (2023).

We use newspapers as proxies for the information environment through which households encounter central bank communication. This choice is motivated by three strands of evidence. First, households rarely obtain central bank information directly and instead rely on intermediated channels, including newspapers (Blinder et al., 2024). Second, media coverage matters for household expectations: inflation reporting affects household inflation forecast disagreement, and recent evidence shows that media coverage of monetary policy communication affects household inflation expectations (Lamla and Maag, 2012; De Fiore et al., 2024). Third, newspapers select and frame central bank messages rather than mechanically reproducing them. Editorial choices reflect both newsworthiness and readers' demand for macroeconomic information (Munday and Brookes, 2021; Fisher et al., 2022). We therefore interpret pre-announcement newspaper narratives as public-facing policy priors, and post-announcement newspaper narratives as evidence of media transmission.

Table 1 reports descriptive statistics for the central bank and newspaper texts used to construct the narrative surprises. For central banks, an observation is a policy announcement document. For newspapers, an observation is the pooled text of all articles published within the relevant event window for a given policy announcement. The newspaper narratives are therefore substantially longer than the central bank documents, with mean lengths ranging from about 6,100 words in Canada to about 9,700 words in the United Kingdom. Our dictionary also picks up a strong tone signal in both sources: at least 88 per cent of central bank documents and 89 per cent of newspaper event-window narratives contain at least one hawkish or dovish match. Overall, the table suggests that the event-level newspaper narratives contain enough text and policy tone variation to proxy for the public information set before central bank announcements.

[13] For the Federal Reserve, the search terms include both 'Federal Reserve Board' and 'Federal Reserve'. For the Bank of Canada, we screen out references to commercial banks whose names include 'Bank of Canada', such as Royal Bank of Canada and National Bank of Canada. We opt for the full institutional name over acronyms or informal references, on the assumption that reporting conventions typically mandate the use of full name at least once in introductory text. To reduce duplication, we excluded articles flagged by Factiva as duplicates. As this process is not exhaustive, we apply further filters to remove articles that (i) share identical date and title; (ii) have the same lead paragraph within the same month; or (iii) are shorter than 100 words.

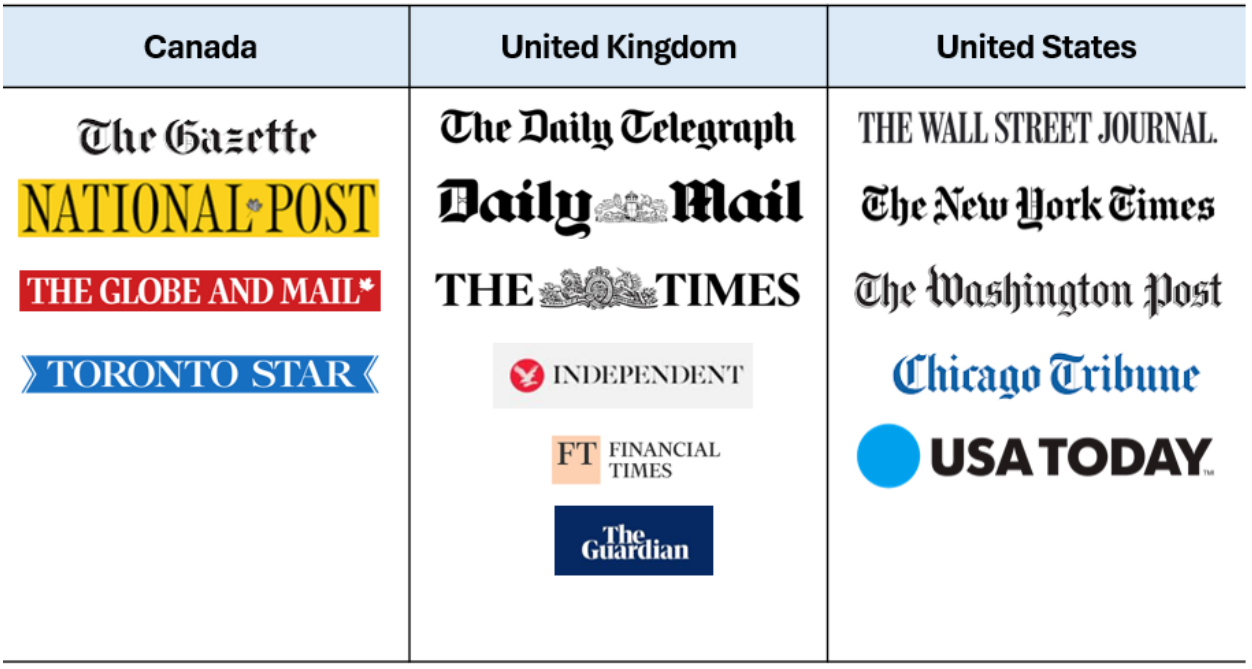

Figure 1: List of newspapers included

Table 1
Descriptive statistics for central bank and newspaper texts

Source	Obs.	Total words	Mean words	Median words	Mean tone matches	Share tone matches (%)
			Canada			
Central bank	208	82,809	398	393	12	88
Newspapers	208	1,261,042	6,063	5,519	36	89
			United Kingdom			
Central bank	123	80,581	655	711	23	93
Newspapers	123	2,609,566	9,665	7,293	80	99
			United States			
Central bank	194	73,212	377	347	18	98
Newspapers	194	1,793,249	9,244	7,882	88	96

Notes: The table reports descriptive statistics for the texts used to construct narrative surprises. For central banks, an observation is a policy announcement document. For newspapers, an observation is the pooled text of all articles published within the relevant event window for a given policy announcement. Word counts are computed after standard text cleaning. Mean tone match is the average number of hawkish plus dovish dictionary matches per observation. Share tone match is the percentage of observations containing at least one hawkish or dovish dictionary match.

A key marker of credible communication is the degree to which public sentiment adjusts in response to central bank messaging. As an initial diagnostic, Figure 2 plots the correlation between two sentiment differentials: (i) the change in hawkish sentiment in newspapers from three days before to five days after a central bank event (vertical axis), and (ii) the difference in hawkish sentiment between the central bank and newspapers over the three days preceding the event (horizontal axis). The correlations exceed 0.4 and are

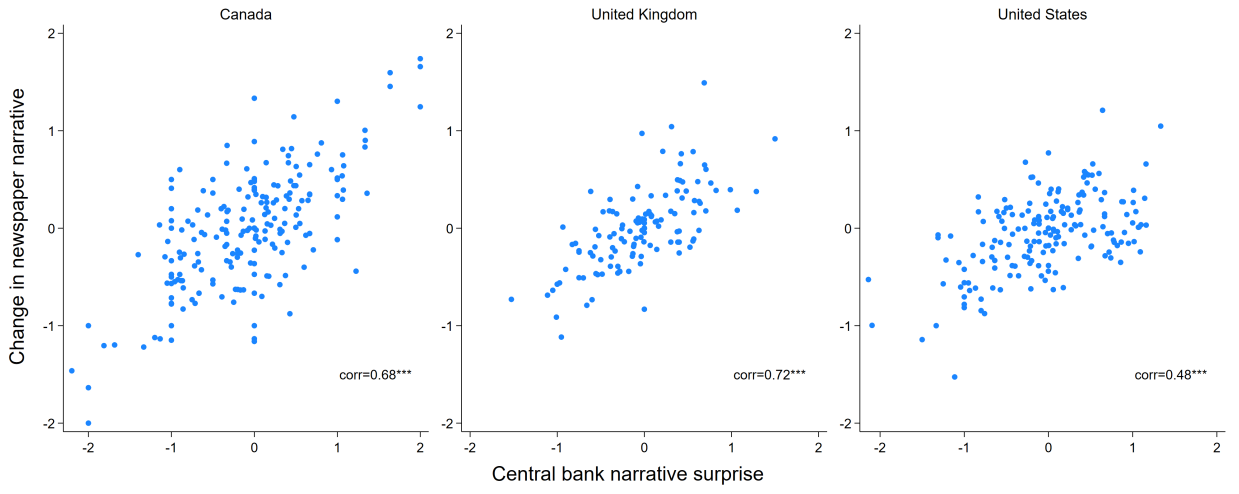

Figure 2: Correlation between public sentiment and central bank communication

Notes: The narrative change in media refers to the difference in newspaper hawkishness between the three days following a central bank event and the three days preceding it.

statistically significant at the 1% level. This provides an initial indication that central bank communication influences public discourse.

To formalise this relationship, Table 2 reports country-specific monthly regressions of the media narrative change on the central bank narrative surprise, controlling for contemporaneous inflation, the policy rate, commodity prices, and industrial production. The coefficient on the central bank narrative surprise is positive and statistically significant at the 1% level in all three countries. The estimated magnitudes are economically significant: a one-unit increase in the central bank narrative surprise is associated with an increase in the media narrative change of 0.58 in Canada, 0.59 in the United Kingdom, and 0.31 in the United States. The adjusted R^2 ranges from 0.3 in the United States to 0.5 in Canada and the United Kingdom, suggesting that the specification explains a substantial share of the monthly variation in media narrative changes. These results indicate that the correlation shown in Figure 2 is not driven solely by contemporaneous macroeconomic conditions. Instead, central bank communication appears to shape public discourse around monetary policy events.

3.3. Dependent variable: subjective inflation expectations

The opening quote of this paper underscores the policy relevance of household expectations.[14] Accordingly, our primary dependent variable is the one-year-ahead expected inflation drawn from household surveys:

- The **Bank of Canada Survey of Consumer Expectations** is a nationally representative quarterly survey of around 2,000 households, conducted online using a rotating panel design. Respondents generally remain in the panel for four quarters, with a similar number joining and leaving each wave,

[14]Their inflation expectations shape savings and consumption decisions via the Euler equation, which links perceived real interest rates to consumption choices (Weber et al., 2022). Inflation expectations also influence wage bargaining, durable investment decisions, and portfolio choices (Bernanke, 2007). As these micro-level decisions aggregate up, inflation expectations play a central role in macroeconomic dynamics. The New Keynesian Phillips curve attributes an important role to inflation expectations in determining realised inflation, and the effectiveness of fiscal and monetary policy in shaping realised inflation hinges on inflation expectations (Sims, 2009; Galí, 2015).

Table 2
Relationship between central bank and media narrative surprises

	(1)	(2)	(3)
	Canada	UK	US
Dependent variable: Media narrative surprise			
Central bank narrative surprise	0.577***	0.588***	0.307***
	(0.042)	(0.059)	(0.047)
Inflation	-0.001	-0.011	-0.042***
	(0.014)	(0.008)	(0.011)
Policy rate	0.015	0.004	-0.000
	(0.011)	(0.008)	(0.006)
Logged commodity price	-0.005	0.159***	0.055
	(0.052)	(0.088)	(0.042)
Logged industrial production	-0.502	0.262*	-0.112
	(0.341)	(0.405)	(0.152)
Observations	331	180	360
R^2	0.510	0.522	0.271
Adjusted R^2	0.503	0.508	0.261

Notes: The table reports country-specific monthly regressions of the media narrative change on the central bank narrative surprise and macroeconomic controls. Robust standard errors are reported in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

which helps reduce variability from changes in composition. The one-year-ahead expectations are elicited by asking: 'What do you expect the rate of [inflation/deflation] to be over the next 12 months? Please give your best guess.' and an interpolated median is used to represent the central tendencies in consumers' expectations.[15]

- The **YouGov/Citigroup inflation survey** is regularly monitored by the Bank of England and is quoted in their Monetary Policy Report (Arioli et al., 2017). It is a monthly, web-based survey of UK adults aged 18 and over, conducted using YouGov's 'active sampling' approach from its online panel. Responses are weighted to match the demographic composition of the UK population. Respondents are asked 'How do you expect consumer prices of goods and services will develop in the next 12 months?', and are asked to select the range of price changes that best summarises their expectations. The use of ranges captures the respondents' uncertainty about future inflation and the reported median figure across respondents mitigates the influence of extreme values (Anderson and Maule, 2014).[16]

[15] See Bank of Canada for survey-design details. To obtain a monthly series before 2014Q4, we estimate a bridge regression on the post-2014 sample, relating the quarterly BoC one-year-ahead expectation to contemporaneous Consensus Economics one-year-ahead CPI forecasts, lagged realised inflation, and lagged consumer confidence. We then use the estimated coefficients and monthly predictors, replacing the quarterly lag with its monthly analogue ($t-3$), to back out monthly fitted values before 2014Q4. For the post-2014 period, we interpolate the quarterly BoC series to monthly frequency using the proportional Denton method with the monthly Consensus series as the indicator. We splice the two series by aligning levels using the mean ratio over a 12-month overlap.

[16] The series begins in October 2005. For 2000–2005, we follow Ramey and Zubairy (2018) and extend coverage by interpolating the quarterly Bank of England/Ipsos Inflation Attitudes Survey using the proportional Denton method, with the monthly ECB's Consumer Expectations Survey for the United Kingdom as the indicator series. We then splice the interpolated series with YouGov/Citigroup series using the average ratio over a 12-month overlap.

Table 3
Search terms for monetary policy attention measures

Country	Country-specific central bank terms	Newspapers
Canada	Bank of Canada; Governing Council	*National Post*; *The Globe and Mail*
United Kingdom	Bank of England; Monetary Policy Committee; MPC	*Financial Times*; *The Times*
United States	Federal Reserve; Federal Open Market Committee; FOMC	*The New York Times*; *The Wall Street Journal*

Notes: The table reports the country-specific central bank search terms and newspaper sources used to construct the monetary policy attention index. An article is classified as monetary-policy-related if it contains at least one country-specific central bank term and at least one common economic term. The common economic terms are: interest rate, monetary, inflation, economy, economic, and unemployment. Searches are case-insensitive. For Canada, articles referring to Royal Bank of Canada or National Bank of Canada are excluded to avoid false matches.

- The **University of Michigan Inflation Expectations** survey interviews around 500 individuals each month, representative of the US population. It has a rotating panel component: each month about 60% of interviewees are first-time respondents, whereas 40% were interviewed six months prior. The survey elicits inflation expectations through a two-step procedure. First, it asks respondents whether they think prices in general will increase, decrease, or stay the same over the next 12 months. Second, the survey asks those who answered 'increase' or 'decrease' 'By about what per cent do you expect prices to go (up/down) on average?'

3.4. State variable: monetary policy attention

We measure attention to monetary policy using newspaper coverage. Following Fisher et al. (2022), we construct a monthly country-level index from articles that refer both to the relevant central bank and to macroeconomic or monetary policy terms. Specifically, for each country and newspaper, we count the number of articles in month t that contain at least one country-specific central bank term and at least one common economic term (Table 3). We then divide this count by the total number of articles published by that newspaper in the same month. The resulting variable is a percentage measure of monetary-policy-related coverage.

