DISCUSSION PAPER SERIES

DP19811

IN THE FED WE TRUST? MEASURING TRUST IN CENTRAL BANKING AND ITS EFFECTS ON THE MACROECONOMY

David Aikman, Francesca Monti and Shunshun Zhang

MONETARY ECONOMICS AND FLUCTUATIONS



ISSN 0265-8003

IN THE FED WE TRUST? MEASURING TRUST IN CENTRAL BANKING AND ITS EFFECTS ON THE MACROECONOMY

David Aikman, Francesca Monti and Shunshun Zhang

Discussion Paper DP19811
Published 23 December 2024
Submitted 22 December 2024

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
187 boulevard Saint-Germain, 75007 Paris, France
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: David Aikman, Francesca Monti and Shunshun Zhang

IN THE FED WE TRUST? MEASURING TRUST IN CENTRAL BANKING AND ITS EFFECTS ON THE MACROECONOMY

Abstract

We develop a novel measure of trust in the Federal Reserve using Generative Artificial Intelligence to analyse millions of tweets about the Fed, its leadership and its policy framework and decisions. Our measure reacts in an intuitive way to various macro-financial variables and indicators of U.S. monetary policy. To study the effects of trust shocks, we use a narrative identification approach based on ethical scandals embroiling some FOMC members, and we study the effects of these shocks using a daily VAR. We find that trust shocks have highly persistent effects on macroeconomic variables despite having short-lived effects on our trust measure: they weaken business conditions, the stock market and news sentiment, while increasing the VIX index. Inflation expectations also increase following a trust shock, worsening the inflation-output trade-off.

JEL Classification: D83, D84, E52, E58

Keywords: Trust, Narrative identification, Central banks

David Aikman - david.aikman@kcl.ac.uk
King's Business School

Francesca Monti - francesca.monti@uclouvain.be
UCLouvain and CEPR

Shunshun Zhang - shunshun.zhang@kcl.ac.uk
King's Business School

In the Fed we trust? Measuring trust in central banking and its effects on the macroeconomy

David Aikman* Francesca Monti† Shunshun Zhang‡

December 23, 2024

Abstract

We develop a novel measure of trust in the Federal Reserve using Generative Artificial Intelligence to analyse millions of tweets about the Fed, its leadership and its policy framework and decisions. Our measure reacts in an intuitive way to various macro-financial variables and indicators of U.S. monetary policy. To study the effects of trust shocks, we use a narrative identification approach based on ethical scandals embroiling some FOMC members, and we study the effects of these shocks using a daily VAR. We find that trust shocks have highly persistent effects on macroeconomic variables despite having short-lived effects on our trust measure: they weaken business conditions, the stock market and news sentiment, while increasing the VIX index. Inflation expectations also increase following a trust shock, worsening the inflation-output trade-off.

*King's Business School. e-mail: david.aikman@kcl.ac.uk

†UC Louvain and CEPR. Email: francesca.monti@uclouvain.be

‡King's Business School. e-mail: shunshun.zhang@kcl.ac.uk

We are grateful to Stephen Cecchetti, Michael Ehrmann, Domenico Giannone, Mikael Juselius, and Michael McMahon for helpful suggestions and discussions. We also thank seminar participants at Ghent University, King's College London, the Bank of England, Banco Central do Brasil, and Université Catholique de Louvain for insightful comments. All errors remain our own. Swarali Deepak Havaldar provided excellent research assistance on this paper.

Why is trust so important? If the public trusts the authorities' actions, they will incorporate them in determining their own behaviour. As a result, it is more likely that the authorities will achieve their objectives. In addition, trust fuels the legitimacy of policies. With trust, the public will be more willing to accept actions that involve short-term costs in exchange for long-term benefits. In sum, trust is vital for policy effectiveness.

