# Central Bank Communication with Public*

Fatih Kansoy **
University of Oxford

Joel Mundy #
Bank of England

March 2026

**Abstract**

Which characteristics of official Bank of England tweets are associated with public engagement? We study the complete record of 9,810 tweets posted by the Bank of England on Twitter/X between July 2011 and July 2022, estimating tweet-level Poisson models with progressively tighter time controls up to week fixed effects. Media format and policy timing dominate. Photo tweets are associated with 181 per cent higher engagement, video tweets with 159 per cent, and Monetary Policy Committee (MPC)-day tweets with 150 per cent. The MPC-day association is substantially larger during the post-2021 high-inflation period. Readability is positively associated with engagement, though the effect is smaller and more sensitive to modelling choices. Reply tweets attract about 89 per cent less engagement than original posts, pointing to a tension between reach and dialogue. Robustness checks across alternative estimators and sample restrictions preserve the core ranking. Estimates are associational throughout.

**Keywords:** Central Bank Communication, Social Media, Public Engagement, Monetary Policy

**JEL Classification:** E52, E58, D83

*The views expressed are those of the authors and do not necessarily reflect the views of any institution.

**Corresponding author: fatih.kansoy@economics.ox.ac.uk
#Contact: joel.mundy@bankofengland.co.uk

# 1. Introduction

Central banks now communicate with the public in crowded digital feeds rather than only through speeches, minutes and press conferences. Yet publication and attention are not the same thing. On social media, an official message competes with breaking news, market commentary and ordinary conversation. Unless it draws engagement, it is easily lost. The economic stakes are concrete. If central bank messages fail to reach households, the transmission channel from communication to expectations is weakened at its very first link, well before any effects on consumption, saving or wage-setting can materialise (Blinder et al. 2024, Masciandaro et al. 2024). We ask which features of official Bank of England tweets are associated with stronger public engagement.

Most of the existing literature begins after exposure. Survey experiments demonstrate that clearer and simpler language can improve comprehension and shift expectations once households are made to read a central bank statement (Bholat et al. 2019, Weber et al. 2025, Ehrmann, Georgarakos, and Kenny 2025). Observational evidence from naturally occurring settings, however, is more mixed. Lamla and Vinogradov (2019) find that FOMC announcements have no measurable direct effect on average household perceptions, though they increase the probability that households hear monetary policy news. Coibion, Gorodnichenko, and Weber (2022) show that information treatments powerfully but transiently shift inflation expectations under forced exposure, raising a natural question about what happens when exposure is voluntary. Neither strand addresses the prior question of whether the message attracts attention in the first place. In a forced-exposure design, every respondent reads the treatment text. On a social media feed, by contrast, no one is compelled to stop scrolling. For institutions that increasingly rely on digital channels to reach the broader public, that earlier margin matters at least as much (Haldane, Macaulay, and McMahon 2020). It has received comparatively little empirical scrutiny, largely because the observational data needed to study it at scale have only recently become available.

We study this question using the complete record of tweets posted by the official Bank of England account (@bankofengland) between July 2011 and July 2022. Our empirical approach is simple. We estimate tweet-level Poisson models under a ladder of tighter timing controls, up to and including week fixed effects, so

that the comparisons run across tweets posted within the same seven-day window. The design is associational rather than experimental, but the progressive addition of controls allows us to separate tweet features whose associations persist after the absorption of broad shifts in platform-wide attention from those that fade once such shifts are accounted for.

The results point in a consistent direction. Media format and policy timing are the strongest and most robust predictors of engagement, holding under tighter time controls up to week fixed effects. Readability carries a positive but more fragile association, and reply tweets attract markedly less engagement than original posts. Section 4 reports these findings in detail.

The paper contributes to a growing literature on public-facing central bank communication. It does so by shifting the focus from comprehension conditional on exposure to the determinants of exposure itself. Whereas most existing studies take readership as given and ask whether simpler language improves understanding (Bholat et al. 2019, McMahon and Naylor 2023), we ask what observable features help an official message compete for attention in an uncontrolled information environment. The contribution is empirical and deliberately narrow. We offer a clean associational account of the tweet-level characteristics that predict public engagement, drawn from the complete posting history of a major central bank.

*Related literature.* Three strands of the literature motivate and frame what follows.

The first is the growing body of work on central bank communication with the general public. Blinder et al. (2008) and Blinder et al. (2024) trace the field from an early focus on financial markets to a broadening interest in whether central bank messages reach households at all; the updated survey concludes that public knowledge of monetary policy remains fragmentary, yet social media channels show "glimmers of hope." Haldane and McMahon (2018) document that mean public understanding of monetary policy in the United Kingdom has flat-lined over seventeen years, while Haldane, Macaulay, and McMahon (2020) formalise a three-part framework of engagement, explanation and education. The present paper speaks to the engagement component. If citizens do not engage with a message, neither explanation nor education can follow.