Table 3 reports the country-specific central bank terms and newspaper sources. For each country, we first demean the newspaper-level coverage series and then average across newspapers to obtain the country-level news attention index. This procedure removes persistent differences in average coverage across newspapers while preserving monthly variation in attention within each country.[17]

From 2004 onwards, we augment the newspaper-based measure with Google Search data. Google Trends provides a complementary measure of public attention. For each country, we obtain the Google Trends series for the relevant central bank topic: Bank of Canada, Bank of England, and Federal Reserve System. We standardise the news attention index and the Google Trends index over their common sample and define the post-2004 composite attention index as their simple average.[18]

4. Results

4.1. Baseline result: evidence that central bank communication reaches the public

Following a tightening stance communication shock, normalised to a one per cent fall in the daily stock price, households' one-year-ahead inflation expectations decline and remain below baseline for

[17] Alternative measures of attention include text-based salience, survey-based measures, and market-implied proxies; see, for example, Song and Stern (2025), Pfäuti (2025), and An et al. (2023).

[18] We use Google Trends topics, which group together related searches and language variants associated with the same concept.

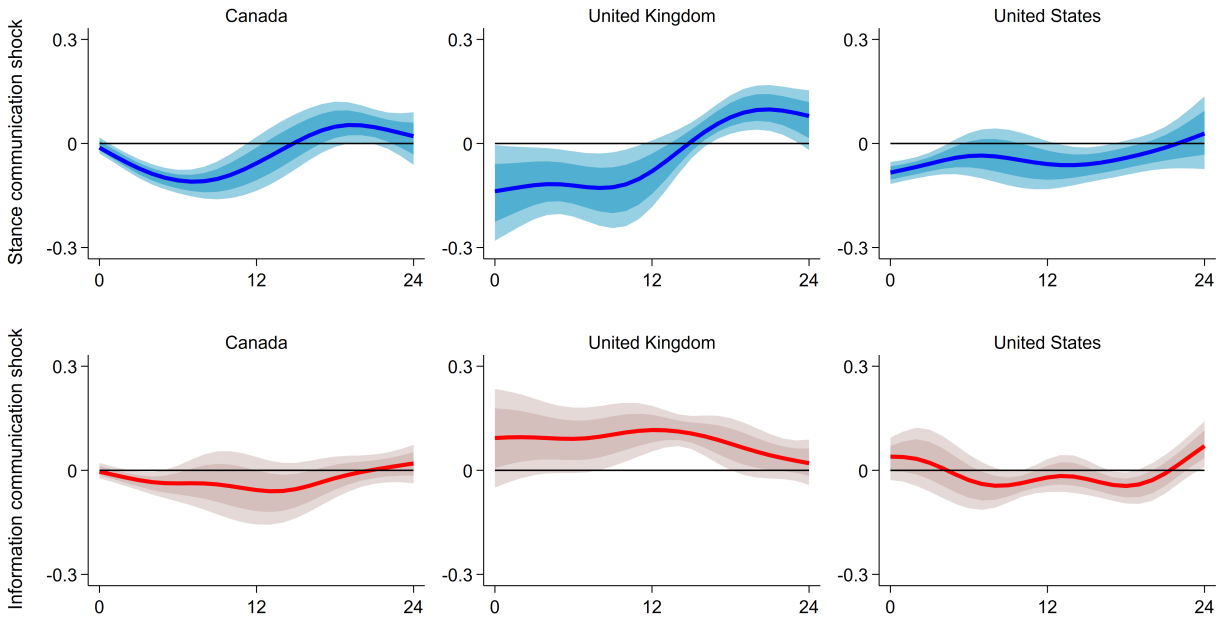

Figure 3: Effects of stance communication shocks on one-year-ahead household expected inflation

Notes: Stance and information communication shocks are normalised to a one per cent fall and rise in the national equity index, respectively. Expected inflation is drawn from household surveys; see Section 3.3 for details. The estimation sample begins in January 1995 for the United States, January 1997 for Canada, and January 2000 for the United Kingdom. Shaded areas denote 68 and 90 per cent confidence bands.

approximately 6–12 months (Figure 3). The estimates are statistically significant at the ten per cent level. The peak response is around −0.1 percentage point relative to baseline—about half the average absolute monthly change in household expectations across the three countries.[19] The relatively rapid adjustment across countries is noteworthy. It contrasts with the predictions of sticky-information models, in which households update expectations only infrequently (Carroll, 2003).

For information communication shocks, the average response is statistically significant only in the United Kingdom. In the United Kingdom, a one-unit information communication shock that raises stock prices by one per cent increases households' one-year-ahead inflation expectations by about 0.1 percentage point. The response is statistically significant between 10 and 18 months after the shock. As shown below, however, the state-dependent results indicate that information communication shocks matter in particular states of the world. The average responses reported here may therefore mask offsetting state-dependent effects.

The decomposition is important for interpreting these baseline responses. Unrotated narrative surprises do not generate the same theoretically consistent responses of household inflation expectations (see Figure 17). Nor are the findings simply a restatement of conventional high-frequency monetary policy shocks. In Section 5.2, we show that our US communication shocks are only weakly correlated with standard asset-price-based monetary policy surprises, and that those benchmark shocks generally produce responses of household inflation expectations that are small, statistically weak, or of the opposite sign.

[19]The average absolute monthly changes are 0.18 percentage point in the United Kingdom and 0.21 percentage point in both Canada and the United States.

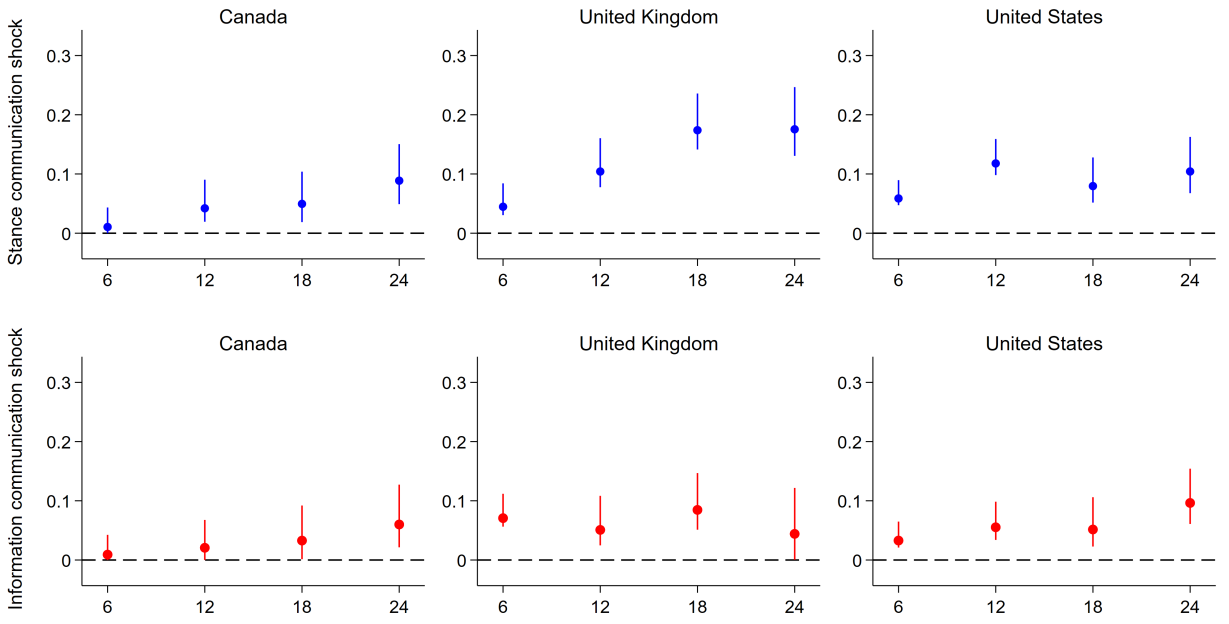

Figure 4: Forecast error variance decomposition of stance and information communication shocks on one-year-ahead households' expected inflation

Notes: The sample begins in January 1995 for the United States, January 1997 for Canada, and January 2000 for the United Kingdom. Shocks are rescaled to correspond to a one per cent fall in the national equity index. The top row reports FEVDs for stance communication shocks and the bottom row for information communication shocks. The vertical axis is expressed in annualised percentage points. Bars denote 90 per cent bootstrap confidence intervals using the block-bootstrap procedure of Kilian and Kim (2011).

Figure 4 reports the forecast error variance decomposition (FEVD). The figures show the share of unexplained variation in households' inflation expectations attributable to each communication shock. At the 12-month horizon, stance and information communication shocks account for around 2–10 per cent of the unexplained variation. This relatively small share is consistent with the view that household inflation expectations are largely driven by supply shocks, gasoline prices, and commodity prices, leaving more limited scope for explanations based on fiscal or monetary policy (Coibion and Gorodnichenko, 2026). Some communication shocks account for a larger share of the variation. In particular, stance communication shocks in the United Kingdom and the United States explain a larger fraction of expected inflation, and similarly for the information communication shocks in the United Kingdom. As these results are conditioned on policy rates and Bu et al. (2021)'s shocks, they can be interpreted as the explanatory power of central bank communication on households' inflation expectations, purged of contemporaneous policy actions.

4.2. State-dependent effects: monetary policy and inflation attention

We next examine whether the effects of central bank communication depend on the attention paid to monetary policy and inflation-related news. The state variable is the country-specific attention index described in Section 3.4. We report impulse responses evaluated at low and high attention, defined as one standard deviation below and above the country-specific mean of the standardised attention index. This normalisation makes the low- and high-attention states comparable across countries.

Figure 5 reports the state-dependent responses to stance communication shocks. In the United Kingdom and the United States, the response of one-year-ahead inflation expectations is larger, and occurs earlier, when

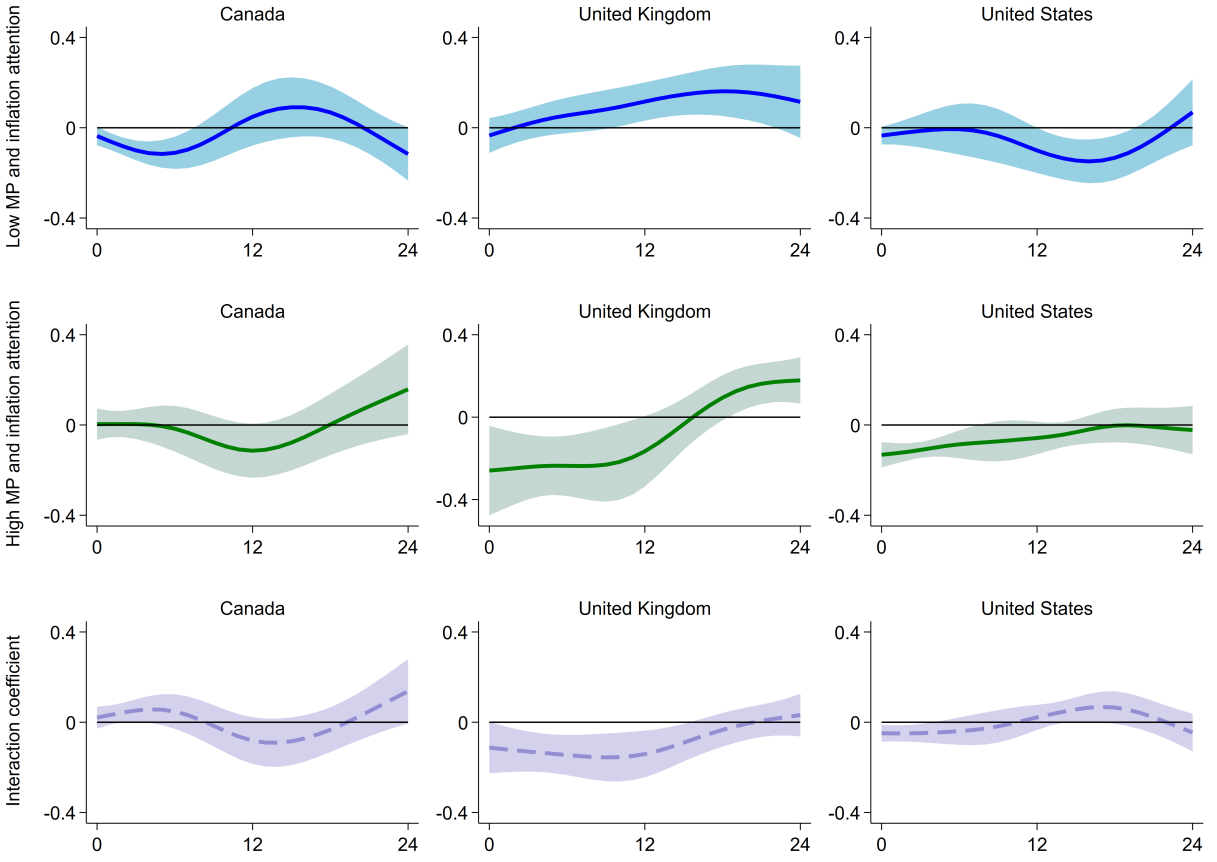

Figure 5: State-dependent effects of stance communication shocks: monetary policy and inflation attention

Notes: The figure reports responses of one-year-ahead household inflation expectations to stance communication shocks, conditional on monetary policy and inflation attention. Low and high attention refer to responses evaluated at one standard deviation below and above the country-specific mean of the standardised attention index, respectively. The bottom row reports the interaction coefficient. Shaded areas denote 90% confidence intervals.

monetary policy and inflation attention is high. The interaction coefficient, $Shock_t \times Att_{t-1}^{MP}$, is negative and significant at the 90% level in both countries, implying that high attention amplifies the disinflationary effect of a stance shock. In contrast, we do not find the same pattern for Canada. The Canadian response is negative when attention is low, but not when attention is high. One possible interpretation is that low-attention periods in Canada coincide with more stable inflation environments in which central bank communication is more credible and expectations are easier to anchor. We therefore interpret the Canadian evidence as indicating that attention is not uniformly transmission-enhancing across countries.

Figure 6 reports the corresponding results for information communication shocks. In Canada and the United Kingdom, the positive effect of information shocks on inflation expectations is concentrated in high-attention states. The interaction coefficients are positive, and the simulated responses indicate that information shocks raise expected inflation only when monetary policy and inflation attention is high. When attention is low, the estimated response is statistically weak and, in Canada, has the opposite sign. We find little evidence in the United States that monetary policy and inflation attention conditions the transmission of information shocks.

