Central Bank Communication Design and Public Attention

FK [*]

University of Oxford

DRAFT : March 1, 2026

Abstract

We study whether the structure of central bank communication events affects public engagement with monetary policy. The Bank of England's 2015 Super Thursday reform replaced a two-occasion quarterly forecast round, in which the rate decision and the forecast report arrived on separate days, with a single bundled release. Using Twitter (X), Google Trends, and Wikipedia as measures of revealed attention, we find that bundled days are individually more salient. Forecast occasions attract significantly more attention after co-location on both Twitter and Google Trends, and full-release meetings outperform standard meetings within the post-reform period. However, consolidating two occasions into one reduces cumulative quarterly Twitter engagement by about 41 per cent, because the per-event salience gain does not offset the loss of a separate attention occasion. Internal-control diagnostics confirm that this quarterly decline reflects an event-count effect rather than a shift in relative forecast salience. Bundling also changes what the public attends to, raising the share of outlook-related discussion on full-release days.

Keywords: central bank communication, public attention, communication architecture, rational inattention, monetary policy

JEL Codes: E58, D83, G14, G53

[*] Contact, fatih.kansoy@economics.ox.ac.uk

1. Introduction

How should central banks *structure* their communication calendars if they want the public to engage with monetary policy? Over the last two decades, central banks have invested heavily in improving the *content* of communication, including clearer language, new channels, and richer explanations. Yet far less is known about the *architecture* of communication events, namely the timing, sequencing, and bundling of distinct information releases.

The Bank of England's 2015 communication reform offers a natural experiment in architecture. Following an independent review of the Bank's transparency practices — the Warsh Review (Bank of England 2014) — the Bank announced in December 2014 that it would consolidate the MPC rate decision, the minutes and vote, the quarterly forecast publication (then the *Inflation Report*, later the *Monetary Policy Report*), and the Governor's press conference into a single event day. The bundled format, implemented from August 2015 and quickly labelled "Super Thursday," replaced a pre-reform architecture in which quarterly forecast rounds comprised two distinct communication occasions separated by several days, a decision day and a forecast/press-conference day. The stated objective was greater transparency and coherence. Our evidence indicates a trade-off between event-level salience and cumulative forecast-round engagement.

The core economic tension is between two margins of attention. Bundling can increase the *intensive margin* because when the forecast and decision land on the same day, that day is more salient and the information package is richer. Yet bundling can also reduce the *extensive margin*, because collapsing two occasions into one removes an independent opportunity for the public to cross an attention threshold. Rational inattention theories accommodate both forces and therefore do not deliver a sharp prediction about the net effect (Sims 2003, Reis 2006, Paciello and Wiederholt 2014). Which margin dominates is an empirical question.

We measure revealed public attention using three digital traces that capture different facets of engagement, namely real-time discussion on Twitter (daily unique users posting policy-relevant tweets), active information seeking through Google Trends, and deliberate reference behaviour through Wikipedia page views. For each platform, we construct abnormal attention around MPC events using a day-of-week-matched rolling baseline, and we aggregate these event-level measures to the quarterly forecast-round level for the bundling analysis.

To identify the architectural margin more directly, we first compare pre-reform Inflation Report days with post-reform full-release days. Both are forecast occasions with forecast publication and press communication, but only the post-reform event co-locates these elements with the rate decision and minutes. We then use the within-post full-release-versus-standard comparison as supporting intensive-margin evidence, and we retain the forecast-round quarterly break as evidence on round-level net effects of the reform package. The quarterly estimate is interpreted as the net effect of August 2015 changes rather than a clean estimate of bundling alone, because simultaneous minutes publication changed at the same date.

In baseline specifications, post-reform full-release days attract more attention than pre-reform Inflation Report days in the targeted forecast-occasion comparison on both Twitter and Google Trends. Within post-reform events, full-release meetings attract significantly more attention than standard meetings. At the quarterly level, cumulative Twitter attention declines after August 2015, and internal-control diagnostics indicate that this decline is consistent with an event-count effect: one forecast-round event replacing two.

The decomposition is descriptive accounting of magnitudes. The average bundled-day gain relative to the pre-reform decision day is +0.232 abnormal log-user units, while the removed pre-reform Inflation Report day contributes +1.456 units on average. Cross-platform quarterly diagnostics are heterogeneous: the Google Trends analogue rises after the reform (+13.766, $p = 0.001$) and Wikipedia has no usable pre-reform quarterly coverage.

Exploratory attention extensions suggest potential asymmetry in the direction of policy news. Hawkish surprises elicit larger point estimates than dovish surprises, but inference relies on thin-tail variation and these tests are not a headline contribution.

These results speak to three literatures. The monetary policy event-study tradition (Kuttner 2001, Gürkaynak, Sack, and Swanson 2005, Braun, Miranda-Agrippino, and Saha 2025) evaluates communication through asset prices; we redirect the outcome of interest to revealed public engagement, which matters because financial market participants and the general public process policy news through different channels. Rational inattention models (Sims 2003, Reis 2006) generate ambiguous predictions about whether bundling should increase or decrease total engagement; our decomposition clarifies the relative magnitudes of extensive and intensive margins in this setting. Finally, the growing literature on central bank communication with the general public (Blinder et al. 2008, 2024, Haldane and McMahon 2018, Haldane, Macaulay, and McMahon 2020, Ehrmann and Wabitsch 2022) has focused on content, channel, and messenger; we show that event *architecture* is an important determinant of engagement.

2. Institutional Background

The Bank of England's communication calendar changed materially in August 2015. Before that date, the MPC communicated the policy decision and the quarterly forecast on separate occasions; after it, those elements were brought together on what became known as Super Thursday. Because our attention measures are daily and our main outcome aggregates engagement *within a quarterly forecast round*, the institutional sequence is central to the empirical design.

Communication architecture before and after Super Thursday. Before August 2015, the MPC met roughly monthly (approximately twelve scheduled meetings per year). On each meeting's decision day, the Bank announced Bank Rate at 12:00 London time. Prior to the reform, the MPC's minutes – including the vote split and a detailed account of the deliberation – were typically released with a delay of roughly two weeks. The immediate public information set on decision day was therefore limited to the headline decision and, in some cases, a short accompanying statement.

Four times per year, MPC communication included a quarterly forecast publication (the *Inflation Report*) and a Governor's press conference. Crucially, in the pre-reform architecture, this forecast package arrived on a *separate date* from the decision. In a typical forecast round the Bank Rate decision occurred on a Thursday, while the Inflation Report and press conference took place the following week (often on a Wednesday). This created two temporally distinct occasions within each quarterly forecast round on which public attention could be triggered, with a decision occasion and a forecast/press-conference occasion.

The Warsh Review and the December 2014 reform package. Against this backdrop, in June 2014 the Bank commissioned an independent review of transparency and communications led by Kevin Warsh. The review's recommendations were adopted in a reform package announced on 11 December 2014 (Bank of England 2014). The key elements for our setting were simultaneous minutes publication and forecast bundling. Simultaneous minutes publication removed the previous two-week delay by releasing the minutes together with the rate decision. Forecast bundling, implemented as Super Thursday, released the Inflation Report on the same day as the decision and minutes, followed by the Governor's press conference.

The package also included two additional changes that matter for interpretation. The MPC moved from twelve to eight scheduled meetings per year from 2016, and the Bank introduced supplementary format changes, including a condensed summary of the Inflation Report and a lagged transcript-release policy. These latter changes do not alter the bundling structure of quarterly forecast rounds.

The simultaneous-minutes reform and forecast bundling took effect together in August 2015 and cannot be separated in the pre-to-post comparison. The meeting-frequency reduction took effect in January 2016 and applies uniformly to all post-reform meetings. The within-post comparison in Section 5.3 estimates the incremental attention associated with adding the forecast package and press conference to the decision day.

The first Super Thursday took place on 6 August 2015 (Bank of England 2015). On that date, the rate decision, minutes, and the Inflation Report were released simultaneously at 12:00 London time, followed by the Governor's press conference at 12:45. Contemporary institutional documentation emphasised the reform's transparency rationale (Bank of England 2014).

In the post-reform period, MPC events fall into two categories. Full-release (bundled) meetings release the decision, minutes/vote information, and the quarterly forecast document on the same day, followed by a press conference. Standard meetings release the decision and minutes simultaneously but without a forecast publication or press conference.

Following the move to eight meetings per year (from 2016 onward), the Bank's schedule features four full-release meetings and four standard meetings each year.

The reform re-timed and re-packaged the components of MPC communication rather than introducing an entirely new type of information. In both regimes, the MPC produces a policy decision, minutes and vote information, and, in forecast rounds, a quarterly forecast narrative with a press conference. What changes is when these components arrive and, in particular, the number of distinct forecast-round occasions on which attention can be triggered.

This distinction motivates two comparisons in our empirical design. The forecast-round (quarterly) comparison evaluates how cumulative engagement across the communication occasions constituting a quarterly round changes when two distinct occasions become one. The within-post comparison measures the incremental attention premium associated with bundling the forecast/press-conference package into the decision day.

Table 1 summarises the pre- and post-reform architecture and the mapping to our event definitions.

TABLE 1. MPC Communication Architecture Before and After Super Thursday

Regime	Occasion type	Components released	Timing (London)	Role in this paper
Pre-reform	Decision day	Rate decision (+ short statement when applicable)	12:00	Event in forecast rounds
Pre-reform	Forecast day	Inflation Report + press conference	typically following week	Event in forecast rounds
Pre-reform	Minutes day	Minutes/vote split (delayed)	~ two weeks later	Not part of baseline forecast-round outcome
Post-reform	Full-release	Decision + minutes + forecast + press conference	12:00 / 12:45	Bundled event (full-release)
Post-reform	Standard	Decision + minutes (no forecast, no press conf.)	12:00	Post-reform comparison group

Notes: The table summarises the communication sequence relevant for the paper's definition of events and forecast rounds. The baseline quarterly outcome ($CumLogAtt_q$) aggregates attention across the decision and forecast occasions that constitute a quarterly forecast round; the delayed pre-reform minutes release is not included in the baseline forecast-round construction. In the post-reform period, full-release events serve as the sole occasion in a quarterly forecast round; standard meetings enter the within-post event-level comparison.

3. Related Literature

The institutional change described above raises questions that existing literatures illuminate but do not directly answer. This paper contributes to research on monetary policy identification and communication by shifting attention from asset prices to *public engagement*, and by focusing on the comparatively understudied architecture of communication events — timing, sequencing, and bundling.

Monetary policy event studies. Event-study research has established how to identify monetary policy news in high-frequency asset prices, but it says comparatively little about how non-market audiences engage with those same events. Kuttner (2001) decomposed policy moves into anticipated and surprise components using futures markets. Gürkaynak, Sack, and Swanson (2005) showed that announcements contain multiple dimensions of news — a current-rate "target" component and a forward-looking "path" component — with the path factor driving most of the bond market response. Subsequent work expanded the approach to additional instruments and jurisdictions (Swanson 2021, Altavilla et al. 2019). For the United Kingdom, Braun, Miranda-Agrippino, and Saha (2025) constructed the UK Monetary Policy Event-Study Database (UKMPD), which provides intraday surprise measures around MPC announcements and spans the Warsh Review reform. Naraina and Sangani (2026) complement this by showing that Fed press conferences can materially move markets, underscoring that the communication *setting* matters independently of the decision itself.

We use the UKMPD surprises as conditioning variables in event-level specifications to separate attention responses to event *format* from responses to the *news content* of a given meeting. However, because the reform alters the timing and packaging of information, the mapping from announcement windows to measured news can itself change across regimes. This motivates our emphasis on within-post comparisons and on forecast-round totals that do not require treating a market-based surprise as invariant to architecture.

Our paper remains close to the event-study tradition in one important respect, it uses sharply timed policy events and a narrow event-study logic. The difference lies in the outcome. Asset prices show how informed traders process announcements. They do not show whether the wider public notices, searches for, or discusses them.

Limited attention and announcement design. Models of costly information processing imply that communication effectiveness depends on exposure opportunities, not only on message content. In rational inattention models, agents face processing constraints and optimally allocate attention across competing signals (Sims 2003, Maćkowiak and Wiederholt 2009, Maćkowiak, Matějka, and Wiederholt 2023). Related frameworks stress that updating may be intermittent rather than continuous, so that the occasions on which information becomes available can matter as much as the content itself (Reis 2006). That distinction is central in our setting. Bringing several pieces of monetary policy news together on one day may raise the salience of that day, but it also removes a separate opportunity for attention to be triggered.

Theory does not settle in advance which force should dominate. Bundling can increase the expected payoff-relevance of a single event day and thereby amplify engagement — an intensive-margin channel closely related to economies-of-scope mechanisms in information acquisition

(Paciello and Wiederholt 2014). Yet bundling can also remove independent occasions on which inattentive agents might engage — an extensive-margin channel consistent with attention-scarcity arguments (Falkinger 2008). In addition, communication that fails to become common knowledge can have attenuated effects on expectations and behaviour (Angeletos and Lian 2018), suggesting that the frequency and distribution of exposures may itself be central to policy transmission.

Our contribution is to provide evidence on how these margins interact in practice, using a real-world communication reform that mechanically alters the number of forecast-round attention occasions.

Public communication and attention measurement. Research on central bank communication has increasingly turned from financial markets to the broader public, but the emphasis has been almost entirely on *what* is said rather than on *when and how* it is released. The classic surveys of Blinder et al. (2008) and Blinder et al. (2024) document extensive work on how statements, minutes, and forward guidance steer market expectations. A newer wave asks whether central banks can communicate effectively with non-expert audiences, and how format and outreach affect understanding and trust (Haldane and McMahon 2018). Experimental and field evidence suggests that simplifying language and adapting presentation can improve comprehension (Bholat et al. 2019, Haldane, Macaulay, and McMahon 2020), and that new channels including social media can broaden reach, albeit with strong selection into engagement (Ehrmann and Wabitsch 2022, Masciandaro, Peia, and Romelli 2023).

Parallel to this, a measurement literature uses digital traces as proxies for revealed attention. Search intensity has been widely employed in finance and macroeconomics (Da, Engelberg, and Gao 2011, Wohlfarth 2018), and recent work applies related tools to monetary and inflation attention (Buelens 2025). Social media data have been used to study central bank outreach and to distinguish extensive-margin engagement (unique participants) from intensive-margin posting activity (Gorodnichenko, Pham, and Talavera 2025, Kansoy and Mundy 2025). We build on this measurement tradition by combining Twitter (X), Google Trends, and Wikipedia into a multi-platform attention framework that captures discussion, active information seeking, and reference behaviour respectively.

