

Ehrmann, Michael; Hubert, Paul

Working Paper

Information acquisition ahead of monetary policy announcements

ECB Working Paper, No. 2770

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Ehrmann, Michael; Hubert, Paul (2023) : Information acquisition ahead of monetary policy announcements, ECB Working Paper, No. 2770, ISBN 978-92-899-5512-6, European Central Bank (ECB), Frankfurt a. M., <https://doi.org/10.2866/45657>

This Version is available at:

<https://hdl.handle.net/10419/278341>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Michael Ehrmann, Paul Hubert

Information acquisition ahead of
monetary policy announcements

No 2770 / February 2023

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

How do financial markets acquire information about upcoming monetary policy decisions, beyond their reaction to central bank signals? This paper hypothesises that sharing information among investors can improve expectations, especially in the presence of disagreement or uncertainty about the economy. To test this hypothesis, the paper studies monetary policy-related content on Twitter during the “quiet period” before European Central Bank announcements, when policymakers refrain from public statements related to monetary policy. Conditional on large disagreement about the economic outlook, higher Twitter traffic is associated with smaller monetary policy surprises, suggesting that exchanging private signals among investors can help improve expectations.

Keywords: Central bank communication, quiet period, Twitter, market expectations, information processing.

JEL Codes: D83, E52, E58, G14.

Non-technical summary

Information provided by the central bank is continuously sought for by financial market participants, and leads them to update their assessment of the state of the economy, the economic outlook and the central bank reaction function. While there is an abundant literature on how financial markets react to central bank communication (which can be seen by everybody, and hence is typically labelled a “public” signal), this paper studies the question how private agents form expectations about upcoming monetary policy announcements through dispersed, individual information, i.e. through “private” signals. Such information acquisition proceeds continuously, whether or not there is a new communication from the central bank. To study this process, this paper builds on an institutional feature of central banks’ communication policies, namely “quiet” periods before policy announcements, during which policymakers abstain from communicating in public about the economic assessment or the outlook for monetary policy.

The ECB’s quiet period is in place for the 7 days preceding the announcement of the ECB’s monetary policy decisions. This policy provides us with a natural experiment to study how information acquisition and the formation of policy expectations proceed, beyond the processing of signals provided by the central bank. The paper develops several hypotheses about this updating process. In particular, we argue that market expectations of upcoming monetary policy decisions can improve if agents can share their views about the economy. This improvement will be stronger, the more views get shared, the more disperse the views are and the larger is the uncertainty about the economy.

To test these hypotheses, the paper analyses ECB-related Twitter traffic in the days before the ECB press conference and whether and how this has a bearing on the magnitude of the monetary policy surprise on the announcement day. Twitter has been shown to be a forum where news about the ECB gets disseminated, but also a platform for discussions about the ECB and its policies – not only, but to a large extent among monetary policy experts. Focusing on Twitter traffic during the ECB’s quiet period thus allows us to better understand information flows among agents in the absence of information supplied by the central bank.

Based on the sample of ECB policy announcements between January 2012 and April 2020, this paper first uncovers that Twitter traffic in the days before Governing Council meetings is higher if the subsequent monetary policy announcement is relatively surprising. Such a pattern can result if attention is triggered, for instance, by market expectations of change decisions (which, in turn, tend to generate larger surprises), or if agents observe a change in market prices. We therefore need to control for various factors that would trigger such increased attention.

When we test our hypotheses that increased information sharing is particularly beneficial if views about the economy are dispersed or if uncertainty is large, we find compelling evidence for the former: the ECB’s monetary policy surprises are larger when there is more disagreement about the economic outlook, but conditional on large disagreement, higher Twitter traffic during the quiet period is associated with lower monetary policy surprises.

In a more disaggregated analysis, we find that the effect stems in particular from tweets that discuss monetary policy, whereas no such pattern is found for tweets on the economy. Interestingly, it is not only tweets from experts, but also those from non-experts that generate our results. We find equivalent results when we use the number of newswire reports about the ECB as an alternative proxy for information exchange. During the quiet period, these

articles often quote analysts' views, and as such their content might be close to the one available on Twitter.

The evidence documented in this paper is in line with the hypothesis that financial market participants and central bank watchers stand to benefit from sharing their views about economic fundamentals when disagreement is large (but not necessarily when uncertainty is large). By doing so, agents can form their expectations about future monetary policy based on a larger information set, and therefore come to more accurate expectations on average. Increased information exchange might therefore serve as a partial substitute for the processing of signals sent by the central bank, suggesting that pausing the information flow from central banks to markets does not pose any immediate concerns even if central bank communication is otherwise dominant.

1. Introduction

Monetary policy has undergone a long journey towards increasing transparency (Issing 2019) and is expected to continue along this path (Blinder 2018). Central banks have become more transparent about their objectives and their reaction function, they publish macroeconomic projections, minutes and even transcripts of their committee meetings. In the course of this journey, central bank decisions have become considerably more predictable (Swanson 2006). Still, as monetary policy decisions are taken under uncertainty, by committees that consist of human beings, and as they take into account a myriad of information that needs to be weighed and assessed each time, they are not – and will never be – perfectly predictable. For instance, Cieslak (2018) shows that investors make large and persistent errors in their forecasts of US short-term interest rates over the business cycle, in large parts because they underestimate how aggressively the US Federal Reserve (Fed) eases monetary policy in economic downturns.

Given the complexity of (forecasting) monetary policy decisions, information that the central bank provides is continuously sought for and leads the private sector to acquire information to update its assessment of the state of the economy, the economic outlook and the central bank reaction function (Byrne et al. 2021). Central bank communication is scrutinized for every word, and speeches by committee members are routinely reported by newswire services and get reflected in financial markets almost instantaneously. While there is an abundant literature on how financial markets *react* to *public* signals, either central bank communication (for a survey, see Blinder et al. 2008) or the release of macroeconomic news (Andersen et al. 2003, Gilbert et al. 2017), the question how private agents acquire information *before* central bank announcements through *private* signals has, to our knowledge, not yet been studied.

Central bank communication is a dominant source of news about monetary policy and private agents' reaction to it might even be too strong relative to the information content of the news, for instance because market participants might give too much weight to public signals (Morris and Shin 2002). Hubert (2014) provides empirical evidence for this mechanism and show that central bank signals acts as a focal point for private expectations. Ehrmann and Sondermann (2012) show that shortly after a release of central bank information, markets react less to macroeconomic announcements from other sources than when the signals from the central bank are less up to date. Similarly, based on a model where agents learn from market signals, Ehrmann et al. (2019) show that central bank communication can potentially lower the informativeness of market signals and, as a consequence, may increase uncertainty, preventing markets from appropriately updating their beliefs.

This begs the question how central bank watchers update their beliefs beyond this channel. Hardly a day passes where commentators would not write about monetary policy, be it in the media or in reports to clients, or without an active discussion about central banking issues in social media. This shows that information acquisition about upcoming policy decisions is going on continuously, whether or not there is a new communication from the central bank, through the exchange of *private* signals. This paper builds on an institutional feature of central banks' communication policies, namely "quiet" periods before policy announcements, to investigate the formation process of policy expectations ahead of policy decisions.

Investors' demand for information about the macroeconomy is particularly strong when uncertainty is large, and market responses to such news are stronger following periods of increased information demand (Benamar et al. 2021). Demand for information about monetary policy is particularly high in the days prior to monetary policy decisions. Ehrmann and

Fratzscher (2009a) show that the responsiveness of short-term interest rates to Fed communication is three to four times as large in the days before policy meetings than otherwise, and that communication during this time window raises rather than reduces market volatility. It is for that reason that many central banks, including the ECB, have adopted a quiet period policy, i.e. they abstain from communicating in the days preceding the announcement of monetary policy decisions. Although the ECB's quiet period, which is in place for the 7 days preceding the announcement of the ECB's monetary policy decisions, is not always adhered to (Rieder and Gnan 2021), the information supply in those days is severely limited, right when information demand peaks. The quiet period policy provides us with a natural experiment to study how information acquisition and the formation of policy expectations proceed, beyond the processing of public signals provided by the central bank.

To do so, we base ourselves on a classic Bayesian updating framework and start from the notion that agents observe private signals about the state of the economy (also) during the quiet period. However, for various reasons not all agents observe such a signal – be it that they are inattentive, only update their information set infrequently, or be it that they are subject to high costs of information acquisition or processing, which prevents them from constantly updating their information set. We show that in such an economy, the market expectation of the economic fundamental is biased, as some agents rely on stale information. If agents can share their private signals, the market expectation of the economic fundamental will be improved. The main hypotheses that emerge from this stylised framework are i) the improvement will be stronger, the more signals get shared; ii) the benefit of sharing private signals is increasing in the share of inattentive agents, and iii) the benefit also increases in the degree of uncertainty about the economy.

The current paper puts these hypotheses to an empirical test. It studies information acquisition about the ECB in the quiet period. It analyses ECB-related Twitter traffic in the days before the ECB press conference and whether and how this has a bearing on the magnitude of the monetary policy surprise on the announcement day. In other words, we take Twitter to be a platform where central bank watchers can exchange information. Twitter has been shown to be a forum where news about the ECB gets disseminated, but also a platform for discussions about the ECB and its policies – not only, but to a large extent among monetary policy experts (Ehrmann and Wabitsch 2022). Focusing on Twitter traffic during the natural experiment that is generated through the ECB's quiet period thus allows us to better understand information flows among agents in the absence of information being supplied by the central bank itself.

Based on the sample of ECB policy announcements between January 2012 and April 2020, this paper first uncovers that Twitter traffic in the days before Governing Council meetings is positively correlated with the magnitude of the subsequent monetary policy surprise, measured using the intraday series of Altavilla et al. (2019). This unconditional correlation can result if attention is endogenous, e.g. if analysts expect a change decision (which, in turn, tend to generate larger surprises), or if agents observe a change in market prices. We therefore aim to control for various factors that would trigger such increased attention.

To test our hypothesis that increased information sharing is particularly beneficial if inattention is high or if uncertainty is large, we use proxy variables that capture information sharing motives and interact them with Twitter traffic. We study correlations that are conditional on the degree to which information sharing might be beneficial. We use disagreement about the economic outlook as our empirical proxy for the degree of inattention. When agents are inattentive, and heterogeneously so, they base their expectations on different information sets and disagreement is large (Andrade and Le Bihan 2013, Giacomini et al. 2020).

Of course, in the absence of disagreement, there is no value in exchanging information. Using this proxy, we find compelling and highly robust results: the ECB's monetary policy surprises are larger when there is more disagreement about the economic outlook, but conditional on large disagreement about the economic outlook, higher Twitter traffic during this quiet period is associated with lower monetary policy surprises.

In a more disaggregated analysis, we find that the effect stems in particular from tweets that discuss monetary policy, whereas no such pattern is found for tweets on the economy. Interestingly, it is not only tweets from experts, but also those from non-experts that generate our results. We find equivalent results when we use the number of newswire reports about the ECB as an alternative proxy for information exchange. During the quiet period, these articles often quote analysts' views, and as such might be close to Twitter.

The evidence documented in this paper is in line with the hypothesis that financial market participants and central bank watchers stand to benefit from sharing their private information about economic fundamentals when disagreement is large (but not necessarily when uncertainty is large). By doing so, agents can form their expectations about future monetary policy based on a larger information set, and therefore come to more accurate expectations on average. Increased information exchange might therefore serve as a partial substitute for the processing of signals sent by the central bank, suggesting that pausing the information flow from central banks to markets does not pose any immediate concerns even if central bank communication is otherwise dominant. It might even contribute to reduce the excess weight put on public signals by financial market participants (Morris and Shin 2002, 2018; Svensson 2006).

The main contribution of this paper is to focus on what happens before policy announcements, while the literature usually investigates the impact of policy announcements on private beliefs. The closest papers to ours in that respect are Lucca and Moench (2015) and Cieslak et al. (2019). The former documents large excess returns on stocks the day ahead of monetary policy announcements by the Federal Open Market Committee (FOMC), the rate-setting committee of the Fed. The latter finds that average excess returns on stocks are statistically significantly higher in the week before FOMC announcements, i.e. during the FOMC's quiet period. Cieslak et al. suggest that monetary policy news reaches financial markets through informal communication by Fed policymakers. Our paper builds on the quiet period policy, measures of potential breaches of this quiet period and Twitter data to specifically investigate how the exchange of private signals, in the absence of public signals, can help financial market participants form their policy expectations.

This paper relates to different strands of the literature. The first focuses on the role of Twitter discussions about monetary policy. Gorodnichenko et al. (2021b) study how tweets by the Fed are received on Twitter, and find a more active engagement among economists and media than among the general public. Based on the 2013 "taper tantrum" episode, Lüdering and Tillmann (2020) find that discussions about Fed policy on social media affect US asset prices. Masciandaro et al. (2022) analyse discussions about monetary policy on Twitter around the time of policy announcements. They show that changes in a similarity measure between the content of tweets and the transcripts of central bank announcements are associated with higher stock market volatility and jumps in sovereign yields. Ehrmann and Wabitsch (2022) study tweets about the ECB after policy announcements and explore whether central bank communication is relayed and discussed by non-experts. Azar and Lo (2016) explore whether social media data contain useful information about future asset prices since anyone can participate in a conversation about asset prices – whether they are informed or not. Using

FOMC meetings, they show that the content of tweets referencing the Fed have some predictive power for future returns. In addition, Bianchi et al. (2022) analyse the effects of President Trump's tweets criticizing the Fed. They find a negative effect on the expected policy rate and on long-term yields, and a positive effect on stock prices. On the same topic, Camous and Matveev (2021) find that financial markets expected the Fed to adjust monetary policy in the direction suggested by President Trump's tweets. Finally, Stiefel and Vivès (2022) use Twitter data to measure the perceived likelihood that the ECB conducts purchases of government bonds following president Draghi's "Whatever it takes" statement. They find that the strong increase in this likelihood is associated with the decreasing sovereign bond yields of the distressed countries.