The second strand uses social media data to study central bank communica-

tion. Gorodnichenko, Pham, and Talavera (2025) analyse Federal Reserve tweets and find that media accounts and economists are the most active interlocutors. Ehrmann and Wabitsch (2022) show that non-experts respond to ECB communication on Twitter but less so than experts. Masciandaro et al. (2024) construct a cross-country index of central bank social media engagement and find that clearer communication is associated with greater public interaction. These papers study the demand side of public discussion or cross-institution comparisons; we study the supply side, asking which features of an individual central bank's own posts predict the engagement those posts receive. Our within-institution, within-week design allows us to control for aggregate attention shifts that cross-country panels cannot account for.

The third strand encompasses experimental work on the effects of communication design. Bholat et al. (2019) show that simplification and relatable information boost understanding of BoE key messages by roughly forty per cent under forced exposure. Ash et al. (2024) demonstrate through a pre-registered experiment that audiovisual formats strengthen household updating toward the ECB's inflation target relative to text alone. Our finding that photo and video tweets attract much higher engagement is consistent with these experimental results, though it operates at the earlier attention margin rather than the comprehension margin they study. The state-dependence of attention that we document, with the MPC-day association strengthening markedly during high inflation, connects to Weber et al. (2025), who show that information treatments have larger effects in environments where inflation is salient and prior inattention is high.

The remainder of the paper proceeds as follows. Section 2 describes the data and variable construction. Section 3 sets out the empirical framework. Section 4 reports the main results and robustness checks. Section 5 concludes.

## 2. Data and Institutional Setting

Our regressions use only tweets posted by the official Bank of England account. They do *not* use the wider universe of tweets about the Bank. That broader conversation matters for context, but the question here is narrower. Once the BoE posts a tweet, which features of that tweet are associated with engagement?

The main sample contains all 9,810 tweets posted by the official account between 11 July 2011 and 20 July 2022. For each tweet we observe likes, retweets, replies and quote tweets at the time of collection. Our main outcome is *total engagement*, defined as the sum of these four counts. Impression data are unavailable for most of the period and are therefore not used, so all measures should be read as cumulative realised interactions, not exposure-adjusted response rates. Appendix Table 4 reports full summary statistics. Engagement is highly skewed: the mean is 32.55, the median is 10, and the maximum reaches 35,189. The Bank is also not a uniformly formatted broadcaster. Of the 9,810 tweets, 7,789 are original posts and 2,021 are replies; photos appear in 24.2 per cent and videos in 1.6 per cent.

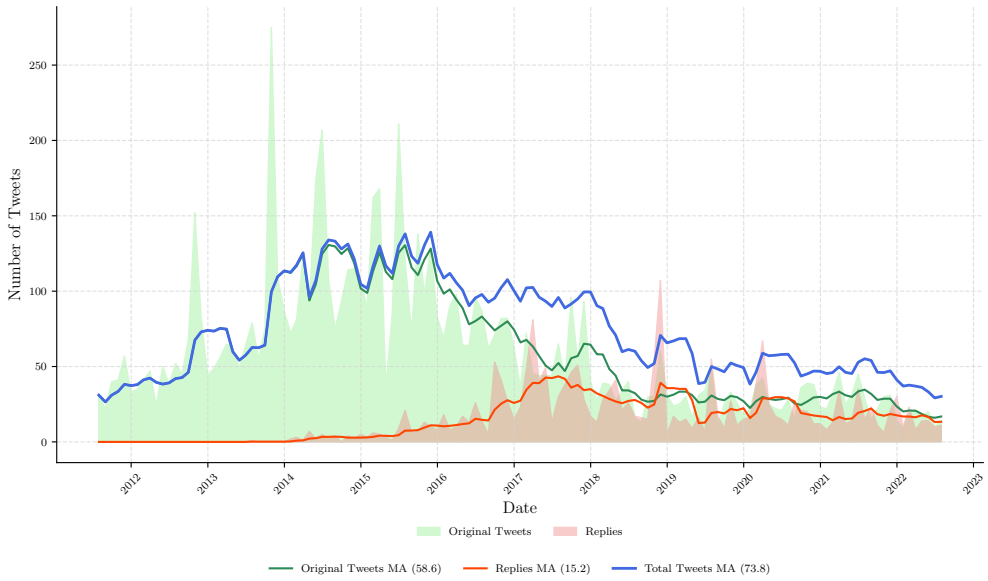


FIGURE 1. Official Bank of England tweet activity, 2011–2022

The figure reports monthly activity for the official BoE account and separates original tweets from replies.

Figure 1 shows the evolution of the official account over time. Monthly activity averages 73.8 tweets, but what changes more noticeably than volume is composition. Replies become more common after 2016, creating substantial within-account variation in communication mode. Periods of heavier posting do predict higher aggregate engagement (Appendix Table 5), but much of the variation in engagement remains unexplained by volume alone, which is why the paper focuses on tweet-level characteristics.