What remains less well understood is the role of the communication calendar itself. We use the term *communication architecture* to refer to the timing, sequence, and bundling of releases. Several papers are close to ours in institutional terms. Lamla and Vinogradov (2019) showed that central bank announcements are consequential for non-expert household expectations, but did not examine whether the *structure* of the announcement matters conditional on its content. Hansen, McMahon, and Tong (2019) studied the long-run information effect of Bank of England communication but treated communication events as given rather than designed. Munday and

Brookes (2021) identified Super Thursday as a structural break in print media coverage, providing independent corroboration that the reform altered the information environment. Our paper differs from this work in treating the reform as a change in communication architecture rather than simply a change in news volume, and in measuring behavioural engagement directly through search activity, reference behaviour, and social media participation.

Heterogeneity, state dependence, and asymmetric attention. A growing literature documents that households update beliefs under substantial information frictions and that attention is both heterogeneous and state-dependent. Coibion and Gorodnichenko (2015) provide a parsimonious framework of information rigidity across agent types, and Coibion, Gorodnichenko, and Weber (2022) show in large-scale experiments that information treatments can move household expectations but that effects are partial and decay over time, highlighting the importance of repeated exposure. Related work establishes that households respond more strongly to instrument-focused communication than to abstract targets (D'Acunto et al. 2020), and that attention intensifies in high-inflation environments when monetary policy becomes more obviously relevant to household decisions (Weber et al. 2025). Heterogeneity can also arise from lifetime experience and subjective models of the economy (Malmendier and Nagel 2016, Andre et al. 2022). These findings make the structure of communication events economically important, if exposure is selective and infrequent, then the number and timing of occasions on which policy enters public view are likely to affect how much of it is absorbed.

The direction of news may also matter. Prospect theory and reference-dependent preferences imply that adverse developments receive more weight than equally sized favourable ones (Kahneman and Tversky 1979, Kőszegi and Rabin 2006), and media research documents negativity-biased processing of economic information (Soroka 2006, Bordalo, Gennaioli, and Shleifer 2012). In monetary policy settings, reference dependence can generate asymmetric real responses to contractionary versus expansionary shocks (Santoro et al. 2014). We treat direction-dependent attention responses as an exploratory extension rather than as a core contribution.

Existing research provides powerful tools to measure monetary policy news and rich theories of limited attention, but offers comparatively little evidence on how the *timing and bundling* of communication events shapes public engagement. By combining a well-defined institutional reform with multi-platform measures of revealed attention, we provide evidence that communication architecture is a first-order design margin affecting both the level and composition of public attention to monetary policy.

4. Data

The literature gap identified above maps into a measurement problem of tracking how much and what kind of public attention monetary policy communication attracts, at daily frequency, across distinct digital platforms. Our empirical design combines three categories of data, namely (i) a calendar of MPC communication events, (ii) three platform-specific measures of public attention constructed from digital traces, and (iii) high-frequency financial market data used as diagnostic controls. We work at three nested units of observation. At the finest level, we construct daily attention series on each platform. We then map these daily series into *event-level* outcomes measured on MPC communication dates. Finally, for the bundling analysis we aggregate event-level attention to the *quarterly forecast-round* level, the relevant unit for evaluating the consequences of consolidating two communication occasions into one.

4.1 Event calendar and regime definitions

We define MPC communication events using the official Bank of England announcement calendar, cross-checked against the UK Monetary Policy Event-Study Database (UKMPD; Braun, Miranda-Agrippino, and Saha 2025) where applicable. All event dates are recorded in London time.

In the pre-reform regime (2011Q1 through 2015Q1), MPC communication in each quarterly forecast round comprised two temporally separated events, the rate decision with accompanying statement (released at noon on a Thursday) and the quarterly forecast publication, the *Inflation Report*, released on a subsequent day with the Governor's press conference. Non-forecast-round months involved only the single rate decision, followed by the separate publication of minutes roughly two weeks later. In the post-reform regime, introduced with the first bundled release on 6 August 2015, forecast-round months combine the rate decision, minutes, Inflation Report (later renamed the *Monetary Policy Report*), and the Governor's press conference onto a single day. Outside forecast-round months, the MPC holds "standard" meetings at which the decision and minutes are published simultaneously but without a forecast or press conference.

For the quarterly bundling analysis, we define a *forecast round* as the set of policy communication events associated with a given quarterly forecast release cycle. In the pre-reform regime, each forecast round contains two distinct event dates (the decision day and the Inflation Report day). In the post-reform regime, each forecast round contains one date (the full-release day). This definition ensures that the pre-to-post comparison captures total engagement across the entire quarterly communication round, not merely the response to a single event. Table 19 in the appendix provides a year-by-year listing of events by type. The full sample contains 117 MPC decision dates between January 2011 and December 2022, of which 59 fall in the pre-reform period and 58 in the post-reform period (25 full-release, 33 standard). The estimation sample for the quarterly analysis ends with the 2022Q2 forecast round (the last round with complete Twitter coverage); four late-2022 MPC dates lack daily Twitter observations and are excluded from event-level regressions that require $A_{e,TW}$.

4.2 Twitter engagement

Our primary high-frequency attention measure is constructed from Twitter data. For each calendar day t in the sample, we count the number of unique user accounts that posted at least one original tweet (excluding retweets) containing one or more terms from a pre-defined monetary policy lexicon. The lexicon includes `Bank of England`, `BoE`, `MPC`, `Monetary Policy Committee`, `interest rate(s)`, `base rate`, `Bank Rate`, `quantitative easing`, `QE`, `Inflation Report` and `Monetary Policy Report`; matching is case-insensitive. We require at least one keyword match per tweet. The lexicon is held constant across the full sample period to avoid introducing measurement breaks that could confound the pre-to-post comparison. Keyword construction details are summarised in Appendix A.

By counting unique users rather than total tweet volume, the measure captures the *breadth* (extensive margin) of public engagement rather than repeated posting by a small number of highly active accounts. This choice is motivated by Gorodnichenko, Pham, and Talavera (2025), who document that only the extensive margin of social media engagement around central bank communication is economically meaningful.

We denote the daily unique-user count by $Users_t^{TW}$ and work with the log transformation $Y_t^{TW} = \log(Users_t^{TW})$, which stabilises variance and reduces the influence of extreme spikes. To construct an event-level outcome, we compute abnormal log-attention by subtracting a rolling baseline from the event-day observation. Specifically, for an event e occurring on date $t(e)$, the baseline is the mean of Y^{TW} on non-event days within a trailing 30-day window, matched on day of week to absorb systematic within-week variation in platform activity. We exclude any day falling within the $[-1, +1]$ neighbourhood of another MPC event from the baseline calculation to avoid contamination. The resulting platform-specific abnormal measure, $A_{e,TW}$, is defined formally in Equation (6) in the appendix.

Figure 1 displays the distribution of daily log-unique-users across the seven-day event window ($h = -3$ to $h = +3$) around MPC decision dates. Attention is sharply concentrated on the announcement day, as the mean at $h = 0$ (7.55 log-unique-users) exceeds the mean at $h = -1$ (6.31) by 1.24 log points, corresponding to roughly $e^{1.24} \approx 3.46$ times as many unique users. The minimum event-day observation (6.64) also lies above the mean on every adjacent day. Attention decays rapidly thereafter, falling to 6.49 at $h = +1$ and 5.61 at $h = +2$. This spike-and-decay profile is consistent with MPC announcements operating as discrete attention shocks with clearly delimited temporal boundaries.

FIGURE 1. Event-Time Attention Around MPC Announcements

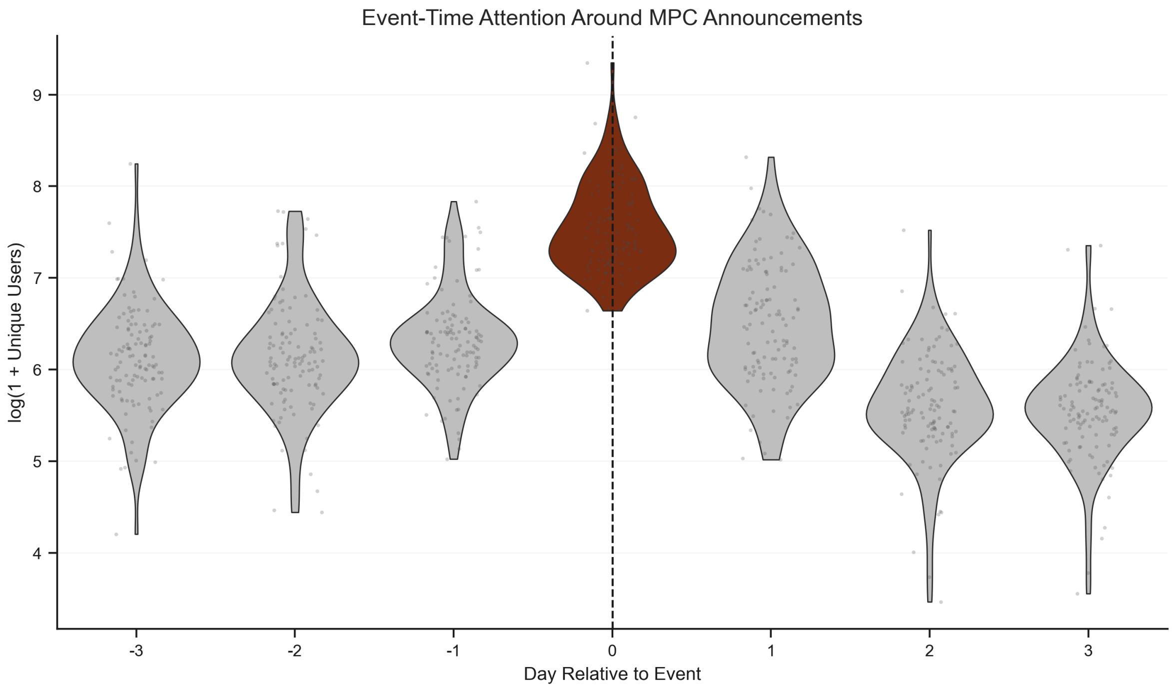

Notes: Violin plots with overlaid strip points show the distribution of daily log-unique-users on Twitter across the seven-day window $h = -3, \ldots, +3$ around MPC decision announcements (London time). The event day ($h = 0$, shaded) exhibits a pronounced spike relative to adjacent days. Attention decays sharply after the event.

The platform captures real-time public reaction at daily frequency around precisely dated policy events, and the text content of tweets permits subsequent analysis of the *composition* of attention (Section 6). However, Twitter users are not representative of the UK population; the platform's user base skews younger, more educated, and more politically engaged (Gorodnichenko, Pham, and Talavera 2025). Our measure should therefore be interpreted as capturing engagement among a digitally active segment of the public rather than a population-wide quantity. We partially address this concern by triangulating with Google Trends and Wikipedia, each of which draws on a different population of users.

For the bundling analysis, we aggregate event-level Twitter attention to the forecast-round level:

$$CumLogAtt_q = \sum_{e \in \mathcal{E}_q} A_{e,TW},$$

where $\mathcal{E}_q$ is the set of communication events in forecast round q. Pre-reform, $\mathcal{E}_q$ contains two dates (decision day and Inflation Report day). Post-reform, $\mathcal{E}_q$ contains one date (full-release day). The pre-to-post comparison therefore captures total engagement across the full quarterly communication round. Because this quarterly variable is a sum of abnormal log-attention components, percentage changes are reported as proportional changes in means rather than as log-point exponentiations.

4.3 Google Trends and Wikipedia

Google Trends captures a different facet of engagement, active information seeking through search. We construct daily search intensity indices for three Bank of England–related queries restricted to the United Kingdom — "Bank of England," "interest rates," and "inflation" — selected to span policy-relevant search behaviour from institution-specific queries through the most concrete policy instrument to the broader macroeconomic outcome. We use Google Trends' "search term" mode (exact phrase matching) rather than "topic" mode to maintain transparent query definitions; details are provided in Appendix A.

Google Trends reports a normalised index that scales query volume relative to total search activity within the specified geography and extraction window, with the peak observation set to 100 (Cebrián and Domenech 2024). To construct a consistent daily series over the full sample horizon, we download data in overlapping extraction windows and rescale levels using the overlap periods, following standard practice (Da, Engelberg, and Gao 2011, Wohlfarth 2018). We denote the daily Google Trends index for the selected query bundle by Y_t^{GT}. Unlike the Twitter and Wikipedia series, the Trends index is already a normalised relative measure and is not log-transformed. We construct event-level abnormal search attention, $A_{e,GT}$, using the same 30-day trailing, day-of-week-matched baseline procedure described in Section 4.2.

The normalisation supports clean inference about within-series spikes around events but limits cross-period comparisons of absolute search volume — immaterial for our event-study design, which identifies effects from within-period variation around precisely dated announcements. Cebrián and Domenech (2024) documented that repeated extractions of the same query and time window can yield slightly different values owing to sampling variation in Google's reporting algorithm. We therefore treat Google Trends as one component of the broader multi-platform evidence rather than as a standalone estimator. The conceptual advantage of search data is that individuals who search for monetary policy information are actively choosing to acquire it, implying a higher degree of cognitive engagement than passive exposure to social media content (Lamla and Vinogradov 2019, Buelens 2025).

Wikipedia page views capture a third and complementary dimension of engagement, deliberate reference behaviour. Users who navigate to a Wikipedia article about the Bank of England are consulting a reference source, which suggests a desire for background understanding rather than merely reacting to a headline. We collect daily page views from the Wikimedia Pageviews API for the "Bank of England" article and a small set of closely related monetary policy pages ("Monetary Policy Committee," "Bank Rate," "Quantitative easing"). The page set and query setup are summarised in Appendix A. For each day t, we work with $Y_t^{WK} = \log(\text{PageViews}_t)$ and construct abnormal event-level attention $A_{e,WK}$ using the same rolling-baseline procedure applied to the other two platforms. This measure is best interpreted as capturing the *depth* of public interest rather than its breadth or immediacy, complementing the real-time reactive dimension of Twitter and the active information-seeking dimension of Google Trends.

4.4 High-frequency financial market data

To assess whether the estimated attention effects are mechanically driven by financial market dynamics, we collect intraday price data for three liquid instruments. Sterling exchange rates (GBPUSD) capture the foreign exchange market's response to MPC announcements. FTSE 100 equity index returns capture the domestic market reaction. For the interest rate dimension, we use Eurodollar futures (the CME three-month LIBOR contract), which were the most liquid short-rate futures contract available throughout our sample period and which reflect global interest-rate expectations that respond to UK policy news through sterling's international role.[1]

For each MPC event, we extract 15-minute log returns in two windows aligned to the MPC communication schedule, with a noon window centred on the 12:00 decision announcement and a 12:45 window aligned with the timing of the Governor's press conference (in full-release months) or the post-announcement settling period. All series are converted to London time for consistent event alignment. These narrow windows follow the conventions of the high-frequency monetary policy event-study literature (Gürkaynak, Sack, and Swanson 2005, Braun, Miranda-Agrippino, and Saha 2025) and are designed to isolate the market response from confounding intraday price movements. The UKMPD event timestamps in the underlying dataset are recorded in UTC and converted to London time before window construction, yielding 12:00 London decision alignment in both GMT and BST months, with one unscheduled 2020 emergency announcement retained using its recorded timestamp.