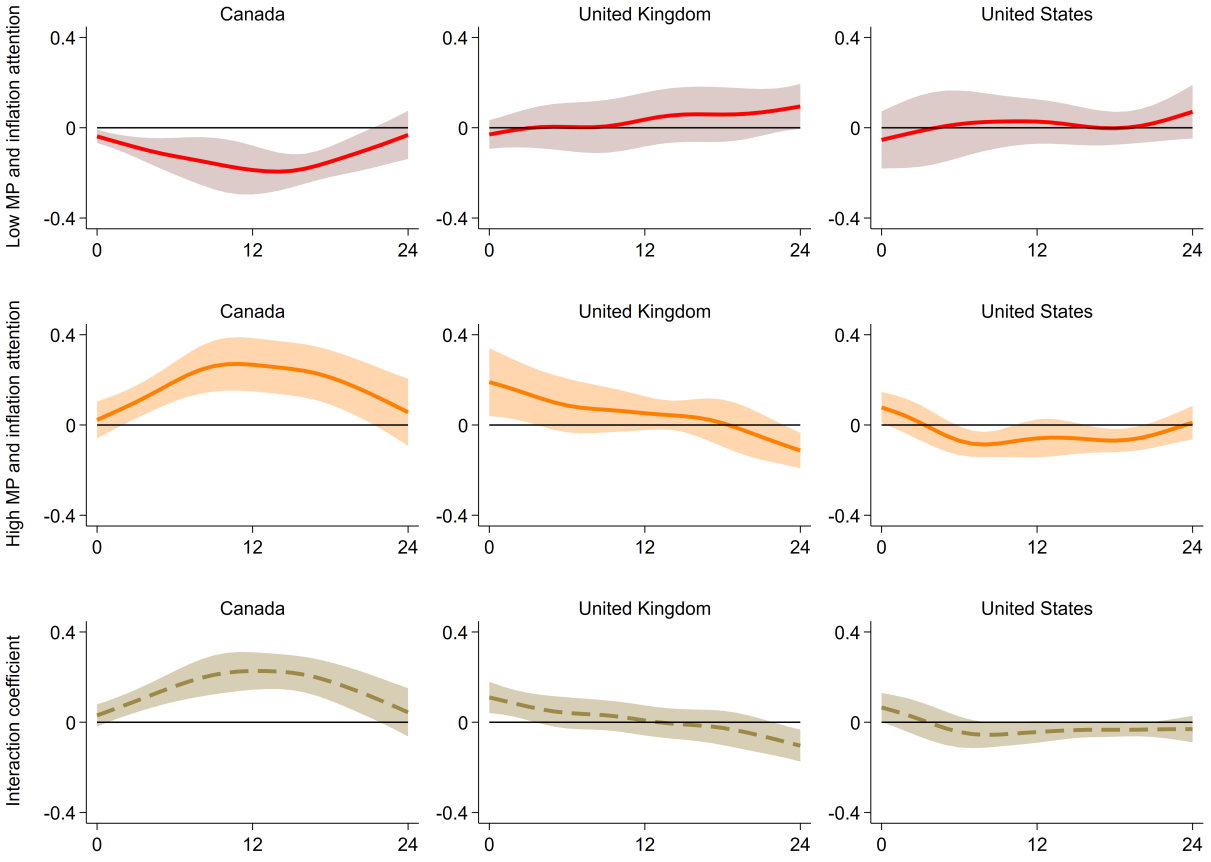

Figure 6: State-dependent effects of information communication shocks: monetary policy and inflation attention

Notes: The figure reports responses of one-year-ahead household inflation expectations to information communication shocks, conditional on monetary policy and inflation attention. Low and high attention refer to responses evaluated at one standard deviation below and above the country-specific mean of the standardised attention index, respectively. The bottom row reports the interaction coefficient. Shaded areas denote 90% confidence intervals.

Taken together, therefore, we find strong evidence that monetary policy attention is a relevant state for the United Kingdom, for both its stance and information communication, and partial evidence for Canada and the United States. Attention brings forward the timing of the impact of stance communication in the United States, and increases the magnitude and persistence of the information communication shocks in Canada.

5. Validation and robustness

Before turning to individual robustness exercises, it is useful to organise them around the main identification concerns. A narrative communication shock should satisfy three requirements. First, it should capture central bank communication that reaches the public information environment, rather than merely reflecting noise in text measurement. Second, it should be distinct from conventional monetary policy shocks and raw narrative tone. Third, its components should have interpretable effects on expectations, with stance and information communication shocks moving beliefs in directions consistent with theory. The exercises in this section are designed to address these requirements.

5.1. Validating the construction of narrative communication shocks

We examine whether the results depend on the construction of the narrative communication shocks.

5.1.1. Alternative media windows

We examine whether the results depend on the window used to measure the pre-announcement newspaper narrative. Specifically, we construct narrative surprises using newspaper coverage over the ten days preceding each central bank announcement. This window follows Ter Ellen et al. (2021), who also use a ten-day window to construct narrative surprises.

Figure 7 shows that the resulting impulse responses are broadly similar to the baseline estimates. This indicates that the main results are not sensitive to the precise length of the pre-announcement newspaper window.

5.1.2. Alternative dictionaries

We next examine whether the results depend on the dictionary used to measure central bank and newspaper narratives. Figure 8 reports responses based on the original dictionary of Apel et al. (2022). Figure 9 instead uses an expanded dictionary that adds topic terms from Aruoba and Drechsel (2025) and monetary policy terms from Cieslak and McMahon (2024). The responses are broadly similar across the two alternatives, although the Apel et al. (2022) dictionary delivers somewhat less significant results. The expanded dictionary produces estimates that are close to the baseline. Overall, the main conclusions are not driven by the precise dictionary used to classify narrative tone, while the additional topic and monetary policy terms appear to improve measurement in this setting.

5.1.3. ChatGPT-based classification

As another measurement check, we replace the dictionary-based narrative scores with classifications generated by a large language model. Korinek (2023) documents the ability of large language models to process context-rich text and approximate patterns in human reasoning. Motivated by this evidence, and following Hansen and Kazinnik (2024), we use GPT-5.4 to classify the hawk–dove policy stance of central bank documents and newspaper articles. Once these classifications are obtained, the remaining steps of the identification procedure follow Section 2.1.

Specifically, we feed the same corpus of central bank communications and news articles to GPT-5.4 and ask it to assign each document to one of five categories: dovish, mostly dovish, neutral, mostly hawkish, and hawkish.[20] To improve reliability, we query each document three times and average the assigned stance across the three responses. This approach provides an alternative to dictionary-based measurement and reduces reliance on researcher choices over keywords, modifiers, and term-level classification rules.

Figure 10 reports the results using GPT-5.4-based narrative classification. The stance responses are broadly consistent with the baseline dictionary-based estimates. In both the United Kingdom and the United States, tightening stance shocks continue to lower household inflation expectations over much of the horizon, with the GPT-based estimates showing somewhat greater persistence but smaller magnitudes in the United States. For information shocks, the United Kingdom continues to show a positive response at short and medium horizons, while the United States remains close to zero. The main difference is Canada, where the GPT-based information shock produces a negative medium-horizon response. The GPT-based estimates generally deliver narrower confidence bands, especially for the United Kingdom and the United States. Overall, the exercise supports the main stance conclusions while showing that some information-shock responses are more sensitive to the measurement of narrative tone.

[20] Table 1 in Hansen and Kazinnik (2024) defines these categories and reports the numerical values we adopt on a -1 to 1 scale, with 0 corresponding to a neutral stance.

5.2. Distinguishing communication shocks from conventional monetary policy shocks

A concern is that the communication shocks may simply repackage standard high-frequency monetary policy surprises. We therefore compare them with existing US monetary policy shock measures. As these series are constructed from high-frequency financial market surprises, the comparison can clarify whether the communication shocks capture variation that is distinct from standard monetary policy surprises.

We compare our shocks with several measures commonly used in the literature:

- The Bauer and Swanson (2023b) surprise is the first principal component of 30-minute changes in money market futures rates over the next four quarters around the FOMC announcements. These surprises reflect changes in expectations for the near-term path of short-term interest rates.
- Orthogonalised Bauer and Swanson surprises are derived as residuals from a regression of the raw surprises on six economic and financial variables that are relevant predictors of the interest rate changes around FOMC announcements. The regression removes predictable variation and isolates the exogenous component of monetary policy announcements.
- Gurkaynak et al. (2005) extract two factors from five high-frequency financial instruments: a target factor, capturing unanticipated changes in the federal funds rate; and a path factor, interpreted as a forward guidance shock, spanning market expectations up to one year ahead.
- Nakamura and Steinsson (2018) construct a shock series based on the first principal component of changes in federal funds and Eurodollar futures at 2- to 4-quarter horizons. These instruments span the liquid segments of the interest rate futures market and are shown to be strong predictors of near-term monetary policy expectations (Gürkaynak et al., 2007).
- Bu et al. (2021) apply the heteroskedasticity-based estimator of Rigobon and Sack (2004) to daily changes in Treasury yields across the full term structure (1 to 30 years) to extract a single-factor summary measure of the monetary policy actions on FOMC days. The two-step Fama and MacBeth (1973) regression is used to mitigate the influence of the Fed's information effect. A notable feature of this series is that it generates non-zero shocks during the effective lower bound (ELB) period.
- Jarociński and Karadi (2018) identify monetary policy shocks as those that simultaneously increase the three-month-ahead Fed Funds futures rate and reduce equity prices within a 30-minute window around FOMC announcements, thereby isolating genuine policy surprises from central bank information effects.

All series are taken from publicly available sources—the Federal Reserve Bank of San Francisco, Dr. Bu's website, Dr. Jarocinski's website, and Dr. Acosta's website.

Table 4 reports pairwise correlations between the US communication shocks and existing monetary policy shock measures. The stance narrative shocks are weakly correlated with standard target, path, and yield-curve-based surprises, including those of Bauer and Swanson (2023b), Nakamura and Steinsson (2018), and Bu et al. (2021). This indicates that the text-based measures are not simply reformulations of conventional monetary policy shocks. At the same time, the policy stance communication shock is most strongly correlated with the Jarociński and Karadi (2018) monetary policy shock, reflecting their shared use of asset-price comovement to separate policy news from central bank information. Conversely, the information communication shock is positively correlated with the Jarociński and Karadi (2018) central bank information shock, but uncorrelated with other conventional shocks. Overall, the table suggests that the communication shocks capture a distinct communication component while preserving the policy-versus-information distinction as in Jarociński and Karadi (2018).

Table 4
Correlations with existing US monetary policy shocks

	Bauer–Swanson	Orth. Bauer–Swanson	GSS target	GSS path	Nakamura–Steinsson	Bu et al. (2021)	JK monetary policy	JK central bank information
stance communication shock	0.155***	0.159***	0.038	0.087	0.089*	0.145***	0.326***	−0.155***
Information communication shock	−0.012	−0.015	0.000	−0.028	−0.020	−0.155***	−0.093*	0.142***

Notes: Correlation is computed over the period of January 1995–December 2023. Bauer–Swanson denotes the monetary policy surprise measure of Bauer and Swanson (2023b); Orth. Bauer–Swanson denotes the corresponding series purged of predictable components. GSS target and GSS path are the target and path factors of Gurkaynak et al. (2005). Nakamura–Steinsson denotes the monetary policy news shock of Nakamura and Steinsson (2018). Bu et al. (2021) denotes the unified monetary policy shock of Bu et al. (2021). JK monetary policy and JK central bank information denote the monetary policy and information components of Jarociński and Karadi (2018). Correlations are computed over the overlapping monthly sample. *** and * denote statistical significance at the 1% and 10% levels, respectively.

As an additional benchmark, Figure 11 reports responses of household inflation expectations to existing US monetary policy shock measures. The specification follows the baseline specification, with two exceptions: the sample is restricted to 1995m1–2023m12, reflecting the availability of the benchmark shock series, and the Bu et al. (2021) shock is excluded from the control set because it is itself one of the shocks considered. The resulting responses are generally small, statistically insignificant, or of the opposite sign. The comparison suggests that narrative communication contains information for household inflation expectations beyond that captured by conventional monetary policy surprises.

5.3. Specification robustness

We assess whether the baseline responses are sensitive to specification choices. We consider three exercises: replacing the Bu et al. (2021) monetary policy shock control with a policy rate surprise, selecting the lag length using an alternative information criterion, and replacing household inflation expectations with professional forecasts.

5.3.1. Alternative policy action control

The baseline specification controls for the Bu et al. (2021) monetary policy shock, which summarises policy news from movements across the yield curve. As an alternative, we use the Bloomberg policy rate surprise as a measure of the immediate policy action. It is defined as the announced policy rate less the pre-announcement Bloomberg median expectation. Bloomberg collects forecasts from professional economists and market economists for scheduled policy decisions. We use the median forecast available before each announcement, which reduces the influence of outliers and provides a real-time measure of the expected policy decision. Figure 12 shows that replacing the Bu et al. (2021) control with this policy rate surprise leaves the main responses largely unchanged.

5.3.2. Alternative lag lengths

The baseline specification selects the lag length using the Bayesian information criterion and includes one additional lag for lag augmentation, following Jordà and Taylor (2025). As an alternative, we select the lag length using the Akaike information criterion, again adding one lag for lag augmentation. Figure 13 shows that the responses are broadly similar to the baseline estimates.

5.3.3. Professional inflation expectations

We replace household inflation expectations with one-year-ahead professional inflation forecasts from Consensus Economics. This exercise is an outcome benchmark rather than a direct robustness check, since professional forecasters differ from households in their expectation formation, expertise, and attention

to monetary policy news.[21] As shown in Figure 14, the responses of professional forecasts to stance communication shocks are broadly similar to the baseline responses for households. The main exception is the response to UK information communication shocks, which is less pronounced than in the household specification. Overall, the exercise suggests that narrative communication shocks contain information relevant for inflation expectations beyond the household survey alone.

5.4. Validating the attention state

Finally, because attention is correlated with inflation salience and monetary policy uncertainty, we test whether the state-dependent results survive alternative attention measures and controls for uncertainty. We consider two robustness exercises. First, we control directly for monetary policy uncertainty. Second, we construct a residualised measure of monetary policy attention that removes variation associated with inflation attention.

5.4.1. Controlling for monetary policy uncertainty

Monetary policy attention and monetary policy uncertainty are related, albeit distinct, state variables. Periods of heightened uncertainty may increase the incentive to acquire information about monetary policy. Conversely, greater attention may either reduce uncertainty, by improving the processing of central bank communication, or increase it, by making disagreement and policy trade-offs more salient. Consistent with this link, Figure 18 shows a positive association between monetary policy uncertainty and monetary policy attention. A natural concern is therefore that the baseline attention estimates partly reflect omitted state dependence associated with monetary policy uncertainty.

To address this concern, we augment Eq. (4) with lagged monetary policy uncertainty. All other elements of the specification are unchanged. Figure 19 and Figure 20 report the results for stance-communication shocks and information-communication shocks, respectively. The main patterns are similar to those in the baseline attention specification. In particular, controlling for uncertainty does not eliminate the state dependence associated with monetary policy attention. We therefore interpret the baseline results as reflecting attention to monetary policy rather than merely variation in monetary policy uncertainty.