To place the official account in its wider platform environment, we also compile a background corpus of 3,126,016 tweets from 719,310 unique users between 2007 and 2022 that contain the terms "Bank of England" or "BoE" (case-insensitive). The corpus spans well before the Bank's own account was created, so it captures public discussion of the institution that predates the official presence on the platform. This corpus is descriptive only and does not enter the regressions. It is assembled by text-based keyword search with regular-expression filters to exclude common mismatches such as "Board of Education"; it does not separately track @-handle mentions of the official account, which may undercount directed replies.[1] The 9,810 official tweets represent roughly 0.3 per cent of this broader corpus, confirming that the Bank's own voice is a small fraction of the public conversation about it.

The tweet-level covariates capture features directly observable to users on the platform. We include indicators for whether a tweet is a reply, contains a link, contains a hashtag, or includes a GIF, photo or video. We also construct an indicator for Monetary Policy Committee announcement days using the official Bank calendar. Some tweets are released on routine days, whilst others coincide with scheduled policy events when attention to the institution is already elevated. To measure linguistic accessibility, we compute the Flesch Reading Ease score after removing URLs, hashtags and user mentions from the tweet text. Higher values indicate easier text. We also include character count so that readability does not merely proxy for brevity. A shorter tweet may attract more interaction simply because it is easier to process quickly, whereas a more readable tweet may do so because it communicates more clearly even at a given length. We therefore interpret readability as an indicator of accessibility, not as a claim about the substantive quality of the underlying message.

Two scope conditions bear emphasis. These data capture observed engagement on one platform, not comprehension, trust or expectation formation; the two questions are complementary, not competing. Twitter users are not representative of the public at large, though Angelico et al. (2022) and Born, Dalal et al. (2025) show that Twitter-derived economic measures correlate well with survey-based and market-based expectations. The paper should accordingly be read as evidence

[1] Manual checks on a random sample indicate that false positives are limited. False negatives are harder to observe directly, so the corpus should be read as a contextual measure of discussion, not a complete census.

on digitally active audiences and on the platform-specific reception of official central bank communication.

# 3. Empirical Framework

The engagement count for tweet $i$ is modelled as

$$\mathbb{E}[Y_i \mid X_i] = \exp(X_i'\beta), \tag{1}$$

where $Y_i$ is total engagement and $X_i$ contains the tweet characteristics and timing controls described in Section 2. We estimate this by Poisson quasi-maximum likelihood throughout, with standard errors clustered by ISO week.

The more important design choice is the fixed-effects structure, not the estimator. We report all results across a five-column ladder that progressively tightens the set of timing controls. Column (1) includes no fixed effects, so it exploits all cross-tweet variation in the sample. Column (2) adds year-quarter, day-of-week and hour-of-day effects, absorbing slow-moving trends in the Bank's follower base and systematic within-week patterns. Columns (3) through (5) replace year-quarter effects with week fixed effects and then layer in day-of-week and hour-of-day controls. The logic of this progression is that if a coefficient survives the move from column (1) to column (5), it is unlikely to be driven simply by changes in platform-wide activity or the broader macroeconomic news cycle. The intermediate columns are also informative: comparing column (3) to column (5) reveals whether within-week estimates are sensitive to the day and hour at which a tweet is posted, which matters because MPC announcements fall on particular weekdays and are typically accompanied by tweets at specific times.

Two identification concerns should be stated plainly. First, format choice is not random. The Bank's communications team may attach a photo or video precisely when it expects a message to be salient. The coefficients therefore reflect both the audience response to format and the sender's selection of format for particular content. We cannot separate these channels. This endogeneity is shared by the broader observational literature on central bank communication design, including the cross-country panel of Masciandaro et al. (2024) and the social media studies of

Gorodnichenko, Pham, and Talavera (2025). Second, even week fixed effects do not absorb all confounders. Within a given week, an MPC announcement day differs from a routine Tuesday not only because the Bank tweets about policy but also because markets, journalists and the public are paying closer attention for reasons unrelated to any individual tweet. The MPC-day coefficient thus captures the joint effect of policy timing and the elevated attention environment that accompanies it, a challenge shared by Ehrmann and Wabitsch (2022) for ECB-related Twitter traffic. We interpret all estimates as conditional associations and avoid causal language accordingly.