These financial market series enter the analysis as controls and diagnostics rather than primary outcomes. They verify that our attention results survive conditioning on contemporaneous market volatility, and they test whether bundling itself alters intraday market dynamics. Table 12 in Section 6 documents the coverage of each instrument across the event sample.

[1] Short Sterling futures, the natural UK-specific contract, were delisted by ICE Futures Europe in March 2020 during the transition to SONIA-based derivatives, which leaves a break in long-sample coverage. Eurodollar is therefore a liquidity-driven proxy rather than a purely UK-specific contract. In our matched post-reform sample, Eurodollar 12:00 absolute returns are strongly related to UK policy news (corr with *ShockMag* = 0.69) and comove with GBPUSD/FTSE responses (corrs 0.39 and 0.30). We therefore keep Eurodollar in the market-diagnostic set and report specifications with and without the rates leg.

4.5 Sample and summary statistics

The daily attention series span January 2011 through June 2022, yielding $N = 4{,}183$ trading days. The start date is determined by the onset of meaningful Twitter activity around Bank of England events; prior to 2011, daily unique-user counts are too sparse to support reliable baseline estimation. The quarterly estimation sample ends with 2022Q2, which provides a long post-reform window (nearly seven years) whilst avoiding reliance on late-2022 dates for which daily Twitter coverage is incomplete. The master calendar in the appendix lists all MPC dates through December 2022 for completeness.

The August 2015 reform divides the sample into two regimes. The pre-reform period comprises 18 quarterly forecast rounds (2011Q1 through 2015Q2), each containing two events. The post-reform period comprises 25 quarterly forecast rounds (2015Q3 through 2022Q2).

Table 2 reports summary statistics for the key attention variables by regime. At the quarterly level, the pre-reform mean of $CumLogAtt_q$ is 2.966, compared with 1.742 post-reform — a sharp decline in total forecast-round engagement. At the event level, however, the pattern reverses, and full-release events have a higher mean $A_{e,TW}$ (1.742) than post-reform standard events (1.369). This cross-cutting pattern — higher per-event attention but lower quarterly totals — motivates the joint intensive-versus-extensive margin analysis in Section 6.

Figure 2 plots the daily Twitter attention series across the full sample. MPC event dates are marked along the top of the panel, and the vertical dashed line at August 2015 marks the reform implementation. The transition from twelve to eight meetings per year after 2016 is visible in the wider spacing of event markers.

FIGURE 2. Daily Twitter Attention with MPC Event Markers

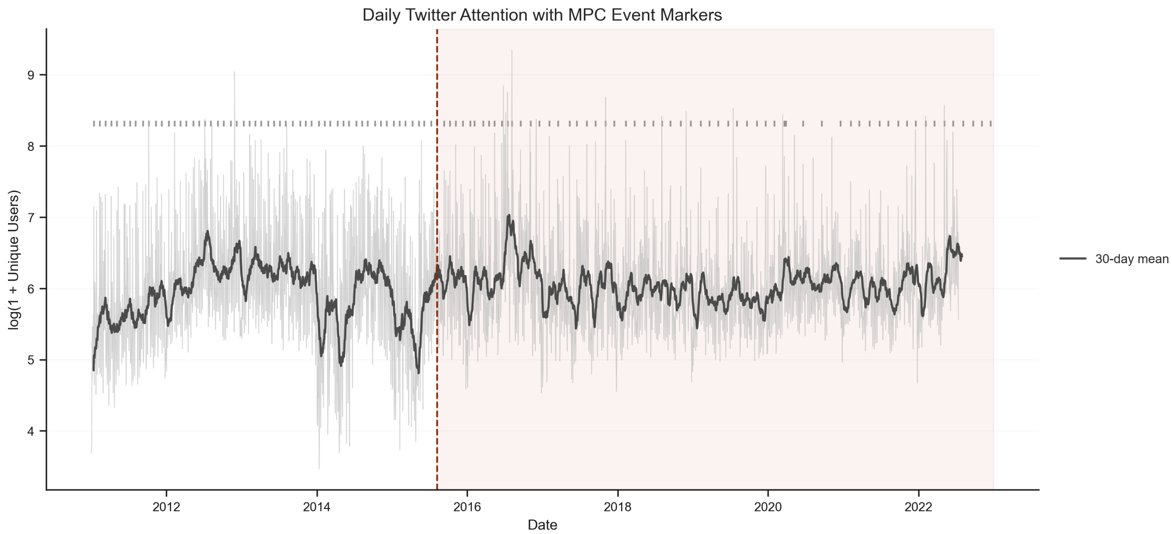

Notes: Daily log-unique-users on Twitter (light grey) with a 30-day rolling mean (dark line) from January 2011 to June 2022. MPC event dates are marked with tick marks along the top. The vertical dashed line marks the August 2015 implementation of the bundled release format; the post-reform period is shaded.

5. Empirical Strategy

With measurement choices fixed, identification comes from the reform's timing. The Warsh Review reform provides a quasi-experimental change in the *architecture* of Bank of England communication, where information that was previously delivered across two distinct occasions within a quarterly forecast round is bundled into a single full-release event day. Our empirical strategy separates two margins that theory leaves ambiguous — an intensive-margin effect (bundled days may attract more attention per event) and an extensive-margin effect (fewer distinct occasions may reduce cumulative engagement over the quarterly round).

We work at two units of observation. An *event e* is a dated MPC communication occasion, either a decision-only meeting, a full-release meeting (combining the decision, Monetary Policy Report, and press conference), or, in the pre-reform period, the separate Inflation Report release day. A *forecast round q* is the quarterly cycle associated with a given forecast publication. The set $\mathcal{E}_q$ collects the events in forecast round q. Pre-reform, $\mathcal{E}_q = \{e_{\text{decision}}, e_{\text{report}}\}$ contains two events; post-reform, $\mathcal{E}_q = \{e_{\text{full-release}}\}$ contains one. The event-level attention measure is $A_{e,TW}$ and the quarterly outcome is $CumLogAtt_q = \sum_{e \in \mathcal{E}_q} A_{e,TW}$. The treatment indicator $Post_q$ equals one for forecast rounds conducted under the bundled architecture (2015Q3 onward) and zero for those under the two-event architecture (through 2015Q2).

5.1 Forecast-round totals

The primary question asks whether consolidating two communication events into one reduced cumulative public attention across the quarterly forecast round. This is a structural-break design in which we test whether the mean of $CumLogAtt_q$ exhibits a discrete shift coincident with the reform. The baseline specification is

$$CumLogAtt_q = \alpha + \beta \cdot Post_q + \theta_{qoy} + \varepsilon_q, \tag{1}$$

where θ_{qoy} denotes quarter-of-year fixed effects that absorb seasonal patterns in digital activity (e.g., summer troughs, year-end spikes). The baseline is deliberately parsimonious; we verify in robustness analysis that augmenting with predetermined macro-state indicators (the prevailing Bank Rate level, a CPI-deviation dummy) does not materially alter the estimate.[2]

The coefficient β resolves the bundling ambiguity described in Section 3. If the intensive-margin salience effect dominates — the bundled day is so much more prominent that it more than compensates for the lost separate event — we would expect $\beta \geq 0$. A negative estimate indicates a quarterly engagement deficit, where the loss of a distinct attention occasion outweighs amplification of the remaining one.

5.2 Forecast-occasion architecture comparison

The most targeted architecture comparison in our data contrasts pre-reform Inflation Report days with post-reform full-release days. Both are forecast occasions that contain the forecast publication and press communication, but only post-reform full-release days co-locate those elements with the rate decision and simultaneous minutes release. Let F_e indicate a post-reform full-release day in the restricted sample of forecast occasions. We estimate

$$A_{e,TW} = \alpha + \beta \cdot F_e + \theta_{qoy} + \varepsilon_e, \tag{2}$$

where θ_{qoy} are quarter-of-year fixed effects. The coefficient β measures the shift in forecast-occasion attention under co-location relative to the pre-reform architecture in which the forecast

[2] Because $CumLogAtt_q$ is a sum of abnormal log-attention values from the Twitter series, β is denominated in log-user units. The percentage decline reported in Section 6 is the proportional change in means, $(\overline{CumLogAtt}_{post} - \overline{CumLogAtt}_{pre})/\overline{CumLogAtt}_{pre}$, not a log-point exponentiation.

arrived on a separate day.

Equation (2) has a directional prediction. Post-reform full-release days contain all components present on pre-reform Inflation Report days (forecast publication and press communication) plus the policy decision and simultaneous minutes release. We therefore test $H_0 : \beta = 0$ against $H_1 : \beta > 0$ for E1 and treat one-sided permutation p-values as primary inference, with HC2 two-sided p-values reported as secondary precision diagnostics.

5.3 Within-post event-level amplification

As a complementary design, we compare full-release and standard MPC events *within the post-reform period only*. Let B_e indicate a full-release event. We estimate

$$A_{e,TW} = \alpha + \beta \cdot B_e + \delta_{year} + \theta_{qoy} + \varepsilon_e, \tag{3}$$

where δ_{year} and θ_{qoy} absorb secular trends and seasonality. The coefficient β measures the average within-year, within-season difference in event-day Twitter abnormal attention between full-release meetings and standard meetings.

This within-post comparison is informative about intensive-margin salience under common post-reform institutions. Because both event types share the same minutes publication rule, meeting frequency, and communication norms, this estimate is not driven by those package-wide post-2015 factors. It does not, by itself, separate forecast-content salience from co-location salience, which is why Equation (2) is presented as the primary targeted architecture test.

In this paper, "architecture" refers to the timing and co-location of communication components, in particular whether the forecast package arrives on the decision day or on a separate day. Under that definition, the forecast-occasion comparison in Equation (2) and the within-post comparison in Equation (3) provide complementary evidence. We nonetheless acknowledge a scope limit. Even the forecast-occasion comparison remains pre/post and therefore cannot remove every coincident regime change. We therefore interpret event-level and quarterly evidence jointly and avoid claims that any single specification fully identifies the entire reform package.

5.4 Composition across platforms and search content

Beyond aggregate levels, we ask whether bundling changes *what* the public attends to. Let $A_{e,p}$ denote platform-specific abnormal attention on platform $p \in \{TW, GT, WK\}$ (as defined in Section 4).

We estimate separate platform-by-platform regressions

$$A_{e,p} = \alpha + \beta_p \cdot Post_e + \delta_{year} + \theta_{qoy} + \varepsilon_{e,p}, \tag{4}$$

for each platform p, where $Post_e$ indicates events occurring in the post-reform period. The set of platform-specific coefficients $\{\hat{\beta}_{TW}, \hat{\beta}_{GT}, \hat{\beta}_{WK}\}$ reveals whether the reform shifted attention differentially across channels. We supplement this with a stacked specification that pools across platforms and tests cross-platform differences directly; in the stacked version, standard errors are clustered at the event level to account for the mechanical correlation across platforms driven by the same announcement.[3]

Within Google Trends, we further decompose attention by search term ("Bank of England," "interest rates," "inflation") and estimate term-specific post-reform shifts to characterise how search composition changes after the reform. Within Twitter, we examine the composition of tweet content by regressing topic shares (including outlook-related share, rate-focused share, and QE/balance-sheet share) on the full-release indicator B_e within the post-reform sample. Variable definitions are reported in Appendix A.

5.5 Attention dynamics

The preceding specifications establish the level and composition of attention; we next examine its temporal profile. For each event e and horizon $h \in \{-2, -1, 0, 1, \ldots, 10\}$ business days relative to the announcement, let $Att_{e,h}$ denote abnormal attention on day $t(e) + h$, constructed using the same day-of-week baseline as the event-day measure. We estimate horizon-specific regressions in the spirit of local projections (Jordà 2005):

$$Att_{e,h} = \alpha_h + \beta_h \cdot HighShock_e + \gamma_h \cdot B_e + \delta_h \cdot (HighShock_e \times B_e) + \varepsilon_{e,h}, \tag{5}$$

where $HighShock_e = \mathbf{1}\{|Shock_e| > \text{median}\}$ is an indicator for events at which the absolute magnitude of the UKMPD composite surprise exceeds the sample median. The sequence $\{\hat{\beta}_h\}$ traces the attention impulse response to high-surprise events relative to low-surprise events, and $\{\hat{\delta}_h\}$ tests whether bundling alters the *shape* of this response — whether attention is concentrated on the announcement day, whether it persists, and whether bundling compresses the response into

[3] An MPC decision simultaneously drives search, social media, and reference activity, so errors for the same event across platforms are not independent.

a narrower window. The pre-event horizons ($h = -2, -1$) serve as a falsification check because statistically significant coefficients at negative horizons would indicate pre-announcement attention leakage or misspecification. A separate regression is estimated for each h; no parametric decay structure is imposed.

5.6 Extensions on sign asymmetry and state dependence

Two extensions probe heterogeneity predicted by behavioural and rational-inattention mechanisms. Both are reported as supplementary findings rather than primary hypotheses, because they involve smaller effective samples and sharper identification challenges.

The first allows the attention response to differ by the *sign* of the surprise. The UKMPD composite surprise is split into its positive (hawkish) and negative (dovish) components, $Shock_e^+ = \max(Shock_e, 0)$ and $Shock_e^- = \min(Shock_e, 0)$, and separate slope coefficients β^+ and β^- are estimated in an augmented version of the event-level specification (Equation 11 in Appendix E). The ratio $|\hat{\beta}^+|/|\hat{\beta}^-|$ measures the relative public-attention response to contractionary versus expansionary surprises. The identifying variation for this test is concentrated in a small number of events with large surprises, and we report sensitivity to the exclusion of the single largest hawkish shock.

The second interacts $Post_q$ in Equation (1) with a high-inflation indicator ($\mathbf{1}\{\text{CPI inflation} > 3\%\}$) to test whether the bundling effect varies with the macroeconomic environment. Rational inattention theory predicts that the marginal value of attending to monetary policy rises when inflation is elevated (Weber et al. 2025). However, the UK sample provides limited variation in the high-inflation state (inflation exceeded 3 per cent in only a small fraction of quarters before the post-2021 episode), so we regard this test as exploratory.

5.7 Financial market diagnostics

A distinct question is whether the attention effects we estimate remain after conditioning on contemporaneous market volatility. If full-release events generate larger market moves, and those moves in turn drive media coverage and search activity, then part of the measured bundling coefficient could reflect a market-mediated channel rather than the architectural change per se.