The present paper also relates to the literature about private agents' demand for information. Benamar et al. (2021) show that financial market participants' demand for information about macroeconomic indicators affects the sensitivity of US Treasury yields to economic news. Tillmann (2022) analyses the demand for information about monetary policy using the Fed's website page views. He finds that macroeconomic news surprises lead financial market participants to revise their policy expectations and actively acquire new information. Our paper also relates to contributions that analyse the influence of uncertainty and disagreement for financial market outcomes (see, e.g., Anderson et al., 2005, 2009, Carlin et al., 2014, Li, 2016, Huang et al., 2021) and the role of disagreement in expectation dynamics (see, e.g., Mankiw et al. 2003, Capistrán and Timmermann 2009, Patton and Timmermann 2010, Andrade et al. 2016). Finally, this paper is linked with the literature using textual analysis to analyse central bank communication. Rosa (2011), Lucca and Trebbi (2009), Hansen and McMahon (2016), Correa et al. (2021), Hubert and Labondance (2021), Gorodnichenko et al. (2021a), among many others, all use various methodologies to measure the content of central bank text data and show that it matters for asset price and macroeconomic dynamics.

The paper proceeds as follows. In Section 2, we develop our hypothesis how agents acquire information about monetary policy in the absence of public signals from the central bank. Section 3 contains a description of the data. In particular, we validate the usefulness of Twitter as a measure of information flow about the ECB. Section 4 reports and discusses the empirical results. Section 5 contains a more detailed analysis on the channels at work, where we study what type of tweets matter, whether there are alternative information channels, and how financial market expectations adjust during our time window. Section 6 concludes.

2. A stylized framework

In this section, we develop a stylised framework of information processing that guides the empirical analysis. It is grounded in a standard Bayesian updating framework. Financial market participants form and update expectations about the upcoming monetary policy decision. They face uncertainty regarding both the central bank reaction function and the state of the economy. Depending on their degree of attentiveness and costs of information acquisition and processing, they observe public and private signals about the central bank reaction function (policymakers' preferences) and the state of the economy.

At the start of the quiet period, the central bank stops disclosing information, such that there is no public signal emitted anymore.² During the quiet period, financial market participants

² In this set-up, the economy is populated with a central bank and financial market participants. One could imagine a set-up with another public agency, the statistical authority, that publishes another public signal. For sake of clarity and for the illustration of the mechanism, we focus on the case of only one public signal. In the empirical

update their beliefs about the state of the economy. We denote the *change* in the fundamental over the quiet period by $\Delta\theta$. Agents form beliefs about this change in the fundamentals, $E_i(\Delta\theta)$, based on noisy private signals, x_i . The private signal for agent i can be written $x_i = \Delta\theta + \varepsilon_i$, where ε_i is normally distributed, independent of $\Delta\theta$, with mean zero and variance $\sigma_{\varepsilon_i}^2$, and $E(\varepsilon_i \varepsilon_j) = 0$ for $i \neq j$. The fact that it is a private signal implies that a signal by one agent is not observable to the others. Let us denote the precision of the private signal with $\beta_i = 1/\sigma_{\varepsilon_i}^2$.³

Let us assume that a share λ of agents observes a signal, whereas a share $(1 - \lambda)$ of agents does not. The latter agents are inattentive, only update their information set infrequently, or are subject to high costs of information acquisition or processing, which they incur only if they think it is worthwhile doing. As these inattentive agents do not observe a signal, they assume that the state of the economy has not changed and set $E_i(\Delta\theta) = 0$.

The expected value of the change in the economic fundamental by attentive agent i , $E_i(\Delta\theta)$, is: $E_i(\Delta\theta) = x_i = \Delta\theta + \varepsilon_i$. In the case where $\lambda = 1$, i.e. all agents observe a signal, the market expectation aggregates the beliefs of all agents (with N the total number of them) to:

$$E_m(\Delta\theta) = \frac{1}{N} \sum_i x_i = \frac{1}{N} \sum_i (\Delta\theta + \varepsilon_i) = \Delta\theta + \frac{1}{N} \sum_i \varepsilon_i = \Delta\theta \quad (1)$$

Assuming that N is sufficiently large such that averaging over ε_i gives zero, the market expectation is unbiased ($E_m(\Delta\theta) - \Delta\theta = 0$). It might be beneficial for individuals to exchange their information, but the aggregate market expectation would not be different.

In the case where $\lambda < 1$, i.e. if there is a share of inattentive agents who do not observe a signal about changes in the state of the economy, the market expectation aggregates the signals of the attentive agents, but furthermore aggregates the beliefs of the inattentive agents that the fundamental has not changed. The market expectation is therefore:⁴

$$E_m(\Delta\theta) = \lambda \frac{1}{\lambda N} \sum_{i=1, \dots, \lambda N} x_i + (1 - \lambda) \frac{1}{(1-\lambda)N} \sum_{j=1, \dots, (1-\lambda)N} 0 = \lambda \frac{1}{\lambda N} \sum_{i=1, \dots, \lambda N} (\Delta\theta + \varepsilon_i) = \lambda \Delta\theta + \frac{1}{\lambda N} \sum_{i=1, \dots, \lambda N} \varepsilon_i = \lambda \Delta\theta \quad (2)$$

In this case, the information frictions imply that the market expectations are inappropriately updated and therefore biased.

Let us introduce a technology (Twitter) whereby agents exchange information about their private signals, thereby reducing information acquisition costs.⁵ There are two possible actions that agents can undertake with this technology. They can share their signals (by posting tweets), or they can observe the signals of others (by reading tweets). We assume that the decision to read tweets is orthogonal to the precision of the signals and orthogonal to being attentive or not. If some market participants do not read tweets, this segments the market. For the segment where agents do not read tweets, the technology is neutral. Hence, it is sufficient

application, we will control for other signals from statistical agencies and other central banks. What is important for our purposes is that these public signals are no longer publicly commented on by the central bank.

³ The notation follows Morris and Shin (2002).

⁴ Let us assume that λN is still sufficiently large such that, on average, the noise in the signals cancels out.

⁵ We leave aside the motivation behind Twitter participation, e.g. reputation, strategic motives, etc. We also abstract from network representation issues, i.e. the fact that some agents have more weight and constitute nodes, and the possibility of "fake news", i.e. untruthful sharing of information. If tweets move markets, there could be an incentive to release wrong signals and thereby influencing the market in a particular direction. However, in a repeated interaction, agents' reputation might be at stake, preventing them from releasing false signals. Also, the issue is ultimately an empirical question. If fake news are dominant, markets should move in the wrong direction. If there is truthful sharing of signals, markets should move in the right direction. The latter is what we find.

for us to show that for the segment of the market where agents read tweets, market expectations improve. For simplification, we assume that all agents read tweets.⁶

In the case where only a fraction κ of attentive agents posts tweets, i.e. shares their signals, the market expectation becomes:⁷

$$E_m(\Delta\theta) = \lambda\Delta\theta + (1 - \lambda)\frac{1}{(1-\lambda)N}\sum_{j=1,\dots,(1-\lambda)N}\left(\frac{1}{\lambda N\kappa}\sum_{j=1,\dots,\lambda N\kappa}x_i\right) \quad (3)$$

The larger κ , the closer the market expectation will be to the true value (as the noise cancels out).⁸ This implies that the market outcome improves, the more signals get shared on Twitter, i.e. the more of the attentive agents engage in the discussion. In the case where all attentive agents post tweets ($\kappa = 1$), the market expectation is unbiased. This is equivalent to the full information case ($E_m(\Delta\theta) = \lambda\Delta\theta + (1 - \lambda)\Delta\theta = \Delta\theta$).

This implies that the Twitter technology leads to an improved market expectation of the underlying fundamental. This, in turn, gets translated into an improved market expectation of the upcoming monetary policy decision (as agents use their updated beliefs about the state of the economy and about the central bank's reaction function to update their expectations about the upcoming decision, see Byrne et al. (2021)).

Another prediction from this simple setup is that the benefit of using the Twitter technology is increasing in the share of inattentive agents. Recall that the expectational error in the absence of Twitter is ($E_m(\Delta\theta) - \Delta\theta = \lambda\Delta\theta - \Delta\theta = (\lambda - 1)\Delta\theta \neq 0$ if $\Delta\theta \neq 0$). The larger the share of inattentive agents, the larger the deviation from the true economic fundamental, i.e. the larger will be the absolute value of the expectation error. In the extreme, if all agents are inattentive, i.e. $\lambda = 0$, the beliefs do not get updated at all.

Finally, the framework also predicts that the benefit of using the Twitter technology is decreasing in the precision of the private signal, β_i . If precision is low, it requires a large κ , i.e. many shared signals, until the beliefs of the inattentive agents are sufficiently precise to lead to an unbiased market expectation. In contrast, if precision is high, the signals do not spread much around $\Delta\theta$, hence sharing signals improves the market outcome relatively quickly.

To summarise, this stylised framework suggests that, *ceteris paribus*, the magnitude of monetary policy surprises should be negatively related to the degree of signal sharing on Twitter, and this effect should increase in the degree of inattention and it should decrease in the precision of the signal. Empirically, we will test these hypotheses as follows:

(i) Our proxy for the degree of inattention, $(1 - \lambda)$, is the cross-sectional disagreement in beliefs about the state of the economy prior to the quiet period. Following Andrade and Le Bihan

⁶ Attentive agents are now in a position to observe the signals of other attentive agents. While they might have an incentive to incorporate that information in their belief, as it improves their individual expectation, this aspect does not change the market outcome, as the market outcome aggregates all the information of attentive agents. In what follows, for simplicity we assume that attentive agents do not update their beliefs with information from Twitter.

⁷ There is a theoretical possibility that the market could be worse off with this technology, namely if very few agents share their signal. Take the example where only one agent shares the signal and this agent has a very bad draw, i.e. a large ε_i , and all the inattentive agents would trade on that signal. We exclude that possibility, i.e. assume that κ is always sufficiently large that the overall signal improves.

⁸ For simplicity, we assume that all agents on Twitter, and more specifically for our purpose, all inattentive agents have access to all signals shared on Twitter. One could otherwise assume that Twitter users follow different accounts and so have access to different signals. While this could matter at the individual level, this is neutral at the market level since all the information is aggregated.

(2013), Giacomini et al. (2020) and Clements (2022), this choice is based on the following reasoning. Over the entire inter-meeting period, agents receive many signals about the state of the economy, as well as signals from the central bank. The smaller the share of attentive agents,⁹ the more agents base their beliefs on stale information sets. This implies that a higher degree of inattention generates more disagreement about the state of the economy.

(ii) To proxy the precision of the signals, β_i , we employ various measures of uncertainty – if precision is low, uncertainty should be high, and vice versa.

(iii) As empirical proxy for κ , we use the number of tweets issued during the quiet period. We will furthermore test for robustness of our results using a concentration measure which indicates to what extent the discussion on Twitter is dominated by few agents, as well as the number of accounts that participate in the ECB-related Twitter discussion on a given day. The number of tweets issued goes beyond the stylised framework in that it allows that individual agents observe and share more than one signal; the number of people participating in the discussion is closer to the formulation of our framework which suggests that it is the number of people who share their (one and only) signal that matters. A limitation of these proxies is that we cannot differentiate whether tweets are posted because Twitter users share their signals, or whether tweets reflect a general discussion about the upcoming decision. In other words, Twitter traffic might increase without additional information content that gets exchanged, e.g. if there is increased attention. This means that Twitter traffic could mix signals about information sharing and increased attention.

Accordingly, we cannot test directly whether exchanging views negatively affects the magnitude of monetary policy surprises (as our stylised framework would suggest). Instead, we rely on interaction terms with proxies for information exchange motives to identify more precisely the role of exchanging views: more Twitter traffic should reduce the magnitude of a monetary policy surprise 1) if inattention, i.e. disagreement about the state of the economy, is large and 2) if macroeconomic uncertainty is large. After all, in the absence of disagreement or uncertainty, there is no benefit from exchanging signals.

3. Data

In this section, we describe the data we use for our empirical analysis. Table A in the Appendix provides summary statistics for all variables.

3.1. ECB-related Twitter traffic

The data regarding Twitter traffic related to the ECB is based on Ehrmann and Wabitsch (2022). It consists of tweets that contain “ecb”, “european central bank”, “draghi” (until the end of his tenure as ECB president) or “lagarde” (since the beginning of her tenure as ECB president) in the text, hashtag or username. These tweets were scraped via Twitter’s Advanced Search using the Python library “GetOldTweets” (Henrique 2016), setting the Twitter Advanced Search language filter to English, given that this is the language spoken in financial markets and the ECB watchers’ community in general. Following this initial collection, the sample is cleaned thoroughly to remove all tweets that are unrelated to the European Central Bank. For instance, tweets related to the English Cricket Board (which also abbreviates as ECB) are removed. The language of the tweets is double-checked using the Python library “langdetect” (Danilak 2015).


⁹ We assume that the degree of inattention is identical for the quiet period and the rest of the intermeeting period, or at least positively correlated.

The data cover the period from January 2012 (when the usage of Twitter in Europe started to stabilise) until April 2020, for a total of 79 Governing Council meetings and press conferences.¹⁰

The underlying data are daily. From there, we construct a measure of Twitter traffic by summing the number of tweets over a period, dividing over the number of days that were aggregated, and taking the logarithm of this metric. Also, we simply count the number of Twitter accounts that post ECB-related tweets on each day and a concentration measure by computing the Herfindahl-Hirschman indicator.¹¹ In our benchmark analysis, we do so for the three days before a press conference (i.e. Monday to Wednesday), but we check for robustness over different time windows.

As shown by Ehrmann and Wabitsch (2022) and Masciandaro et al. (2022), Twitter traffic is highly responsive to the ECB’s press conference – it increases already several days before the press conference, and remains elevated for several days after. In addition, more Twitter users join the discussion about the ECB around the press conference, such that the Herfindahl-Hirschman indicator falls substantially and significantly. Figure 1 shows that the outcome of the Governing Council meetings has an effect on Twitter traffic. In line with the Bank of England results in Haldane et al. (2021), we find that whenever there is a “change” decision, i.e. policy rates or QE purchase amounts are altered, liquidity operations are announced or the ECB changes its forward guidance, Twitter traffic is higher than otherwise.