Carstens (2023) – Bank for International Settlements

1 Introduction

Public trust in the institution of central banking has undoubtedly been under question of late. In part, this reflects the large unforeseen increase in inflation experienced by many countries since 2022, the first spike in inflation for a generation. But other factors are also likely to have been at play, including: the turmoil experienced in banking and financial systems in late 2022 and early 2023; controversial policies such as Quantitative Easing that are viewed by some as "currency debasement" and others as a "gift to the banking sector"; large losses on many central banks' asset purchase programs as a result of the rise in interest rates; overt criticism of central bank's policy stance from various publicly elected officials, and President Trump in particular; and, last but not least, the fact that a number of prominent senior central bankers have been forced to resign their posts after being embroiled in financial scandals.

In this paper, we propose a new method for measuring trust in central banking. Our method involves collecting data from Twitter/X and using Generative Artificial Intelligence (GAI) to assess whether the tweets reflect criticism or support of the central bank, its policies, and its leadership. Although our method could be applied to any public media, it has been argued that in some respects Twitter/X has replaced conventional media as a public sphere for discussion and debate. We focus our analysis on the US Federal Reserve—as the world's leading central bank, the Fed's actions

attract significant public discourse, and it therefore provides a perfect case study for our method.

What do we mean by trust in the central bank? The economics literature has long emphasised the importance of the related concept of credibility in macroeconomic policymaking. *Credibility* here refers to the likelihood that the central bank will fulfil its stated commitments, e.g., commitments to reduce inflation from an undesirably high level, or to maintain an exchange rate peg. Kydland and Prescott (1977), Barro and Gordon (1983) and others taught us that credibility hinges critically on making commitments that align with the incentives of policymakers. In a repeat-game context, credibility can also be fostered by the central bank's track record of honouring past commitments.

By *trust*, we have in mind a related but broader concept. Trust captures the public's belief in the goodwill or integrity of the central bank, its leadership and its staff. It also captures an expectation that the central bank will act in the public interest rather than prioritising special interest groups. And it encompasses the idea that the central bank is perceived to have the technical competency to meet the goals it has been set—related to, but distinct from, whether those goals are incentive compatible. These factors are summarised as ability, benevolence, and integrity in the management organisation literature (see Mayer et al. (1995)). While credibility can be established through formal mechanisms, trust relies on informal norms and social capital, making it more fragile.

Public trust is critical to the mission of central banking. At a fundamental level, money is a social convention, the value of which depends on trust in the money-issuing institution (e.g., Schnabel (2020)). Trust also underpins the transmission of monetary policy to the economy—Ehrmann et al. (2013) and Christelis et al. (2020) present empirical evidence that higher trust in the central bank translates into a better ability to anchor household inflation expectations around its target. During financial crises, a central bank that has widespread public trust as a technocratic, competent

organisation is likely to find that its actions and communications are viewed as credible and reassuring, increasing their effectiveness. Finally, a central bank that is trusted as having high integrity and whose policies are supported by the public is likely to be less prone to short-term political influence. Indeed as Blinder et al. (2024) note, when explaining the rationale for their communications with the public central banks increasingly refer to public trust. However, despite the importance of trust to central banks' success, quantitative metrics of the public's trust in these institutions are still few and far between.

Most existing trust measures are based on surveys. For example, Brouwer and de Haan (2022) and Bursian and Fürth (2015) use Eurobarometer, a biannual survey used to gauge public opinion throughout the European Union on a variety of topics, including respondents' level of trust in the European Central Bank (ECB). Similar surveys exist in the United States, including an annual Gallup survey that assesses the confidence of the U.S. public in economic leaders including the Chair of the Federal Reserve, and the now-discontinued Chicago Booth Financial Trust Index, which polled U.S. citizens' confidence in financial institutions including the Federal Reserve. Survey-based methods have various shortcomings as metrics for understanding the evolution of trust in central banking. They tend to be relatively infrequent, and the questions tend to be closed-end (e.g., paraphrasing Eurobarometer: do you tend to trust or not to trust the ECB), and as such they offer limited insight into why responses might have changed. Respondents might also interpret the survey question in different ways, with responses capturing views about related institutions or general gut reactions. And there are various well-known problems that affect all surveys including social desirability bias, non-response bias, and the effects of question framing. For a discussion of these problems, see Ehrmann (2024).