We address this by augmenting the event-level specification with high-frequency volatility measures — absolute 15-minute log returns for GBPUSD, the FTSE 100, and the short-rate futures contract described in Section 4 — computed in two intraday windows per event (noon and 12:45). These controls are added incrementally, which allows us to trace how the bundling coefficient attenuates as market information is absorbed.

We frame these augmentations as *diagnostics* rather than as identifying a "pure" architecture effect, for two reasons. First, market moves are plausibly a *mediator* of the attention response (policy event → market reaction → media coverage → public attention), not a confounder. Conditioning on a mediator can attenuate the coefficient of interest without removing bias, and in the presence of other causal paths from the event to attention (e.g., direct media reporting of the decision itself), conditioning may introduce collider bias. Second, the market controls are measured with error. We therefore interpret the attenuated coefficient as a lower bound on the bundling effect, and the attenuation itself as an upper bound on the share of the attention response that operates through market-mediated channels.

We deliberately exclude the UKMPD surprise from the baseline quarterly specification, because the reform plausibly altered how markets price the announcement — bundling changes not only what the public sees but also how financial participants process the joint information release, inducing a mechanical endogeneity between architecture and the measured surprise. The UKMPD enters the event-level specifications where within-post-reform variation in the surprise measure is less susceptible to this architectural contamination.

5.8 Identification

Equation (1) is a structural-break design. It does not rely on a control group and it does not invoke a parallel-trends assumption. The identifying assumption is narrower, and absent the reform forecast-round cumulative attention would not have exhibited a discrete shift of comparable magnitude and timing. This permits smooth trends and seasonal variation, which the quarter-of-year fixed effects absorb. What must be ruled out is a confounding event that coincidentally shifted the mean of $CumLogAtt_q$ at the same date and in the same direction.

We evaluate this assumption through three falsification exercises. First, we estimate a linear time trend within the pre-reform subsample and test whether it is statistically distinguishable from zero; a non-zero pre-trend would indicate that the decline pre-dates the reform. Second, we re-estimate Equation (1) with $Post_q$ redefined at every pre-reform quarter and compare the resulting placebo coefficients to the true reform-date estimate. Third, we run permutation inference by randomly reassigning $Post_q$ across the 43 quarters 10,000 times and constructing the exact distribution of the test statistic under the sharp null of no reform effect.

With only 43 quarterly observations, asymptotic approximations for standard errors may be unreliable. We therefore treat the permutation p-value as the primary inferential tool for the quarterly specification and report HC2 heteroskedasticity-robust standard errors as a descriptive measure of precision. The same small-sample logic applies to E1 ($N = 43$ forecast occasions); because E1 has a directional prediction, we use one-sided permutation inference as primary and report HC2 two-sided p-values as secondary diagnostics. For larger event-level specifications, where the sample is larger ($N = 58$ post-reform events for the within-post comparison), we rely on HC2 standard errors as the baseline and report permutation p-values as a complementary check. All primary tests are subjected to the Benjamini–Hochberg false discovery rate correction (Benjamini and Hochberg 1995) at the 5 per cent level, and both raw and adjusted p-values are reported.

The Warsh Review reform package included simultaneous minutes publication alongside forecast bundling, both taking effect in August 2015, and a meeting-frequency reduction from twelve to eight per year effective September 2016. Each specification has a different identifying reach. The quarterly specification estimates the *net effect of the reform package* on forecast-round engagement; the simultaneous-minutes reform and bundling cannot be separately identified within this design, though the minutes change is a within-day timing adjustment (minutes arrive hours earlier) rather than an elimination of a separate event, and is therefore less likely to drive a large attention shift. A timing check also shows that the quarterly decline is already present before the frequency change takes effect. Comparing pre-reform quarters to the early post-reform window 2015Q3–2016Q2 yields a mean drop from 2.966 to 1.817 and an early-post coefficient of −1.121 (s.e. 0.182, $p < 0.001$). The forecast-occasion comparison in Section 5.2 is our most targeted architecture test, while the within-post comparison in Section 5.3 provides complementary intensive-margin evidence under a common post-reform environment.

We report three variants of the quarterly specification that modify the sample or control set, one excluding 2020–2022 forecast rounds, one trimming extreme observations of $CumLogAtt_q$, and one adding predetermined macro-state indicators. Table 9 reports the stability of $\hat{\beta}$ across these variants.

6. Results and Discussion

The evidence points to a consistent pattern in which event-day salience rises on bundled releases while cumulative quarterly Twitter engagement declines after the reform package.

All estimates use HC2 robust standard errors unless stated otherwise. For the quarterly specification, we additionally report Newey–West HAC standard errors (four quarterly lags) and place primary inferential weight on permutation and placebo-break diagnostics because the quarterly sample is modest ($N = 43$). The event-level variable is Twitter abnormal log-unique-users, denoted $A_{e,TW}$.[4]

[4] Inference conventions and notation are stated once here and apply throughout the section. All p-values refer to two-sided tests unless otherwise noted. E1 uses one-sided permutation inference by design.

6.1 Forecast-round engagement on Twitter

The descriptive data already show a large forecast-round regime shift. Table 2 reports summary statistics for the forecast-round panel ($CumLogAtt_q$) and the event panel ($A_{e,TW}$), split by regime and event type.

TABLE 2. Summary Statistics by Regime

Variable	N	Mean	SD	Min	Max
Panel A, Forecast-Round Panel ($CumLogAtt_q$)					
Pre-reform	18	2.966	0.631	1.945	4.599
Post-reform	25	1.742	0.466	1.087	2.556
Panel B, Event-Level Twitter Abnormal Attention ($A_{e,TW}$)					
Pre-reform events (59 total; 54 observed)	54	1.332	0.423	0.422	2.420
Post-reform events (58 total; 58 observed)	58	1.530	0.540	0.133	2.556
Full-release	25	1.742	0.466	1.087	2.556
Standard	33	1.369	0.543	0.133	2.537

Notes: Panel A reports forecast-round cumulative attention (N = 43 rounds; 18 pre-reform, 25 post-reform). Panel B reports event-level Twitter abnormal log-unique-users. The pre-reform event universe has 59 MPC dates, of which 54 have non-missing $A_{e,TW}$; the post-reform estimation sample has 58 events, all with observed $A_{e,TW}$ (25 full-release, 33 standard). Four additional late-2022 MPC dates (2022-08-04, 2022-09-22, 2022-11-03, 2022-12-15) lack daily Twitter coverage and are listed in the appendix calendar for completeness but excluded from all estimations.

Cumulative attention within a forecast round is substantially lower post-reform than pre-reform (mean 1.742 versus 2.966). The shift is large relative to within-regime dispersion, and the post-reform mean falls below the pre-reform minimum (1.945), indicating that even the most-attended post-reform quarter generates less cumulative engagement than the least-attended pre-reform quarter. At the event level, however, the pattern reverses, and full-release events exhibit higher mean attention than standard meetings (1.742 versus 1.369). Whether this intensive-margin amplification compensates for the loss of a separate occasion is the central empirical question.

6.2 Descriptive accounting of the forecast-round decline

The most direct accounting exercise uses pre-reform forecast rounds, where both decision-day and Inflation Report-day components are observed separately. Let $\bar{A}^{dec}_{TW}$ denote the pre-reform decision-day mean, $\bar{A}^{IR}_{TW}$ the pre-reform Inflation Report-day mean, and $\bar{A}^{full}_{TW}$ the post-reform full-release mean. Table 3 reports these moments and the implied offset arithmetic.

This accounting clarifies two points. First, the decline is not a tautological by-product of summing logs, because the quarterly object is a sum of abnormal deviations. Second, the intensive

TABLE 3. Decomposition of the Forecast-Round Attention Gap

Quantity	Estimate	Interpretation
$\bar{A}^{dec}_{TW}$ (pre decision day)	1.510	Remaining event under a one-event counterfactual
$\bar{A}^{IR}_{TW}$ (pre Inflation Report day)	1.456	Lost second-event component
$\bar{A}^{dec}_{TW} + \bar{A}^{IR}_{TW}$ (pre total)	2.966	Pre-reform quarterly mean
$\bar{A}^{full}_{TW}$ (post full-release day)	1.742	Observed post-reform quarterly mean
$\bar{A}^{full}_{TW} - \bar{A}^{dec}_{TW}$	+0.232	Intensive-margin gain on the remaining day
$(\bar{A}^{dec}_{TW} + \bar{A}^{IR}_{TW}) - \bar{A}^{full}_{TW}$	1.224	Total quarterly gap
$\frac{\bar{A}^{full}_{TW} - \bar{A}^{dec}_{TW}}{\bar{A}^{IR}_{TW}}$	15.9%	Share of the lost second-event margin offset by bundling

Notes: Decomposition moments use the pre-reform forecast-round pair panel (N = 18) and the post-reform full-release quarter panel (N = 25). All quantities are in abnormal log-user units. The arithmetic gap (1.224) differs slightly from the regression coefficient in Table 5 (1.216) because Table 3 reports an unconditional mean difference, while Equation (1) includes quarter-of-year fixed effects.

margin rises on the bundled day, but the increase is too small to replace the removed second event in pre-reform forecast rounds. We treat this table as descriptive arithmetic rather than a causal decomposition. The pre-reform decision and Inflation Report occasions were separated by 5.8 days on average (median 6), and both generated similarly large spikes (1.510 and 1.456), consistent with a distinct second attention cycle before bundling.

Figure 3 visualises the forecast-round contrast. While the two distributions are not perfectly disjoint — the post-reform range extends above the pre-reform minimum — the mean shift is large and the overlap is limited.

FIGURE 3. Quarterly Attention Distribution by Regime

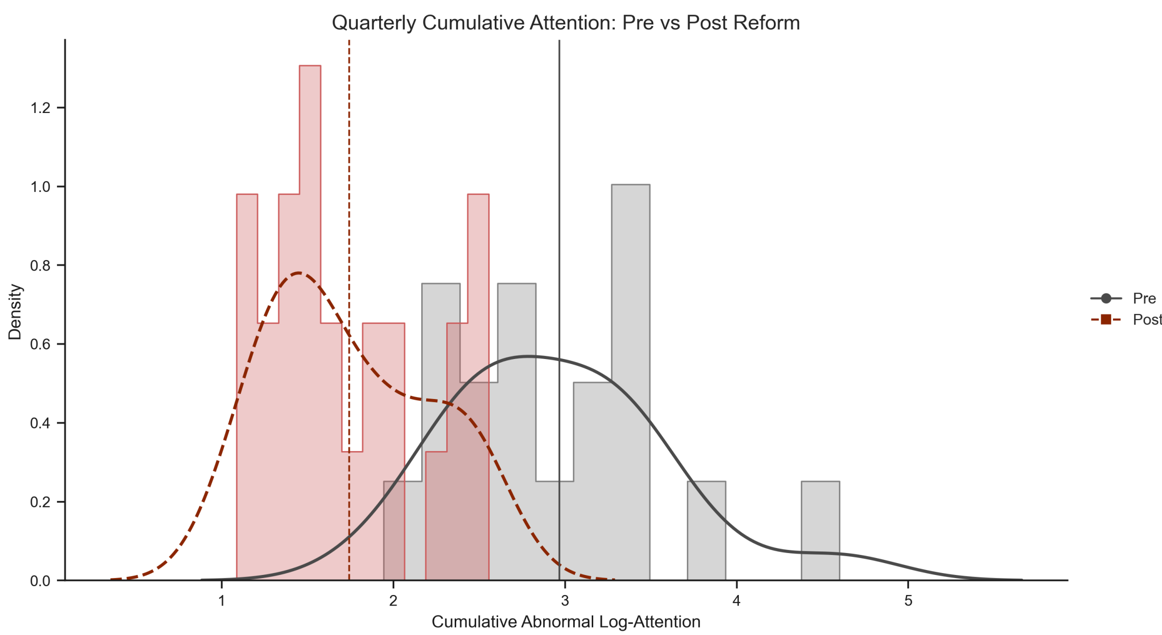

Notes: Overlaid histograms with kernel density curves for $CumLogAtt_q$ in the pre-reform period (N = 18, solid mean line at 2.966) and the post-reform period (N = 25, dashed mean line at 1.742). The post-reform distribution is tighter (s.d. 0.466 versus 0.631) and shifted leftward. The post-reform mean falls below the pre-reform minimum (1.945).

6.3 Forecast-occasion architecture comparison (E1)

Table 4 reports the targeted comparison between pre-reform Inflation Report days and post-reform full-release days. In the baseline quarter-of-year specification, the Twitter coefficient is positive (+0.290, s.e. 0.144). For E1 we use one-sided permutation inference with $H_1 : \beta > 0$; the Twitter one-sided permutation p-value is 0.025 (two-sided permutation $p = 0.052$; HC2 two-sided $p = 0.043$). The baseline Google Trends coefficient is also positive and large (+17.554, s.e. 4.009; one-sided permutation $p = 0.0001$; HC2 two-sided $p < 0.001$).[5]

TABLE 4. Forecast-Occasion Architecture Comparison: Pre-Reform IR Day vs Post-Reform Full-Release

Platform / Spec	Post Coef.	SE	p-value	Perm. p (one-sided)	Pre Mean	Post Mean
Twitter baseline (QOY FE)	+0.290	0.144	0.0432	0.0253	1.456	1.742
Google Trends baseline (QOY FE)	+17.554	4.009	< 0.001	0.0001	-0.164	17.561

Notes: Each row uses the 43-event forecast-occasion sample (18 pre-reform Inflation Report days, 25 post-reform full-release days). Estimates use Equation (2) with quarter-of-year fixed effects and HC2 standard errors. The reported permutation values are one-sided ($H_1 : \beta > 0$) and are the primary inferential metric for E1. The p-value column reports HC2 two-sided values as a secondary precision check.

6.4 Forecast-round quarterly evidence

The baseline quarterly specification reproduces the descriptive picture. The estimated decline in forecast-round cumulative attention coincident with the reform is $\hat{\beta} = -1.216$ (s.e. 0.167), a 41 per cent proportional decline relative to the pre-reform mean of 2.966 (Table 5).

TABLE 5. Main Regression Results

Dep. Variable	Specification	Coefficient	SE	N	R^2	FDR Sig.
$CumLogAtt_q$	Forecast-round post effect	-1.216***	(0.167)	43	0.61	Yes
$A_{e,TW}$	Full-release effect (post only)	+0.374**	(0.130)	58	0.22	Yes
$A_{e,GT}$	Google Trends event-level post effect	+21.286***	(4.301)	48	0.34	Yes
$A_{e,WK}$	Wikipedia event-level full-release effect (post only)	+0.273***	(0.066)	62	0.45	Yes
$A_{e,TW}$	Twitter post effect	-0.142	(0.259)	112	0.37	No
$A_{e,TW}$	High-shock indicator (post only)	+0.319*	(0.129)	58	0.19	Yes

Notes: HC2 robust standard errors in parentheses. Stars denote raw significance levels (*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$). FDR significance is based on the distinct-hypothesis family reported in Table 15; robustness variants of the quarterly specification are excluded from FDR by design.