Figure 1 - Post-Governing Council Twitter traffic and ECB decisions



Note: the figure shows a box plot of the number of ECB-related tweets that are issued on Twitter per day in the three days after a press conference, separately for 59 press conference days when no policy change was announced and 19 days when a policy change was announced. The box plot reports the median (line in the box), 25th and 75th percentiles (Q1 and Q3, border of box), adjacent values (Q3+1.5*(Q3-Q1) and Q1-1.5*(Q3-Q1), whiskers) and outliers (dots).

The difference is large, and highly statistically significant. The average number of tweets per day over the three-day time window following a no-change decision is 999, the median stands at 889. This compares to a mean of 2592 and a median of 2072 following the announcement of a change decision. Still, there can be substantially higher Twitter traffic also after no-change announcements, as shown by the outliers in Figure 1. The press conference on 22 October 2015 is such an example. On that occasion, no policy change was made, but there were hints that

¹⁰ Meetings always take place on Thursdays, at a monthly frequency until 2014, and 8 times a year since 2015.

¹¹ The indicator is constructed as $HHI_t = \sum_{i=1}^{N_t} s_{i,t}^2$, where $s_{i,t}$ is the “market share” of a tweeting user i in the “tweet market” during time interval t ($s_{i,t} = \text{tweets}_{i,t} / \sum \text{tweets}_t$), and N_t is the number of users during time interval t . Hence, the larger the indicator, the more concentrated is Twitter traffic, as the “market share” of the participating accounts in a given time interval is larger.

the Governing Council would consider further monetary easing at its subsequent meeting in December (when substantial easing measures were indeed announced). This pushed stock markets higher and the euro lower,¹² and substantially raised Twitter traffic. In Figure 1, it is the maximum observations among the no-change decisions, with 2967 tweets per day over the 3-day window following the press conference. This example illustrates that it is important to understand to what extent policy decisions have been anticipated, and to focus on the unanticipated component of policy decisions. We will now explain how this is measured.

3.2. Monetary policy surprises

Following Gürkaynak et al. (2005), Campbell et al. (2012) and Hanson and Stein (2015), the literature has started differentiating surprises that relate to the *current* policy stance and surprises that relate to news about the expected path of future policy and policymakers' views of the future state of the economy. A simple and transparent way to capture these different types of surprises is to use the change in Overnight Index Swaps (OIS) rates at different maturities, where the change in shorter maturities reflects surprises about the current policy stance, and changes in longer maturities are a proxy for surprises related to the future outlook. Based on Altavilla et al. (2019), we measure monetary surprises in a tight window around ECB announcements on Governing Council meeting days starting around 13.30 and ending around 15.45. Such a narrow time window facilitates identification, as any market move is likely to be dominated by the news about the ECB. It also enables us to limit the possibility that a “pre-meeting drift” could bias monetary surprise measures (Lucca and Moench, 2015). Following Hanson and Stein (2015), we consider the change in the OIS rates at the 2-year maturity as our benchmark surprise measure.¹³ This maturity has the advantage that it is responsive to news about standard policy rates but also about unconventional policy tools such as extended liquidity provisions, forward guidance or asset purchases. For robustness purposes, we also consider various other horizons.

As a first test of the relationship between monetary policy surprises and Twitter traffic, we test how Twitter traffic after policy announcements responds to surprises contained therein. Our hypothesis is that for larger *absolute* monetary surprises, Twitter traffic increases relatively more, given that the news content and the need (or desire) to discuss the ECB's decisions increases. We estimate Equation (4) over 79 policy announcements using OLS and heteroscedasticity-robust standard errors:

$$T_x = \alpha + \beta |MPS_t| + \varepsilon_t \quad (4)$$

where T_x denotes one of various Twitter-related measures in time window x , and $|MPS_t|$ is the absolute monetary policy surprise generated on the announcement day t , i.e. $|MPS_t| = |OIS_{t,15:45} - OIS_{t,13:30}|$. The different columns in Table 1 report results for various Twitter-related measures. The first set of columns estimates the responsiveness of the log number of tweets on the day after the press conference and of the average number of tweets per day over the three days following the press conference. These effects are positive and highly statistically significant. They are furthermore large (if the absolute surprise increases by 1 basis point, Twitter traffic increases by around 14 percentage points on the policy announcement day), and

¹² See, e.g., the article in the Guardian on 22 October 2015: “Markets lifted by Draghi hints on more stimulus measures - as it happened”, <https://www.theguardian.com/business/blog/live/2015/oct/22/markets-expect-ecb-draghi-to-hint-on-more-qe-business-live>.

¹³ For the measure of monetary surprises and in the rest of the paper, we consider spot OIS rates when not indicated otherwise. When we compare forward and spot OIS rates, we stress explicitly the difference.

persistently so. The effect is estimated to be significant for the three days after.¹⁴ The number of retweets also increases.

The even columns in Table 1 furthermore control for whether or not the ECB has changed policy at a given meeting, including changes to policy rates, asset purchases, liquidity operations or forward guidance. The results show that Twitter traffic responds disproportionately if there has been some action.

Table 1 - Twitter traffic after ECB meetings and ECB monetary surprises

Tweets _x	Tweets _{t+1}		Tweets _{3d}		Retweets _{3d}	
	(1)	(2)	(3)	(4)	(5)	(6)
MPS _{2y}	0.142*** (6.07)	0.079** (2.41)	0.135*** (6.23)	0.076** (2.43)	0.098** (2.08)	0.005 (0.13)
Action		0.714*** (3.30)		0.670*** (3.21)		1.040*** (5.42)
Constant	7.450*** (87.03)	7.398*** (104.79)	6.743*** (82.22)	6.694*** (99.29)	6.013*** (46.90)	5.937*** (47.28)
Obs.	78	78	78	78	78	78
R ²	0.21	0.38	0.21	0.38	0.05	0.24

Note: the table shows OLS parameters of Equation (4) estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variables correspond to various measures of Twitter traffic. The explanatory variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level. Even columns also control for whether or not the ECB has changed policy at a given meeting. Twitter data are available until April 30, 2020, which coincides with the date of a press conference. Accordingly, only 78 observations are available. Columns (1) and (2) report results for Twitter traffic on the day after the press conference. Columns (3) and (4): Twitter traffic per day over the three days after the press conference; columns (5) and (6): Number of retweets over the three days after the press conference.

To summarise, the fact that our various measures of Twitter traffic are highly responsive to the ECB's policy announcements and the magnitude of monetary surprises makes us confident that our data constitute a good proxy for information flow about the ECB on social media.

3.3. Motives to share information

The hypotheses developed in Section 2 suggest that there are instances when exchanging views on Twitter, or sharing private information, might be particularly valuable. This could be the case when disagreement among central bank watchers is large or when aggregate uncertainty is high.

Our first proxy is an agnostic summary measure based on Google trend, the number of Google searches for the ECB at a given day.¹⁵ This measure is likely driven by different potential motives for information demand, which might imply that an exchange of views is more beneficial to market participants. Its advantage is that it is a parsimonious summary measure, but it does therefore not allow us to understand the underlying drivers. We will try to get at this through our more specialised proxies.

¹⁴ Estimating these results for each subsequent day (not reported here for brevity) shows that the effect is largest the day after the press conference, but is still present and large (nearly 10%) three days later (which is a Sunday).

¹⁵ Google trends provides daily data, but only for shorter time periods than required here, whereas monthly data are available for the full period. To get a daily measure over the entire sample, we calculated the average in a given month and mapped this to the monthly data in that particular month.

Our proxies for disagreement are based on surveys among professional forecasters conducted by Consensus Economics. They relate to disagreement about one-year ahead euro area HICP inflation and GDP growth (i.e. about the macroeconomic outlook) and about one-year ahead 3-month euro area interest rates (i.e. about the future path of monetary policy). For HICP inflation and GDP growth, the one-year ahead forecasts are constructed from current calendar-year forecasts and next calendar-year forecasts following Dovern et al. (2012). In line with the earlier literature that has employed these data (e.g., Dovern et al. 2012 or Mankiw et al. 2003), we use the inter-decile and the inter-quartile range as our measure of disagreement, i.e. measures that are robust to outliers. The survey has a monthly frequency, whereas we require the data to be at the frequency of Governing Council meetings. We match the frequencies to ensure that for each meeting, the uncertainty measure is based on the most recent survey conducted before the beginning of the quiet period. We also use, for a robustness test, a measure of disagreement derived from the ECB's Survey of Professional Forecasters (SPF). We compute the cross-sectional standard deviation of point estimates of 2-year ahead inflation and GDP forecasts. One limitation of this database is that it is available at the quarterly frequency only.

As to aggregate uncertainty, we use several proxies, covering uncertainty about the state of the economy and the macroeconomic outlook, financial market uncertainty and policy uncertainty. Uncertainty about the macroeconomy is measured in different ways. First, we obtain an uncertainty measure that is based on Twitter from Baker et al. (2021).¹⁶ This measure consists of the total number of English-language tweets containing keywords related to uncertainty and related to the economy. From this index, which has a daily frequency, we compute the average over the three days preceding the Governing Council meeting.

Second, financial market uncertainty is proxied through the uncertainty component contained in the VSTOXX. In analogy to its U.S. counterpart VIX, the VSTOXX measures the 30-day implied volatility of the EUROSTOXX50, and is meant to reflect investor sentiment and overall economic and financial uncertainty. The uncertainty component contained in this measure is taken from Bekaert et al. (2021).¹⁷ Starting from daily data, we compute the average over the three days preceding the Governing Council meeting.

Third, the last type of uncertainty that we aim to measure relates to policy uncertainty, where we use the European version of the Baker et al. (2016) economic policy uncertainty (EPU) index, which is based on a word count of the terms “uncertain” or “uncertainty” and “economic” or “economy” in 10 European newspaper. We compute the average of this index over the three days preceding the press conference.

We also use, for robustness, measures of uncertainty derived from the ECB's SPF. We compute the mean of the standard deviation of forecasters' individual probability distribution of inflation and output forecasts and the standard deviation of the aggregate probability distribution.¹⁸ These measures capture the uncertainty around the outlook for inflation and output 2 years ahead.

¹⁶ The data, as well as those for the economic policy uncertainty index described below, are available on <https://www.policyuncertainty.com>.

¹⁷ Based on Bekaert and Hoerova (2014), they decompose the squared VIX index into the conditional variance of stock returns (the uncertainty measure) and a variance risk premium measuring risk aversion.

¹⁸ See <https://www.ecb.europa.eu/pub/pdf/scpops/ecbocp59.pdf> for more details about the ECB's SPF and these uncertainty measures.

3.4. News flow controls and controls for expectations regarding upcoming decisions

While our sample covers the ECB's quiet period and is therefore characterised by limited news flows from the ECB, it is not entirely free from news that might affect ECB-related Twitter traffic as well as the predictability of the ECB's decision. Accordingly, it is important to control for such news.

The first proxy relates to breaches of the ECB's quiet period, which occasionally happen. Speeches by Governing Council members during the quiet period might well lead to increased traffic on Twitter, and might possibly also lead to a pre-meeting drift in financial markets (Lucca and Moench, 2015) and affect expectations of the upcoming decision. We control for such breaches using the dataset by Rieder and Gnan (2021), in the form of a dummy variable that is equal to one if there is at least one quiet period breach in the three days before the press conference. This is the case in 10 out of the 79 observations. For 8 of these instances, one breach was observed; for the remaining two instances, two breaches were recorded. Given the small number of double breaches, we do not differentiate them further from single breaches.

The next set of proxies takes into account that even if the ECB is in its quiet period, other central banks are typically not, so news emanating from other central banks might also be important to assess the outlook for the ECB. We test for this in two ways. First, by including a dummy variable that is equal to one for quiet periods during which the Fed announces its policy decision, which is the case for 7 of our 79 observations. Second, by controlling for speeches given by members of the FOMC. This variable is based on FOMC Speak, a speech dataset provided by the Federal Reserve Bank of St Louis,¹⁹ and is set to one for each of the 58 quiet periods that contain at least one speech by a FOMC member.

Two proxies summarise the surprise component contained in macroeconomic data releases during the time window we study. The first is the index of macroeconomic surprises developed by Scotti (2016). It is based on the surprise component contained in the releases of GDP, industrial production, the unemployment rate, retail sales and the PMI. By taking the difference of the index over the three-day time window that we study (i.e. subtracting the surprise index four days prior to the press conference from the index on the day preceding the press conference), and by taking the absolute value of this difference, our proxy measures the magnitude of surprising announcements of all relevant macroeconomic indicators that were made in the time window of our analysis. For 28 out of our 79 observations, no relevant news arrived, such that our measure is equal to zero. The second proxy is the Citigroup Economic Surprise Index.²⁰ As this covers more macroeconomic releases, the absolute difference over the three-day window is different from zero for all observations in our sample.

In addition, we will control for expectations regarding the upcoming monetary policy decision. The more financial market participants expect a change in monetary policy in the upcoming meeting, the stronger will be the incentive to pay attention to the meeting (Boguth et al. 2019). Similarly, changes in market prices can be observed by all market participants and might generate higher attention. We proxy for this in two ways. First, by computing the absolute change in 2-year OIS rates between the previous Governing Council press conference and the beginning of our Twitter time window (so three days before the Governing Council). Specifically, we take the absolute difference of the opening price on the day after the

¹⁹ <https://www.stlouisfed.org/fomcspeak/viewbydate>.

²⁰ To construct the index, Citi measures the surprises in units of standard deviation, weights them according to their importance in terms of their previous impact on market prices, and then constructs a moving average of the surprises over the past ninety days using a roughly exponential weighting scheme.


Governing Council meeting and the closing price on the Friday before the upcoming Governing Council meeting, i.e. $|\Delta OIS_{pre\ QP}| = |OIS_{t-6,close} - OIS_{\tau+1,open}|$, where t denotes the day of the subsequent press conference and τ the day of the preceding Governing Council meeting. This variable controls for the change in the monetary stance that has already been priced into markets *before* the start of the Twitter window, and helps us to control for the possibility of a “pre-meeting drift” in financial markets (Lucca and Moench, 2015) which, in turn, could affect the monetary policy surprise as well as Twitter traffic by generating more attention to the upcoming decision.