5.4.2. Residualising inflation attention

The baseline attention index is designed to capture attention to monetary policy and related macroeconomic news. Because monetary policy coverage often coincides with inflation coverage, we next examine whether the results are driven by attention to inflation rather than by attention to monetary policy per se. To do so, we construct a residualised, 'pure', monetary policy attention index.

The residualisation is conducted at the component level, as follows.

1. For each country, we regress the monetary policy attention series on the inflation attention series at the newspaper level. The inflation attention measure is constructed using the country-specific search terms in Table 6.
2. We retain the residuals from these newspaper-level regressions, average the residuals within country, and standardise the resulting series.
3. From 2004 onwards, we apply the same procedure to Google Trends data: the central bank search index is regressed on the inflation search index, and the residual is standardised.
4. Before 2004, the residualised attention index is the standardised residualised newspaper measure. From 2004 onwards, it is the average of the standardised residualised newspaper and Google Trends measures.

[21] Professional forecasts are from Consensus Economics, which surveys private sector forecasters at monthly frequency. We construct fixed-horizon one-year-ahead expectations following the interpolation approach in Mehrotra and Yetman (2018) and Yetman (2018).

For the Google Trends series, we use topics rather than literal search strings where available. Google recommends using topics when possible because they are language agnostic and account for spelling variants, misspellings, and multiple names for the same concept (Google, 2021).

Figure 21 and Figure 22 report the state-dependent responses using this residualised attention measure. The results show that the baseline findings are not solely driven by inflation attention. Some estimates become weaker after residualisation, especially for Canada, where the low-attention response to stance-communication shocks is attenuated. This is useful because it suggests that the counterintuitive Canadian baseline result is partly associated with inflation-related attention. In other cases, the evidence becomes sharper. For the United Kingdom, for example, the state dependence of stance-communication shocks is more pronounced after purging inflation attention.

Overall, the robustness exercises suggest two conclusions. First, monetary policy attention and inflation attention are correlated, making it difficult to disentangle them in practice. Second, the main attention results are not explained solely by inflation attention or by monetary policy uncertainty. The evidence therefore supports the interpretation that attention to central bank communication is an independent state variable for the transmission of communication shocks.

5.4.3. Binary state specifications

To assess whether the main state-dependent results are robust to a discrete classification of the information environment, we replace the continuous attention interaction in the baseline state-dependent specification with a binary high-attention indicator. In particular, we estimate separate impulse responses in months of high and low monetary policy and inflation attention à la Ascari and Haber (2022).

Let $D(Att^{MP}_{t-1})$ be an indicator equal to one when the standardised monetary policy and inflation attention index is above zero in the month before the shock, and equal to zero otherwise. We estimate

$$\begin{aligned} Y_{t+h} = \tau_h t + D(Att^{MP}_{t-1}) \left(\alpha^{HI}_h + \beta^{HI}_h Shock_t + \sum_{k=1}^{K} \gamma^{HI}_{h,k} X_{t,k} \right) \\ + \left(1 - D(Att^{MP}_{t-1})\right) \left(\alpha^{LO}_h + \beta^{LO}_h Shock_t + \sum_{k=1}^{K} \gamma^{LO}_{h,k} X_{t,k} \right) + u_{t+h}, \end{aligned} \tag{5}$$

where $Shock_t$ denotes either the stance communication shock or the information communication shock. The coefficients β^{HI}_h and β^{LO}_h trace the impulse responses in high- and low-attention states, respectively.

Figure 15 and Figure 16 report the results. The binary specification confirms the main conclusions from the continuous state-dependent analysis. Stance communication shocks have stronger effects in the United Kingdom and the United States, while information communication shocks matter primarily in Canada and the United Kingdom. The estimated responses are generally larger than in the headline specification, suggesting that the continuous interaction may understate the contrast between high- and low-attention regimes. Overall, the results indicate that our main findings are not driven by the particular functional form used to model attention. Rather, they continue to hold when attention is measured using a simple high-versus-low regime split.

Taken together, these exercises support the interpretation of the estimated shocks as communication shocks rather than as repackaged macroeconomic news, conventional monetary policy surprises, or artefacts of a particular text-classification method. The results are strongest when narrative surprises are decomposed into stance and information components, consistent with the view that central bank communication contains both policy-stance news and information about the central bank's economic outlook.

6. Conclusion

This paper studies whether central bank communication reaches the general public. We identify communication shocks from the texts in which central bank messages originate and the newspaper narratives through which households encounter them. The approach departs from the standard identification via asset prices, which reflect the views of financial market participants rather than those of households. By comparing central bank narratives with pre-announcement newspaper narratives, and by separating stance communication shocks from information communication shocks, we obtain shocks that generate impulse responses consistent with theory.

Applying the framework to the Bank of Canada, the Bank of England, and the Federal Reserve System, we find that central bank narratives pass through to media coverage and affect households' one-year-ahead inflation expectations. Stance communication shocks lower inflation expectations, while information communication shocks raise them, especially when households' attention is high. Communication shocks account for a meaningful share of the unexplained variation in inflation expectations. These findings qualify the view that central bank communication rarely reaches the public and point to conditions under which communication can serve as an additional policy tool alongside the policy rate. For policymakers seeking to manage inflation expectations, the implication is that communication policy must be designed not only around what is said, but also around how messages are received and perceived.

A. Robustness figures

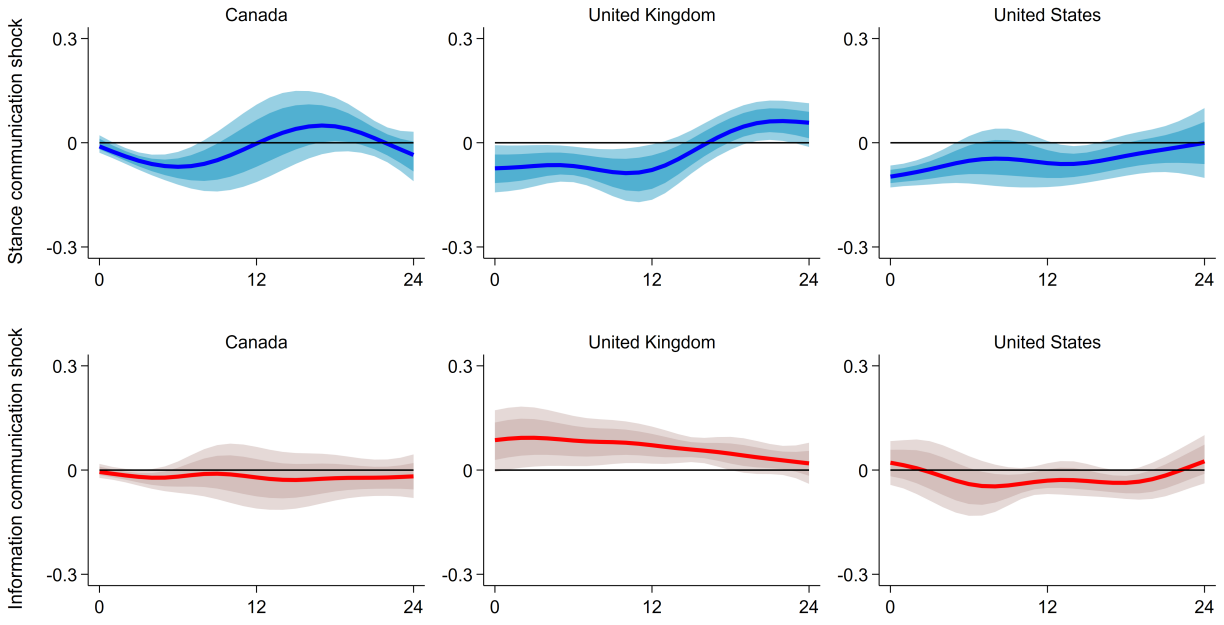

Figure 7: Robustness check 1: ten-day newspaper window

Notes: This figure reports impulse responses to communication shocks constructed using newspaper narratives measured over the ten days preceding each central bank announcement. The resulting narrative surprise is decomposed into policy-stance and information narrative shocks using the same procedure as in the baseline specification.

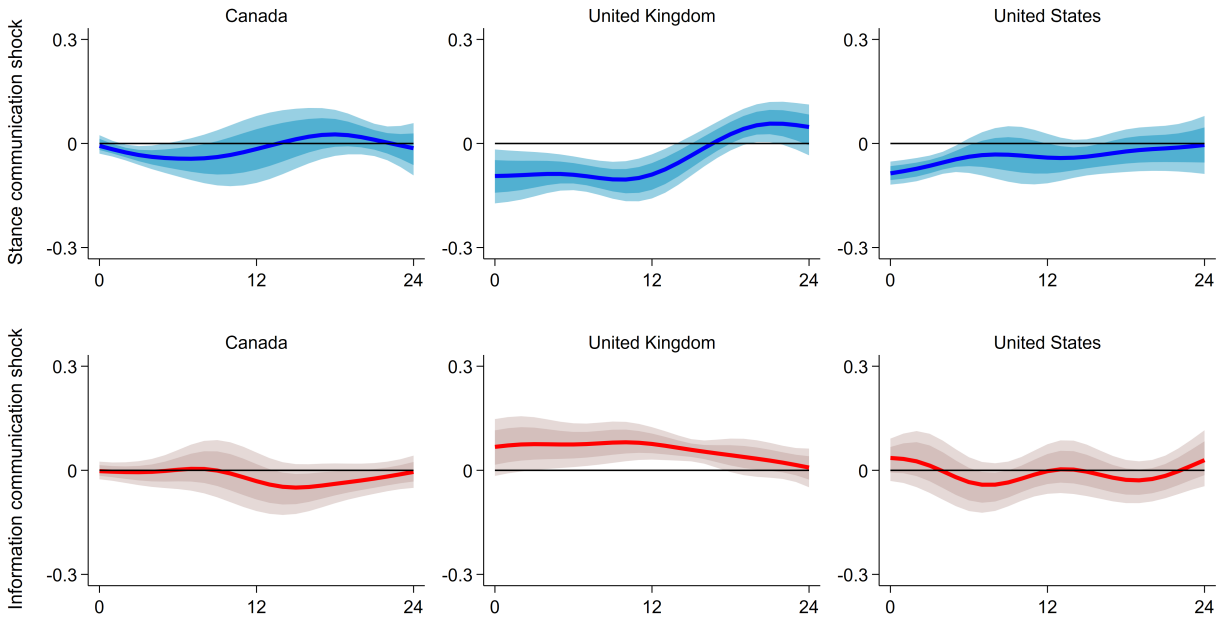

Figure 8: Robustness check 2: original Apel–Grimaldi dictionary

Notes: This figure reports impulse responses to communication shocks constructed using the original dictionary of Apel et al. (2022). The resulting narrative surprises are decomposed into policy-stance and information narrative shocks using the same procedure as in the baseline specification.

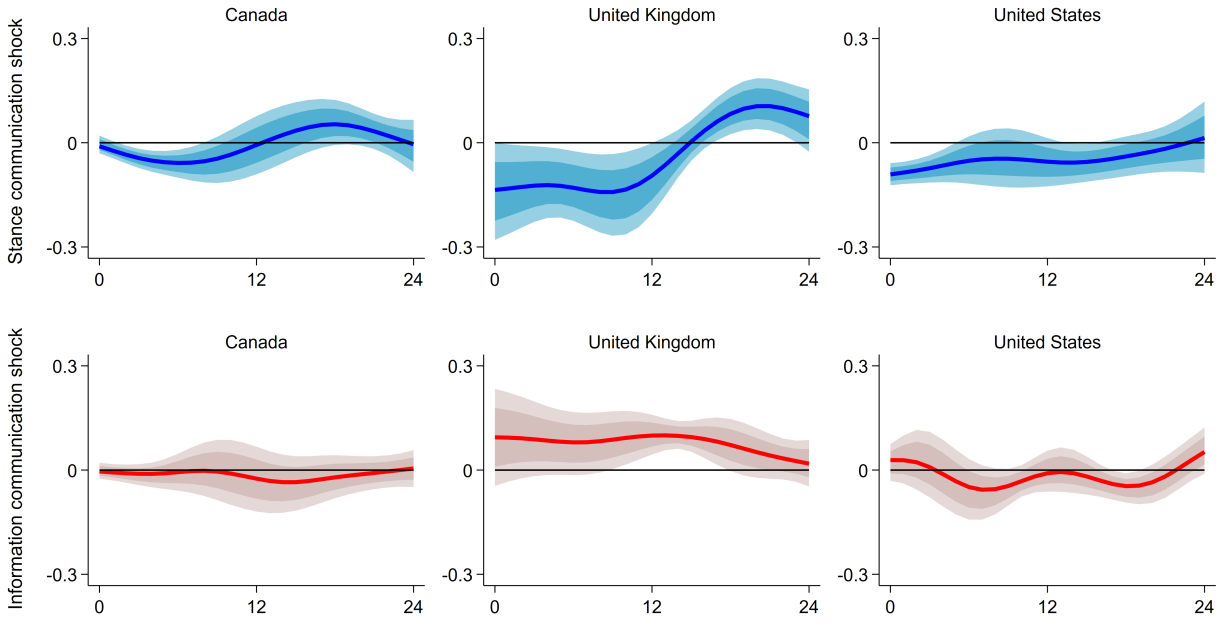

Figure 9: Robustness check 3: expanded central-bank dictionary

Notes: This figure reports impulse responses to communication shocks constructed using an expanded version of the Apel et al. (2022) dictionary. The expanded dictionary adds topic terms from Aruoba and Drechsel (2025) and monetary-policy terms from Cieslak and McMahon (2024). The resulting narrative surprises are decomposed into policy-stance and information narrative shocks using the same procedure as in the baseline specification.