# 4. Results

Table 1 reports the five-column fixed-effects ladder for total engagement. What matters most in the table is how the estimates behave as controls are tightened. Consider first the media variables. Photo and video coefficients both attenuate sharply between columns (1) and (2), which is expected. Raw correlations between format and engagement partly reflect time-composition confounding: photos and videos became more common in later years when the follower base was larger and engagement was generally higher. Once year-quarter effects remove that trend, the coefficients fall. They then stabilise across the week-FE ladder. The photo coefficient moves from 0.928 in column (3) to 1.034 in column (5); the video coefficient moves from 1.029 to 0.950. Neither shift is large relative to the standard errors. The media associations are therefore robust to the choice of time controls once broad trends are accounted for. In the full specification, photo tweets are associated with 181.3 per cent higher engagement and video tweets with 158.7 per cent higher engagement. Against a sample mean of 32.55, these imply predicted engagement levels of roughly 91 and 84 respectively.

TABLE 1. Main Results: Tweet-Level Poisson Models for Total Engagement

| Variable | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | | | Week FE ladder | | |
| | No FE | YQ + Day + Hour FE | Week only | Week + Day | Week + Day + Hour |
| MPC day | 0.877*** | 0.816*** | 0.755*** | 0.765*** | 0.916*** |
| | (0.235) | (0.211) | (0.169) | (0.223) | (0.254) |
| Reply tweet | -1.367*** | -2.122*** | -2.164*** | -2.215*** | -2.245*** |
| | (0.180) | (0.204) | (0.171) | (0.167) | (0.182) |
| Link | 0.130 | -0.118 | -0.037 | -0.033 | 0.032 |
| | (0.186) | (0.221) | (0.289) | (0.279) | (0.325) |
| Hashtag | 0.408** | 0.314 | 0.363** | 0.357** | 0.326* |
| | (0.193) | (0.209) | (0.168) | (0.169) | (0.176) |
| GIF | 0.840*** | 0.468** | 0.411 | 0.367 | 0.411 |
| | (0.220) | (0.219) | (0.289) | (0.373) | (0.309) |
| Photo | 1.344*** | 0.863*** | 0.928*** | 0.960*** | 1.034*** |
| | (0.159) | (0.158) | (0.193) | (0.238) | (0.259) |
| Video | 2.026*** | 1.170*** | 1.029*** | 1.015** | 0.950*** |
| | (0.202) | (0.228) | (0.378) | (0.399) | (0.265) |
| Readability (Flesch) | 0.006* | 0.005** | 0.003 | 0.002 | 0.003* |
| | (0.003) | (0.002) | (0.002) | (0.002) | (0.002) |
| Character count | 0.003*** | -0.003*** | -0.003 | -0.002* | -0.002 |
| | (0.001) | (0.001) | (0.002) | (0.001) | (0.001) |
| Year-quarter FE | No | Yes | No | No | No |
| Week FE | No | No | Yes | Yes | Yes |
| Day-of-week FE | No | Yes | No | Yes | Yes |
| Hour-of-day FE | No | Yes | No | No | Yes |
| Observations | 9810 | 9810 | 9810 | 9810 | 9810 |
| Pseudo-$R^2$ | 0.296 | 0.521 | 0.607 | 0.620 | 0.657 |

Notes: Poisson QMLE with week-clustered standard errors. Outcome is total engagement (likes + retweets + replies + quote tweets). Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

These format effects align with experimental evidence that audiovisual communication strengthens the reception of central bank messages. Ash et al. (2024) find that video and audio formats improve household updating toward the ECB's inflation target relative to text alone. Our estimates capture a different margin—attention, not comprehension—but visual presentation is associated with stronger reception at both stages of the communication pipeline. The MPC-day coefficient follows a different and more instructive path. It attenuates modestly from 0.877

in the baseline to 0.755 with week fixed effects alone, but then recovers to 0.916 once day-of-week and hour-of-day controls are added in column (5). The recovery matters. MPC announcements fall on specific weekdays and the Bank's tweets about them are posted at particular hours. Without controlling for these regularities, week fixed effects partly capture the systematic day-and-hour component of MPC-day engagement. Once that component is modelled separately, the pure within-week MPC premium re-emerges. In the full specification, MPC-day tweets are associated with 150.0 per cent higher engagement, implying a predicted level of about 81.

Readability retains a positive association throughout, but it is the most fragile of the main covariates. The coefficient is significant at the five per cent level in column (2) and at the ten per cent level in column (5), but it dips below conventional significance thresholds in columns (3) and (4). A one-standard-deviation increase in Flesch Reading Ease (27.03 points) is associated with approximately 8.4 per cent higher engagement in the full specification. This is meaningful but modest, and it should be read with the caveat that the estimate is sensitive to the exact configuration of time controls. We return to this fragility in the robustness discussion below.