The second proxy variable measures the change in the monetary stance that is expected by markets in between the beginning of our time window and the subsequent press conference, i.e. what markets expect to happen *after* the start of the Twitter window. We do so by taking the absolute difference of the one-week ahead OIS 2-year forward rate and the OIS 2-year spot rate at market close on the Friday before the upcoming Governing Council meeting, i.e. $|OIS_{for} - OIS_{spot}| = |OIS_{t-6,close}^{for,t+7} - OIS_{t-6,close}^{spot}|$.

4. Twitter traffic before policy meetings and monetary surprises

In the previous section, we have shown that Twitter traffic *after* the Governing Council meetings responds to the magnitude of monetary policy surprises. We now move our focus on what happens *before* Governing Council meetings. Market expectations about policy in the days before Governing Council meetings move continuously. To study this evolution, we look at the absolute value of the difference between 2-year OIS rates h days before a Governing Council meeting and 2-year OIS rates after the press conference on the day of policy announcements, averaged over our sample of 79 meetings. Even though financial markets move most in the uprun to the quiet period, they still show considerable movement during the quiet period. Specifically, market expectations improve, on average, by 1.1 basis points in the week before the quiet period (compared to a 0.8 basis points improvement over the two preceding weeks), and still record an improvement of 0.5 basis points during the four trading days in the quiet period. In other words, during the quiet period, market expectations move, on average, by nearly half of what we see in the week before (where Governing Council members still provide speeches), and the movement is more than one fourth of the average resulting monetary policy surprise (which amounts to 1.87 basis points). This shows clearly that market participants continue to update their assessment also in the absence of central bank communication.

Figure 2 - Pre-Governing Council Twitter traffic and ECB decisions



Note: the figure shows a box plot of the number of ECB-related tweets that are issued on Twitter per day in the three days before a press conference, separately for 59 press conference days when no policy change was announced and 19 days when a policy change was announced. The box plot reports the median (line in the box), 25th and 75th percentiles (Q1 and Q3, border of box), adjacent values ($Q3+1.5(Q3-Q1)$ and $Q1-1.5*(Q3-Q1)$, whiskers) and outliers (dots).

Ehrmann and Wabitsch (2022) also show that Twitter traffic already intensifies in the days *before* Governing Council meetings. Figure 2 shows that there is more Twitter traffic prior to change decisions. This is probably not overly surprising – often, change decisions are broadly anticipated in the sense that central bank watchers expect some change with high probability, but might be less certain about the details of the decision (which tool is adjusted and by how much). As can be seen in figure 2, there is substantial variation in Twitter traffic both ahead of no-change and before change decisions. This suggests, once more, that it is important to go beyond the change/no change distinction and study how Twitter traffic relates to the subsequent monetary policy surprise. To get at this, we estimate the following relationship:

$$|MPS_t| = \alpha + \beta T_x + \varepsilon_t \quad (5)$$

which is the mirror image of Equation (4). Previously, we correlated post-Governing Council Twitter traffic with the magnitude of monetary policy surprises. Now, we write the correlation the other way, i.e. we link the magnitude of monetary policy surprises with Twitter traffic *ahead* of Governing Council meetings. Although we acknowledge that we are only estimating correlations that do not allow for causal inference, the timing restriction – the left-hand side in both equations is measured later in time than the right-hand side – limits the possibility for a reverse causation. Equation (5) is estimated over 79 policy announcements using OLS and heteroscedasticity-robust standard errors.


Table 2 – Twitter traffic before ECB meetings and ECB monetary surprises

$ MPS_{2y} $	Tweets _{t-3}	Tweets _{s,3d}	Retweets _{s,3d}	Tweets _{p95}	MPS _{p95}
	(1)	(2)	(3)	(4)	(5)
Tweets _x	1.091** (2.43)	1.251** (2.30)	0.314 (0.84)	2.460*** (3.66)	0.625* (1.86)
Constant	-5.630* (-1.89)	-6.931* (-1.87)	-0.100 (-0.04)	-15.235*** (-3.33)	-2.816 (-1.21)
Obs.	79	79	79	75	75
R ²	0.11	0.12	0.01	0.26	0.06

Note: the table shows OLS coefficients of Equation (5) estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level. The explanatory variables listed in the column headers correspond to various measures of Twitter traffic. Column (1) reports results for Twitter traffic three days preceding the press conference. Column (2) to (3): Twitter traffic per day over the three days preceding the press conference, for the number of original tweets and the number of retweets. Columns (4) and (5) remove observations above the 95th percentile of Twitter traffic and monetary policy surprises, respectively.

Table 2 shows estimates of Equation (5) and uncovers a positive correlation – the larger the subsequent monetary policy surprise, the higher is Twitter traffic in the preceding days. This relationship is very robust.²¹ It already exists at the beginning of our time window three days preceding the press conference, and is larger in magnitude if estimated over all three days preceding the press conference.²² Figure 3 shows how this correlation evolves in the uprun to the Governing Council meetings. It becomes statistically significantly positive just on the first day of the quiet period, and remains so on all business days during the quiet period (note that days -5 and -4 are weekends).

Figure 3 – Twitter traffic and its correlation with monetary policy surprises



Note: the figure shows the average number of ECB-related tweets (blue line, right axis) as well as the estimated correlation between Twitter traffic and the absolute monetary policy surprise (following equation 5, left axis). Point estimates are provided as orange dots, along with ± 1 and 2 standard deviation confidence bands. Days -12, -11, -5 and -4 are weekends.

²¹ Table B in the Appendix shows that the effect arises from surprises around the press release and the Q&A session. In addition, Table C in the Appendix shows that this effect is at work for various horizons at the short and medium end of the maturity spectrum.

²² It is also estimated significantly for each of the three days preceding the press conference separately, and does not result from a few outliers, as shown in columns (6) and (7), where we remove the most extreme values for Twitter traffic and for monetary policy surprises, respectively. The correlation becomes even stronger when we remove the most extreme observations for Twitter traffic.

This result seems to go against the hypothesis tested in this paper, that exchanging views among central bank watchers should help improve knowledge about the ECB’s monetary policy, and hence reduce the magnitude of the subsequent monetary policy surprises. However, estimating this correlation – in Equation (5) – falls short of testing our hypothesis for various reasons. First, as already discussed in the hypothesis section, Twitter volumes mix different signals and can only partially serve as a proxy for the exchange of views. For instance, a tweet might raise questions about the upcoming decision, or raise awareness about it, without providing an own view. In such a case, increased Twitter traffic ahead of large surprises might simply reflect increased attention. Second, larger surprises might be correlated with expectations about an upcoming change decision. In such a case, Twitter participation would increase with the probability of an upcoming change decision. Third, it might be important to control for the (limited) news flow during the quiet period. These three types of factors would suggest that the previous result is driven by an omitted variable bias.

Table 3 – Controlling for news flow and expectations of upcoming decisions

MPS _{2y}	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Tweets _{.3d}	1.251** (2.30)	1.539** (2.47)	1.303*** (2.69)	1.254** (2.30)	1.251** (2.29)	1.266** (2.29)	1.287** (2.39)	0.803 (1.58)	1.296** (2.45)	0.985* (1.75)
Tweets _{pre QP}		-0.412 (-1.03)								
Quiet period breach			-0.310 (-0.32)							
FOMC meeting				-0.272 (-0.44)						
FOMC speech					-0.010 (-0.02)					
Surprise _{Scotti}						-12.096 (-0.67)				
Surprise _{Citi}							0.045 (0.95)			
Δ OIS _{pre QP}								0.213*** (3.48)		0.141** (2.20)
OIS _{for} - OIS _{spot}									0.134*** (2.82)	0.090* (1.68)
Constant	-6.931* (-1.87)	-6.159 (-1.65)	-7.255** (-2.19)	-6.922* (-1.87)	-6.921* (-1.85)	-6.953* (-1.83)	-7.497* (-1.98)	-4.717 (-1.35)	-8.244** (-2.18)	-6.347 (-1.59)
Obs.	79	79	79	79	79	79	79	79	79	79
R ²	0.12	0.12	0.12	0.12	0.12	0.12	0.13	0.26	0.26	0.31

Note: the table shows OLS coefficients of Equation (5) with the addition of the listed control variables, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level. “Tweets” denotes Twitter traffic three days preceding the press conference.

Table 3 does therefore repeat the earlier analysis (for the benchmark variant that studies Twitter traffic for the three days preceding the press conference), but adding one by one our various controls for news flow and policy expectations. It turns out that none of the information flow controls is estimated to be statistically significant. In contrast, both the backward-looking and the forward-looking measure of market expectations are significant (when entered individually and jointly): Twitter traffic is larger when 2-year OIS rates moved more between the previous press conference and the beginning of our Twitter time window, and when the market expects a larger change in interest rates between now and the subsequent

Governing Council meeting. Controlling for the backward-looking attention effect, we find that the correlation of Twitter traffic with the policy surprises becomes insignificant.²³

This is in line with the idea that attention is endogenous – if agents see market prices move, they pay more attention to the upcoming decision, which is reflected in more Twitter traffic, but does not imply that more signals are being shared.²⁴ Hence, to test our hypotheses related to the role of information exchange, we move beyond the unconditional correlation and estimate the extended relationship:

$$|MPS_t| = \alpha + \beta T_x + \gamma \Omega_{pre-qp} + \delta T_x \Omega_{pre-qp} + \varepsilon_t \quad (6)$$

where Ω_{pre-qp} denotes the proxies for instances where an exchange of information during the days before the policy meeting could be particularly valuable (measured prior to the quiet period, hence denoted with subscript *pre-qp*). While the parameter β is not well identified for the reasons outlined above, the hypothesis of interest is that if $\gamma > 0$, then we also find that $\delta < 0$, implying that increased Twitter traffic mitigates the effect of, for instance, uncertainty on the magnitude of monetary policy surprises. As before, Equation (6) is estimated over 79 policy announcements using OLS and heteroscedasticity-robust standard errors.

Table 4 – The role of information exchange metrics

MPS _{2y}	Info demand	Disagreement			Uncertainty		
	Google trends	HICP _{IDR}	GDP _{IDR}	3-month rates _{IDR}	Uncert _{Econ}	Uncert _{VSTOXX}	Uncert _{Ec.Policy}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Tweets _{.3d}	2.993*** (3.13)	4.715** (2.19)	3.860** (2.13)	1.372 (1.53)	1.085 (1.03)	2.116** (2.34)	0.500 (0.27)
Proxy	0.269** (2.29)	40.929 (1.60)	28.667* (1.91)	12.572 (0.40)	-0.010 (-0.18)	0.291 (1.51)	-0.028 (-0.44)
Tweets · Proxy	-0.034** (-2.45)	-6.543* (-1.75)	-3.929* (-1.81)	-1.413 (-0.31)	0.002 (0.19)	-0.037 (-1.38)	0.004 (0.42)
Constant	-20.704*** (-2.95)	-28.815* (-1.98)	-25.842** (-2.09)	-8.152 (-1.34)	-5.865 (-0.82)	-13.624** (-2.18)	-1.501 (-0.12)
Obs.	79	79	79	79	79	79	79
R ²	0.16	0.25	0.18	0.14	0.12	0.20	0.12

Note: the table shows OLS coefficients of Equation (6) estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. The “Tweets” explanatory variable is the daily average of tweets over the 3 preceding days for each Governing Council meeting. The “Proxy” explanatory variable covers ECB-related Google trend measure (column 1), disagreement (columns 2 to 4), and uncertainty (columns 5 to 7). The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

Table 4 reports the corresponding estimates. We start by using the daily Google trend measure, i.e. a general proxy for information demand that synthesizes all potential motives for information exchange. The limitation of this measure is that we cannot give it an economic interpretation, but it allows for a test whether higher Twitter traffic is associated with smaller monetary policy surprises if demand for information related to monetary policy is high. As can be seen in column (1), higher information demand in the days ahead of Governing Council

²³ We have also controlled for monetary surprises or whether there was a change decision at the previous Governing Council meeting, as this may affect expectations about the subsequent policy decision. Neither of these variables affects the correlation between Twitter traffic and the magnitude of monetary surprises.

²⁴ Table D in the Appendix shows to what extent Twitter traffic is responsive to the various information exchange motives proxies. With the exception of economic policy uncertainty, we generally find this to be the case, in the expected direction: Twitter traffic increases if disagreement or uncertainty are high. To test whether the patterns identified here are related to Twitter or to disagreement, Table E in the Appendix shows that the pattern is also present if we relate it to the part of Twitter traffic that is not correlated with disagreement.

meetings positively correlates with the magnitude of monetary surprises. But we find that the interaction term between Twitter traffic and Google trend data is negative, meaning that higher Twitter traffic is associated with a smaller correlation between information demand and the magnitude of monetary surprises. This is in line with the notion that Twitter traffic matters in influencing the correlation between at least one of the economic factors driving information demand and monetary surprises.

To understand which factors are driving this result, we will now proceed by using the various proxies for disagreement and uncertainty. We start by studying one proxy at a time, and subsequently will estimate a joint model with all significant proxies.

The first proxies relate to disagreement (columns (2) to (4)). We find the expected effect for disagreement about the economic outlook, i.e. about inflation and GDP growth. We find that $\delta < 0$, which provides evidence in line with the notion that exchanging views on Twitter can moderate the effect of disagreement on the magnitude by which agents are surprised by the policy decision.²⁵ More concretely, while the effect of a 1 percentage point increase in inflation disagreement yields a 41 basis points increase in the magnitude of monetary surprises (although only significant at the 11% level, we show later that this effect becomes significant at the 5% level in more complete specifications), a 1 standard-deviation (SD) increase in Twitter traffic (0.58%) is associated with a reduction by 3.8 basis points, i.e. by 9%. Similarly, the effect of a 1 percentage point increase in output disagreement yields a 29 basis points increase in the magnitude of monetary surprises, and a 1 SD increase in Twitter traffic is associated with a reduction by 2.3 basis points, i.e. by 8%. In contrast, no significant relationship is found for disagreement regarding the outlook for interest rates.