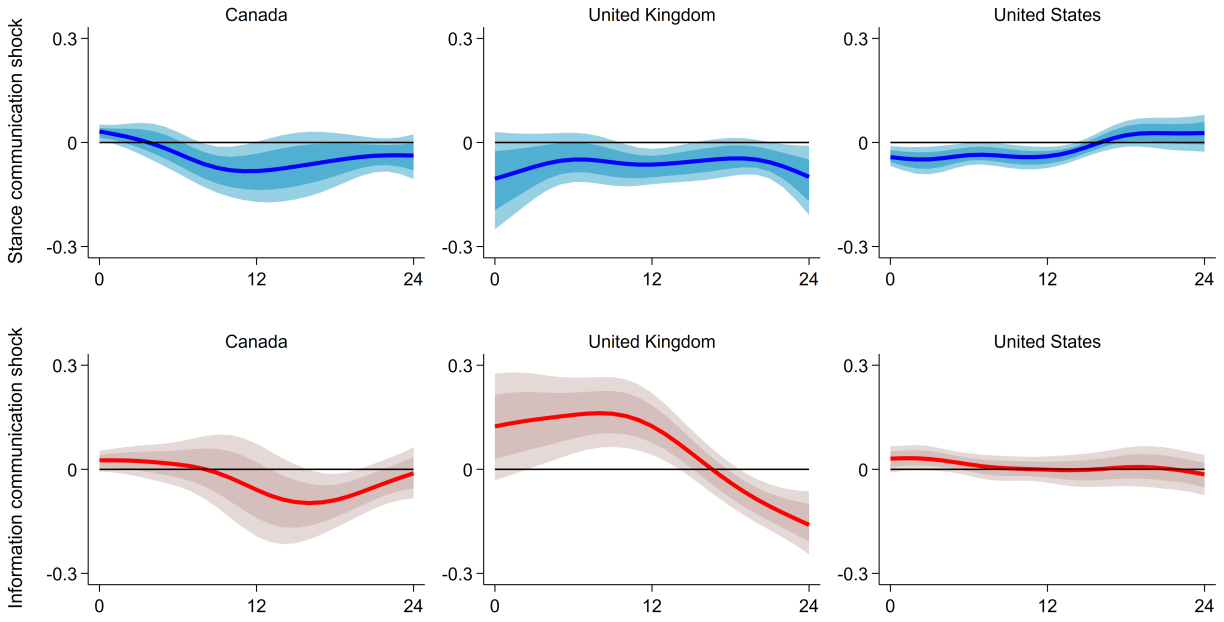

Figure 10: Robustness check 4: ChatGPT-based narrative classification

Notes: This figure reports impulse responses to communication shocks constructed from ChatGPT-based classifications of central-bank communications and newspaper articles. Each document is assigned to one of five policy-stance categories: dovish, mostly dovish, neutral, mostly hawkish, and hawkish. The resulting numerical scores are used to construct narrative surprises, which are then decomposed into policy-stance and information narrative shocks using the same procedure as in the baseline specification.

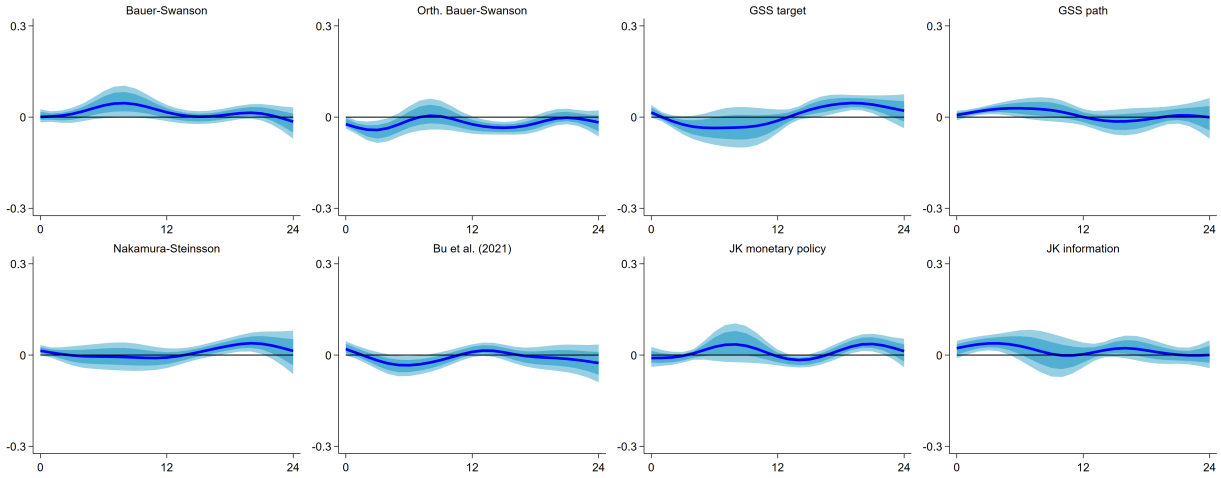

Figure 11: Robustness check 5: responses to existing US monetary-policy shocks

Notes: This figure reports impulse responses of one-year-ahead household inflation expectations to existing US monetary policy shock measures. Responses are estimated over January 1995–December 2023. The specification is the same, except that we no longer control for Bu et al. (2021) shock given it's one of the main shocks. Bauer-Swanson denotes the monetary-policy surprise measure of Bauer and Swanson (2023b); Orth. Bauer-Swanson denotes the corresponding series purged of predictable components. GSS target and GSS path are the target and path factors of Gurkaynak et al. (2005). Nakamura-Steinsson denotes the monetary policy news shock of Nakamura and Steinsson (2018). Bu et al. (2021) denotes the unified monetary policy shock in their paper. JK monetary policy and JK information denote the monetary policy and information components of Jarociński and Karadi (2018). Shaded areas denote 68% and 90% confidence intervals.

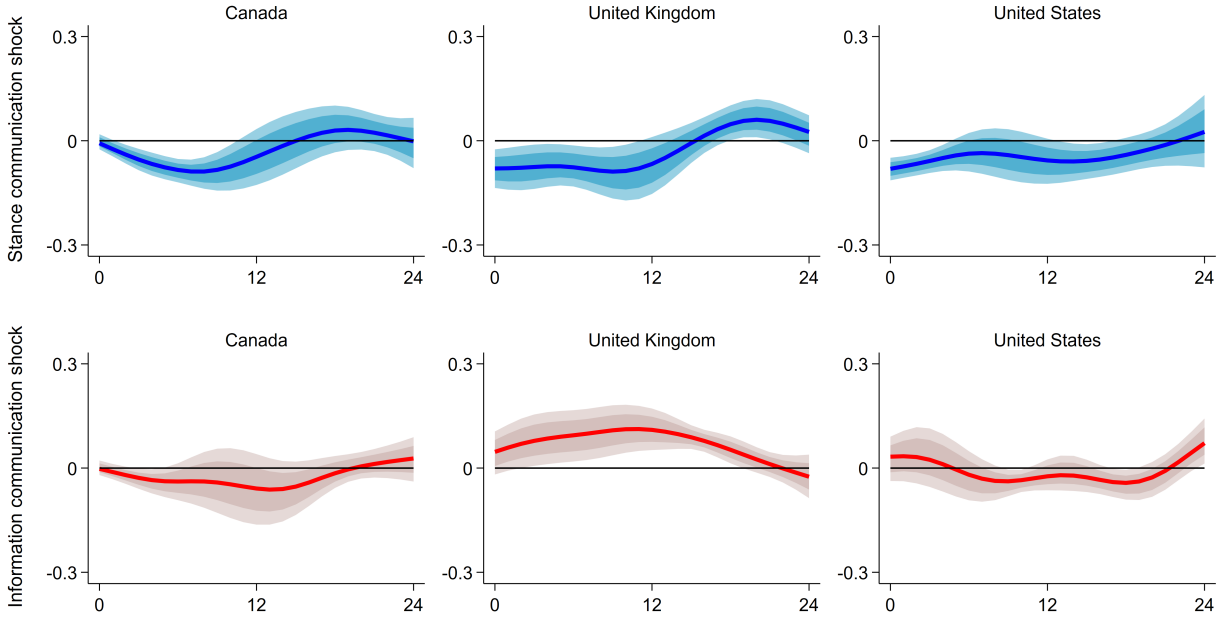

Figure 12: Robustness check 6: Bloomberg policy-rate surprise as policy-action control

Notes: This figure reports impulse responses of one-year-ahead household inflation expectations to narrative communication shocks. The specification is the same as in the baseline, except that the Bu et al. (2021) monetary-policy shock control is replaced by a Bloomberg policy-rate surprise, measured as the announced policy rate less the pre-announcement Bloomberg median expectation. Shaded areas denote 68% and 90% confidence intervals.

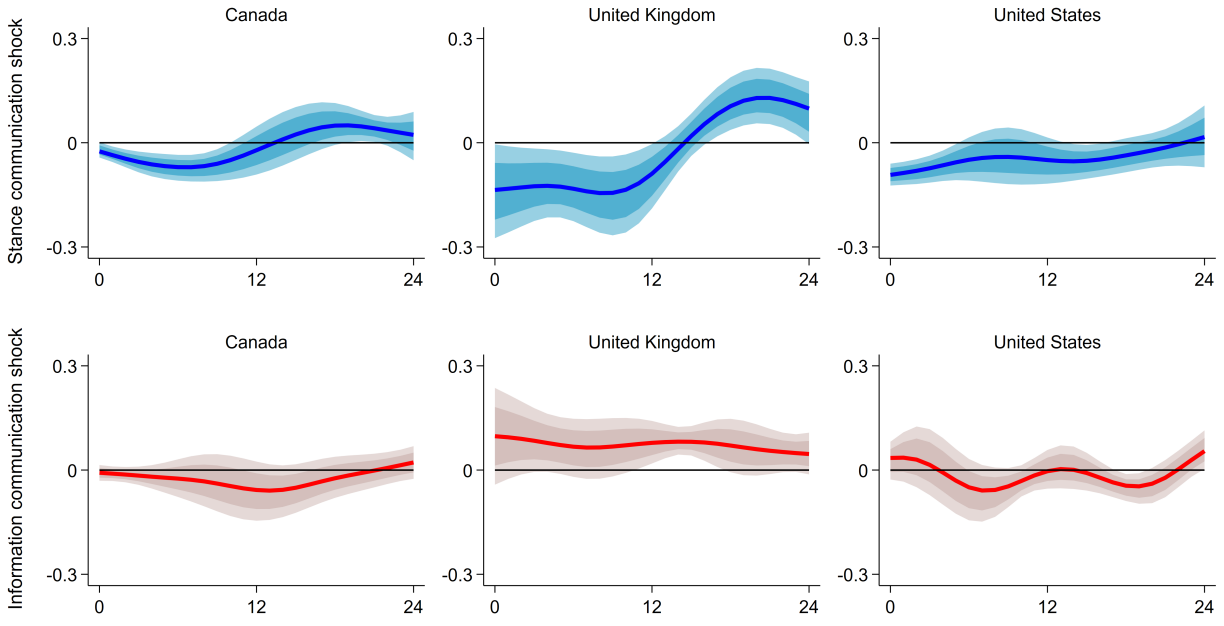

Figure 13: Robustness check 7: lag length selected by the Akaike information criterion

Notes: This figure reports impulse responses of one-year-ahead household inflation expectations to narrative communication shocks. The lag length is selected using the Akaike information criterion in the one-step-ahead projection, with one additional lag included for lag augmentation. All other aspects of the specification are unchanged. Shaded areas denote 68% and 90% confidence intervals.

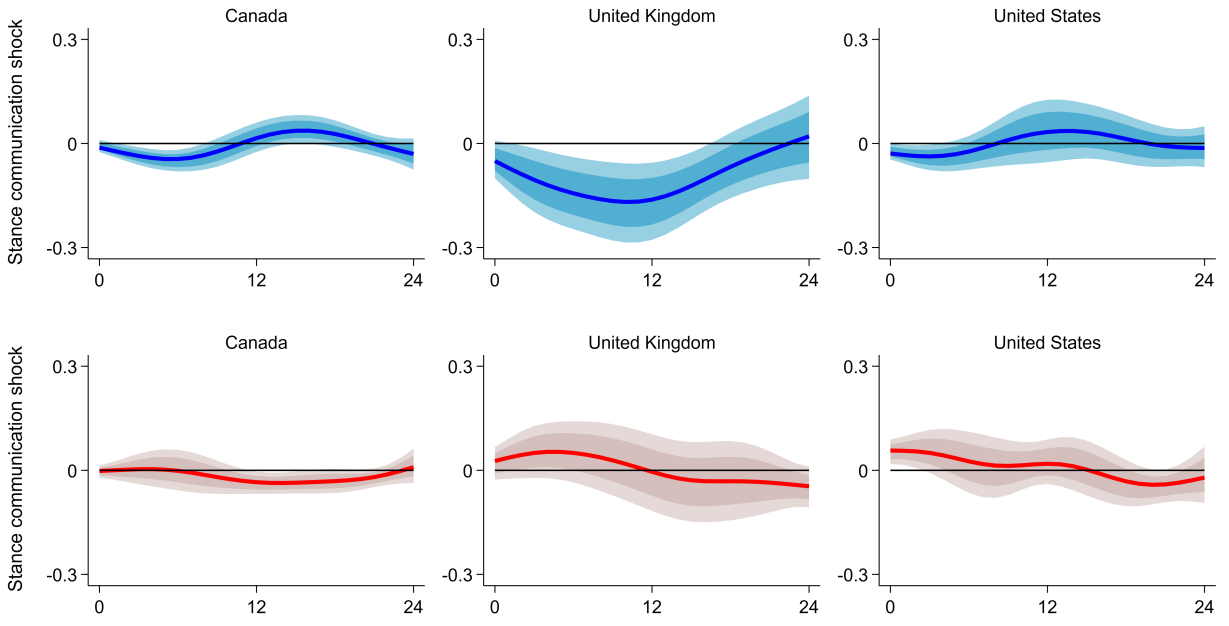

Figure 14: Robustness check 8: professional inflation expectations

Notes: This figure reports impulse responses of one-year-ahead professional inflation expectations to narrative communication shocks. Professional forecasts are drawn from Consensus Economics. The specification is otherwise the same as in the baseline. Shaded areas denote 68% and 90% confidence intervals.

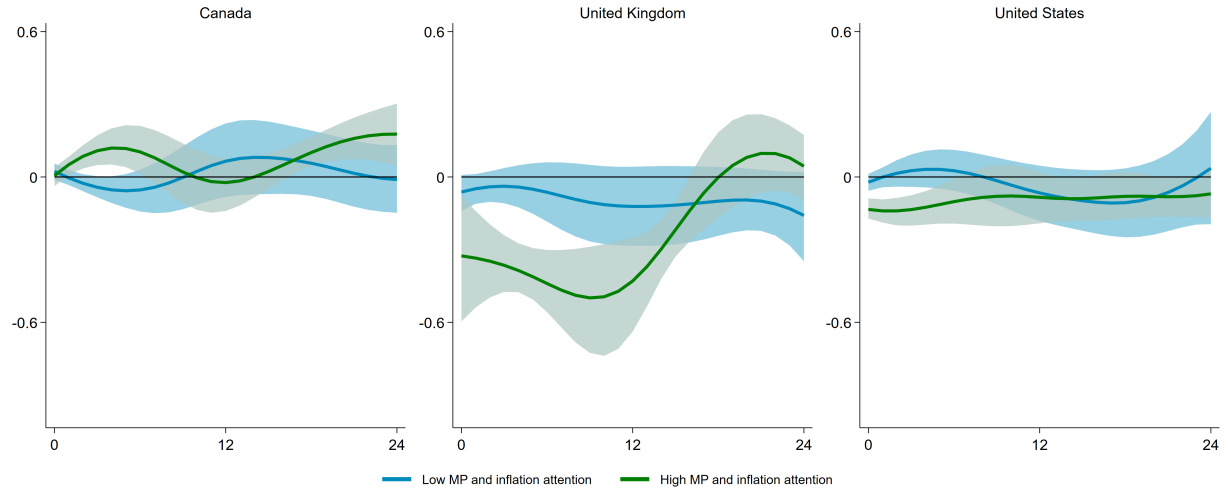

Figure 15: Robustness check 9a: binary attention states, stance communication shock

Notes: This figure reports impulse responses of one-year-ahead household inflation expectations to stance communication shocks, allowing the response to differ across low- and high-attention states. The attention state is based on the standardised monetary-policy and inflation attention index. High attention is defined as months in which the index is above zero in the month before the shock; low attention is defined as months in which the index is below zero. Shocks are normalised as in the baseline specification. Shaded areas denote confidence intervals.