For our new measure of public trust, we collect relevant data about the Federal Reserve using the Twitter/X API. Specifically, we download all tweets mentioning the Federal Reserve and its senior leadership that also include unigrams related to

policy, competence, or ethics/governance—these words help us to filter out tweets that solely provide objective commentary about the Fed. In total, this procedure generates around 4 million tweets, providing a rich source of daily information about Twitter/X participants' view the Federal Reserve. Our data cover the period 2010-2023, a period that includes the aftermath of the global financial crisis, public criticism of the Fed by various senior politicians, three Fed Chairs, the Covid-19 crisis and its aftermath, various quantitative easing programmes, a banking crisis, and the jump and decline in inflation.

We then use GAI (ChatGPT) to analyse the information contained in these tweets, categorising them as supportive, critical, neutral or unrelated to the Fed, its leadership or its policy framework and decisions. We also explored other methods for performing this categorisation, including off-the-shelf dictionary-based methods such as VADER, and machine learning algorithms trained on a sample of tweets categorised by hand, but we found the GAI approach to be far superior. It took significant trial and error to produce a prompt that worked well for our purposes. Our public trust index in the Fed is based on the difference between the number of positive and negative tweets (normalised appropriately). The index is quite volatile at daily frequency, but its four-week moving average has a pronounced cyclical pattern.

Which variables explain our trust index? Macro-financial factors such as the yield curve, inflation expectations, and market volatility all have statistically significant effects on our index, with the largest impact coming from the Fed Funds rate—a 1 standard deviation increase in the Fed Funds rate reduces public trust in the Fed by 13-15 percentage points. Monetary policy-related indicators and events also matter, with the semi-annual Monetary Policy Report, the Jackson Hole speech by the Fed Chair, and key dates in recent leadership transitions all entering the regression with large positive coefficients. Events that are unrelated to monetary and financial policy also matter. Notably, news of ethical scandals affecting senior FOMC members has the largest measured (negative) effect on the trust index, with each event reducing

our index by about 20 percentage points. Tweets by President Trump perceived as threatening the Fed's independence (as per Bianchi et al. (2023)) also have large negative impacts. Our index has only a weak relationship with newspaper sentiment.

We then explore the effects of trust shocks on the economy. To do this, we use a narrative identification approach that exploits the news of alleged ethical scandals embroiling FOMC officials, which considerably reduce the trust index. As we argue in Section 3, the release of this news is exogenous to the state of the economy and to the events themselves. We also have an example of a positive shock that increases our trust index, but is unrelated to the state of the economy: the Twitter trend #popyourcollar, launched to celebrate Janet Yellen as a role model for women in STEM a few days ahead of the end of her tenure as Fed Chair. Adopting the insight by Plagborg-Møller and Wolf (2021), we implement these shock sign restrictions by recasting them as a relevant and valid discrete-valued proxy for the structural shock, and then use them to identify impulse responses to that shock in a proxy-Structural Vector Autoregression as discussed in Mertens and Ravn (2013) and Stock and Watson (2018). We also compare the responses to a trust shock identified using scandals (and the #popyourcollar trend) with those obtained assuming that tweets by President Trump threatening the independence of the Fed represent negative shocks to trust, and with a simple Cholesky identification approach that treats all variation in the trust index that cannot be explained by macro-financial factors as a "shock".

We find that shocks that reduce public trust in the Federal Reserve have a persistent contractionary effect on the economy, despite only having a short-lived impact on our trust index. In particular, a negative trust shock reduces the business conditions index, news sentiment, and the stock market, while increasing VIX index. These effects are robust across a range of model specifications. The response of inflation expectations and the Fed Funds rate to trust shocks, however, depends critically on the identification strategy. In our main narrative identification, which focuses purely on news of ethical scandals and the #popyourcollar trend (i.e., the narrative proxy), we find that short-

term inflation expectations and the Fed Funds rate increase following a trust shock. The "supply shock" nature of these responses aligns with the findings of Bursian and Faia (2018), who develop a fully-fledged macroeconomic model featuring trust in the monetary authority (discussed below). If instead we construct the narrative proxy based on tweets by President Trump identified by Bianchi et al. (2023) as threats to central bank independence, or by using a simple Cholesky identification, we find that inflation expectations fall and the effects resemble those of a confidence shock.