Columns (5) to (7) show that in contrast to disagreement, we cannot identify any relevant correlation for the various uncertainty indices. Taken together, these results are in line with our hypothesis that exchanging information via Twitter can improve the understanding of the upcoming monetary policy decision and thereby reduce the surprise component therein. However, the effect is not universal; while we find it for disagreement, it does not show up for the various uncertainty measures.

Table 5 reports results if all variables that yield significant results in isolation are studied jointly, i.e. two disagreement proxies and the two proxies for market expectations about upcoming decisions are added. As shown in column (1), in this specification, the absolute change in 2-year OIS rates is no longer significant, such that we drop it in column (2). The earlier individual relationships with the disagreement measures hold and are even sharpened, as the statistical significance of the estimated coefficients increases in most cases, and the economic effects tend to be similar or larger.²⁶

²⁵ Figures H and I in the Appendix show how the estimated coefficients for these interaction terms evolve in the uprun to the Governing Council meetings.

²⁶ In Appendix Table F, we show that these results are robust to varying the time window over which we measure Twitter traffic, and by varying the measure of Twitter traffic itself.

Table 5 – Combining increased attention and information exchange motives

$ \text{MPS}_{2y} $	IDR (1)	IDR (2)	IQR (3)	SPF _{disagreement} (4)	SPF _{ind.uncert} (5)	SPF _{agg.uncert} (6)
Tweets _{.3d}	7.301*** (2.93)	7.566*** (3.39)	5.349** (2.55)	10.053*** (3.10)	1.392 (0.12)	2.804 (0.32)
HICP _x	37.567** (2.03)	37.959** (2.08)	59.217 (1.53)	126.519*** (2.77)	-83.065 (-0.46)	2.120 (0.02)
GDP _x	32.092** (2.25)	33.491*** (2.66)	37.633* (1.80)	90.258* (1.93)	77.684 (0.57)	11.882 (0.16)
Tweets * HICP _x	-5.985** (-2.21)	-6.106** (-2.31)	-8.890 (-1.60)	-18.402*** (-2.80)	11.790 (0.45)	-0.205 (-0.02)
Tweets * GDP _x	-4.407** (-2.18)	-4.584** (-2.53)	-5.142* (-1.75)	-12.989* (-1.92)	-11.077 (-0.56)	-1.721 (-0.16)
$ \Delta \text{OIS}_{\text{pre QP}} $	0.035 (0.39)
$ \text{OIS}_{\text{for}} - \text{OIS}_{\text{spot}} $	0.091* (1.98)	0.101** (2.59)	0.130*** (2.91)	0.135*** (3.16)	0.136*** (2.88)	0.135*** (2.75)
Constant	-48.709*** (-2.92)	-50.368*** (-3.34)	-36.326** (-2.48)	-68.687*** (-3.08)	-8.696 (-0.10)	-19.167 (-0.32)
Obs.	79	79	79	79	79	79
R ²	0.43	0.43	0.33	0.36	0.27	0.26

Note: the table shows OLS coefficients of Equation (6) allowing for several Ω_{t-h} variables with the addition of the listed control variables, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. Column (3) replaces the interdecile range among inflation and GDP growth forecasts by the interquartile range. Columns (4) to (6) use disagreement and individual and aggregate uncertainty based on the SPF. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

Both disagreement measures remain in the specification jointly, whereas uncertainty does not seem to matter. It has been shown repeatedly that disagreement and uncertainty are two different concepts and do not always comove (for the United States, see Rich and Tracy (2010), for the UK, Boero et al. (2008), and for the euro area, Abel et al. (2016)). In columns (3) and (4), we explore further the disagreement result. In particular, we are interested whether Twitter exchanges are associated with smaller surprises if disagreement arises due to divergent views in the tails of the distribution, or if disagreement is a more general phenomenon. To do so, we modify the measure of disagreement, by using the interquartile range as opposed to the interdecile range. This modification replicates the earlier results qualitatively, but the statistical fit is worse, as can be seen by the reduction in R². The interpretation of this result is that what matters is the disagreement in the tails. The interdecile range includes a broader range of views, as it measures disagreement between the 10th and the 90th percentile of the distribution. If this measure of disagreement is large, Twitter traffic is more strongly associated with a better understanding of monetary policy than if the disagreement between the 25th and the 75th percentile of the distribution is large.²⁷

Column (4) is based on disagreement in the SPF, and shows that the results are not unique to the use of Consensus Economics. The SPF allows us to directly compare disagreement and uncertainty, as it not only contains forecasts of different individuals, but also density forecasts for output growth and inflation. Columns (5) and (6) show that neither individual uncertainty nor aggregate uncertainty generate the pattern that we obtain for disagreement.

²⁷ The result that the heterogeneity of beliefs matters for macroeconomic dynamics is consistent with Meeks and Monti (2022). They show that the cross-sectional distribution of expectations summarized by three factors that capture disagreement, skew and shape is key for inflation dynamics.

5. Understanding the mechanism

The evidence so far reveals a clear pattern whereby increased Twitter traffic is associated with relatively smaller monetary policy surprises when disagreement about the economic outlook is large. But what could generate such a pattern? In this section, we try to dissect the channel that could be at work, along three dimensions. First, we study which tweets are the most relevant. Second, we check whether there are alternative channels for information exchange beyond Twitter. Third, we analyse how market expectations evolve during the time window under study here.

The results so far considered all tweets related to the ECB, regardless of the source or the impact of the tweets. A more granular analysis might be warranted at this stage. To do so, we will use the distinction of tweets posted by experts and non-experts following Ehrmann and Wabitsch (2022), and differentiate according to the content of the tweet. We will also look at tweets that get retweeted (as these get distributed more widely), and at long vs short tweets (as short tweets might contain less information than long tweets, and are less costly to produce, hence might be more noisy measures).

The differentiation of experts and non-experts in Ehrmann and Wabitsch (2022) is based on the accounts from which tweets are posted. If the account owners post tweets about the ECB regularly (at least on the occasion of every second press conference), they are classified as experts. If they are irregular and furthermore have a low level of ECB centrality (below the 25th percentile of the distribution across all accounts), i.e. they post tweets about many topics and only irregularly about the ECB, they are classified as non-experts. As columns (1) and (2) of Table 6 show, the patterns are identical for experts and non-experts. To interpret this result, it is important to note that there are relatively few experts which, however, post many tweets, implying that there is relatively little diversity in views, whereas there are many non-experts who tend to post relatively fewer tweets about the ECB, but reflect a large heterogeneity of views. Both of these aspects seem to be helpful.

We also differentiate based on content. On the one hand, we characterise tweets as “economic” if they contain the words inflation, cost, price, wage, oil, employ*, labor, labour, output, growth, econom*, cpi, hicp, forecast or projection. “Monetary policy” tweets need to contain the terms monetary, policy, interest, rate, purchase, forward, guidance, fg, liquidity, decision, action, path, hik*, decreas*, APP, PSPP, PEPP, eas* or tight*. Note that a tweet can be characterised both as economic and monetary policy. Columns (3) and (4) in Table 6 show that the statistical fit is much better for monetary policy-related tweets. While we find negative interaction terms for both types of tweets, they are not or less statistically significant for the economic tweets, leading also to a considerably lower R².

Column (5) only includes tweets by experts that get retweeted at least once (which are only 7% of tweets), and shows that our relationship holds and is highly statistically significant. In contrast, as shown in column (6), the effect for tweets by experts that do not get retweeted does not yield significant results.²⁸ Similarly, the length of a tweet seems to matter – for tweets that have more than the median number of characters, we find the usual pattern, whereas this is not the case for short tweets (see columns (7) and (8)).

²⁸ We report the distinction for tweets by experts here – for tweets by experts, there is no difference between retweeted and not retweeted tweets.

Columns (9) and (10) use alternative proxies for κ , the degree to which agents share their signals. It is apparent that the relationship also depends on how concentrated the discussion on Twitter is and how many users participate. The results suggest that a reduction in the concentration measure (i.e. a discussion that has a more equal distribution of tweets across the participating Twitter accounts) and an increased number of users are associated with a smaller surprise when disagreement is large. This corroborates the earlier findings with the number of tweets and provides further support towards our hypothesis – the more financial market participants and central bank watchers exchange views and share their private signals about the state of the economy, the stronger the effect.

Table 6 – Differentiating tweets and participation in Twitter traffic

MPS _{2y}	Differentiating tweets								Participation	
	Experts' tweets	Non-experts' tweets	Mon pol tweets	Econ tweets	Retwt experts' tweets	Not retwt experts' tweets	Long tweets	Short tweets	h-stat _{3d}	TwAcc _{3d}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Tweets _{x,-3d}	6.131*** (2.65)	6.618*** (3.29)	7.339*** (2.93)	5.311** (2.38)	8.137*** (3.45)	4.993** (2.33)	8.615*** (3.53)	4.452*** (2.70)	-22.082*** (-2.75)	0.010*** (3.18)
HICP _{IDR}	21.773 (1.51)	26.160** (2.15)	32.374** (2.13)	23.424* (1.69)	20.025** (2.00)	14.032 (1.10)	39.919** (2.43)	13.377 (0.97)	-8.771** (-2.37)	2.436 (0.87)
GDP _{IDR}	20.913** (2.16)	21.274** (2.61)	21.719* (1.78)	12.750* (1.70)	25.181*** (3.72)	16.244* (1.80)	36.085*** (2.80)	18.826** (2.02)	-3.894* (-1.80)	6.165*** (4.01)
Tweets · HICP _{IDR}	-4.851* (-1.80)	-5.905** (-2.47)	-6.487** (-2.40)	-5.278* (-1.90)	-6.032** (-2.41)	-3.651 (-1.46)	-6.942*** (-2.66)	-2.973 (-1.35)	10.438 (0.98)	-0.009** (-2.38)
Tweets · GDP _{IDR}	-3.531* (-1.95)	-3.930** (-2.40)	-3.653* (-1.67)	-2.174 (-1.39)	-5.676*** (-3.36)	-2.856 (-1.60)	-5.423** (-2.63)	-2.827* (-1.91)	16.187*** (2.80)	-0.006*** (-3.03)
OIS _{for} - OIS _{spot}	0.100** (2.44)	0.118*** (2.84)	0.084** (2.26)	0.138*** (3.26)	0.101** (2.60)	0.101** (2.37)	0.096** (2.60)	0.112** (2.52)	0.118** (2.55)	0.087** (2.28)
Constant	-31.258** (-2.58)	-31.420*** (-3.16)	-38.840*** (-2.84)	-25.624** (-2.32)	-31.391*** (-3.37)	-23.425** (-2.23)	-52.559*** (-3.47)	-24.952** (-2.56)	10.088*** (3.16)	-5.413** (-2.61)
Obs.	79	79	79	79	79	79	79	79	79	79
R ²	0.39	0.39	0.43	0.38	0.42	0.36	0.42	0.37	0.41	0.42

Note: the table shows OLS coefficients of Equation (6) allowing for several $\Omega_{\tau-h}$ variables with the addition of the listed control variable, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. Columns (1) to (8) replicate the results for subsets of tweets as denoted by the column header. Columns (9) and (10) are based on a measure of user concentration and the number of accounts that participate in the ECB-related Twitter discussion, respectively. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

Our next step is to consider whether there are alternative ways of measuring information exchange. Table 7 shows results for several possible alternatives. We obtain the same pattern when we use the number of newswire articles in English that mention the ECB (sourced from Dow Jones' Factiva DNA service). Note that during the ECB's quiet period, these articles mostly quote analysts' views (an example is provided in Appendix G), and as such might be relatively close to Twitter. Still, the Twitter specification dominates the one with newswire articles in terms of R². In contrast, results with newspaper articles (also sourced from Dow Jones' Factiva DNA service, covering articles in English that mention the ECB) are not statistically significant, likely because they are catering for a less specialised audience, do not go into as much detail as the newswire reports and thus do not provide as broad an overview of analysts' views. The same holds true for Bloomberg's News Trend for the topic "European Central Bank", which covers newswire reports on the topic, but is much broader, as it also contains articles in newspapers and over 90,000 internet sources. It therefore seems that Twitter traffic captures information exchange better than the other proxies, likely because it captures the views of relatively more people.

Table 7 – Alternative proxies for information exchange

$ \text{MPS}_{2y} $	Twitter (1)	Newswires (2)	Newspapers (3)	Bloomberg (4)
Information channel	7.566*** (3.39)	5.290*** (3.18)	0.958 (0.29)	6.152** (2.03)
HICP _{IDR}	37.959** (2.08)	7.375 (1.13)	-2.324 (-0.15)	28.150 (1.15)
GDP _{IDR}	33.491*** (2.66)	12.568*** (3.33)	4.160 (0.48)	26.748* (1.67)
Info channel · HICP _{IDR}	-6.106** (-2.31)	-4.181* (-1.91)	-0.246 (-0.06)	-5.044 (-1.33)
Info channel · GDP _{IDR}	-4.584** (-2.53)	-3.623*** (-2.74)	-0.702 (-0.32)	-3.665 (-1.45)
$ \text{OIS}_{\text{for}} - \text{OIS}_{\text{spot}} $	0.101** (2.59)	0.108** (2.34)	0.123* (1.91)	0.093** (2.18)
Constant	-50.368*** (-3.34)	-13.301*** (-2.87)	-2.002 (-0.17)	-38.536** (-2.00)
Obs.	79	79	79	74
R ²	0.43	0.32	0.19	0.34

Note: the table shows OLS coefficients of Equation (6) allowing for several Ω_{t-h} variables with the addition of the listed control variable, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. Column(1) replicates the results for all tweets, all remaining columns include subsets of tweets. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

But how does Twitter traffic relate to smaller subsequent monetary policy surprises? One necessary condition is that Twitter traffic during the quiet period affects the formation of policy expectations during that window, so it can be linked to the smaller subsequent surprises. Table E in the Appendix shows that Twitter traffic during the quiet period is positively correlated with absolute movements in 2-year OIS rates. It also documents that the movements in 2-year OIS rates that can be predicted from Twitter traffic also generate the pattern whereby larger movements lead to smaller monetary policy surprises when disagreement is large. In addition, in Table 8, we inspect a possible channel that could be at work. If more information exchange on Twitter was causal for improving market expectations of upcoming decision, we should expect that the market expectation moves closer to the actual outcome during the Twitter time window – in particular when the benefit to information exchange is high, i.e. when disagreement is large. Hence, Table 8 reports results where we replace the left-hand side variable of our regressions with a variable that measures how much closer market expectations move to the actual outcome. This measure, which we label “improvement”, is constructed as the absolute difference between the hypothetical absolute surprise at the beginning of the time window and the actual absolute surprise. We do so in two ways. First, we start from the spot price at the beginning of the Twitter time window: $Improv_{spot,2y} = |OIS_{t,15:45} - OIS_{t-3,open}^{spot}| - |OIS_{t,15:45} - OIS_{t,13:30}|$. Second, we use as a starting point the 7-day ahead forward price on the Friday before the Governing Council meeting: $Improv_{for,2y} = |OIS_{t,15:45} - OIS_{t-6,clos}^{for,t+7}| - |OIS_{t,15:45} - OIS_{t,13:30}|$.