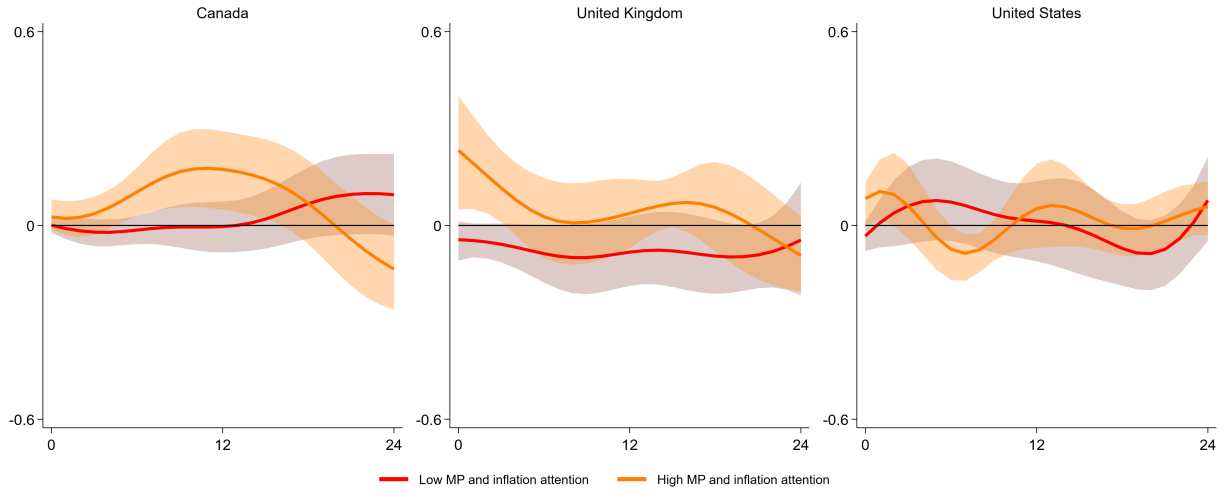

Figure 16: Robustness check 9b: binary attention states, information communication shock

Notes: This figure reports impulse responses of one-year-ahead household inflation expectations to information communication shocks, allowing the response to differ across low- and high-attention states. The attention state is based on the standardised monetary-policy and inflation attention index. High attention is defined as months in which the index is above zero in the month before the shock; low attention is defined as months in which the index is below zero. Shocks are normalised as in the baseline specification. Shaded areas denote confidence intervals.

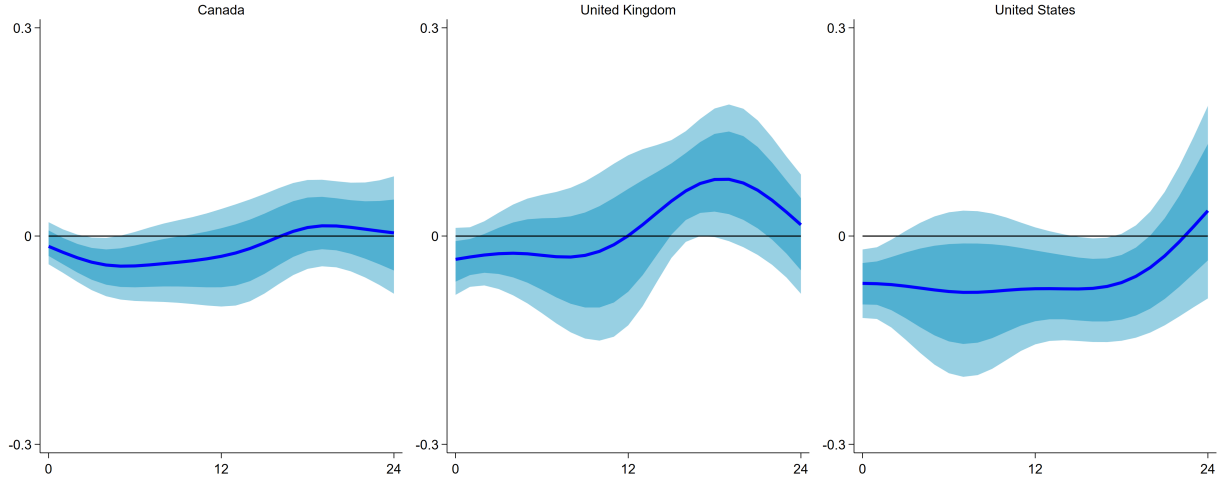

Figure 17: Robustness check 10: responses to unrotated narrative surprises

Notes: This figure reports impulse responses of one-year-ahead household inflation expectations to unrotated narrative surprises. The narrative surprise is defined as the difference between the central bank narrative and the expected narrative reflected in pre-announcement newspaper coverage. Unlike the baseline specification, the narrative surprise is not decomposed into stance communication shocks and information communication shocks. The specification is otherwise the same as in the baseline. Shaded areas denote 68% and 90% confidence intervals.

B. Baseline dictionary

B.1. Dictionary construction

Our baseline dictionary builds on the central bank communication dictionary of Apel et al. (2022), which classifies hawkish and dovish tone by searching for modifier words around terms related to inflation, economic activity, and employment. We use their modifier dictionaries as the starting point and retain their search logic: a modifier is counted when it appears within seven tokens of a topical term in the same sentence, with negation terms reversing the polarity of the modifier. We adapt the topical terms by adding relevant synonyms from Aruoba and Drechsel (2025). In the implementation, spelling variants are adjusted for Canadian, British, and American English where relevant to the country-specific corpus. For expositional consistency, the dictionary is documented below using British English.

B.2. Topical categories and trigger words

Each topical phrase is treated as a trigger term. Multi-word phrases are collapsed into single tokens connected by underscores before tokenisation.

Consumer prices

- 'consumer prices', 'consumer price', 'cpi'

Inflation

- 'inflation'

Inflation pressures

- 'inflation pressure', 'inflation pressures', 'price pressure', 'price pressures'
- 'inflationary pressure', 'inflationary pressures'
- 'wage pressure', 'wage pressures'
- 'employment cost', 'employment costs'
- 'unit labour cost', 'unit labour costs'

Consumer spending

- 'consumer spending', 'consumption', 'household spending'

Economic activity

- 'economic activity', 'output', 'gdp', 'gross domestic product', 'industrial production'

Economic growth

- 'economic growth'

Resource utilisation

- 'utilisation'

Employment

- 'employment', 'hiring', 'hires', 'jobs', 'payrolls'
- 'labour force participation', 'hours'

Labour markets

- 'labour market', 'labour markets'

Unemployment

- 'unemployment'

B.3. Implementation

The implementation follows Apel et al. (2022) closely. Text is converted to lower case, non-standard apostrophes and whitespace are normalised, and sentences are identified using the NLTK sentence tokenizer with a rule-based fallback. Multi-word trigger phrases are first linked with underscores, so that phrases such as 'labour market' are treated as single tokens. Modifier expressions are implemented as regular expressions, allowing stemmed forms such as `increas\w*`. For each trigger term, we search up to seven tokens on either side within the same sentence. When a negation term appears within seven tokens of the modifier, the modifier's polarity is reversed.

C. FEVD algorithm

Formally, the h-step-ahead forecast error of the endogenous variable is defined as:

$$f_{t+h|t-1} \equiv y_{t+h} - P\left[\, y_{t+h} \,\middle|\, \Omega_{t-1} \right], \tag{6}$$

where $P\left[\, y_{t+h} \,\middle|\, \Omega_{t-1} \right]$ denotes the projection of y_{t+h} on the information set Ω_{t-1}, which includes lags of the dependent variable, the identified shocks, and additional controls.

We then decompose the forecast error into contributions from the identified shocks and other sources of variation:

$$f_{t+h|t-1} = \theta_0 Shocks_{t+h} + \cdots + \theta_h Shocks_t + \nu_{t+h|t-1}, \tag{7}$$

where $\nu_{t+h|t-1}$ captures innovations orthogonal to $\{Shocks_t, \ldots, Shocks_{t+h}\}$ and the information set Ω_{t-1}.

The R^2 statistic is then used to measure the share of the variance of forecast errors explained by current and future shocks:

$$s_h = \frac{\operatorname{var}\left(\theta_0 Shocks_{t+h} + \cdots + \theta_h Shocks_t\right)}{\operatorname{var}\left(f_{t+h|t-1}\right)}. \tag{8}$$

By construction, $\hat{s}_h$ lies between zero and one. We also ensure that any serial correlation and predictable variation in $\{Shocks_t, \ldots, Shocks_{t+h}\}$ is removed. Confidence bands are obtained using the block bootstrap for local projections proposed by Kilian and Kim (2011).

Table 5
Adaptation of Husted's monetary policy uncertainty index

Country	Monetary-policy terms	Newspapers
Canada	'monetary policy' or 'monetary policies'; 'interest rate' or 'interest rates'; 'overnight rate'; 'target for the overnight rate'; 'target overnight rate'; 'policy interest rate'; 'key policy interest rate'; 'key interest rate'; 'bank rate'	*National Post*; *The Globe and Mail*; *Toronto Star*
United Kingdom	'monetary policy' or 'monetary policies'; 'interest rate' or 'interest rates'; 'bank rate'; 'official bank rate'	*The Daily Telegraph*; *Financial Times*; *The Times*; *Daily Mail*; *The Independent*; *The Guardian*

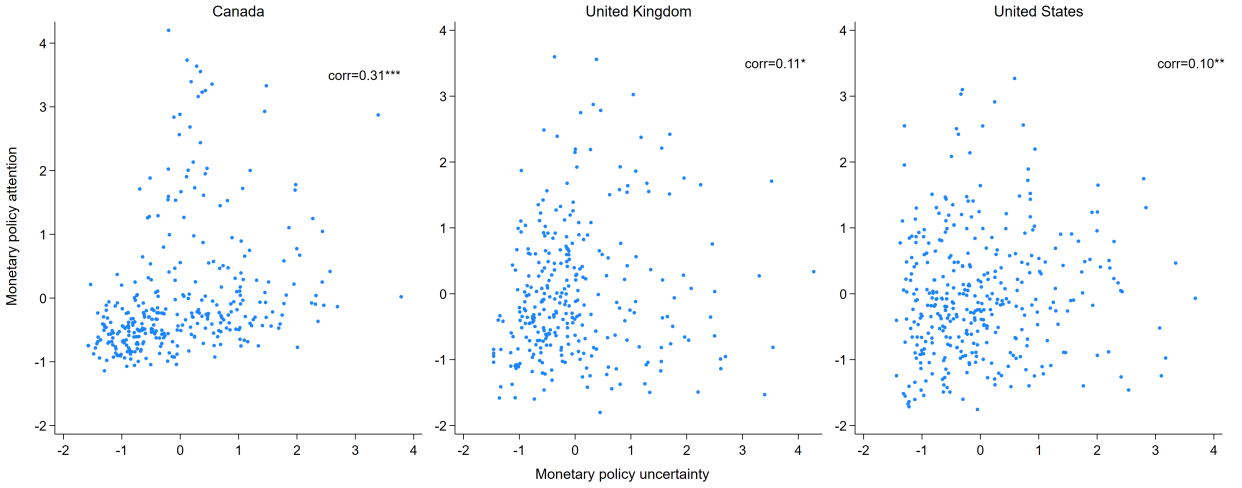

Figure 18: Correlation between monetary policy uncertainty and attention

D. Monetary policy uncertainty

We follow Husted et al. (2020), who identify articles containing terms from each of three categories: (i) 'uncertainty' or 'uncertain'; (ii) monetary-policy terms, such as 'monetary policy', 'interest rate', 'Federal funds rate', or 'Fed funds rate'; and (iii) central-bank terms, such as 'Federal Reserve', 'the Fed', 'Federal Open Market Committee', or 'FOMC'. For each newspaper and period, the number of articles satisfying the three-part rule is divided by the total number of articles containing category-(iii) terms. The resulting newspaper-level shares are standardised to have unit standard deviation over the sample. The aggregate index is obtained by summing the standardised newspaper-level series and rescaling the sum to have mean 100.

We construct analogous MPU indices for Canada and the United Kingdom. The adaptation preserves the structure of Husted et al. (2020): we retain the uncertainty component, replace the U.S.-specific monetary-policy and central-bank terminology with country-specific equivalents, and use major national newspapers in each country. Table 5 reports the monetary-policy terms and newspapers used in the adaptation.