The reply coefficient behaves differently from every other variable in the table. Rather than shrinking as controls are added, it grows monotonically larger in absolute value, from −1.37 in column (1) to −2.25 in column (5). This implies that within a given week, conditional on day and hour, reply tweets are associated with roughly 89.4 per cent lower engagement than original posts. The pattern is conceptually distinct from the media and readability results, which describe how presentation is associated with the reception of broadcast content. The reply penalty instead reflects a structural difference in communication mode. Reply tweets are directed at individual users and may receive less engagement because they are narrower in audience, lower in salience, or less prominently surfaced on the platform; the design cannot distinguish these channels. In practice, replies and original posts are not comparable objects, and any evaluation of a central bank's social media performance that pools the two will conflate reach with dialogue. This tension echoes the dilemma identified by Haldane, Macaulay, and McMahon (2020), who argue that engagement, explanation and education can work at cross-purposes. A reply may achieve the Bank's educational goal while generating far

less measurable interaction than a broadcast post.

Table 2 examines whether the MPC-day association varies with the macroeconomic environment. We interact the MPC-day indicator with a dummy for the high-inflation period beginning in January 2021 and report the results across the same five-column ladder.

TABLE 2. MPC-Day Engagement in the High-Inflation Period

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Variable | No FE | YQ + Day + Hour FE | Week FE ladder | | |
| | | | Week only | Week + Day | Week + Day + Hour |
| MPC day | 0.654*** | 0.593*** | 0.582*** | 0.558** | 0.648** |
| | (0.227) | (0.206) | (0.202) | (0.246) | (0.255) |
| MPC × High inflation (2021–2022) | 1.107*** | 1.037*** | 0.622* | 0.743** | 1.171*** |
| | (0.257) | (0.285) | (0.335) | (0.330) | (0.306) |
| Reply tweet | -1.510*** | -2.199*** | -2.191*** | -2.247*** | -2.287*** |
| | (0.198) | (0.189) | (0.170) | (0.166) | (0.178) |
| Photo | 1.318*** | 0.862*** | 0.924*** | 0.958*** | 1.036*** |
| | (0.163) | (0.164) | (0.194) | (0.240) | (0.262) |
| Video | 2.013*** | 1.178*** | 1.027*** | 1.012** | 0.958*** |
| | (0.199) | (0.227) | (0.376) | (0.398) | (0.259) |
| Readability (Flesch) | 0.007* | 0.005** | 0.003* | 0.002 | 0.004** |
| | (0.003) | (0.002) | (0.002) | (0.002) | (0.002) |
| Character count | 0.003*** | -0.003*** | -0.003 | -0.002* | -0.002 |
| | (0.001) | (0.001) | (0.002) | (0.001) | (0.001) |
| Year-quarter FE | No | Yes | No | No | No |
| Week FE | No | No | Yes | Yes | Yes |
| Day-of-week FE | No | Yes | No | Yes | Yes |
| Hour-of-day FE | No | Yes | No | No | Yes |
| Observations | 9810 | 9810 | 9810 | 9810 | 9810 |
| Pseudo-$R^2$ | 0.303 | 0.526 | 0.608 | 0.622 | 0.660 |

Notes: Poisson QMLE with week-clustered standard errors. All columns include reply, link, hashtag, GIF, photo, video, readability, and character-count controls. High-inflation indicator equals one for dates on or after 1 January 2021.

The interaction term is positive in every column, but its path across the ladder mirrors the pattern observed for the MPC-day main effect in Table 1. Without fixed effects the interaction is large (1.107) and precisely estimated. It remains so with year-quarter controls. When week fixed effects replace year-quarter effects in column (3), however, the coefficient drops to 0.622 and significance weakens to the ten per cent level. This drop is expected for the same reason as before. Week

fixed effects sweep up much of the variation that distinguishes the high-inflation period from earlier years, leaving the interaction with less aggregate time-series variation to exploit. As day-of-week and then hour-of-day controls are layered back in through columns (4) and (5), the interaction recovers, reaching 1.171 and regaining significance at the one per cent level. The recovery mirrors what we observed for the MPC main effect. Once systematic day-and-hour variation in engagement is modelled separately instead of left for the week dummies to absorb, the within-week MPC premium sharpens.

The estimates point to state-dependence in public attention to monetary policy decisions. During the high-inflation period, the combined MPC-day effect in the full specification is 0.648 + 1.171 = 1.819, which implies roughly 516 per cent higher engagement on announcement days relative to non-MPC days in the same week. In the pre-2021 period, the corresponding effect is $\exp(0.648) - 1 \approx 91$ per cent. The difference is large. It does not tell us whether the heightened response reflects greater public concern about price stability, increased media amplification, or some combination. But it does suggest that the engagement premium on announcement days is far from constant, and that the value of policy-timed communication rises sharply when inflation is salient. Rational inattention models (Maćkowiak, Matějka, and Wiederholt 2023) predict exactly this pattern, holding that agents allocate more processing capacity to information when the stakes of attending to it are higher. It also aligns with the experimental evidence of Weber et al. (2025), who find that information treatments have larger effects in high-inflation settings. Our result complements their finding by showing that the same pattern appears at the attention margin. People pay more attention to central bank communication when inflation is high, not only learn more from it.