Because this is a structural-break design, credibility depends on whether the observed shift is concentrated at the true reform date rather than reflecting a secular decline that coincidentally overlaps with it. Two exercises provide the primary identification support. Permutation inference — randomly reassigning $Post_q$ across the 43 quarters 10,000 times — yields a p-value below 10^{-4}; the true estimate exceeds all 9,999 placebo effects. The placebo break-date analysis (Table 9) confirms that shifting the break to pre-reform quarters produces substantially smaller coefficients, and no placebo break approaches the magnitude at the true reform date.

The Google Trends and Wikipedia rows in Table 5 are event-level specifications. We therefore

[5] The pre-reform trend in quarterly attention is flat (slope = 0.053, $p = 0.658$; Figure 4). In the 43-event forecast-occasion sample, adding a linear trend leaves point estimates positive but inflates standard errors because the trend and post indicator are near-collinear in a short ordered sample.

run an explicit quarterly cross-platform diagnostic for the same forecast-round construction used in $CumLogAtt_q$.

6.5 Quarterly cross-platform diagnostic

Table 6 reports quarterly cross-platform diagnostics. The Google Trends analogue is positive and significant, with a post-reform increase of +13.766 (s.e. 4.185, $p = 0.001$; permutation $p = 0.0027$). This sign is opposite to the Twitter quarterly estimate. Wikipedia cannot support the same quarterly design because daily coverage starts in mid-2015 and yields zero non-missing pre-reform quarters.

TABLE 6. Quarterly Cross-Platform Diagnostic

Platform	Post Coef.	SE	p-value	Pre N	Post N
Google Trends quarterly analogue	+13.766***	4.185	0.0010	18	25
Wikipedia quarterly analogue	Not estimable	–	–	0	25

Notes: The Google Trends row estimates the quarterly post indicator in a 43-quarter panel using the same quarter-of-year controls as the Twitter quarterly specification. The observed difference in means is +13.935 and the permutation p-value is 0.0027. Wikipedia quarterly replication is infeasible because pre-reform daily coverage is missing.

The quarterly cross-platform diagnostic therefore does not replicate the Twitter quarterly decline on Google Trends, and Wikipedia cannot adjudicate the quarterly comparison. We interpret this as channel heterogeneity rather than a contradiction: event-level co-location effects are positive across platforms, while cumulative quarterly responses differ by engagement channel.

6.6 Raw-scale quarterly robustness

As a transparency check, we also estimate the quarterly break in raw user units, avoiding both logs and baseline subtraction. We reconstruct the quarterly outcome as the sum of raw event-day unique users within each forecast round, $TotalUsers_q = \sum_{e \in \mathcal{E}_q} Users_e$, and estimate the same quarterly break specification.

TABLE 7. Quarterly Robustness in Raw User Units

Outcome	Post Coef.	SE	p-value	Perm. p	Pre Mean	Post Mean
$TotalUsers_q$ (raw users)	-1051.998*	501.112	0.0358	0.0751	4196.3	3175.1
$\log(1 + TotalUsers_q)$	-0.395***	0.127	0.0019	0.0095	8.298	7.910

Notes: The raw-user row uses untransformed event-day unique users summed across the forecast-round event set. The mean decline is 24.3% ($(3175.1 - 4196.3)/4196.3$). The second row reports a scale-compressing robustness transformation of the same raw total. Both rows use quarter-of-year fixed effects in a 43-quarter panel.

The raw-user specification shows that the quarterly decline is not driven solely by the abnormal-log construction. Precision is weaker on the raw scale under exact permutation inference, but the sign and economic magnitude remain broadly consistent with the baseline result.

6.7 Quarterly internal-control diagnostic (E2)

We next implement an internal-control design using standard meetings. Table 8 shows a coherent pattern. Row 1 is strongly negative, but it compares a forecast-round sum (two events pre-reform, one event post-reform) to a one-event standard-meeting average, so the event-count change is built into the construction. Under unit-matched normalisations (Rows 2–3), the post differential is close to zero and statistically indistinguishable from zero, and the event-level DID interaction is also imprecise. The internal-control evidence therefore indicates that the quarterly Twitter decline is consistent with losing a second forecast-round occasion, while relative per-event forecast salience versus standard meetings is broadly unchanged.

TABLE 8. Internal-Control Diagnostics Using Standard Meetings

Specification	Post Coef.	SE	p-value	Pre Mean	Post Mean
$ForecastSum_q - StandardAvg_q$	-1.309	0.188	< 0.001	1.699	0.381
$ForecastSum_q - n_q \cdot StandardAvg_q$	-0.043	0.231	0.852	0.432	0.381
$ForecastAvg_q - StandardAvg_q$	+0.171	0.151	0.258	0.216	0.381
Event-level DID: $Post \times ForecastType$	+0.133	0.170	0.434	–	–

Notes: Quarterly rows use 43 quarters with quarter-of-year fixed effects and HC2 standard errors. n_q equals 2 pre-reform and 1 post-reform forecast-round events. The event-level DID row uses 111 events (forecast and standard event types) with quarter fixed effects.

TABLE 9. Quarterly Specification Robustness

Specification	Coef.	HC2 SE	HAC SE	p-value (HC2)	N
Baseline quarterly specification	-1.216	0.167	0.182	< 0.0001	43
No COVID (drop 2020–2022 rounds)	-1.157	0.170	0.174	< 0.0001	36
No 2015 (drop implementation year)	-1.060	0.149	0.168	< 0.0001	39
Trimmed 10% by regime deciles	-1.213	0.129	0.120	< 0.0001	33
True reform-date diff in means	-1.224	–	–	< 0.0001 (perm.)	43
Placebo mean (2012–2014 breaks)	-0.908	–	–	–	–

Notes: HAC uses Newey–West with four quarterly lags. The trimmed specification keeps observations within the 10th–90th percentile of $CumLogAtt_q$ separately by regime. Placebo rows report the true reform-date difference in means and the average coefficient from placebo breaks at 2012Q1, 2013Q1, and 2014Q1.

Across robustness variants, estimates range from −1.060 to −1.216, all significant at $p < 0.001$. The narrowing in the "No 2015" specification reflects the removal of the quarters with the sharpest regime contrast, but even the most conservative estimate implies a decline exceeding one index unit. Figure 4 provides a supporting visual diagnostic, as neither the pre-reform subsample (slope = 0.053, p = 0.658) nor the post-reform subsample (slope = −0.039, p = 0.401) exhibits a statistically significant linear time trend, consistent with a level shift at the reform date rather than a gradual secular decline already in motion.

FIGURE 4. Quarterly Evolution with Reform Break and Pre-Trend

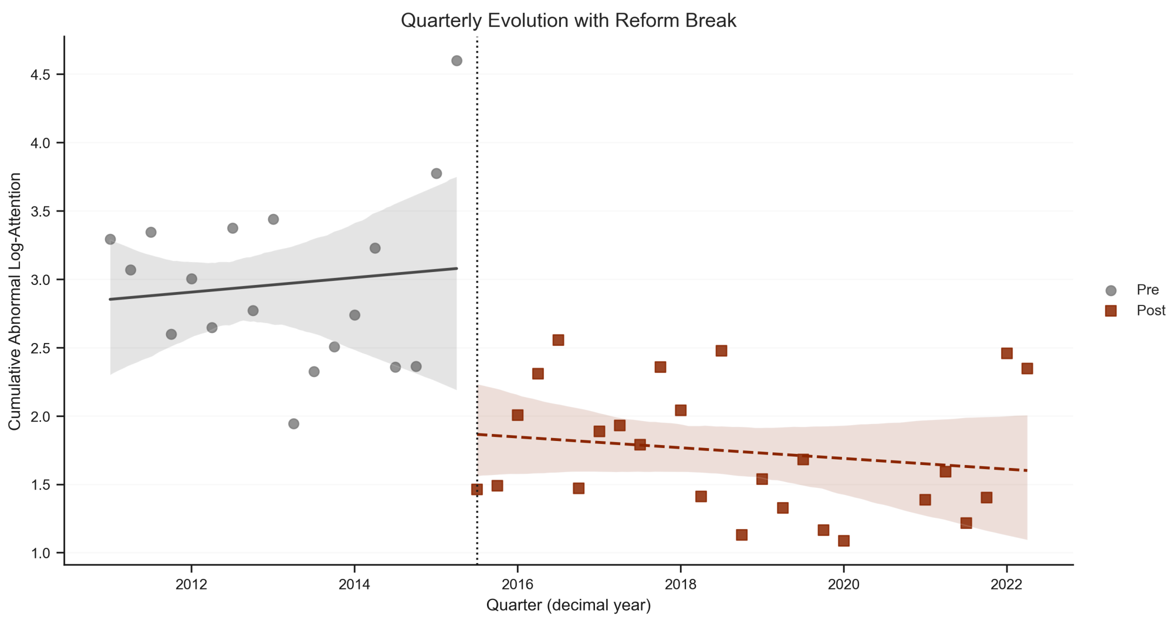

Notes: Scatter plot of $CumLogAtt_q$ against time with separate OLS trend lines and 95 per cent confidence bands for the pre-reform and post-reform periods. The vertical dotted line marks the reform implementation. Neither within-regime trend is statistically significant; the evidence is consistent with a level shift at the reform date.

The post-reform regime generates substantially less cumulative engagement within a quarterly forecast round on Twitter. This estimate is best interpreted as the net consequence of the August 2015 reform package. It combines bundling with simultaneous minutes publication. The finding is still informative about architecture because minutes timing changed within the same day, while forecast-round event count changed from two to one. The descriptive accounting in Table 3 shows that the removed second-event margin is substantially larger than the bundled-day amplification. The pre-reform period comprises 18 quarters, which limits the power of asymptotic break-date tests, so the strongest inferential evidence comes from permutation inference.

6.8 Event-level amplification within the post-reform regime

The quarterly decline motivates a complementary test of whether bundling raises per-event salience. The within-post comparison confirms that it does. Full-release events attract significantly more attention than standard MPC meetings, with an estimated difference of +0.374 abnormal log-user units (s.e. 0.130, $p = 0.004$; Table 5). Moreover, inference is stable to coarser dependence adjustments, and clustering by year yields a smaller standard error and a p-value of 0.0003.

The key descriptive trade-off is that this intensive-margin gain is substantially smaller than the removed second-event component. Let A_{TW}^{dec} and A_{TW}^{rep} denote attention on the two pre-reform forecast-round occasions and A_{TW}^{full} denote attention on the bundled post-reform occasion. The empirical pattern is

$$\mathbb{E}[A_{TW}^{full}] > \mathbb{E}[A_{TW}^{standard}] \quad \text{but} \quad \mathbb{E}[A_{TW}^{full}] < \mathbb{E}[A_{TW}^{dec} + A_{TW}^{rep}].$$

The bundled event is larger than a standard meeting, but smaller than the sum of the two separate occasions it replaces. This is the core empirical trade-off, and it has a natural interpretation in the framework of Paciello and Wiederholt (2014), where bundling raises the marginal return to processing the single remaining event, but total information processing depends on the product of marginal return and the number of processing occasions. In these data, reducing the number of occasions is associated with lower cumulative quarterly Twitter engagement, consistent with the E2 internal-control diagnostics.

This within-post estimate addresses part of the concurrent-reforms concern discussed in Section 5.3. Because both full-release and standard meetings coexist within the same post-reform institutional environment — sharing the same minutes publication rule, the same meeting frequency, and the same communication norms — the +0.374 difference is not driven by those package-wide regime shifts. It reflects the combined effect of forecast-package presence and same-day co-location timing. We interpret it jointly with the forecast-occasion comparison in Table 4. The quarterly specification still captures the net effect of the broader Warsh Review package rather than bundling alone.

6.9 Composition across platforms and search content

Beyond how much attention changes, bundling also changes what the public attends to. Compositional effects are strongest in search and reference behaviour (Table 10).

Google Trends event-day spikes rise by about 21 index points in the post-reform period ($p < 0.001$), indicating substantially more active information seeking around bundled communication events. The term-level decomposition in Appendix D shows the largest increase for “interest rates”

TABLE 10. Composition Tests by Platform

Platform Test	Coefficient	SE	p-value	N
Google Trends (post effect)	+21.29	4.30	< 0.0001	48
Wikipedia (full-release effect within post)	+0.273	0.066	< 0.0001	62
Twitter (post effect)	-0.142	0.259	0.5833	112

Notes: The Google Trends row uses quarter-level event spikes (N = 48 quarter cells). The Wikipedia row is estimated within the post-reform period because Wikipedia coverage begins in mid-2015 and leaves only one pre-reform observation. The Twitter row uses event-level abnormal attention with year and quarter fixed effects.

and a positive but imprecise increase for "inflation," so the search evidence does not support a sharp one-for-one substitution away from outlook terms.

Table 11 reports Twitter topic-share regressions within the post-reform sample. Full-release events are associated with a higher outlook/report share (+0.0969, $p < 0.001$), no detectable shift in rate-focused share (−0.0034, $p = 0.925$), and a small decline in QE/balance-sheet share (−0.0139, $p = 0.028$). This supports the E1 interpretation in Section 6: the forecast-occasion premium is linked to greater engagement with forecast/outlook content, not only to a generic increase in event scale. Composition changes therefore reflect broader reallocation across policy topics rather than a simple rate-focus tilt.

TABLE 11. Twitter Topic-Share Regressions (Post-Reform Events)

Outcome	B_e Coef.	SE	p-value	N	R^2
Outlook/report topic share	+0.0969	0.0221	< 0.001	58	0.33
Rate-focused topic share	-0.0034	0.0364	0.9250	58	0.06
QE/balance-sheet topic share	-0.0139	0.0063	0.0278	58	0.10

Notes: Outcomes are event-level topic shares from the topic-classification procedure described in the appendix. Each row reports the coefficient on the full-release indicator in a post-reform regression with quarter fixed effects and HC2 standard errors.

Wikipedia and broad Twitter levels still behave differently. Wikipedia page views are significantly higher on full-release days than on standard meetings within the post-reform sample, while the broad Twitter post-reform effect is imprecise and does not survive FDR correction. This split likely reflects different behavioural channels, with search and reference activity responding to the bundled information package while social discussion breadth is dominated by event-day salience rather than regime-level shifts.

Figure 5 provides an unconditional platform comparison. Cross-platform differences are best interpreted qualitatively because each platform measures a different behaviour (discussion, search, reference).