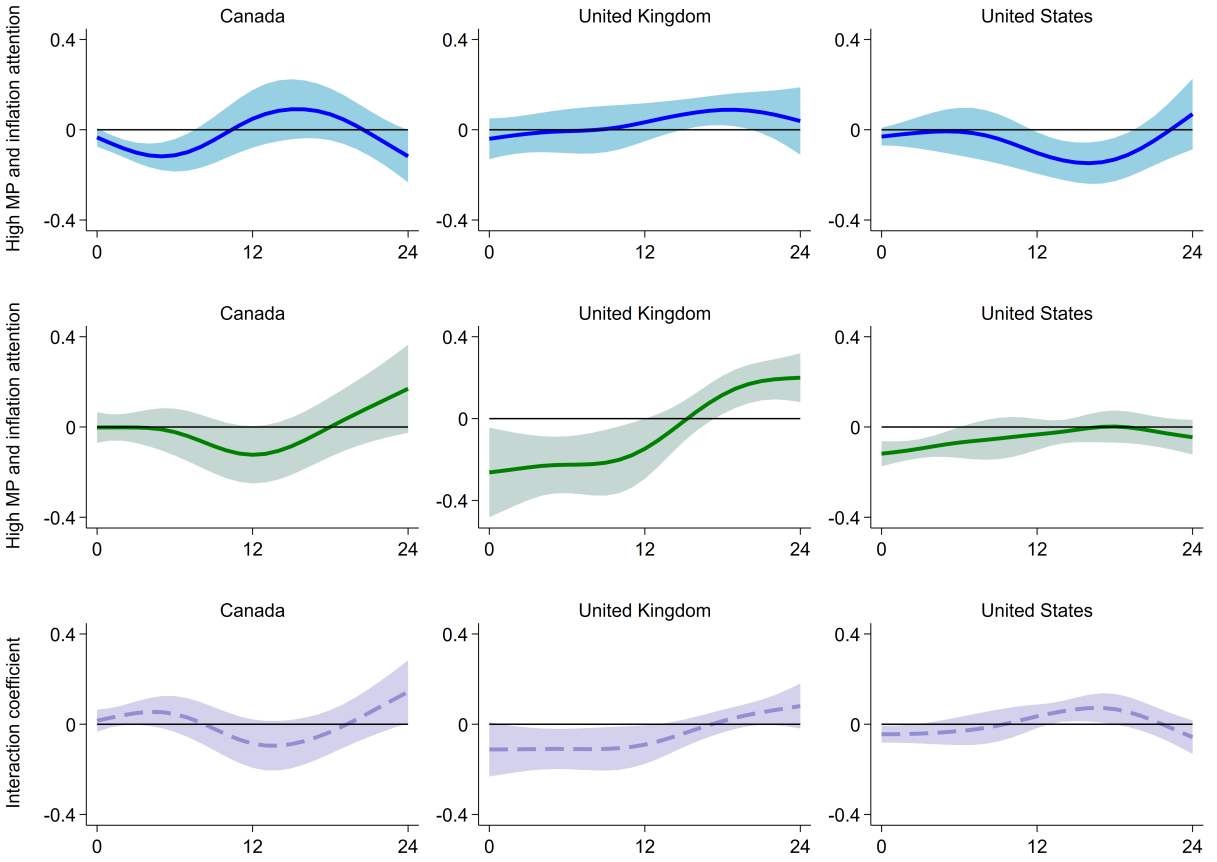

Figure 19: Robustness: state-dependent effects of policy-stance communication shocks controlling for monetary policy uncertainty

Notes: The figure reports responses of one-year-ahead household inflation expectations to policy-stance narrative shocks, conditional on monetary-policy and inflation attention, controlling for monetary policy uncertainty. Low and high attention refer to responses evaluated at one standard deviation below and above the country-specific mean of the standardised attention index, respectively. The bottom row reports the interaction coefficient. Shaded areas denote 90% confidence intervals.

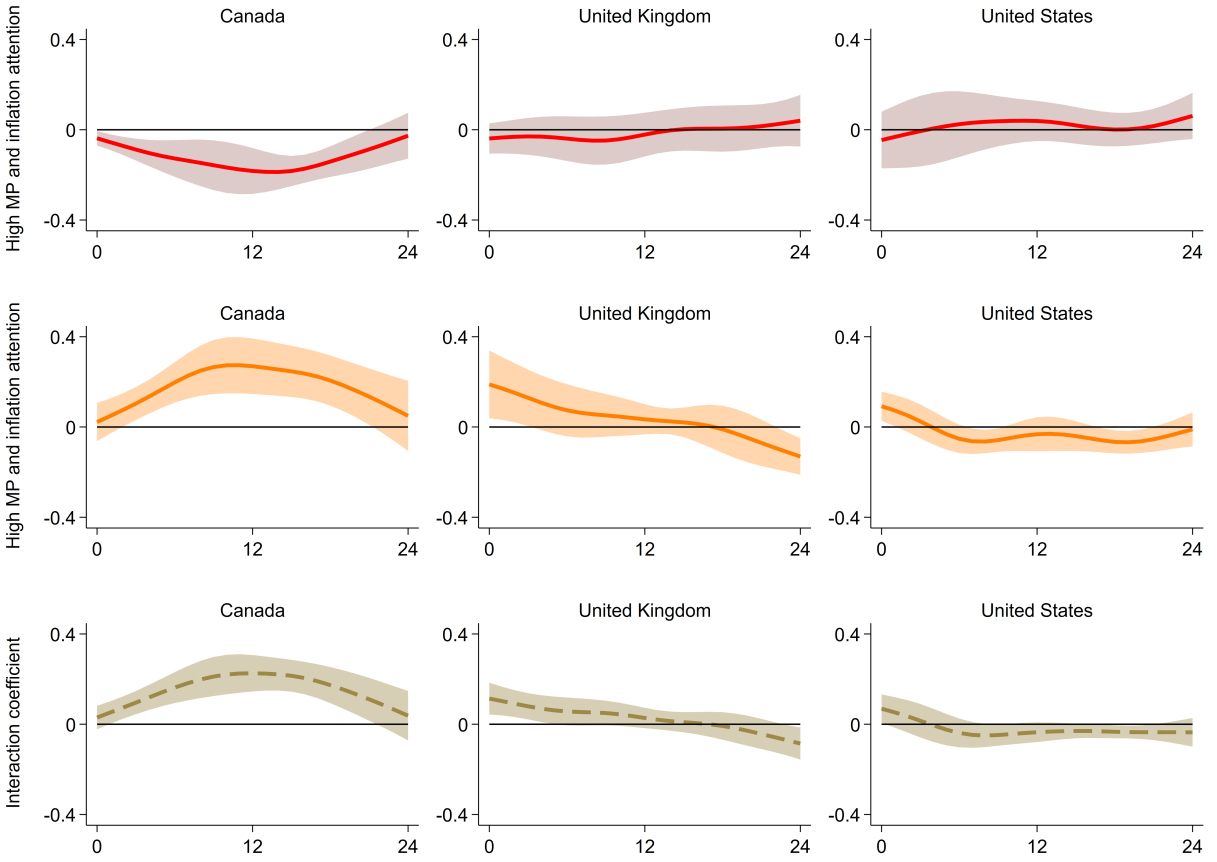

Figure 20: Robustness: state-dependent effects of information communication shocks controlling for monetary policy uncertainty

Notes: The figure reports responses of one-year-ahead household inflation expectations to information narrative shocks, conditional on monetary-policy and inflation attention, controlling for monetary policy uncertainty. Low and high attention refer to responses evaluated at one standard deviation below and above the country-specific mean of the standardised attention index, respectively. The bottom row reports the interaction coefficient. Shaded areas denote 90% confidence intervals.

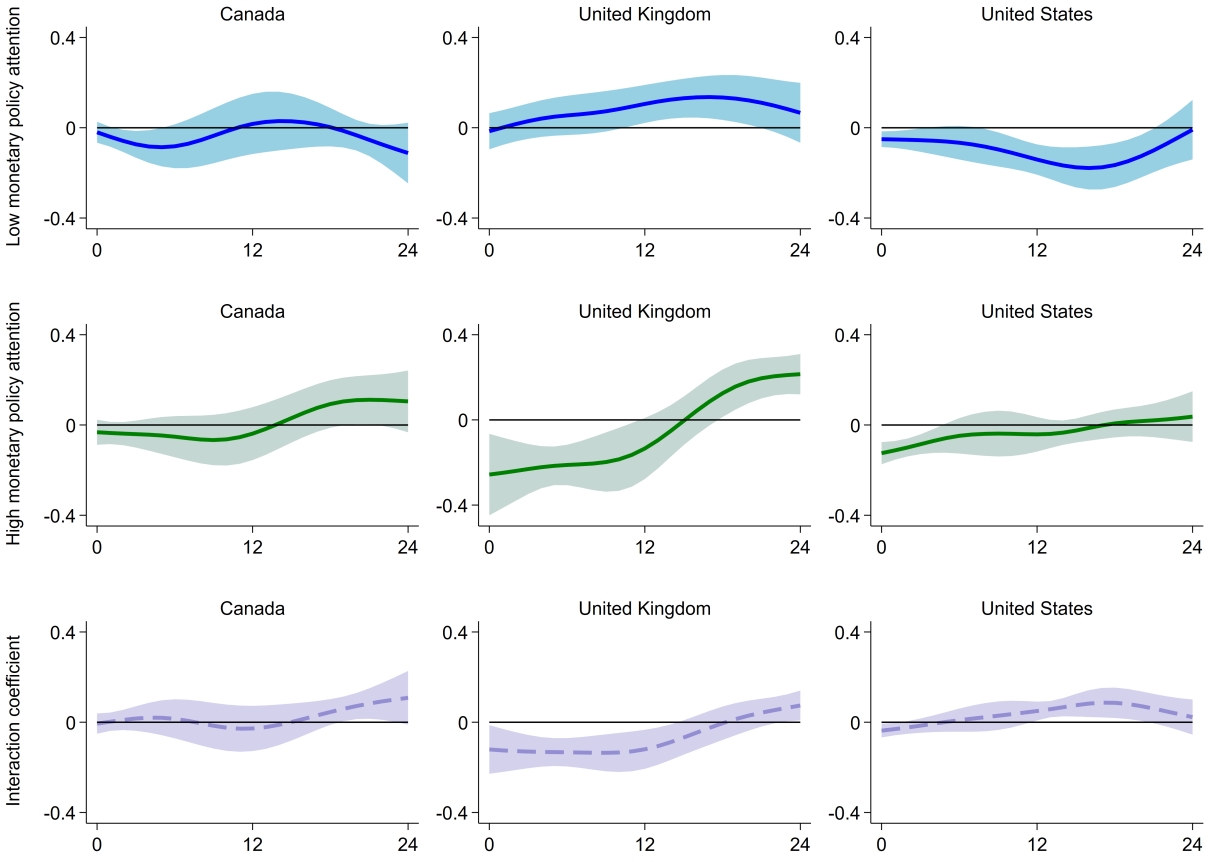

Figure 21: State-dependent effects of policy-stance communication shocks: monetary-policy attention

Notes: The figure reports responses of one-year-ahead household inflation expectations to policy-stance narrative shocks, conditional on monetary-policy attention. Low and high attention refer to responses evaluated at one standard deviation below and above the country-specific mean of the standardised attention index, respectively. The bottom row reports the interaction coefficient. Shaded areas denote 90% confidence intervals.

E. Pure monetary policy attention

F. Predictors used to purge narrative surprises

This appendix describes the predictors used to purge narrative surprises of predictable components. The purpose of this step is to ensure that the communication shocks are identified from the unexpected component of central bank narratives rather than from variation that could have been anticipated from the pre-announcement information set.

Bauer and Swanson (2023a), Bauer and Swanson (2023b), and Swanson (2024) use different predictor sets. To be conservative, we use the union of their predictors as the predictor set. The predictor set includes the following variables:

- *Macroeconomic surprises*: unemployment rate, quarterly GDP growth, and annualised CPI inflation surprises, measured as the released value less the Bloomberg median expectation.

Table 6
Search Terms for Inflation Attention Measures

Country	Country-name variants	Inflation terms
Canada	Canada	inflation; consumer price index; industrial product price index; producer price index
United Kingdom	U.K.; United Kingdom; UK	inflation; consumer price index; consumer prices index; producer price index; producer prices index
United States	U.S.; United States	inflation; consumer price index; producer price index

Notes: The table reports the country-name variants and inflation terms used to construct the inflation attention index. An article is classified as inflation-related if it contains at least one inflation term and at least one country-name variant. Searches are case-insensitive. The newspaper sources are *National Post* and *The Globe and Mail* for Canada, *Financial Times* and *The Times* for the United Kingdom, and *The New York Times* and *The Wall Street Journal* for the United States.

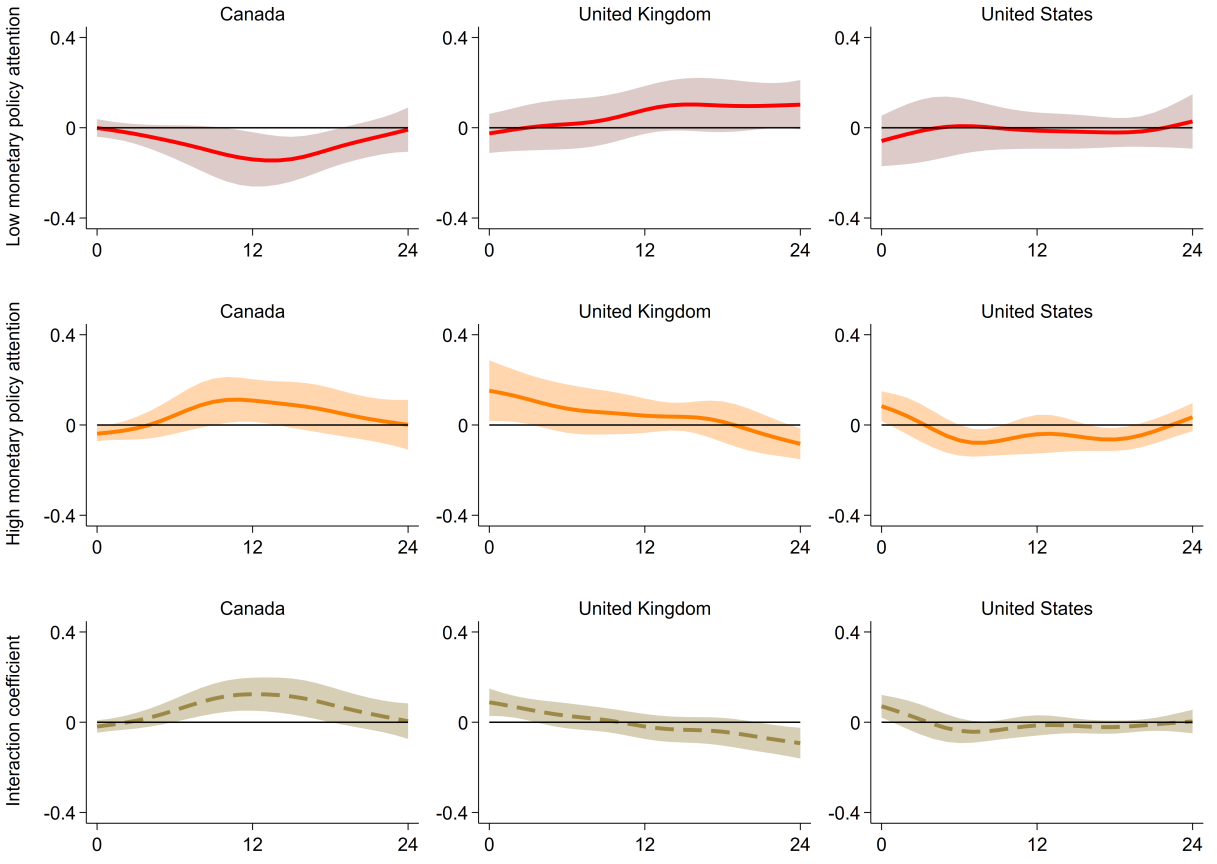

Figure 22: State-dependent effects of information communication shocks: monetary-policy attention

Notes: The figure reports responses of one-year-ahead household inflation expectations to information narrative shocks, conditional on monetary-policy attention. Low and high attention refer to responses evaluated at one standard deviation below and above the country-specific mean of the standardised attention index, respectively. The bottom row reports the interaction coefficient. Shaded areas denote 90% confidence intervals.