The remaining coefficients in Table 2 are stable in ways that reinforce the main results. The reply penalty ranges from −1.51 to −2.29 and grows as controls tighten. Photo and video associations fall from baseline to the fixed-effects columns but remain large and significant throughout. Readability is positive in four of the five columns, though it dips below significance in column (4), again underlining its relative fragility.

Table 3 subjects the full specification to a series of alternative estimators and sample restrictions. Column (3) replaces the Poisson with a negative binomial model. Column (4) uses OLS on $\ln(1 + Y)$. Column (5) drops the five highest-

engagement tweets. Column (6) restricts the sample to original posts only.

TABLE 3. Robustness Checks: Total Engagement

| | (1) Pref | (2) Week | (3) NegBin | (4) OLS | (5) Excl top 5 | (6) Original only |
|---|---|---|---|---|---|---|
| MPC day | 0.816*** | 0.916*** | 0.895*** | 0.657*** | 0.880*** | 0.893*** |
| | (0.211) | (0.254) | (0.111) | (0.076) | (0.197) | (0.279) |
| Reply tweet | -2.122*** | -2.245*** | -2.259*** | -1.707*** | -1.933*** | |
| | (0.204) | (0.182) | (0.083) | (0.061) | (0.187) | |
| Photo | 0.863*** | 1.034*** | 0.982*** | 0.966*** | 0.721*** | 0.925*** |
| | (0.158) | (0.259) | (0.056) | (0.052) | (0.088) | (0.286) |
| Video | 1.170*** | 0.950*** | 1.574*** | 1.470*** | 1.286*** | 0.827*** |
| | (0.228) | (0.265) | (0.165) | (0.095) | (0.247) | (0.253) |
| Readability | 0.005** | 0.003* | -0.000 | -0.001*** | 0.002 | 0.005** |
| | (0.002) | (0.002) | (0.001) | (0.000) | (0.001) | (0.002) |
| Character count | -0.003*** | -0.002 | 0.001 | 0.000 | 0.001 | -0.002 |
| | (0.001) | (0.001) | (0.001) | (0.000) | (0.001) | (0.002) |
| Observations | 9810 | 9810 | 9810 | 9810 | 9805 | 7789 |
| Model fit stat | 0.521 | 0.657 | 0.653 | 0.677 | 0.627 | 0.658 |
| NegBin $\alpha$ | | | 2.9533 | | | |

Notes: Column (4) uses OLS with dependent variable $\ln(1 + Y)$. Model fit stat is pseudo-$R^2$ for count models and $R^2$ for OLS. NegBin $\alpha$ is estimated from an auxiliary NB model without high-dimensional fixed effects.

The core results for MPC timing, photos and videos are stable in sign and significance across all six columns; the reply penalty remains large and significant in every specification where it is estimated (it is undefined in the original-tweets-only column, which excludes replies by construction). Magnitudes shift, as one would expect when moving between count and log-linear models, but the ranking of associations is preserved. Excluding the five most-engaged tweets reduces the photo and video coefficients somewhat, confirming that the upper tail contributes to the estimates, but the effects remain large and significant. This check matters because the extreme skewness of the engagement distribution raises the possibility that a handful of viral tweets could drive the results; the stability of the ranking after exclusion suggests they do not.

Readability, however, is more fragile. The coefficient is positive and significant in the Poisson specifications and in the original-tweets-only sample, but it turns negative in the OLS log specification and loses significance when the top observations are excluded. We regard the readability result as suggestive rather than settled. The fragility is itself informative. McMahon and Naylor (2023) show that the Bank of England's syntactic complexity has decreased over time while its conceptual complexity has increased. If the Flesch score captures only the syntactic dimension, it may miss the conceptual barriers that matter more for whether a reader engages with a message. The cross-country evidence from Masciandaro et al. (2024), who find a positive association between communication clarity and social engagement, is based on broader clarity measures that may capture both dimensions more completely. Additional checks are reported in the appendix. Table 6 decomposes total engagement into its four components and shows that the photo, video and reply associations hold across likes, retweets, replies and quote tweets individually. Table 7 replaces week fixed effects with date fixed effects on the subsample of days with multiple official tweets; the reply, photo and video results hold under this considerably more demanding specification.

# 5. Conclusion

This paper asks which characteristics of official Bank of England tweets are associated with public engagement. The answer, drawn from the complete record of 9,810 tweets over eleven years, is that tweet-level characteristics and policy timing add substantial explanatory power beyond posting frequency alone. The two most robust associations are with media format and MPC timing, both of which survive the full battery of time controls. The MPC-day premium is itself state-dependent. The combined coefficient rises from 0.65 before 2021 to 1.82 during the high-inflation period, implying a shift from roughly 91 to 516 per cent higher engagement on announcement days. Reply tweets attract much less engagement than original posts, highlighting a tension between reach and dialogue.

These findings speak to two open questions in the literature. The first is whether central bank communication with the general public is "promise or false hope," in the framing of Blinder et al. (2024). At the attention margin, there is real promise.