Table 8 – Market expectations and Twitter traffic

	Improv _{spot,2y}				Improv _{for,2y}			
	HICP _{IDR} & GDP _{IDR}	HICP _{IQR} & GDP _{IQR}	HICP _{IDR} & GDP _{IQR}	HICP _{IQR} & GDP _{IDR}	HICP _{IDR} & GDP _{IDR}	HICP _{IQR} & GDP _{IQR}	HICP _{IDR} & GDP _{IQR}	HICP _{IQR} & GDP _{IDR}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tweets _{3d}	-4.364*** (-3.12)	-4.205*** (-3.42)	-4.053** (-2.41)	-4.801*** (-4.04)	-6.329*** (-2.93)	-10.159*** (-3.29)	-6.835*** (-3.12)	-11.273*** (-3.62)
HICP _x	-21.061 (-1.17)	-66.448* (-1.96)	-19.992 (-1.02)	-69.885** (-2.40)	-31.851 (-1.03)	-164.711*** (-2.94)	-39.673 (-1.45)	-178.476*** (-3.13)
GDP _x	-22.701*** (-3.03)	-28.395 (-1.49)	-44.401*** (-2.72)	-17.212** (-2.27)	-48.098** (-2.26)	-88.038** (-2.09)	-96.843* (-1.98)	-50.582*** (-2.71)
Tweets + HICP _x	3.305 (1.30)	9.819** (2.01)	3.148 (1.12)	10.329** (2.48)	3.885 (0.91)	23.091*** (2.95)	4.942 (1.33)	25.024*** (3.12)
Tweets + GDP _x	3.448*** (3.33)	4.232 (1.58)	6.534*** (2.83)	2.652** (2.51)	6.846** (2.28)	12.545** (2.11)	13.817** (1.99)	7.148*** (2.72)
OIS _{for} - OIS _{spot}	0.049 (1.31)	0.029 (0.70)	0.041 (1.02)	0.037 (0.95)	-0.002 (-0.03)	0.017 (0.26)	-0.008 (-0.11)	0.024 (0.37)
Constant	28.532*** (3.07)	28.600*** (3.32)	26.973** (2.37)	32.170*** (3.89)	47.259*** (3.03)	72.513*** (3.31)	51.006*** (3.20)	80.518*** (3.65)
Obs.	79	79	79	79	79	79	79	79
R ²	0.18	0.14	0.12	0.20	0.18	0.22	0.17	0.26

Note: the table shows OLS coefficients of Equation (6) allowing for several $\Omega_{\tau-h}$ variables with the addition of the listed control variable, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is one of the two variables that measure the improvement of market expectations during the Twitter time window. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

In this analysis, we would expect to see the opposite sign (market expectations should move closer to the actual outcome, hence the effect on “improvement” should be positive if disagreement is large. Starting with the spot-based improvement measure, we find this to be the case for disagreement about GDP growth, whereas the effect related to disagreement about inflation is positive, but not statistically significant. Instead, when we measure disagreement by means of the interquartile range, i.e. focus more on information closer to the centre of the distribution, the coefficient on inflation disagreement turns significant, whereas output growth disagreement turns insignificant. When mixing the two concepts accordingly, i.e. measuring disagreement about output growth by means of the interdecile range and disagreement about inflation through the interquartile range, both variables show the expected effect - more Twitter traffic is associated with financial market participants improving their expectations when disagreement is large (see column (4)). For the forward-based measure, we find equivalent results.

6. Conclusion

This paper studies how financial market participants and central bank watchers update their expectations about upcoming monetary policy decisions in the absence of – otherwise dominant – central bank signals. Given the importance of monetary policy decisions for the economy overall, financial market participants have a strong desire to anticipate upcoming decisions as precisely as possible. In the absence of central bank signals, they are left to sharing their own views about the economy and the conduct of monetary policy by the central bank. The question we ask is to what extent such a sharing of information is beneficial and helps them anticipate the upcoming policy decisions.

Using Twitter traffic as a proxy for information exchange among financial market participants and central bank watchers, we find that conditional on large disagreement about the economic outlook, higher Twitter traffic is associated with lower monetary policy surprises. The results are supportive of the hypothesis that financial market participants stand to benefit from sharing their private information about economic fundamentals. By doing so, all agents can form their expectations about future monetary policy based on a larger information set, and therefore come to more accurate expectations on average.

The paper has implications for central bank communication. It is well known that signals sent by the central bank receive a lot of attention in financial markets. Investors might even give too much weight to such signals (Morris and Shin 2002, Svensson 2006), up to a point that the central bank might lose a valuable source of information, namely independent pricing signals stemming from market participants' assessment of the economic situation (Morris and Shin 2018). Given the large and increasing number of communications that central banks provide throughout the intermeeting period (Blinder et al. 2017), this begs the question how information acquisition proceeds in the absence of such communications. The evidence shown in this paper suggests that information exchange among private agents can help the formation of policy expectations, also when the central bank is in its quiet period. Increased information exchange might therefore serve as a partial substitute for the processing of signals sent by the central bank directly, suggesting that pausing the information flow from central banks to markets does not pose any immediate concerns.

References

- Abel, J., R. Rich, J. Song and J. Tracy (2016). The Measurement and Behavior of Uncertainty: Evidence from the ECB Survey of Professional Forecasters. *Journal of Applied Econometrics* 31, 533-550.
- Altavilla, C., L. Brugnolini, R. Gürkaynak, R. Motto and G. Ragusa (2019). Measuring euro area monetary policy, *Journal of Monetary Economics* 108, 162-179.
- Andersen, T., T. Bollerslev, F. Diebold and C. Vega (2003). Micro effects of macro announcements: real-time price discovery in foreign exchange. *American Economic Review* 93(1), 38-62.
- Anderson, E., E. Ghysels and J. Juergens (2005). Do heterogeneous beliefs matter for asset pricing? *Review of Financial Studies* 18, 875-924.
- Anderson, E., E. Ghysels and J. Juergens (2009). The impact of risk and uncertainty on expected returns. *Journal of Financial Economics* 94, 233-263.
- Andrade, P. and H. Le Bihan (2013). Inattentive professional forecasters. *Journal of Monetary Economics* 60, 967-982.
- Andrade, P., R. Crump, S. Eusepi, and E. Moench (2016). Fundamental disagreement. *Journal of Monetary Economics* 83, 106-128.
- Azar, P. and A. Lo (2016). The wisdom of Twitter crowds: Predicting stock market reactions to FOMC meetings via Twitter feeds. *Journal of Portfolio Management* 42(5), 123-134.
- Benamar, H., T. Foucault and C. Vega (2021). Demand for Information, Uncertainty, and the Response of U.S. Treasury Securities to News. *Review of Financial Studies* 34, 3403-3455.
- Baker, S., N. Bloom and S. Davis (2016). Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics* 131 1593-1636.
- Baker, S., N. Bloom, S. Davis and T. Renault (2021). Twitter-Derived Measures of Economic Uncertainty. *Manuscript*.
- Bekaert, G. and M. Hoerova (2014). The VIX, the variance premium and stock market volatility. *Journal of Econometrics* 183, 181-192.

- Bekaert, G., M. Hoerova and N.R. Xu (2021). Risk, Monetary Policy and Asset Prices in a Global World. *Manuscript*.
- Bianchi, F., R. Gomez-Cram, T. Kind and H. Kung (2022). Threats to Central Bank Independence: High-Frequency Identification with Twitter, *Manuscript*.
- Blinder, A., M. Ehrmann, J. de Haan and D.-J. Jansen (2017). Necessity as the Mother of Invention: Monetary Policy after the Crisis. *Economic Policy* 32, 707-755.
- Blinder, A., M. Ehrmann, M. Fratzscher, J. De Haan and D.-J. Jansen (2008). Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence. *Journal of Economic Literature* 46(4), 910-945.
- Blinder, A. (2018). Through a Crystal Ball Darkly: The Future of Monetary Policy Communication. *American Economic Association Papers and Proceedings* 108, 567-571.
- Boguth, O., V. Gregoire and C. Martineau (2019). Shaping expectations and coordinating attention: the unintended consequences of FOMC press conferences. *Journal of Financial and Quantitative Analysis* 54(6), 2327-2353
- Boero, G., J. Smith and K. Wallis (2008). Uncertainty and disagreement in economic prediction: the Bank of England survey of external forecasters. *Economic Journal* 118, 1107-1127.
- Byrne, D., R. Goodhead, M. McMahon and C. Parle (2021). The Central Bank Crystal Ball: Temporal information in monetary policy communication. Mimeo, Oxford University.
- Camous, A. and D. Matveev (2021). Furor over the Fed: A President's Tweets and Central Bank Independence, *CESifo Economic Studies* 67(1), 106-127.
- Campbell, J., C. Evans, J. Fisher and A. Justiniano (2012). Macroeconomic Effects of Federal Reserve Forward Guidance, *Brookings Papers on Economic Activity* 43(1), 1-80.
- Capistrán, C. and A. Timmermann (2009). Disagreement and biases in inflation expectations. *Journal of Money, Credit and Banking* 41(2-3), 365-396.
- Carlin, B., F. Longstaff and K. Matoba (2014). Disagreement and asset prices. *Journal of Financial Economics* 114, 226-238.
- Cieslak, A. (2018). Short-rate expectations and unexpected returns in Treasury bonds. *Review of Financial Studies* 31(9), 3265-3306.
- Cieslak, A., A. Morse and A. Vissing-Jorgensen (2019). Stock returns over the FOMC cycle. *Journal of Finance* 74(5), 2201-2248.
- Clements, M. (2022). Forecaster Efficiency, Accuracy, and Disagreement: Evidence Using Individual-Level Survey Data. *Journal of Money, Credit and Banking* 54(2-3), 537-568.
- Correa, R., K. Garud, J. Londono and N. Mislav (2021). Sentiment in central banks' financial stability reports. *Review of Finance* 25(1), 85-120.
- Dovern, J., Fritsche, U., and Slacalek, J. (2012). Disagreement among Forecasters in G7 Countries, *The Review of Economics and Statistics* 94(4), 1081-1096.
- Ehrmann, M. and M. Fratzscher (2009a). Purdah: On the rationale for central bank silence around policy meetings. *Journal of Money, Credit and Banking* 41 (2/3), 517-528.
- Ehrmann, M. and M. Fratzscher (2009b). Explaining Monetary Policy in Press Conferences. *International Journal of Central Banking* 5, 41-84.
- Ehrmann, M., G. Gaballo, P. Hoffmann and G. Strasser (2019). Can More Public Information Raise Uncertainty? The International Evidence on Forward Guidance. *Journal of Monetary Economics* 108, 93-112.
- Ehrmann, M. and D. Sondermann (2012). The News Content of Macroeconomic Announcements: What if Central Bank Communication Becomes Stale? *International Journal of Central Banking* 8, 1-53.
- Ehrmann, M. and A. Wabitsch (2022). Central Bank Communication with Non-Experts - A Road to Nowhere? *Journal of Monetary Economics* 127, 69-85.
- Giacomini, R., V. Skreta and J. Turen (2020). Heterogeneity, Inattention, and Bayesian Updates. *American Economic Journal: Macroeconomics* 12, 282-309.