- *Expected inflation*: the one-year-ahead Consensus Forecasts mean forecast of year-on-year CPI inflation.
- *Recent inflation dynamics*: the change in six-month core CPI inflation relative to six months earlier.
- *Labour-market conditions*: annual employment growth, following Cieslak (2018).
- *Equity-market performance*: the local stock-market return from 65 trading days before the central bank event to the day before the event.
- *Yield-curve movements*: the change in the second principal component of country-specific zero-coupon yields over the three months before the event.
- *Commodity prices*: the return on the local commodity price index over the 65 trading days before the event.[22]
- *Government bond yields*: the three-month changes in two-year and ten-year government bond yields.
- *Financial conditions*: the one-month change in the Chicago Fed National Financial Conditions Index for the United States and the Goldman Sachs Financial Conditions Index for Canada and the United Kingdom.

All predictors are measured using information available before the monetary policy event. The residualisation is conducted separately by country, so that the orthogonalised narrative surprises are purged only of predictable variation in each country's own pre-event macroeconomic and financial information set.

[22] We use the Bank of Canada Commodity Price Index for Canada and the Bloomberg Commodity Spot Price Index for the United Kingdom and the United States.

Acknowledgements

We are grateful to David Elliott, Andrew Gimber, Simon Lloyd, Ricardo Reis, Leif Thorsrud, and colleagues at the Bank of England for comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily represent those of the Bank of England, the Reserve Bank of New Zealand, or any of their committees.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used OpenAI's ChatGPT to support editing and LaTeX formatting. The authors reviewed and edited the output as needed and take full responsibility for the content of the published article.

References

Acosta, M. (2023). The Perceived Causes of Monetary Policy Surprises. https://www.acostamiguel.com/papers/acosta_mp.pdf

An, Z., Abo-Zaid, S., and Sheng, X. S. (2023). Inattention and the Impact of Monetary Policy. *Journal of Applied Econometrics*, 38(4), 623–643.

Anderson, G., & Maule, B. (2014). Assessing the Risk to Inflation from Inflation Expectations. *Bank of England Quarterly Bulletin*, 2014 Q2. https://www.bankofengland.co.uk/-/media/boe/files/quarterly-bulletin/2014/assessing-the-risk-to-inflation-from-inflation-expectations.pdf

Apel, M., Blix Grimaldi, M., Hull, I., & Lindé, J. (2022). How much information do monetary policy committees disclose? Evidence from the FOMC and the Riksbank. *Journal of Money, Credit and Banking*, *54*(6), 1491–1533. doi.org/10.1111/jmcb.12885

Arioli, R., Brugiavini, A., D'Acunto, F., & Weber, M. (2017). EU consumers' quantitative inflation perceptions and expectations. ECB Occasional Paper No. 186. https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op186.en.pdf

Aruoba, B., & Drechsel, T. (2025). Identifying Monetary Policy Shocks: A Natural Language Approach. *American Economic Journal: Macroeconomics*. doi:10.3386/w32417

Ascari, G., & Haber, T. (2022). Non-linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy. *The Economic Journal*, *132*(641), 37–57. doi:10.1093/ej/ueab049

Baker, S., Bloom, N., & Davis, S. (2016). Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics*, *131*(4), 1593–1636. doi:10.1093/qje/qjw024

Barnichon, R., & Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *Review of Economics and Statistics*, *101*(3), 522–530. doi:10.1162/rest_a_00778

Bauer, M., & Swanson, E. (2023a). An Alternative Explanation for the "Fed Information Effect." *American Economic Review*, *113*(3), 664–700. doi:10.1257/aer.20201220

Bauer, M., & Swanson, E. (2023b). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual*, *37*(1), 77–138. doi:10.1086/723574

Bernanke, B. S., & Kuttner, K. N. (2005). What Explains the Stock Market's Reaction to Federal Reserve Policy? *Journal of Finance*, *60*(3), 1221–1257. doi:10.1111/j.1540-6261.2005.00760.x

Bernanke, B. S. (2007). Inflation Expectations and Inflation Forecasting. Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts, July 10. Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm

Blinder, A., Ehrmann, M., de Haan, J., & Jansen, D.-J. (2024). Central bank communication with the general public: Promise or false hope? *Journal of Economic Literature*, *62*(2), 425–457. doi:10.1257/jel.20231683

Brennan, C. M., Jacobson, M. M., Matthes, C., & Walker, T. B. (2024). Monetary Policy Shocks: Data or Methods? FEDS Working Paper No. 2024-011R1. doi:10.17016/FEDS.2024.011r1

Bu, C., Rogers, J., & Wu, W. (2021). A Unified Measure of Fed Monetary Policy Shocks. *Journal of Monetary Economics*, *118*, 331–349. doi:10.1016/j.jmoneco.2020.11.002

Bullard, J. (2016). Inflation Expectations Are Important to Central Bankers, Too. *The Regional Economist*, Federal Reserve Bank of St. Louis, April. https://www.stlouisfed.org/publications/regional-economist/april-2016/inflation-expectations-are-important-to-central-bankers-too

Caldara, D., & Iacoviello, M. (2022). Measuring Geopolitical Risk. *American Economic Review*, *112*(4), 1194–1225. doi:10.1257/aer.20191823

Carroll, C. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics*, *118*(1), 269–298. doi:10.1162/00335530360535207

Cieslak, A. (2018). Short-Rate Expectations and Unexpected Returns in Treasury Bonds. *Review of Financial Studies*, *31*(9), 3265–3306. doi:10.1093/rfs/hhy051

Cieslak, A., & McMahon, M. (2024). Tough Talk: The Fed and the Risk Premium. *CEPR Discussion Papers*, No. 19313. cepr.org/publications/dp19313

Coibion, O., & Gorodnichenko, Y. (2026). Inflation Expectations and Monetary Policy: What Have We Learned and to What End? *The Econometrics Journal*, utag004. doi:10.1093/ectj/utag004

Coibion, O., Gorodnichenko, Y., & Kamdar, R. (2018). The Formation of Expectations, Inflation, and the Phillips Curve. *Journal of Economic Literature*, *56*(4), 1447–1491. doi:10.1257/jel.20171300

Coibion, O., Gorodnichenko, Y., & Weber, M. (2022). Monetary Policy Communications and Their Effects on Household Inflation Expectations. *Journal of Political Economy*, *130*(6), 1537–1584. doi:10.1086/718982

D'Acunto, F., Charalambakis, E., Georgarakos, D., Kenny, G., Meyer, J., & Weber, M. (2024). Household Inflation Expectations: An Overview of Recent Insights for Monetary Policy. *NBER Working Paper*, No. 32488. National Bureau of Economic Research. https://www.nber.org/papers/w32488

De Fiore, F., Maurin, A., Mijakovic, A., & Sandri, D. (2024). Monetary Policy in the News: Communication Pass-Through and Inflation Expectations. *BIS Working Papers*, No. 1231.

Desroches, B., Kozicki, S., & Simon, L. (2024). Monetary Policy Governance: Bank of Canada Practices to Support Policy Effectiveness. *Bank of Canada Staff Discussion Paper*, 2024–14. doi:10.34989/sdp-2024-14

Fama, E. F., & MacBeth, J. D. (1973). Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy*, *81*(3), 607–636. doi:10.1086/260061

Fisher, A., Martineau, C., and Sheng, J. (2022). Macroeconomic Attention and Announcement Risk Premia. *Review of Financial Studies*, 35(11), 5057–5093.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications* (2nd ed.). Princeton University Press.

Georgiadis, G., & Jarociński, M. (2025). Global Spillovers from Multi-Dimensional US Monetary Policy. ECB Working Paper No. 2881. https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2881 a2c5c8070a.en.pdf

Gertler, M., & Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, *7*(1), 44–76. https://doi.org/10.1257/mac.20130329

Gorodnichenko, Y., & Lee, B. (2020). Forecast Error Variance Decompositions with Local Projections. *Journal of Business & Economic Statistics*, *38*(4), 921–933. doi:10.1080/07350015.2019.1610661

Gürkaynak, R. S., Sack, B., & 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.

Gürkaynak, R. S., Sack, B. P., & Swanson, E. T. (2007). Market-Based Measures of Monetary Policy Expectations. *Journal of Business & Economic Statistics*, *25*(2), 201–212. https://www.jstor.org/stable/27638925

Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., & Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour*, *5*, 529–538. doi:10.1038/s41562-021-01079-8

Hansen, A. L., & Kazinnik, S. (2024). Can ChatGPT Decipher Fedspeak? SSRN Working Paper. https://ssrn.com/abstract=4399406

Hubert, P., & Labondance, F. (2021). The Signaling Effects of Central Bank Tone. *European Economic Review*, *133(C)*, 103684. doi:10.1016/j.euroecorev.2021.103684

Husted, L., Rogers, J., & Sun, B. (2020). Monetary Policy Uncertainty. *Journal of Monetary Economics*, *115*, 20–36. doi:10.1016/j.jmoneco.2019.07.009

Jarociński, M. (2022). Central Bank Information Effects and Transatlantic Spillovers. *Journal of International Economics*, *139*. doi:10.1016/j.jinteco.2022.103683

Jarociński, M., & Karadi, P. (2018). Deconstructing Monetary Policy Surprises: The Role of Information Shocks. ECB Working Paper No. 2133. https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2133.en.pdf

Jarociński, M., & Karadi, P. (2020). Deconstructing Monetary Policy Surprises—The Role of Information Shocks. *American Economic Journal: Macroeconomics*, *12*(2), 1–43. doi:10.1257/mac.20180090

Jordà, Ò., & Taylor, A. (2025). Local Projections. *Journal of Economic Literature*, *63*(1), 59–110. doi:10.1257/jel.20241521

Kalamara, E., Turrell, A., Redl, C., Kapetanios, G., & Theodoridis, K. (2022). Making text count: Economic forecasting using newspaper text. *Journal of Applied Econometrics*, *37*(5), 896–919. doi:10.1002/jae.2907

Kilian, L., & Kim, Y. J. (2011). How Reliable Are Local Projection Estimators of Impulse Responses? *Review of Economics and Statistics*, *93*(4), 1460–1466. doi:10.1162/REST_a_00143

Koop, C., & Scotto di Vettimo, M. (2023). How do the media scrutinise central banking? Evidence from the Bank of England. *European Journal of Political Economy*, *77*, 102296. doi:10.1016/j.ejpoleco.2022.102296

Korinek, A. (2023). Generative AI for economic research: Use cases and implications for economists. *Journal of Economic Literature*, 61(4), 1281–1317. doi:10.1257/jel.20231736

Kuttner, K. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, *47*(3), 523–544. doi:10.1016/S0304-3932(01)00055-1

Lamla, M. J., & Maag, T. (2012). The Role of Media for Inflation Forecast Disagreement of Households and Professional Forecasters. *Journal of Money, Credit and Banking*, *44*(7), 1325–1350.

Malmendier, U., Nagel, S., & Yan, Z. (2021). The Making of Hawks and Doves: Inflation Experiences on the FOMC. *Journal of Monetary Economics*, *117*, 19–42. doi:10.1016/j.jmoneco.2020.04.002

Mehrotra, A., & Yetman, J. (2018). Decaying Expectations: What Inflation Forecasts Tell Us about the Anchoring of Inflation Expectations. *International Journal of Central Banking*, *14*(5), 55–101. http://www.ijcb.org/journal/ijcb18q4a2.pdf

Munday, T., & Brookes, J. (2021). Mark My Words: The Transmission of Central Bank Communication to the General Public via the Print Media. *Bank of England Staff Working Paper*, No. 944.

Nakamura, E., & Steinsson, J. (2018). High-Frequency Identification of Monetary Non-Neutrality: The Information Effect. *Quarterly Journal of Economics*, *133*(3), 1283–1330. jstor.org/stable/26864975

Pfäuti, O. (2025). Inflation—Who Cares? Monetary Policy in Times of Low Attention. *Journal of Money, Credit and Banking*, 57(5), 1211–1240.

Plagborg-Møller, M., & Wolf, C. K. (2017). Instrumental Variable Identification of Dynamic Variance Decompositions. Manuscript. economics.mit.edu/sites/default/files/publications/decomp_iv.pdf

Plagborg-Møller, M., & Wolf, C. K. (2022). Instrumental Variable Identification of Dynamic Variance Decompositions. *Journal of Political Economy*, *130*(8), 2164–2202. doi:10.1086/720141

Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. In J. B. Taylor & H. Uhlig (Eds.), *Handbook of Macroeconomics*, Vol. 2 (pp. 71–162). Elsevier. doi:10.1016/bs.hesmac.2016.03.003

Ramey, V. A., & Zubairy, S. (2018). Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data. *Journal of Political Economy*, *126*(2), 850–901. doi:10.1086/696277

Rigobon, R., & Sack, B. (2004). The impact of monetary policy on asset prices. *Journal of Monetary Economics*, *51*(8), 1553–1575. doi:10.1016/j.jmoneco.2004.02.004

Shapiro, A., & Wilson, D. (2022). Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives using Text Analysis. *Review of Economic Studies*, *89*(5), 2768–2805. https://doi.org/10.1093/restud/rdab094

Sims, C. A. (2009). Inflation Expectations, Uncertainty and Monetary Policy. *BIS Working Paper No. 275*. Bank for International Settlements. https://www.bis.org/publ/work275.htm

Song, W., and Stern, S. (2025). Firm Inattention and the Efficacy of Monetary Policy: A Text-Based Approach. *Review of Economic Studies*, 92(5), 3438–3469.

Swanson, E. T. (2024). The Macroeconomic Effects of the Federal Reserve's Conventional and Unconventional Monetary Policies. *IMF Economic Review*, *72*(3), 1152–1184. doi:10.1057/s41308-024-00250-1

Swanson, E. T., & Jayawickrema, V. (2024). Speeches by the Fed Chair Are More Important Than FOMC Announcements: An Improved High-Frequency Measure of U.S. Monetary Policy Shocks. Working paper, version 2.0, March 27.

Ter Ellen, S., Larsen, V., & Thorsrud, L. (2021). Narrative monetary policy surprises and the media. *Journal of Money, Credit and Banking*, *54*(5), 1525–1549. doi:10.1111/jmcb.12868

Weber, M., D'Acunto, F., Gorodnichenko, Y., & Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, *36*(3), 157–184. doi:10.1257/jep.36.3.157

Yetman, J. (2018). The Perils of Approximating Fixed-Horizon Inflation Forecasts with Fixed-Event Forecasts. *BIS Working Paper No. 700*. Bank for International Settlements, Monetary and Economic Department. https://www.bis.org/publ/work700.pdf