Official messages posted at policy-relevant moments and presented in visual formats do attract measurable public interaction. But the promise is conditional, depending heavily on format, timing and macroeconomic context, and does not extend uniformly across communication modes. The reply penalty suggests that the conversational engagement Haldane, Macaulay, and McMahon (2020) advocate comes at a quantifiable cost to visible reach. The second is whether attention itself varies with the macroeconomic environment. Our evidence complements the experimental finding that communication treatments are more effective when inflation is salient (Weber et al. 2025, Coibion, Gorodnichenko, and Weber 2022) by showing that a similar pattern appears at the pre-comprehension stage.

The findings do not identify causal format effects. The Bank's communications team selects formats endogenously, and within-week confounders that we cannot observe may still bias the estimates. Twitter users are not representative of the general public, and the engagement measures we use capture visible interaction, not comprehension, trust or expectation formation. These limitations are real, and the estimates should be read with them in mind.

The paper's contribution is deliberately narrow. It documents the tweet-level associations that predict whether an official central bank message receives any visible public response at all. Three questions follow naturally from these findings. Does the architecture of communication events, specifically the timing and bundling of releases, shape the quantity of public engagement beyond what tweet-level characteristics explain? Does the textual content of those releases, net of what informed observers could have anticipated, shift how the public interprets policy direction among those already engaged? And does a change in interpretation among attentive users eventually propagate into the broader household beliefs that matter for monetary policy transmission? Answering these questions would connect the attention margin studied here to the expectation formation that Ash et al. (2024) and related experimental work have examined under forced exposure, closing a gap that neither strand has been able to bridge alone.

# References

Angelico, Cristina, Juri Marcucci, Marcello Miccoli, and Filippo Quarta. 2022. "Can we measure inflation expectations using Twitter?" *Journal of Econometrics* 228 (2): 259–277.

Ash, Elliott, Heiner Mikosch, Alexis Perakis, and Samad Sarferaz. 2024. "Seeing and Hearing Is Believing: The Role of Audiovisual Communication in Shaping Inflation Expectations." KOF Working Paper 515, KOF Swiss Economic Institute.

Bholat, David, Nida Broughton, Janna Ter Meer, and Ed Walczak. 2019. "Enhancing central bank communications using simple and relatable information." *Journal of Monetary Economics* 108: 1–15.

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.

Blinder, Alan S., Michael Ehrmann, David-Jan Jansen, and Jakob de Haan. 2024. "Central Bank Communication with the General Public: Promise or False Hope?." *Journal of Economic Literature* 62 (1): 3–44.

Born, Benjamin, Hrishbh Dalal et al. 2025. "Inflation Expectations in the Social Media Age." University of Frankfurt Working Paper.

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.

Ehrmann, Michael, Dimitris Georgarakos, and Geoff Kenny. 2025. "Credibility Gains from Central Bank Communication with the Public." *European Economic Review* 177.

Ehrmann, Michael, and Alena Wabitsch. 2022. "Central banks on social media–the reception of ECB communication among experts and non-experts on Twitter." *ECB Working Paper Series* 2716.

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

Haldane, Andrew G., 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 G, and Michael McMahon. 2018. "Central bank communications and the general public." *AEA Papers and Proceedings* 108: 578–83.

Lamla, Michael J, and Dmitri V Vinogradov. 2019. "Central bank announcements: Big news for little people?" *Journal of Monetary Economics* 108: 21–38.

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

Masciandaro, Donato, Davide Romelli, Gaia Rubera, and Stefano Vena. 2024. "Central bank communication and social media: From silence to Twitter." *Journal of Economic Surveys* 38 (2): 365–388.

McMahon, Michael, and Matthew Naylor. 2023. "Getting Through: Communicating Complex Information." Staff Working Paper 1047, Bank of England.

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

# Appendix

This appendix reports supplementary tables referenced in the main text. Table 4 provides full summary statistics for the official-tweet sample. Table 5 presents an aggregate weekly regression of total engagement on posting volume. Table 6 decomposes total engagement by type. Table 7 reports a date-fixed-effects specification on the subsample of days with multiple official tweets.