- Gilbert, T., C. Scotti, G. Strasser and C. Vega (2017). Is the intrinsic value of a macroeconomic news announcement related to its asset price impact? *Journal of Monetary Economics* 92, 78-95.
- Gorodnichenko, Y., T. Pham and O. Talavera (2021a). The Voice of Monetary Policy. *NBER Working Paper*, No. 28592.
- Gorodnichenko, Y., T. Pham and O. Talavera (2021b). ‘Liked’, ‘Shared’, ‘Commented’: Central Bank Communication on Facebook and Twitter. University of Birmingham Discussion Paper No. 21-05.
- Gürkaynak, R., B. Sack and E. 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, A., A. Macaulay and M. McMahon (2021). The 3 E’s of Central Bank Communication with the Public. In: Ernesto Pastén and Ricardo Reis (eds), *Independence, Credibility, and Communication of Central Banking*, pp. 279–342. Santiago: Banco Central de Chile.
- Hansen, S. and M. McMahon (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics* 99, S114–S133.
- Hanson, S. and J. Stein (2015). Monetary policy and long-term real rates, *Journal of Financial Economics* 115(3), 429-448.
- Henrique, J. (2016). Getoldtweets-python. Available at: <https://github.com/Jefferson-Henrique/GetOldTweets-python>.
- Huang, D., L. Jiangyuan and L. Wang (2021). Are disagreements agreeable? Evidence from information aggregation. *Journal of Financial Economics* 141, 83–101.
- Hubert, P. (2014). FOMC Forecasts as a Focal Point for Private Expectations. *Journal of Money, Credit and Banking* 46(7), 1381-1420.
- Hubert, P. and F. Labondance (2021). The Signaling Effects of Central Bank Tone. *European Economic Review* 133, 103684.
- Issing, O. (2019). The Long Journey of Central Bank Communication. In: Karl Brunner Distinguished Lecture Series, MIT Press.
- Jarocinski, M. and P. Karadi (2020). Deconstructing Monetary Policy Surprises – The Role of Information Shocks. *American Economic Journal: Macroeconomics* 12, 1-43.
- Li, F. (2016). Macro disagreement and the cross-section of stock returns. *Review of Asset Pricing Studies* 6(1), 1-45.
- Lucca, D. and E. Moench (2015). The Pre-FOMC Announcement Drift. *Journal of Finance*, 70(1), 329-371.
- Lucca, D. and F. Trebbi (2009). Measuring central bank communication: an automated approach with application to FOMC statements. *NBER Working Paper*, No. 15367.
- Lüdering, J., and P. Tillmann (2020). Monetary policy on twitter and asset prices: Evidence from computational text analysis. *North American Journal of Economics and Finance* 51, 100875.
- Mankiw, N.G., R. Reis and J. Wolfers (2003). Disagreement on Inflation Expectations, in M. Gertler and K. Rogoff (eds.), *NBER Macroeconomics Annual*, Cambridge, MA: MIT Press, 209–248.
- Masciandaro, D., D. Romelli and G. Rubera (2022). Monetary policy and financial markets: evidence from Twitter traffic. *Bocconi Working Paper*, No. 160.
- Meeks, R. and F. Monti (2022). Heterogeneous beliefs and the Phillips curve. *Manuscript*.
- Morris, S. and H. S. Shin (2002). Social Value of Public Information. *American Economic Review* 92(5), 1521-1534.
- Morris, S. and H. S. Shin (2018). Central Bank Forward Guidance and the Signal Value of Market Prices. *AEA Papers and Proceedings* 108, 572-577.
- Patton, A. and A. Timmermann (2010). Why do forecasters disagree? Lessons from the term structure of cross-section dispersion. *Journal of Monetary Economics* 57, 803–820.

- Rich, R. and J. Tracy (2010). The Relationships among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts. *Review of Economics and Statistics* 92, 200–207.
- Rieder, K. and P. Gnan (2021). Who Talks During Monetary Policy Quiet Periods, and Why? Evidence from the European Central Bank's Governing Council. *CEPR Discussion Paper*, No. 15735.
- Rosa, C. (2011). Words that shake traders: The stock market's reaction to central bank communication in real time. *Journal of Empirical Finance* 18 (5), 915-934.
- Scotti, C. (2016). Surprise and uncertainty indexes: Real-time aggregation of real-activity macro-surprises. *Journal of Monetary Economics* 82, 1–19.
- Stiefel, M. and R. Vivès (2022). 'Whatever it takes' to change belief: evidence from Twitter. *Review of World Economics* 158, 715-747.
- Svensson, L.E.O. (2006). Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con. *American Economic Review* 96, 448–452.
- Swanson, E. (2006). Have Increases in Federal Reserve Transparency Improved Private Sector Interest Rate Forecasts? *Journal of Money, Credit, and Banking* 38, 791-819.
- Tillmann, P. (2022), Macroeconomic Surprises and the Demand for Information about Monetary Policy, *International Journal of Central Banking*, forthcoming.

Appendix

Table A - Descriptive statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
<i>Tweets after press conference</i>					
Tweets _t	79	8906	5792	2994	32698
Tweets _{t+1}	78	2919	2637	749	14288
Tweets _{t+2}	78	667	598	161	3937
Tweets _{t+3}	78	576	536	127	3289
Tweets _{3d}	78	1387	1224	390	7166
Retweets _{3d}	78	743	775	63	4648
All tweets _{3d}	78	2130	1748	453	10576
h-stat _{3d}	78	0.61	0.55	0.09	4.67
<i>Tweets before press conference</i>					
Tweets _{t-1}	79	1765	1613	460	10379
Tweets _{t-2}	79	1287	1220	374	7670
Tweets _{t-3}	79	1254	1093	213	7174
Tweets _{-3d}	79	1435	1245	483	7388
Retweets _{-3d}	79	790	633	129	3610
All tweets _{-3d}	79	2226	1637	619	10280
h-stat _{-3d}	79	0.34	0.15	0.09	0.92
TwAcc _{-3d}	79	863	565	348	3308
Tweets _{QP}	79	1091	909	412	6280
<i>Dependent variables</i>					
MPS _{2y}	79	1.91	2.11	0.03	11.57
Improv _{spot,2y}	79	0.38	1.63	-5.35	5.05
Improv _{for,2y}	79	0.89	2.60	-7.73	12.43
<i>Proxies for increased attention</i>					
Tweets _{pre QP}	79	1103	846	324	5135
Surprise _{Scotti}	79	0.01	0.01	0.00	0.06
Surprise _{Citi}	79	7.08	5.55	0.10	37.60
Δ OIS _{pre QP}	79	4.47	3.94	0.10	16.20
OIS _{for} - OIS _{spot}	79	0.07	0.06	0.00	0.26
<i>Proxies for instances where an exchange of information is particularly valuable</i>					
Google trends	79	52.94	15.63	30.21	143.67
HICP _{IDR}	79	0.47	0.10	0.23	0.76
HICP _{IQR}	79	0.24	0.06	0.10	0.43
GDP _{IDR}	79	0.66	0.25	0.21	1.80
GDP _{IQR}	79	0.32	0.14	0.12	1.30
3-month rates _{IDR}	79	0.15	0.12	0.00	0.60
3-month rates _{IQR}	79	0.07	0.06	0.00	0.25
Uncert _{Econ}	79	130.69	101.46	29.87	595.57
Uncert _{vSTOXX}	79	20.20	6.53	11.77	53.89
Uncert _{Ec.Policy}	79	201.08	53.96	111.80	424.38

(to be continued)

Table A (continued) - Descriptive statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
<i>Alternative proxies for information exchange</i>					
Bloomberg	74	833	439	274	2488
Newswires	79	26	22	3	141
Newspapers	79	44	27	17	158
<i>Differentiating tweets</i>					
Experts' tweets	79	311	236	73	1518
Non-experts' tweets	79	199	180	54	930
Mon pol tweets	79	364	285	103	1477
Econ tweets	79	200	166	25	996
Retweeted experts' tweets	79	74	52	15	371
Not retweeted experts' tweets	79	237	187	59	1147
Long tweets	79	703	553	202	3247
Short tweets	79	733	710	158	4141
<i>Decomposing 2-year yields monetary policy surprises</i>					
$ MPS_{2y} _{\text{Press release}}$	79	0.85	1.46	0.00	6.85
$ MPS_{2y} _{\text{Press conference}}$	79	1.38	1.53	0.00	6.29
$ MPS_{2y} _{\text{Introductory statement}}$	79	0.86	1.02	0.00	5.43
$ MPS_{2y} _{\text{Q\&A}}$	79	1.08	1.14	0.00	6.52
<i>Alternative monetary policy surprises</i>					
$ MPS_{3m} $	79	1.08	1.87	0.00	10.35
$ MPS_{6m} $	79	1.14	1.81	0.00	8.61
$ MPS_{1y} $	79	1.44	1.92	0.00	9.88
$ MPS_{5y} $	79	2.45	2.32	0.05	12.96
$ MPS_{\text{targ}} $	79	1.06	1.78	0.02	9.90
$ MPS_{\text{path}} $	79	1.39	1.32	0.02	5.29
$ MP_{JK} $	70	0.00	0.03	-0.07	0.14
$ CBI_{JK} $	70	0.00	0.02	-0.10	0.06

Note: the table shows summary statistics of the various variables.

Table B - Decomposing the segments of the ECB monetary surprises

$ \text{MPS}_{2y} _x$	Overall	Press release	Press conference	Introductory statement	Q&A
	(1)	(2)	(3)	(4)	(5)
Tweets _{3d}	7.566*** (3.39)	4.022** (2.50)	3.995*** (2.66)	0.514 (0.41)	3.417*** (3.63)
HICP _{IDR}	37.959** (2.08)	29.848** (2.56)	10.658 (0.74)	-14.531 (-1.24)	17.213* (1.91)
GDP _{IDR}	33.491*** (2.66)	6.499 (0.79)	27.346*** (3.39)	14.355** (2.41)	18.434*** (4.12)
Tweets * HICP _{IDR}	-6.106** (-2.31)	-4.620** (-2.57)	-1.838 (-0.90)	1.907 (1.16)	-2.583** (-2.03)
Tweets * GDP _{IDR}	-4.584** (-2.53)	-1.009 (-0.84)	-3.707*** (-3.30)	-1.850** (-2.21)	-2.496*** (-3.94)
$ \text{OIS}_{\text{for}} - \text{OIS}_{\text{spot}} $	0.101** (2.59)	0.038 (1.47)	0.070** (2.17)	0.061*** (2.79)	0.047 (1.62)
Constant	-50.368*** (-3.34)	-26.027** (-2.47)	-26.898** (-2.63)	-3.552 (-0.41)	-23.348*** (-3.63)
Obs.	79	79	79	79	79
R ²	0.43	0.28	0.28	0.28	0.26

Note: the table shows OLS coefficients of Equation (6) allowing for several Ω_t variables with the addition of the listed control variable, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. Column (1) reports results for the benchmark specification, measuring the monetary policy surprise over the entire event window capturing the press release at 13:45 and the press conference beginning at 14:30. Columns (2) and (3) differentiate the time window into the press release and the press conference. Columns (4) and (5) differentiate the press conference time window into the reading of the prepared introductory statement and the Q&A session. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

In Table B, we vary the timing over which the monetary policy surprise is measured. The ECB's communication on Governing Council meeting days lends itself nicely to a disaggregated analysis, as there is a very clear time sequence. At 13:45 CET, a press release is issued that enumerates the decisions taken, without providing any explanation. The background explanation is then provided in the press conference, which starts at 14:30 on the same day. At the beginning, the ECB president reads an introductory statement,²⁹ which is then followed by a question and answer session with the media. The introductory statement contains the economic rationale for the decisions and discusses the current economic situation and the economic outlook, for instance by releasing some results of the quarterly macroeconomic projections prepared by the ECB and Eurosystem staff. The surprise measure considered so far takes the entire time window covering the press release and the entire press conference. The database provided by Altavilla et al. (2019) contains surprise measures that separate the press release from the press conference window. In addition, we further disentangle the press conference into the introductory statement and the Q&A part.³⁰

²⁹ As of July 2021, i.e. after the end of our sample period, the "introductory statement" is called the "monetary policy statement". While the structure and the language used has changed considerably at that point in time, the content of the introductory statement has been stable during our time sample.

³⁰ Altavilla et al. (2019) define the surprise measure for the press release, released at 13:45, as the difference between the median price in the 13:25–13:35 interval and the median price in the 14:00–14:10 interval. The same is done for the press conference, using the intervals 14:15–14:25 and 15:40–15:50 before and after the press conference. Based on Ehrmann and Fratzscher (2009b), who show that the introductory statement takes on average 12 minutes to read, plus allowing for a possible slight delay in the start of the press conference, we define the end of the introductory statement by using the median price in the 14:42–14:45 interval.

We find that disagreement about inflation tends to be associated with larger press release surprises, and that higher Twitter exchanges moderate this effect, whereas this is not the case for disagreement about GDP growth. This is intuitive, as the ECB does not have (in contrast to the Fed, for instance) a dual mandate; its single mandate is to maintain price stability. In that sense, the outlook for inflation is bound to be at the core of uncertainty about the actual decisions. In contrast, uncertainty about the real economy, or disagreement about economic forecasts, will likely matter more for the economic outlook and therefore for future monetary policy decisions. In that sense, they are more likely associated with surprises in the ECB's communication during the press conference. Indeed, we find that disagreement about GDP growth is associated with the surprise component in the press conference, i.e. in the parts where the economic outlook is discussed, and that higher Twitter exchanges moderate this effect.

Table C – Differentiating the horizons of the ECB monetary surprises

$ \text{MPS}_x $	$ \text{MPS}_{3m} $	$ \text{MPS}_{6m} $	$ \text{MPS}_{1y} $	$ \text{MPS}_{2y} $	$ \text{MPS}_{5y} $	$ \text{MPS}_{\text{targ}} $	$ \text{MPS}_{\text{path}} $	$ \text{MP}_{JK} $	$ \text{CBI}_{JK} $
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Tweets _{-3d}	4.411** (2.32)	5.478*** (2.83)	6.572*** (3.27)	7.566*** (3.39)	5.020 (1.65)	4.043** (2.20)	4.177*** (3.93)	0.050 (0.99)	0.010 (0.24)
HICP _{IDR}	33.165** (2.37)	34.272** (2.58)	41.517*** (2.96)	37.959** (2.08)	6.450 (0.25)	29.395** (2.19)	16.012 (1.31)	0.714* (1.69)	-0.284 (-0.78)
GDP _{IDR}	6.129 (0.52)	15.162 (1.33)	22.226* (1.98)	33.491*** (2.66)	32.550** (2.19)	5.912 (0.51)	24.076*** (3.66)	-0.043 (-0.19)	0.302 (1.29)
Tweets * HICP _{IDR}	-5.027** (-2.37)	-5.364*** (-2.65)	-6.347*** (-3.05)	-6.106** (-2.31)	-1.793 (-0.47)	-4.485** (-2.20)	-2.719 (-1.59)	-0.102* (-1.67)	0.039 (0.75)
Tweets * GDP _{IDR}	-0.928 (-0.53)	-2.129 (-1.27)	-3.092* (-1.88)	-4.584** (-2.53)	-4.428** (-2.08)	-0.893 (-0.52)	-3.302*** (-3.61)	0.007 (0.21)	-0.044 (-1.30)
$ \text{OIS}_{\text{for}} - \text{OIS}_{\text{spot}} $	0.066** (2.09)	0.066** (2.10)	0.104*** (3.11)	0.101** (2.59)	0.081 (1.67)	0.064** (2.17)	0.062** (2.43)	0.001 (0.91)	-0.000 (-0.88)
Constant	-29.089** (-2.34)	-36.293*** (-2.84)	-44.276*** (-3.29)	-50.368*** (-3.34)	-31.382 (-1.52)	-26.550** (-2.20)	-27.441*** (-3.80)	-0.356 (-1.03)	-0.059 (-0.20)
Obs.	79	79	79	79	79	79	79	70	70
R ²	0.25	0.34	0.41	0.43	0.25	0.25	0.35	0.12	0.14

Note: the table shows OLS coefficients of Equation (6) allowing for several $\Omega_{\tau-h}$ variables with the addition of the listed control variable, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in OIS rates around policy announcements. Columns (1) to (5) define the monetary policy surprise via the change in the 3-month, 6-month, 1-year, 2-year and 5-year OIS. Columns (6) and (7) differentiate the target and path factors of monetary surprises measured the fitted values and the residuals, respectively, of the OLS regression of the intraday change in 2-year OIS on the intraday change in 3-month OIS. Columns (8) and (9) report results for the monetary policy and the central bank information shocks identified by Jarocinski and Karadi (2020). The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

In our next exercise, we vary the maturity over which the monetary policy surprises are defined, ranging from 3-month to 5-year OIS rates. Movements at the short end of the maturity spectrum are more closely related to the current stance of monetary policy, whereas movements at longer horizons reflect the outlook for monetary policy. Another way to get at this distinction is by separating so-called “target factor” surprises from the “path factor” surprises – following the language of Gürkaynak et al. (2005), with the latter corresponding more to movements at the short end of the yield curve, and the latter to movements at the longer end. The path factor of ECB monetary policy surprises is measured as the residuals of the OLS regression of the intraday change in 2-year OIS rates on the intraday change in 3-month OIS rates, whereas the fitted values of that regression corresponds to the target factor.