While we are the first to use Twitter/X data to assess trust in central banking, other recent papers have made use of such data for related purposes. Ehrmann and Wabitsch (2022) use Twitter data to study how experts and non-experts respond to ECB communication, while Angelico et al. (2022) use a dictionary-based approach to construct a Twitter-based proxy for consumer inflation expectations. Bianchi et al. (2023) examine threats to central bank independence by demonstrating how President Trump's tweets about the Fed had affected monetary policy, the stock market, bond premia, and the macroeconomy.

Our paper also relates to the nascent literature analysing the determinants of and role of public trust in central banking and macroeconomic institutions. Ehrmann (2024) provides a recent comprehensive review of this literature, organised around ten statements covering the reasons why trust matters, its determinants, means to generate and maintain trust, and future challenges to do so. While most of this literature is based on survey methods, two recent papers attempt to integrate a role for trust in the central bank into macroeconomic models. In Bursian and Faia (2018), trust emerges as an equilibrium outcome in a game between betrayal-averse agents and the central bank. When trust declines, consumers increase precautionary savings, reducing demand. But the monetary transmission mechanism is also less effective, increasing the slope of the Phillips curve. As agents perceive a worsening in the inflation-output trade-off, they increase their inflation expectations in anticipation, despite the fall in aggregate demand. De Grauwe and Ji (2024) study how trust impinges on the

transmission of negative demand and supply shocks in the context of a behavioral macroeconomic model. In their setting trust is defined as the degree to which people believe that the central bank can keep inflation close to the announced target and maintain macroeconomic stability. They find that when trust is low and large shocks occur, a bad equilibrium can occur in which trust erodes further and output declines very persistently.

The rest of this paper is organised as follows. In section 2 we describe our Twitter/X dataset and our approach to analysing tweet sentiment. In section 3 we present our trust index and examine its determinants. Section 4 presents a proxy SVAR to study the effects of exogenous shocks to our trust index on the economy. Section 5 concludes with a summary of the paper's findings and its implications for policymakers. Appendix A provides further detail on the approach used to compile and analyse our dataset. Appendix B presents additional robustness results.

2 Data and methodology

In this section, we first describe our approach to locating relevant tweets about the Fed, its leadership and its policy framework and decisions. The detailed queries we use are listed in Appendix A.1. We then describe our use of GAI (ChatGPT) to analyse the sentiment contained in these data, with the detailed prompts described in Appendix A.2.

2.1 Twitter data

We use data from Twitter/X to construct our index of public trust in the Federal Reserve. Our dataset begins on January 1, 2007 and runs until December 31, 2023. Twitter/X has a vast user base, with hundreds of millions of active users, including diverse groups such as the government, the media, academics and the general pub-

lic.[1] This platform therefore provides an important real-time forum to gauge public attitudes to the Federal Reserve.

Given the now high financial costs of downloading historical data from Twitter/X,[2] we implement a series of filtering procedures to ensure the accurate capture of public attitudes toward the Federal Reserve. In total, there are 61,922,464 tweets containing the terms "the fed" or "federal reserve" or the names of key Federal Reserve officials over our sample period. Many of these, however, are cases of homonymy. For instance, "fed" is also used to refer to Roger Federer, FBI agents (referred to as "the feds"), and the federal government (fed govt). To address these issues, we exclude tweets containing the terms "fed gov", "fed govt", "fed government", "federal gov", "federal govt", "federal government", "the fed election", "the feds", "the fed ex cup", and "fed ex". This reduces the tweet count to 14,104,900.