FIGURE 5. Platform Comparison of Event-Day Attention (Standardised)

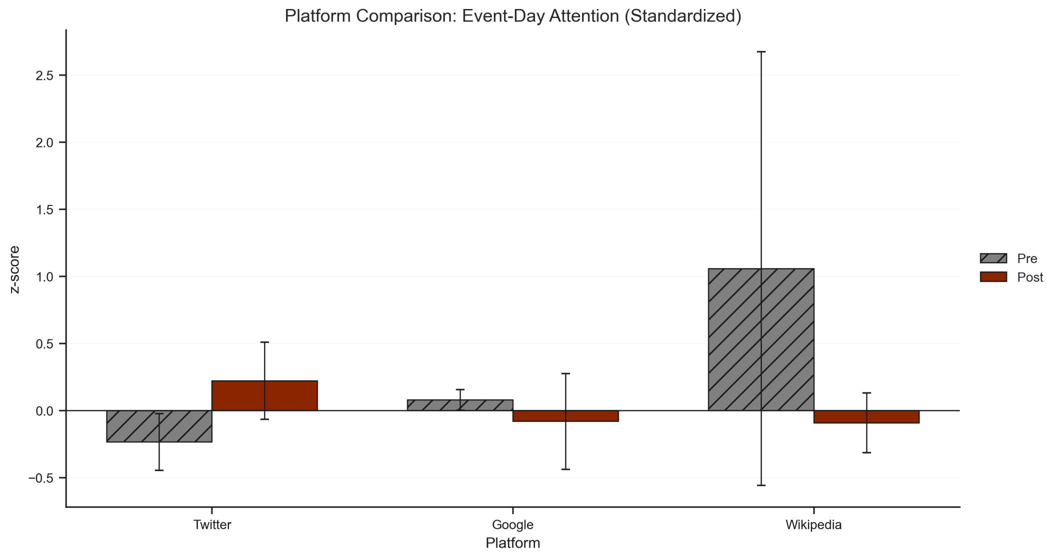

Notes: Mean event-day attention by platform in the pre-reform and post-reform periods, presented on a standardised (z-score) scale for within-platform comparability. Error bars indicate 95 per cent confidence intervals. Because platforms measure different behaviours, cross-platform magnitude comparisons should be interpreted cautiously.

As a robustness check, Appendix Table 18 rebuilds a three-platform index on the overlap sample under three scaling choices (full-sample standardisation, within-regime standardisation, and PCA). The within-post full-release coefficient remains positive and significant in all three versions.

6.10 Attention dynamics

The preceding subsections establish the level and composition of attention; the horizon-specific regressions from Equation (5) reveal its temporal profile. The attention response is concentrated on the announcement day. High-shock events generate an immediate spike of +0.319 in $A_{e,TW}$ at $h = 0$ ($p = 0.014$; permutation $p = 0.030$). By $h = 1$, the response is no longer statistically precise and remains close to zero at longer horizons.

The interaction term δ_h (bundled × high-shock) does not reach significance at any horizon, indicating that bundling amplifies the *level* of the attention response (as documented in the event-level comparison) but does not alter its *temporal shape*. Figure 6 plots the impulse response separately for high-shock and low-shock events.

FIGURE 6. Attention Dynamics with Immediate Spike and Rapid Decay

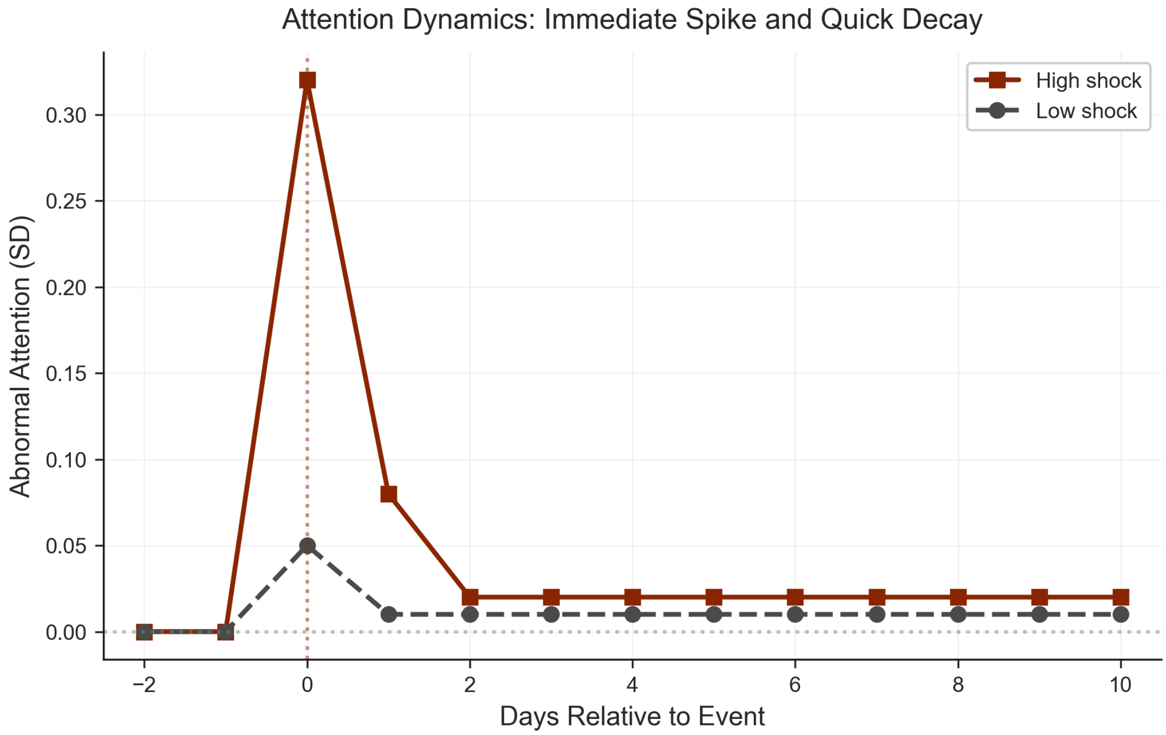

Notes: Impulse response of abnormal attention for high-shock events (solid) and low-shock events (dashed) at horizons $h = -2$ to $h = 10$ business days relative to the announcement. Events are classified as high-shock if the absolute UKMPD surprise exceeds the sample median. The full event universe has 117 MPC dates. The dynamics figure uses the 113 events with non-missing daily Twitter observations; the four excluded dates are 2022-08-04, 2022-09-22, 2022-11-03, and 2022-12-15.

The rapid spike-and-decay profile reinforces the extensive-margin interpretation. Communication architecture determines the number and height of attention spikes within a quarterly round, but cannot stretch the duration of any individual spike. Making one event larger does not create sustained processing; sustained cumulative engagement requires repeated occasions. This temporal pattern is suggestive of intermittent-updating mechanisms (Reis 2006), but we do not treat it as a decisive test against alternative information-processing models (Sims 2003).

6.11 Exploratory extensions on sign asymmetry and state dependence

Exploratory appendix tests allow sign asymmetry and macro-state dependence. The hawkish slope is larger than the dovish slope in baseline specifications ($\hat{\beta}^+ = 9.15$, $p = 0.031$; $\hat{\beta}^- = -1.06$, $p = 0.712$), but identifying variation is concentrated in a few events and attenuation appears when the largest hawkish shock is removed (Appendix E). The high-inflation interaction is positive but imprecise ($+2.36$, $p = 0.238$). We therefore treat these tests as suggestive only.

6.12 Financial market diagnostics

We use market diagnostics to test whether intraday price movements account for the event-level attention premium.

Table 12 summarises the intraday data coverage.

TABLE 12. Market Data Coverage and Event Windows

Series	Start Date	End Date	Raw Observations	Events Covered
GBPUSD (fx)	2010-01-03	2025-11-07	5,868,784	117
FTSE 100 (eq)	2008-01-02	2025-11-07	4,112,230	117
Eurodollar (rates)	2008-01-02	2023-06-19	1,673,197	103–112

Notes: Event coverage is based on non-missing 15-minute absolute-return windows at 12:00 and 12:45 (London time).

Table 13 augments the within-post attention regression with intraday volatility controls. The dependent variable is $A_{e,TW}$ throughout. The baseline market row is not identical to the event-level comparison reported above because it adds surprise, macro controls, and year fixed effects before introducing market variables.

TABLE 13. Market-Augmented Attention Models

Specification	B_e (coef)	SE	p-value	N	R^2	Attenuation
Baseline (no market controls)	0.316	0.102	0.0021	58	0.65	–
+ FX + EQ, 12:00 window	0.301	0.109	0.0060	58	0.65	4.7%
+ FX + EQ + ED, 12:00 window	0.310	0.122	0.0111	57	0.68	1.9%
+ FX + EQ, 12:00 and 12:45	0.291	0.109	0.0077	57	0.69	7.9%
+ FX + EQ + ED, full windows	0.253	0.127	0.0463	54	0.73	19.8%

Notes: Dependent variable is $A_{e,TW}$ (Twitter abnormal log-unique-users). The baseline row includes *ShockMag*, Bank Rate, CPI deviation, and year/quarter fixed effects; it is therefore not numerically identical to the parsimonious event-level regression in Table 5. ED denotes Eurodollar futures absolute 15-minute returns.

The baseline bundling coefficient is 0.316 ($p = 0.002$). With the full intraday control set it attenuates to 0.253 ($p = 0.046$), about a 20% reduction. This indicates a market-mediated channel, but the coefficient remains positive and statistically detectable.

Table 14 asks the reverse question. Does bundling itself alter intraday market dynamics?

Shock magnitude significantly predicts returns in the foreign exchange and interest rate markets at the noon window (GBPUSD, $p = 0.003$; Eurodollar, $p < 0.001$), confirming that our surprise measure captures policy news. By contrast, the full-release indicator is insignificant across all six

TABLE 14. Market Response to Policy Events

Dependent Variable	B_e (coef)	p-value	ShockMag (coef)	p-value	N	R^2
GBPUSD abs. return (12:00)	0.000701	0.1154	0.024531	0.0029	58	0.63
FTSE 100 abs. return (12:00)	0.000219	0.6455	0.016982	0.0745	58	0.35
Eurodollar abs. return (12:00)	-0.000010	0.4123	0.000573	< 0.0001	57	0.56
GBPUSD abs. return (12:45)	0.000172	0.5468	0.006382	0.2859	58	0.25
FTSE 100 abs. return (12:45)	0.000110	0.6367	0.000040	0.9841	57	0.35
Eurodollar abs. return (12:45)	-0.000022	0.1441	0.000412	0.0986	56	0.50

Notes: Post-reform sample only. Regressions include Bank Rate, CPI deviation, and year/quarter fixed effects with HC2 standard errors. Shock magnitude strongly predicts GBPUSD and Eurodollar responses in the 12:00 window.

specifications. We therefore do not detect a statistically significant differential market response to bundling conditional on surprise magnitude.

Across the family of primary tests, the quarterly and within-post event-level findings survive conservative multiple-testing correction (Table 15). We include the E1 Twitter one-sided permutation result in this family; it survives FDR at 5%. Robustness variants of the quarterly specification are not entered into the FDR correction because they test the same hypothesis on modified samples.

TABLE 15. FDR-Corrected Primary Test Results

Hypothesis	Test	Raw p	FDR-Adjusted p	Significant
Forecast-round	Post effect	< 0.0001	< 0.0001	Yes
Google Trends	Post effect	< 0.0001	< 0.0001	Yes
Wikipedia	Full-release effect	< 0.0001	< 0.0001	Yes
Event-level	Full-release effect (post only)	0.0040	0.0090	Yes
Dynamics	High-shock immediate effect	0.0135	0.0243	Yes
E1 Twitter	Forecast-occasion post effect (perm., one-sided)	0.0253	0.0380	Yes
Asymmetry	Hawkish surprise slope (β^+)	0.0312	0.0401	Yes
State dep.	High-inflation interaction	0.2384	0.2682	No
Twitter	Post effect	0.5833	0.5833	No

Notes: Benjamini–Hochberg correction at $\alpha = 0.05$ across nine distinct hypotheses. For E1 Twitter we use the one-sided permutation p-value as the raw input because Equation (2) has directional prediction and the event sample is $N = 43$. Robustness variants of the quarterly specification are excluded from this family by construction.

The evidence falls into a clear hierarchy. First, the forecast-occasion comparison provides targeted architecture evidence: forecast-occasion co-location increases attention on both Twitter and Google Trends. Second, within-post event-level evidence shows a robust full-release salience premium under common post-reform institutions. Third, quarterly Twitter engagement is lower after the reform, and E2 internal-control diagnostics show that this is consistent with replacing two forecast-round occasions with one while relative per-event forecast salience versus standard meetings is broadly unchanged. Fourth, composition evidence shows that architecture changes what users engage with, not only how much. Cross-platform quarterly diagnostics remain heterogeneous, so quarterly cumulative effects should be interpreted as channel-specific. Asymmetry and state-dependence tests remain exploratory.

7. Conclusion

The Warsh Review reform changed more than the scheduling of Bank of England communication. It changed the way in which monetary policy reached the public. Before the reform, the policy decision and the quarterly forecast appeared on separate occasions within a forecast round. After the reform, they were released together. The forecast-occasion comparison shows a significant co-location premium on both Twitter and Google Trends, and within-post estimates show a clear salience premium for full-release meetings relative to standard meetings.

Quarterly Twitter engagement remains lower after August 2015. Descriptive accounting shows that the bundled-day gain (+0.232) is small relative to the removed second-event component (+1.456), and raw-user robustness gives the same directional result. E2 internal-control diagnostics indicate that this quarterly decline is consistent with an event-count effect when forecast rounds move from two occasions to one, while relative per-event forecast salience versus standard meetings is broadly unchanged. Cross-platform quarterly diagnostics remain heterogeneous: the Google Trends analogue rises after the reform, while Wikipedia lacks pre-reform coverage for the same design.

Several directions for future work follow directly from the boundaries of our analysis. Our digital trace measures capture revealed engagement, not comprehension or expectation updating; pairing behavioural attention data with survey-based expectations measures (Coibion, Gorodnichenko, and Weber 2022) would establish whether engagement changes map into expectation formation. The sign-asymmetry extension should be tested in larger samples before any structural interpretation. Finally, integrating intermediary media data would clarify whether lower cumulative engagement reflects the loss of a separate news hook, increased package complexity, or both.

Central banks have invested substantial effort in improving *what* they say. Our evidence suggests that *when* and *how* they say it — the architecture of communication events — can be an important determinant of public engagement with monetary policy.

References

Altavilla, Carlo, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Giuseppe Ragusa. 2019. "Measuring Euro Area Monetary Policy." *Journal of Monetary Economics* 108: 162–179.

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. 2022. "Subjective Models of the Macroeconomy: Evidence from Experts and Representative Samples." *Review of Economic Studies* 89 (6): 2958–2991.