TABLE 4. Summary Statistics: Official BoE Tweet-Level Sample

| Variable | N | Mean | SD | Median | P25 | P75 | Max |
|---|---|---|---|---|---|---|---|
| Likes | 9810 | 14.92 | 241.42 | 3.00 | 1.00 | 11.00 | 22775.00 |
| Retweets | 9810 | 13.12 | 96.20 | 6.00 | 1.00 | 13.00 | 8875.00 |
| Replies | 9810 | 2.45 | 17.94 | 1.00 | 0.00 | 2.00 | 1154.00 |
| Quote tweets | 9810 | 2.06 | 33.45 | 0.00 | 0.00 | 0.00 | 2917.00 |
| Total engagement | 9810 | 32.55 | 377.40 | 10.00 | 3.00 | 28.00 | 35189.00 |
| Reply tweet (dummy) | 9810 | 0.21 | 0.40 | 0.00 | 0.00 | 0.00 | 1.00 |
| Link (dummy) | 9810 | 0.76 | 0.42 | 1.00 | 1.00 | 1.00 | 1.00 |
| Hashtag (dummy) | 9810 | 0.33 | 0.47 | 0.00 | 0.00 | 1.00 | 1.00 |
| GIF (dummy) | 9810 | 0.02 | 0.14 | 0.00 | 0.00 | 0.00 | 1.00 |
| Photo (dummy) | 9810 | 0.24 | 0.43 | 0.00 | 0.00 | 0.00 | 1.00 |
| Video (dummy) | 9810 | 0.02 | 0.13 | 0.00 | 0.00 | 0.00 | 1.00 |
| Readability (Flesch Reading Ease) | 9810 | 53.79 | 27.03 | 56.98 | 40.09 | 72.32 | 121.22 |
| Tweet character count (URL-stripped) | 9810 | 111.17 | 58.04 | 98.00 | 71.00 | 130.00 | 297.00 |
| MPC announcement day (dummy) | 9810 | 0.06 | 0.23 | 0.00 | 0.00 | 0.00 | 1.00 |

The heavy right skew of total engagement (standard deviation more than eleven times the mean) motivates the use of Poisson and negative binomial specifications rather than OLS in levels.

TABLE 5. Aggregate Weekly Regression: Engagement on Posting Volume

| Regressor | Coefficient | SE | p-value | $R^2$ | N |
|---|---|---|---|---|---|
| $\ln(1 + \text{weekly volume})$ | 1.0784 | 0.0667 | 0.0000 | 0.3411 | 577 |

The near-unit elasticity and $R^2$ of 0.34 show that posting volume predicts aggregate engagement but leaves more than two-thirds of the weekly variation unexplained, motivating the main text's focus on tweet-level characteristics.

TABLE 6. Outcome Decomposition: Week-FE Poisson

| | Likes | Retweets | Replies | Quote tweets |
|---|---|---|---|---|
| MPC day | 0.642** | 1.171*** | 0.498 | 1.238*** |
| | (0.298) | (0.179) | (0.321) | (0.466) |
| Reply tweet | -2.267*** | -2.416*** | -1.478*** | -2.180*** |
| | (0.167) | (0.241) | (0.187) | (0.378) |
| Photo | 1.255*** | 0.841*** | 0.747*** | 1.403*** |
| | (0.329) | (0.168) | (0.192) | (0.399) |
| Video | 0.913*** | 0.979*** | 1.080** | 1.565*** |
| | (0.216) | (0.258) | (0.424) | (0.540) |
| Readability | 0.004** | 0.002 | 0.009** | 0.007 |
| | (0.002) | (0.001) | (0.004) | (0.005) |
| Character count | -0.003* | -0.001 | -0.000 | -0.002 |
| | (0.001) | (0.001) | (0.002) | (0.002) |
| Observations | 9810 | 9810 | 9810 | 9810 |

The photo, video and reply associations hold across all four engagement components. The MPC-day effect is concentrated in retweets and quote tweets, consistent with policy announcements prompting amplification over simple liking. Readability is significant for likes and replies but not for retweets or quote tweets, suggesting that the accessibility premium operates primarily through lower-effort engagement actions.

TABLE 7. Date-FE Poisson on Days with Multiple Official Tweets

| | Date-FE specification |
|---|---|
| MPC day | -1.806* |
| | (1.001) |
| Reply tweet | -1.840*** |
| | (0.176) |
| Link | -0.191** |
| | (0.085) |
| Hashtag | 0.127* |
| | (0.068) |
| GIF | 0.731*** |
| | (0.220) |
| Photo | 0.836*** |
| | (0.070) |
| Video | 1.205*** |
| | (0.194) |
| Readability (Flesch) | 0.000 |
| | (0.001) |
| Character count | 0.002** |
| | (0.001) |
| Date FE | Yes |
| Hour-of-day FE | Yes |
| Observations | 9117 |
| Pseudo-$R^2$ | 0.752 |

Table 7 pushes the fixed-effects structure further by replacing week effects with date effects and restricting the sample to days on which the Bank posted more than one tweet ($N$ = 9,117). Identification now comes entirely from within-day variation. The reply, photo and video coefficients all persist under this considerably more demanding specification. The MPC-day coefficient, however, turns negative, which is mechanical: on an MPC announcement day, the date fixed effect absorbs the day-level attention premium, leaving the MPC indicator to capture only residual within-day variation that may run in either direction depending on the composition of that day's tweets. This result is therefore not in tension with the main findings but rather illustrates the limit of what date-level controls can identify.