We also encounter cases of homonymy when downloading tweets relating to Fed officials. We focus on the Fed Chairs who served during the sample period—Ben Bernanke, Janet Yellen, and Jerome Powell—along with several other key officials including Donald Kohn, Stanley Fischer, Richard Clarida, Lael Brainard, Philip Jefferson, Mary Daly, Eric Rosengren, Robert Kaplan, Raphael Bostic, James Bullard, and Jeffrey Lacker. To eliminate confusion with celebrities sharing these surnames, we refine our search query by only including tweets that also contain the terms "market", "economy", "business", "chair", "president", or references to specific regional Federal Reserve Banks. We also include in our search query widely used nicknames for these individuals. For example, for Jerome Powell, our query is: "powell AND (fed OR jerome OR jay OR jaypow OR chair OR chairman OR market OR business OR economy)".[3] The number of tweets about Fed officials totals 6,586,843.

[1] According to Statista, in 2023 Twitter/X had 237.8 million daily active users and 541.6 million monthly active users https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/.

[2] At the time of writing, the subscription cost to access the "Pro" Twitter/X API was US$6k per month, permitting up to 1,000,000 downloads of historical tweets.

[3] Table A.1 in Appendix A lists the details of the query for each Fed official.

We refine our search query one step further by focusing on tweets that relate specifically to confidence and trust in the Fed, its officials or its policy framework and decisions; we seek to differentiate these from tweets that provide objective, neutral reporting about the Fed or those that simply share factual information about the institution or its policy decisions. We do this by constructing a lexicon of context-specific synonyms for 'confidence' and 'trust', which we develop by analysing the language used in a sample of several tens of thousands of tweets about Fed officials. We first used ChatGPT to identify from this set of tweets those that expressed criticism of or support for the Fed or its officials as opposed to objective information about upcoming meetings, testimony etc.[4] Next, we computed the term frequency of the identified tweets, focusing on words that either specifically relate to the policy framework and actions of the Fed, are thematically related to ethical and governance issues, or express sentiment or judgement about the competence of the Fed as an institution or its leadership. The resulting lexicon is reported in Table 1.

Bringing this all together, the query we use in the Twitter API returns a tweet if it either (a) contains one element from the set {"the fed", "federal reserve", "thefed", "fedreserve", "federalreserve", "federalreservebank", "usfed", "fomc", "#Fed", "#FOMC"}, or (b) contains the name of a Fed Chair or key Fed official and one element from {"market", "economy", "business", "chair", "president"} or references to specific regional Federal Reserve Banks, and (c) contains one of the words listed in the lexicon in Table 1. This procedure aims to capture the public's subjective views about the Fed and its leadership, while minimising noise in the dataset.

In total, this yields 3,798,214 individual tweets. For each tweet, we record the number of retweets and likes it receives, as these data can be seen as an indicator of additional approval are will be used in constructing our index. Table 2 reports summary statistics for these data. Evidently, tweets about the Fed were relatively rare in the early days of Twitter, and became rather more prevalent from 2010 onward. We

[4]For details, see Appendix A.1.

Table 1: Lexicon of context-specific synonyms for confidence and trust in the Fed

Category	**Words**
Policy	aggressive, authority, bubble, credibility, credible, crisis, decision, dissent, dots, easy, fedspeak, inflation, loose, overheat, policy, qe, regulate, regulation, regulator, remit, soundness, stability, strategy, surprise, taper, tapering, testimony, threat, tight, volatility
Ethics/Governance	accountability, bitcoin, controversy, corrupt, corruption, criminal, disclosure, disruption, ethic, ethical, ethics, gold, governance, independence, influence, insider, integrity, investigation, jail, regulatebitcoin, reliability, scandal, scrutiny, trading, unethical, violate, violation
Competence	abuse, assurance, blame, bogus, bombshell, buffoon, bulls**t, catastrophic, cautious, clear, clown, clumsy, comfort, competent, composed, confident, conflict, confuse, consistent, crazy, critic, deception, denial, disgrace, disgraceful, distrust, doubt, erratic, excuse, experience, experienced, fake, fear, fickle, flop, fool, foolish, fraud, gaffe, genius, honest, idiot, ignorant, innocent, knowledgeable, lie, logical, loonie, mistake, muddle, muddy, nonsense, nuts, outrage, outrageous, overreact, panic, predictable, rational, rookie, sensible, smart, socialist, sophisticated, stable, steady, tantrum, trouble, trump, trust, trustworthiness, trustworthy, unpredictable, vague, wacky, wise
Hashtags	#auditthefed, #endthefed, #popyourcollar

therefore focus our analysis on tweets from 2010 to 2023 in what follows.