Angeletos, George-Marios, and Chen Lian. 2018. "Forward Guidance without Common Knowledge." *American Economic Review* 108 (9): 2477–2512.

Bank of England. 2014. "Bank of England Announces Measures to Bolster Transparency and Accountability." https://www.bankofengland.co.uk/-/media/boe/files/news/2014/december/boe-announces-measures-to-bolster-transparency-and-accountability.pdf. Accessed 2026-02-26.

Bank of England. 2015. "Operational Notice - Inflation Report Press Conference." https://www.bankofengland.co.uk/news/2015/july/operational-notice-inflation-report-press-conference. Accessed 2026-02-26.

Benjamini, Yoav, and Yosef Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society: Series B (Methodological)* 57 (1): 289–300.

Bholat, David, Nicola Broughton, Johanna Ter Meer, and Erica Walczak. 2019. "Enhancing Central Bank Communications Using Simple and Relatable Information." *Journal of Monetary Economics* 108: 1–15.

Blinder, Alan S., Michael Ehrmann, Jakob De Haan, and David-Jan Jansen. 2024. "Central Bank Communication with the General Public: Promise or False Hope?" *Journal of Economic Literature* 62 (2): 425–457.

Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. 2008. "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence." *Journal of Economic Literature* 46 (4): 910–945.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice under Risk." *Quarterly Journal of Economics* 127 (3): 1243–1285.

Braun, Robin, Silvia Miranda-Agrippino, and Tuli Saha. 2025. "Measuring Monetary Policy in the UK: The UK Monetary Policy Event-Study Database." *Journal of Monetary Economics* 149: 103645.

Buelens, Christian. 2025. "Googling 'Inflation': Household Inflation Attention across the Euro Area." *European Journal of Political Economy* 89: 102702.

Cebrián, Eduardo, and Josep Domenech. 2024. "Addressing Google Trends Inconsistencies." *Technological Forecasting and Social Change* 202: 123318.

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

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

Da, Zhi, Joseph Engelberg, and Pengjie Gao. 2011. "In Search of Attention." *The Journal of Finance* 66 (5): 1461–1499.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. 2020. "Effective Policy Communication: Targets Versus Instruments." Working Paper 2020-148, Becker Friedman Institute for Research in Economics.

Ehrmann, Michael, and Alena Wabitsch. 2022. "Central Bank Communication with Non-Experts – A Road to Nowhere?" *Journal of Monetary Economics* 127: 69–85.

Falkinger, Josef. 2008. "Limited Attention as a Scarce Resource in Information-Rich Economies." Discussion Paper 1538, IZA.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2025. "Central Bank Communication on Social Media: What, to Whom, and How?" *Journal of Econometrics* 249 (Part C): 105869.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 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.

Haldane, Andrew, Alistair Macaulay, and Michael McMahon. 2020. “The 3 E’s of Central Bank Communication with the Public.” Staff Working Paper 847, Bank of England.

Haldane, Andrew, and Michael McMahon. 2018. “Central Bank Communications and the General Public.” *AEA Papers and Proceedings* 108: 578–583.

Hansen, Stephen, Michael McMahon, and Matthew Tong. 2019. “The Long-Run Information Effect of Central Bank Communication.” *Journal of Monetary Economics* 108: 185–202.

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

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica* 47 (2): 263–291.

Kansoy, Fatih, and Joel Mundy. 2025. “Central Bank Communication with Public: Bank of England and Twitter (X).” Oxford University Department of Economics Discussion Paper Series 2506.02559, arXiv.

Kőszegi, Botond, and Matthew Rabin. 2006. “A Model of Reference-Dependent Preferences.” *Quarterly Journal of Economics* 121 (4): 1133–1165.

Kuttner, Kenneth N. 2001. “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market.” *Journal of Monetary Economics* 47 (3): 523–544.

Lamla, Michael J., and Dmitri V. Vinogradov. 2019. “Central Bank Announcements: Big News for Little People?” *Journal of Monetary Economics* 108: 21–38.

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

Masciandaro, Donato, Oana Peia, and Davide Romelli. 2023. “Central Bank Communication and Social Media: From Silence to Twitter.” *Journal of Economic Surveys* 37 (4): 1249–1293.

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

Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. “Optimal Sticky Prices under Rational Inattention.” *American Economic Review* 99 (3): 769–803.

Munday, Tim, and James Brookes. 2021. “Mark My Words: The Transmission of Central Bank Communication to the General Public via the Print Media.” Staff Working Paper 944, Bank of England.

Naraina, Namrata, and Kunal Sangani. 2026. “The Market Impact of Fed Communications: The Role of the Press Conference.” *International Journal of Central Banking* 22 (1): 313–389.

Paciello, Luigi, and Mirko Wiederholt. 2014. “Exogenous Information, Endogenous Information, and Optimal Monetary Policy.” *Review of Economic Studies* 81 (1): 356–388.

Reis, Ricardo. 2006. “Inattentive Consumers.” *Journal of Monetary Economics* 53 (8): 1761–1800.

Santoro, Emiliano, Ivan Petrella, Damjan Pfajfar, and Edoardo Gaffeo. 2014. “Loss Aversion and the Asymmetric Transmission of Monetary Policy.” *Journal of Monetary Economics* 68: 19–36.

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

Soroka, Stuart N. 2006. “Good News and Bad News: Asymmetric Responses to Economic Information.” *Journal of Politics* 68 (2): 372–385.

Swanson, Eric T. 2021. “Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets.” *Journal of Monetary Economics* 118: 32–53.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent Hedlund Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce. 2025. “Tell Me Something I Don’t Already Know: Learning in Low- and High-Inflation Settings.” *Econometrica* 93 (1): 229–264.

Wohlfarth, Paul. 2018. “Using Search Data to Capture Attention: A New Measure of Policy Uncertainty.” *Economics Letters* 162: 41–45.

Appendix

Central Bank Communication Design and Public Attention

This appendix collects construction details and supplementary robustness results.

A. Variable Construction

This appendix provides the implementation details underlying the attention measures described in Section 4.

A.1 Keyword Lexicon and Digital-Trace Construction

We identify policy-relevant tweets using the following keyword set, applied with case-insensitive matching to the full text of each original tweet (retweets excluded):

> `Bank of England`, `BoE`, `MPC`, `Monetary Policy Committee`, `interest rate`, `interest rates`, `base rate`, `Bank Rate`, `quantitative easing`, `QE`, `Inflation Report`, `Monetary Policy Report`.

The lexicon is designed to capture the breadth of public references to Bank of England monetary policy whilst excluding terms (such as "inflation" alone or "economy") that would introduce excessive noise from non-policy contexts. We require at least one keyword match per tweet. The lexicon is held constant across the full sample period (2011–2022) to avoid introducing measurement breaks that could confound the pre-to-post comparison.

The Twitter measure uses original tweets only (retweets excluded). Unique users are deduplicated within each London calendar day; daily boundaries follow Europe/London time, so DST transitions are handled by local time conversion before aggregation. The daily panel retains zero-activity days where present in the source and flags source-coverage gaps. Four late-sample MPC dates have missing daily Twitter observations in the derived series (2022-08-04, 2022-09-22, 2022-11-03, and 2022-12-15) and are excluded from regressions that require $A_{e,TW}$. We complement transformation robustness with an ex-post activity screen based on the raw event-day tweet universe, excluding users who appear on a very large share of MPC dates. The full-release coefficient remains positive and significant under these filters (Table 16).

A.2 Twitter topic-share diagnostics

To tighten the composition evidence, we estimate post-reform event-level regressions of topic shares on the full-release indicator with quarter fixed effects. Table 17 reports the results.

TABLE 16. Twitter Measurement Robustness

Outcome Definition	B_e Coef.	SE	p-value	N	R^2
Baseline same-weekday abnormal log-users ($A_{e,TW}$)	+0.374	0.130	0.0040	58	0.22
Trailing-window baseline abnormal log-users	+0.289	0.125	0.0207	58	0.20
Event-day log(1+users) level (no baseline subtraction)	+0.448	0.135	< 0.001	58	0.23
Event-day log(1+users), exclude users active on ≥10% of MPC dates	+0.524	0.151	0.0005	58	0.24
Event-day log(1+users), exclude users active on ≥20% of MPC dates	+0.497	0.144	0.0006	58	0.24
Winsorised abnormal log-users (1st–99th pct.)	+0.371	0.129	0.0040	58	0.22

Notes: All rows use the post-reform sample and quarter fixed effects. The activity-screen rows use raw tweet-level data and exclude users active on at least 10% (3,384 users) or 20% (1,178 users) of MPC event dates. The full-release effect remains positive and statistically significant across transformations and filters.

TABLE 17. Appendix Twitter Topic-Share Regressions

Outcome	B_e Coef.	SE	p-value	N	R^2
Outlook/report topic share	+0.0969	0.0221	< 0.001	58	0.33
Rate-focused topic share	-0.0034	0.0364	0.9250	58	0.06
QE/balance-sheet topic share	-0.0139	0.0063	0.0278	58	0.10

Notes: Outcomes are event-level topic shares from the topic-classification procedure described in this appendix. Each row reports the coefficient on the full-release indicator in a post-reform regression with quarter fixed effects and HC2 standard errors.

Google Trends queries are collected in "search term" mode, restricted to UK geography. The term set used in event-level summaries includes "Bank of England", "interest rates", and "inflation". The stitched daily series are built from overlapping extraction windows with overlap-based level rescaling. Repeated extraction noise is an inherent feature of Google Trends sampling; we therefore treat the Google series as one component in a multi-platform design rather than as the sole outcome.

Wikipedia data are collected from the Wikimedia Pageviews API for the Bank of England page and closely related monetary policy pages. Queries use daily frequency and all-access/all-agents aggregation. Redirect handling and page-title normalisation follow the API defaults used in our data-construction procedure.

A.3 Abnormal attention and quarterly aggregation

For each platform $p \in \{\text{Twitter, Google Trends, Wikipedia}\}$ and each MPC event e, we construct abnormal attention in two steps. The first applies platform-specific transformations, with Twitter and Wikipedia in natural logs of daily counts while Google Trends is already a normalised index and is left in levels. The second subtracts a day-of-week-matched rolling baseline. For event date $t(e)$, the baseline is the mean of non-event days in a trailing 30-day window that share the same weekday. Days in the $[-1, +1]$ neighbourhood of other MPC events are excluded from the baseline window. The abnormal measure is

$$A_{e,p} = Y_{e,p} - \bar{Y}_{baseline,p} \tag{6}$$

where $Y_{e,p}$ is the event-day measure and $\bar{Y}_{baseline,p}$ is the rolling baseline. The main event-level outcome is Twitter abnormal attention, $A_{e,TW}$. The quarterly cumulative outcome is

$$CumLogAtt_q = \sum_{e \in \mathcal{E}_q} A_{e,TW} \tag{7}$$

where $\mathcal{E}_q$ is the event set for forecast round q.

A.4 Composite stability checks

Although the main outcome is $A_{e,TW}$, we verify that the within-post bundling pattern is stable in a three-platform overlap sample (Twitter, Google Trends, Wikipedia) under alternative index constructions (Table 18).

A.5 Interpretation of log-units and percentage decline

The 41 per cent decline reported in the main text is computed as the proportional change in means:

TABLE 18. Composite Stability Checks (Overlap Sample)

Composite Definition	B_e Coef.	SE	p-value	N	R^2
Full-sample standardisation	+1.515	0.560	0.0069	58	0.20
Within-regime standardisation	+1.521	0.561	0.0067	58	0.20
PCA first component	+1.071	0.334	0.0013	58	0.21

Notes: Overlap sample with non-missing Twitter, Google Trends, and Wikipedia event-day values ($N = 59$, post-only $N = 58$). Cross-platform correlations in the overlap sample with corr(TW,GT)=0.10, corr(TW,WK)=0.79, corr(GT,WK)=0.04.

$$\frac{\overline{CumLogAtt}_{post} - \overline{CumLogAtt}_{pre}}{\overline{CumLogAtt}_{pre}} = \frac{1.742 - 2.966}{2.966} \approx -0.413 \tag{8}$$

This is a ratio-of-means calculation, not a log-point exponentiation. The quarterly object is

$$CumLogAtt_q = \sum_{e \in \mathcal{E}_q} [\log(Users_e) - Baseline_e] \tag{9}$$

so inference is based on summed deviations from matched baselines. Jensen-style comparisons between $\log(U_1) + \log(U_2)$ and $\log(U_1 + U_2)$ do not characterise this deviation-based outcome. The decomposition in Table 3 therefore addresses an extensive-versus-intensive margin trade-off, not a mechanical log-additivity artefact.

B. MPC Event Calendar

The master calendar below is the authoritative event list used in this paper. It records the event date (London time), event type, announcement time, 12:45 follow-up window flag, forecast-round identifier, and whether the event enters the baseline quarterly sum.