Table C reports the corresponding results, and shows that Twitter exchange is associated with a moderation of the surprises at the short end of the maturity spectrum that are related to disagreement about inflation, whereas the Twitter effect regarding the surprises at the long end of the maturity spectrum is more related to disagreement about the real economic outlook and uncertainty about the current economic situation. These results are well aligned with those reported in Table D – the path surprise is related to the communication about the outlook during the press conference and less so to the release of current decisions.

Albeit not statistically significant, we find related results when using the monetary policy and the central bank information shocks identified by Jarocinski and Karadi (2020). The interaction term is negative for inflation disagreement and the monetary policy shock, as well as for GDP growth disagreement and the central bank information shock.

Table D – Twitter response to information sharing motives

Tweets _{s-3d}	Info demand	Disagreement			Uncertainty		
	Google trends (1)	HICP _{IDR} (2)	GDP _{IDR} (3)	3-month rates _{IDR} (4)	Uncert _{Econ} (5)	Uncert _{VSTOXX} (6)	Uncert _{Ec.Policy} (7)
Proxy	0.027*** (6.34)	1.359** (2.03)	0.462* (1.80)	0.966* (1.97)	-0.001 (-1.09)	0.004* (1.68)	-0.001 (-0.90)
Constant	5.654*** (26.95)	6.427*** (20.63)	6.763*** (39.04)	6.921*** (75.75)	7.145*** (70.85)	6.964*** (92.61)	7.252*** (31.61)
Obs.	79	79	79	79	79	79	79
R ²	0.52	0.06	0.04	0.04	0.01	0.03	0.01

Note: the table shows OLS coefficient estimates for the relation between Twitter traffic pre-Governing Council meetings and the various proxies for instances where an exchange of information is particularly valuable, estimated as $T_{t-h} = \alpha + \beta \Omega_{t-h} + \varepsilon_{t-h}$. The dependent variable is the daily average of tweets over the 3 preceding days for each Governing Council meeting. The X explanatory variable covers ECB-related Google trend measure (column 1), policy expectations (column 2), disagreement (columns 3 to 5), and uncertainty (columns 6 to 9). The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

Table E – Predicted OIS changes and ECB monetary surprises; controlling for the correlation between Twitter traffic and disagreement

	First stage		Second stage		
	(1)	(2)	(3)	(4)	
	Tweets _{-3d}	ΔOIS_{QP}	Tweets _{-3d, residual}	$\Delta OIS_{QP, predicted}$	
			$ MPS_{2y} $	$ MPS_{2y} $	
HICP _{IDR}	1.272*		Variable (column header)	7.547***	19.508***
	(1.80)			(3.05)	(3.39)
GDP _{IDR}	0.414*		HICP _{IDR}	-3.203**	16.974*
	(1.67)			(-2.40)	(1.84)
Tweets _{-3d}		0.388**	GDP _{IDR}	1.261	17.735***
		(2.03)		(1.42)	(2.78)
			Variable · HICP _{IDR}	-5.895*	-15.741**
				(-1.97)	(-2.31)
			Variable · GDP _{IDR}	-4.645**	-11.819**
				(-2.36)	(-2.53)
			$ OIS_{for} - OIS_{spot} $	0.110***	0.101**
				(2.85)	(2.59)
Constant	6.197***	-1.333	Constant	1.775*	-24.362***
	(18.69)	(-0.99)		(1.89)	(-3.29)
Obs.	79	79	Obs.	79	79
R ²	0.09	0.04	R ²	0.42	0.43

Note: Column (1) relates Twitter traffic to disagreement, in order to obtain the residual of regression $T_{-3d} = \alpha + \beta \Omega_{t-h} + \varepsilon_{-3d}$, where Ω_{t-h} contains disagreement about both, inflation and output growth. Column (2) relates the absolute movements in 2-year OIS rates over the Twitter time window to Twitter traffic during that window, in order to obtain the predicted value of regression $|\Delta OIS_{2y, QP}| = \alpha + \beta T_{-3d} + \varepsilon_{QP}$. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level. Columns (3) and (4) then shows OLS coefficients of Equation (6) using the unexplained Twitter traffic after controlling for disagreement (from column 1) and using the predicted value of the absolute movements in 2-year OIS rates over the Twitter time window by Twitter traffic during that window, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in OIS rates around policy announcements.

Table F - Alternative Twitter metrics

$ \text{MPS}_{2y} $	Tweets _{s,3d}	Tweets _{t-1}	Tweets _{t-2}	Tweets _{t-3}	Retweets _{s,3d}	Tweets _{QP}
	(1)	(2)	(3)	(4)	(5)	(6)
Tweets _x	7.566*** (3.39)	6.925*** (3.83)	5.779** (2.52)	7.021*** (3.13)	4.267** (2.46)	8.996*** (3.31)
HICP _{IDR}	37.959** (2.08)	39.171** (2.39)	21.424 (1.10)	28.971* (1.95)	16.841 (1.35)	42.492** (2.07)
GDP _{IDR}	33.491*** (2.66)	29.077** (2.89)	28.465** (2.38)	35.437*** (2.78)	22.181** (2.27)	41.029*** (2.97)
Tweets · HICP _{IDR}	-6.106** (-2.31)	-6.107*** (-2.68)	-3.779 (-1.31)	-4.898** (-2.22)	-3.109 (-1.60)	-7.028** (-2.27)
Tweets · GDP _{IDR}	-4.584** (-2.53)	-3.839*** (-2.75)	-3.954** (-2.22)	-4.938** (-2.62)	-3.213** (-2.07)	-5.900*** (-2.86)
$ \text{OIS}_{\text{for}} - \text{OIS}_{\text{spot}} $	0.101** (2.59)	0.090** (2.39)	0.113** (2.59)	0.110*** (2.71)	0.124** (2.60)	0.095** (2.42)
Constant	-50.368*** (-3.34)	-47.170*** (-3.77)	-37.259** (-2.47)	-45.840*** (-3.06)	-26.187** (-2.39)	-57.771*** (-3.27)
Obs.	79	79	79	79	79	79
R ²	0.43	0.42	0.37	0.41	0.29	0.42

Note: the table shows OLS coefficients of Equation (6) allowing for several Ω_{t-h} variables with the addition of the listed control variable, estimated over the 79 ECB meetings from January 2012 to April 2020. The dependent variable is the absolute monetary policy surprise measured by the intraday change in the 2-year OIS rates around policy announcements. Column(1) replicates the results for all tweets, all remaining columns include different tweet measures. The regression allows for robust standard errors. Numbers in brackets are t-statistics. ***/**/* denote statistical significance at the 1%/5%/10% level.

In Table F, we check for the robustness of our benchmark result by varying the time window over which we measure Twitter traffic, and by varying the measure of Twitter traffic itself. As before, column (1) replicates the benchmark result (as shown in column (2) of Table 5). The subsequent columns differentiate Twitter traffic by day. Our results are qualitatively robust on each of the three days under analysis here, with only slight nuances. Consistent with the initial result shown in Table 2, the effect is predominantly driven by original tweets. Finally, as the quiet period lasts for a full week, we measure Twitter traffic over the entire quiet period and again find similar results.

Appendix G: Newswire report during an ECB quiet period

RTRS - Rising Italy bond yields to test an ECB set to pull rate hike trigger again
02-Sep-2022, 02:50:52 PM

LONDON, Sept 2 (Reuters) - Rising borrowing costs in highly-indebted Italy are again testing the European Central Bank's resolve to contain bond market strain. Just days before the ECB is tipped to deliver a second big interest rate hike to curb record-high inflation, worries about a more aggressive move have unnerved investors. Italy's 10-year bond yield on Thursday topped 4% for the first time since mid-June, when a sharp move above that level pushed the closely-watched spread to German debt to around 250 bps and prompted the ECB to hold an emergency meeting to discuss how to contain bond stress as it withdraws stimulus. This level is generally viewed as one where concern about Italy's ability to service its debt sets in. At around 150% of gross domestic product (GDP), Italy has the second-highest debt to-GDP ratio in the euro area.

Societe Generale believes that Italian yields are entering a "danger zone" with 10-year borrowing costs now above its estimate of the level at which the debt ratio would remain stable. Alert to the dangers of tightening policy against a backdrop of sharp rises in borrowing costs, the ECB unveiled its Transmission Protection Instrument (TPI) in July. It is a new bond purchase scheme to help more indebted euro zone states and prevent a divergence of borrowing costs from benchmark issuer Germany it sees as happening through no fault of their own.

"Everyone in the market knows that 4% is tricky for debt sustainability in Italy and the growth outlook has deteriorated" said Pictet Wealth Management fixed income strategist Laureline Renaud-Chatelain. Yet, analysts suspect the new tool was unlikely to be used soon - especially as a snap Italian election looms on Sept. 25. "The yield level is going to be a problem. But I don't think the ECB will activate the new tool before an election" Renaud-Chatelain said. Renewed Italian political instability has contributed to the bond selloff, while the recent yield surge is in line with peers. Italian and German 10-year yields rose around 70 bps each in August as fears about higher inflation and rates took hold. The risk premium over Germany, at around 235 bps, has widened but is below recent peaks, supported perhaps by the ECB skewing reinvestments from maturing bonds it bought for its pandemic purchase scheme at Italy.

"The spread remains orderly and more than the level it is the behavior of the spread (and more generally the periphery) that the ECB would be concerned about" said Peter Schaffrik, global macro strategist at RBC Capital Markets.


PAIN THRESHOLD?

Still, markets were expected to keep pushing yields higher to test just where the ECB's tolerance level for pain in Italian bond markets lies. In the past, analysts have viewed the 250-300 bps area in the spread as a danger zone for the ECB and some analysts expect the spread to reach this area in coming months. UBS, for instance, reckons the spread could test 300 bps. Italian borrowing costs meanwhile climbed to new multi-year highs at a Tuesday auction. "No one knows when the ECB will start intervening and they obviously won't tell us" said Mike Riddell, senior fixed income portfolio manager at Allianz Global Investors. "My assumption since the announcements of potential support for the periphery and namely Italy, is that markets will test (the ECB) given the trajectory for economic growth and rates."

Battered by soaring energy prices, many economists expect the euro zone economy to slip into a recession - a challenging backdrop for the ECB as inflation nears double digits. On the plus side, Italy's election noise so far hasn't alarmed investors. The Italian rightist alliance's ambitious spending plans will respect European Union budget rules and not blow a hole in Italy's finances, according to Giorgia Meloni, who heads the Brothers of Italy party topping the polls.


“The worrying thing for the market was that you'd have Italy's yields rising and a considerably worse growth outlook after (former prime minister Mario) Draghi left,” said Mizuho rates strategist Peter McCallum. “It now seems like politics isn't going to be so much of a shock as some people have been fearing.”

Figure H - Interaction of Twitter traffic and inflation disagreement



Note: the figure shows the coefficient estimates for the interaction of Twitter traffic and inflation disagreement as estimated in equation (6). Point estimates are provided as orange dots, along with ± 1 and 2 standard deviation confidence bands. Days -12, -11, -5 and -4 are weekends.

Figure I - Interaction of Twitter traffic and GDP growth disagreement



Note: the figure shows the coefficient estimates for the interaction of Twitter traffic and GDP growth disagreement as estimated in equation (6). Point estimates are provided as orange dots, along with ± 1 and 2 standard deviation confidence bands. Days -12, -11, -5 and -4 are weekends.

Acknowledgements

The authors would like to thank Christophe Blot, Erwan Gautier, Fabien Labondance, Michael McMahon, Sarah Mouabbi, Alena Wabitsch, and participants at the conference “New Challenges for Monetary Policy” and at seminars at the Banque de France, ECB, Norges Bank, HEC Paris and St. Gallen University for comments, and Giada Bozzelli, Catherine Le Grand and Justus Meyer for their help with the data collection.

This paper presents the authors’ personal opinions and does not necessarily reflect the views of Banque de France, the European Central Bank or the Eurosystem.

Michael Ehrmann

European Central Bank, Frankfurt am Main, Germany; Centre for Economic Policy Research, London, United Kingdom;
email: michael.ehrmann@ecb.europa.eu

Paul Hubert

Banque de France, Paris, France; Sciences Po – OFCE, Paris, France; email: paul.hubert@sciencespo.fr

© European Central Bank, 2023

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB’s website](#).

PDF

ISBN 978-92-899-5512-6

ISSN 1725-2806

doi:10.2866/45657

QB-AR-23-007-EN-N