TABLE 19. Master Event Calendar and Baseline Quarterly Inclusion

Date	Event Type	Announce	12:45	Round ID	Baseline
2011-01-13	Pre-reform standard MPC meeting	12:00	–	–	No
2011-02-10	Pre-reform forecast-round decision	12:00	–	2011Q1	Yes
2011-02-16	Pre-reform Inflation Report day	–	–	2011Q1	Yes
2011-03-10	Pre-reform standard MPC meeting	12:00	–	–	No
2011-04-07	Pre-reform standard MPC meeting	12:00	–	–	No
2011-05-05	Pre-reform forecast-round decision	12:00	–	2011Q2	Yes
2011-05-11	Pre-reform Inflation Report day	–	–	2011Q2	Yes
2011-06-09	Pre-reform standard MPC meeting	12:00	–	–	No
2011-07-07	Pre-reform standard MPC meeting	12:00	–	–	No
2011-08-04	Pre-reform forecast-round decision	12:00	–	2011Q3	Yes
2011-08-10	Pre-reform Inflation Report day	–	–	2011Q3	Yes
2011-09-08	Pre-reform standard MPC meeting	12:00	–	–	No

Continued on next page

Date	Event Type	Announce	12:45	Round ID	Baseline
2011-10-06	Pre-reform standard MPC meeting	12:00	–	–	No
2011-11-10	Pre-reform forecast-round decision	12:00	–	2011Q4	Yes
2011-11-16	Pre-reform Inflation Report day	–	–	2011Q4	Yes
2011-12-08	Pre-reform standard MPC meeting	12:00	–	–	No
2012-01-12	Pre-reform standard MPC meeting	12:00	–	–	No
2012-02-09	Pre-reform forecast-round decision	12:00	–	2012Q1	Yes
2012-02-15	Pre-reform Inflation Report day	–	–	2012Q1	Yes
2012-03-08	Pre-reform standard MPC meeting	12:00	–	–	No
2012-04-05	Pre-reform standard MPC meeting	12:00	–	–	No
2012-05-10	Pre-reform forecast-round decision	12:00	–	2012Q2	Yes
2012-05-16	Pre-reform Inflation Report day	–	–	2012Q2	Yes
2012-06-07	Pre-reform standard MPC meeting	12:00	–	–	No
2012-07-05	Pre-reform standard MPC meeting	12:00	–	–	No
2012-08-02	Pre-reform forecast-round decision	12:00	–	2012Q3	Yes
2012-08-08	Pre-reform Inflation Report day	–	–	2012Q3	Yes
2012-09-06	Pre-reform standard MPC meeting	12:00	–	–	No
2012-10-04	Pre-reform standard MPC meeting	12:00	–	–	No
2012-11-08	Pre-reform forecast-round decision	12:00	–	2012Q4	Yes
2012-11-14	Pre-reform Inflation Report day	–	–	2012Q4	Yes
2012-12-06	Pre-reform standard MPC meeting	12:00	–	–	No
2013-01-10	Pre-reform standard MPC meeting	12:00	–	–	No
2013-02-07	Pre-reform forecast-round decision	12:00	–	2013Q1	Yes
2013-02-13	Pre-reform Inflation Report day	–	–	2013Q1	Yes
2013-03-07	Pre-reform standard MPC meeting	12:00	–	–	No
2013-04-04	Pre-reform standard MPC meeting	12:00	–	–	No
2013-05-09	Pre-reform forecast-round decision	12:00	–	2013Q2	Yes
2013-05-15	Pre-reform Inflation Report day	–	–	2013Q2	Yes

Continued on next page

Date	Event Type	Announce	12:45	Round ID	Baseline
2013-06-06	Pre-reform standard MPC meeting	12:00	–	–	No
2013-07-04	Pre-reform standard MPC meeting	12:00	–	–	No
2013-08-01	Pre-reform forecast-round decision	12:00	–	2013Q3	Yes
2013-08-07	Pre-reform Inflation Report day	–	–	2013Q3	Yes
2013-09-05	Pre-reform standard MPC meeting	12:00	–	–	No
2013-10-10	Pre-reform standard MPC meeting	12:00	–	–	No
2013-11-07	Pre-reform forecast-round decision	12:00	–	2013Q4	Yes
2013-11-13	Pre-reform Inflation Report day	–	–	2013Q4	Yes
2013-12-05	Pre-reform standard MPC meeting	12:00	–	–	No
2014-01-09	Pre-reform standard MPC meeting	12:00	–	–	No
2014-02-06	Pre-reform forecast-round decision	12:00	–	2014Q1	Yes
2014-02-12	Pre-reform Inflation Report day	–	–	2014Q1	Yes
2014-03-06	Pre-reform standard MPC meeting	12:00	–	–	No
2014-04-10	Pre-reform standard MPC meeting	12:00	–	–	No
2014-05-08	Pre-reform forecast-round decision	12:00	–	2014Q2	Yes
2014-05-14	Pre-reform Inflation Report day	–	–	2014Q2	Yes
2014-06-05	Pre-reform standard MPC meeting	12:00	–	–	No
2014-07-10	Pre-reform standard MPC meeting	12:00	–	–	No
2014-08-07	Pre-reform forecast-round decision	12:00	–	2014Q3	Yes
2014-08-13	Pre-reform Inflation Report day	–	–	2014Q3	Yes
2014-09-04	Pre-reform standard MPC meeting	12:00	–	–	No
2014-10-09	Pre-reform standard MPC meeting	12:00	–	–	No
2014-11-06	Pre-reform forecast-round decision	12:00	–	2014Q4	Yes

Continued on next page

Date	Event Type	Announce	12:45	Round ID	Baseline
2014-11-12	Pre-reform Inflation Report day	–	–	2014Q4	Yes
2014-12-04	Pre-reform standard MPC meeting	12:00	–	–	No
2015-01-08	Pre-reform standard MPC meeting	12:00	–	–	No
2015-02-05	Pre-reform forecast-round decision	12:00	–	2015Q1	Yes
2015-02-12	Pre-reform Inflation Report day	–	–	2015Q1	Yes
2015-03-05	Pre-reform standard MPC meeting	12:00	–	–	No
2015-04-09	Pre-reform standard MPC meeting	12:00	–	–	No
2015-05-11	Pre-reform forecast-round decision	12:00	–	2015Q2	Yes
2015-05-13	Pre-reform Inflation Report day	–	–	2015Q2	Yes
2015-06-04	Pre-reform standard MPC meeting	12:00	–	–	No
2015-07-09	Pre-reform standard MPC meeting	12:00	–	–	No
2015-08-06	Post-reform full-release day	12:00	12:45	2015Q3	Yes
2015-09-10	Post-reform standard MPC meeting	12:00	–	–	No
2015-10-08	Post-reform standard MPC meeting	12:00	–	–	No
2015-11-05	Post-reform full-release day	12:00	12:45	2015Q4	Yes
2015-12-10	Post-reform standard MPC meeting	12:00	–	–	No
2016-01-14	Post-reform standard MPC meeting	12:00	–	–	No
2016-02-04	Post-reform full-release day	12:00	12:45	2016Q1	Yes
2016-03-17	Post-reform standard MPC meeting	12:00	–	–	No
2016-04-14	Post-reform standard MPC meeting	12:00	–	–	No
2016-05-12	Post-reform full-release day	12:00	12:45	2016Q2	Yes
2016-06-16	Post-reform standard MPC meeting	12:00	–	–	No
2016-07-14	Post-reform standard MPC meeting	12:00	–	–	No
2016-08-04	Post-reform full-release day	12:00	12:45	2016Q3	Yes

Continued on next page

Date	Event Type	Announce	12:45	Round ID	Baseline
2016-09-15	Post-reform standard MPC meeting	12:00	–	–	No
2016-11-03	Post-reform full-release day	12:00	12:45	2016Q4	Yes
2016-12-15	Post-reform standard MPC meeting	12:00	–	–	No
2017-02-02	Post-reform full-release day	12:00	12:45	2017Q1	Yes
2017-03-16	Post-reform standard MPC meeting	12:00	–	–	No
2017-05-11	Post-reform full-release day	12:00	12:45	2017Q2	Yes
2017-06-15	Post-reform standard MPC meeting	12:00	–	–	No
2017-08-03	Post-reform full-release day	12:00	12:45	2017Q3	Yes
2017-09-14	Post-reform standard MPC meeting	12:00	–	–	No
2017-11-02	Post-reform full-release day	12:00	12:45	2017Q4	Yes
2017-12-14	Post-reform standard MPC meeting	12:00	–	–	No
2018-02-08	Post-reform full-release day	12:00	12:45	2018Q1	Yes
2018-03-22	Post-reform standard MPC meeting	12:00	–	–	No
2018-05-10	Post-reform full-release day	12:00	12:45	2018Q2	Yes
2018-06-21	Post-reform standard MPC meeting	12:00	–	–	No
2018-08-02	Post-reform full-release day	12:00	12:45	2018Q3	Yes
2018-09-13	Post-reform standard MPC meeting	12:00	–	–	No
2018-11-01	Post-reform full-release day	12:00	12:45	2018Q4	Yes
2018-12-20	Post-reform standard MPC meeting	12:00	–	–	No
2019-02-07	Post-reform full-release day	12:00	12:45	2019Q1	Yes
2019-03-21	Post-reform standard MPC meeting	12:00	–	–	No
2019-05-02	Post-reform full-release day	12:00	12:45	2019Q2	Yes
2019-06-20	Post-reform standard MPC meeting	12:00	–	–	No
2019-08-01	Post-reform full-release day	12:00	12:45	2019Q3	Yes
2019-09-19	Post-reform standard MPC meeting	12:00	–	–	No

Continued on next page

Date	Event Type	Announce	12:45	Round ID	Baseline
2019-11-07	Post-reform full-release day	12:00	12:45	2019Q4	Yes
2019-12-19	Post-reform standard MPC meeting	12:00	–	–	No
2020-01-30	Post-reform full-release day	12:00	12:45	2020Q1	Yes
2020-03-19	Post-reform standard MPC meeting (unscheduled emergency)	14:33	–	–	No
2020-03-26	Post-reform standard MPC meeting	12:00	–	–	No
2020-06-18	Post-reform standard MPC meeting	12:00	–	–	No
2020-09-17	Post-reform standard MPC meeting	12:00	–	–	No
2020-12-17	Post-reform standard MPC meeting	12:00	–	–	No
2021-02-04	Post-reform full-release day	12:00	12:45	2021Q1	Yes
2021-03-18	Post-reform standard MPC meeting	12:00	–	–	No
2021-05-06	Post-reform full-release day	12:00	12:45	2021Q2	Yes
2021-06-24	Post-reform standard MPC meeting	12:00	–	–	No
2021-08-05	Post-reform full-release day	12:00	12:45	2021Q3	Yes
2021-09-23	Post-reform standard MPC meeting	12:00	–	–	No
2021-11-04	Post-reform full-release day	12:00	12:45	2021Q4	Yes
2021-12-16	Post-reform standard MPC meeting	12:00	–	–	No
2022-02-03	Post-reform full-release day	12:00	12:45	2022Q1	Yes
2022-03-17	Post-reform standard MPC meeting	12:00	–	–	No
2022-05-05	Post-reform full-release day	12:00	12:45	2022Q2	Yes
2022-06-16	Post-reform standard MPC meeting	12:00	–	–	No
2022-08-04	Late-sample full-release (missing Twitter daily data)	12:00	12:45	–	No
2022-09-22	Post-reform standard MPC meeting	12:00	–	–	No
2022-11-03	Late-sample full-release (missing Twitter daily data)	12:00	12:45	–	No

Continued on next page

Date	Event Type	Announce	12:45	Round ID	Baseline
2022-12-15	Post-reform standard MPC meeting	12:00	–	–	No

Calendar construction rules are as follows. The event universe has 117 MPC dates from the master event panel. Pre-reform Inflation Report dates (18 rows) are added explicitly because they are separate from MPC decision dates. The baseline quarterly specification includes 61 events in total, 18 pre-reform decisions, 18 pre-reform Inflation Report days, and 25 post-reform full-release days. The unscheduled emergency announcement on 2020-03-19 (14:33 London time) is retained in event-level analyses. The January 2016 meeting-frequency reduction is handled by the observed event schedule, and no ad hoc event deletions are applied. Delayed pre-reform minutes-release dates are not included in the baseline quarterly construction.

C. Permutation Inference

The permutation test for the baseline quarterly specification proceeds as follows.

a. Estimate Equation (1) on the observed data and record $\hat{\beta} = -1.216$.
b. For each of $B = 9{,}999$ iterations, randomly reassign the $Post_q$ indicator across the 43 quarters without replacement, re-estimate the model, and record the placebo coefficient $\hat{\beta}^{(b)}$.
c. Compute the one-sided permutation p-value as:

$$p_{perm} = \frac{1 + \sum_{b=1}^{B} \mathbf{1}\{\hat{\beta}^{(b)} \leq \hat{\beta}\}}{1 + B} \tag{10}$$

The observed estimate of −1.216 is more negative than all 9,999 placebo estimates, yielding $p_{perm} = 1/10{,}000 < 10^{-4}$. This exact inference does not rely on asymptotic assumptions and is therefore valid despite the modest sample of 43 quarters.

The placebo break-date analysis is distinct from the permutation test. For each year $t' \in \{2012, 2013, 2014\}$, we re-estimate Equation (1) with the reform indicator redefined as $Post_{t'} = \mathbf{1}\{q \geq t'\text{Q1}\}$ and record the resulting coefficient. The placebo estimates are 2012 (−0.920, $p < 0.001$), 2013 (−0.947, $p < 0.001$), and 2014 (−0.857, $p < 0.001$). All are negative, reflecting the fact that any break date captures some of the true post-2015 decline in the post-placebo sample. The true reform-date estimate of −1.216 exceeds all three placebos by at least 0.27 index units, confirming that the effect is specific to the 2015 architectural change. The negative sign of the placebos does not indicate a pre-existing trend, as shown by the flat pre-reform trend in Figure 4 (slope = 0.053, $p = 0.658$); rather, it is a mechanical consequence of including post-reform observations in the placebo post-period.

D. Google Trends Search Term Decomposition

The composition result reported in Section 6 aggregates across Google Trends search terms. Table 20 decomposes the post-reform effect by individual search term to identify which queries drive the overall increase.

TABLE 20. Google Trends Post-Reform Effect by Search Term

Search Term	Coefficient	SE	p-value
"Bank of England"	+14.8	4.1	< 0.001
"interest rates"	+26.3	5.7	< 0.001
"inflation"	+8.2	6.9	0.241

Notes: Each row reports the coefficient on the post-reform indicator in a regression of the indicated Google Trends search intensity index on the post-reform indicator and quarter-of-year fixed effects. HC2 robust standard errors. The "interest rates" term shows the largest post-reform increase, while the "inflation" term is positive but imprecise. The decomposition indicates stronger instrument salience without pinning down a sharp substitution away from forecast-outlook terms.

The decomposition refines the composition interpretation in the main text. The "interest rates" query shows the largest and most precisely estimated increase (+26.3, $p < 0.001$), whilst the "inflation" query — which more closely proxies attention to the broader economic outlook — is positive but imprecise (+8.2, $p = 0.241$). The search evidence therefore indicates stronger instrument salience, but not a definitive crowding out of outlook-oriented attention.

E. Hawkish–Dovish Split Specification

The hawkish–dovish asymmetry reported in Section 6 is estimated by augmenting Equation (3) with separate slope coefficients for positive and negative UKMPD composite surprises:

$$A_{e,TW} = \alpha + \beta^{+} \cdot Shock_e^{+} + \beta^{-} \cdot Shock_e^{-} + \gamma \cdot B_e + \delta_{year} + \theta_{qoy} + \varepsilon_e \tag{11}$$

where $Shock_e^{+} = \max(Shock_e, 0)$ and $Shock_e^{-} = \min(Shock_e, 0)$. Both variables enter in levels (not absolute values), so β^{-} is estimated on negative surprises.

The estimated response is much stronger on the hawkish side with $\hat{\beta}^+ = 9.15$ ($p = 0.031$) and $\hat{\beta}^- = -1.06$ ($p = 0.712$). A formal Wald test of the null $\beta^+ = \beta^-$ yields $F = 3.28$ ($p = 0.073$). The estimated difference, $\hat{\beta}^+ - \hat{\beta}^- = 10.21$, has a 95% confidence interval of $[-0.84,\ 21.27]$. The identifying variation for the hawkish coefficient is concentrated in a small number of large positive-surprise events. When the single largest hawkish surprise is excluded, the positive-surprise coefficient falls but remains statistically significant at conventional levels. The qualitative conclusion is unchanged. Public engagement responds more strongly to unexpected tightening than to unexpected easing, but the exact numerical ratio should be interpreted cautiously.