Table 2: Summary statistics of daily tweets about the Fed

	min	mean	median	max	25%	75%	<10	std	count
2007	0	0.07	0	17	0	0	364	0.90	25
2008	0	0.95	0	45	0	1	360	3.57	349
2009	0	13.05	5	257	2	14	238	23.62	4764
2010	3	103.93	51	793	26	120	22	136.75	37,936
2011	12	183.73	97	2,338	57	184	0	259.51	67,062
2012	23	300.37	160	3,543	98	321.25	0	408.46	109,934
2013	38	437.58	276	5,546	168	442	0	586.24	159,715
2014	39	476.39	268	9,131	164	420	0	806.35	173,881
2015	56	532.00	321	7,878	182	540	0	818.90	194,180
2016	55	525.55	280.5	6,368	156	486	0	803.75	192,352
2017	65	533.24	304	9,176	193	524	0	872.29	194,634
2018	55	481.96	260	6,062	164	451	0	716.22	175,914
2019	123	736.82	507	9,088	349	838	0	757.05	268,940
2020	207	903.35	707	10,569	510	996.75	0	825.40	330,625
2021	263	1,100.16	870	6,310	594	1,288	0	795.74	401,557
2022	446	2,275.02	1,796	10,504	1,344	2,723	0	1,507.08	830,382
2023	111	1,779.47	1,333	11,593	973	1,921	0	1,507.28	649,507

We complete our description of the data by noting the potential impact of automated tweet generation and bot activity, which could affect the validity of our index. Existing methods to identify bot accounts often rely on specialised tools, such as Botometer X,[5] or relatively crude metrics such as the follower-to-following ratio (Tabassum et al., 2023). However, these methods have limitations: Botometer X has not been updated since May 2023 and does not cover all Twitter/X users, and it is unclear what threshold for the follower-to-following ratio should be employed to reliably identify bot accounts. It is also unclear *a priori* whether bot activity should be excluded or not—after all, bots rapidly amplify certain topics, which in turn can influence public attitudes. We take an agnostic approach to this issue in what follows. We first note that tweets generated by bots tend to be identical or highly similar to others in the dataset.[6] Recognising this, in Appendix A.2 we present a version of the

[5] https://botometer.osome.iu.edu/

[6] For example, on June 10, 2020, c. 8k tweets were posted, each of which contained the text:

index that removes duplicate tweets entirely. While this affects specific values of the index over certain time periods, the overall index is relatively unaffected by whether duplicate tweets are included or not and we choose to retain these in our baseline approach.

2.2 Sentiment analysis

We experimented with a variety of methods to assess the sentiment expressed in our tweets, including lexicon-based methods and supervised machine learning models. For our purposes these methods, and particularly the lexicon-based methods like VADER, performed poorly. Focusing on the sentiment of each word, they are unable to distinguish sentiment about the economy from sentiment regarding the Fed and its policies, and therefore generated results that were highly counterintuitive, even where the true sentiment of the tweet was clear. Instead, we found Generative Artificial Intelligence (GAI), and specifically ChatGPT, to significantly outperform these other methods, with its analysis according well with our intuition.[7]

GAI has a number of advantages over these alternative methods. First, it excels at understanding contextual meaning, which is critical in our context—e.g., some tweets may express sentiment directed at subjects other than the Fed and as GAI evaluates the full tweet it is better able to determine the true target of the sentiment. Second, it can detect implicit sentiment and interpret complex linguistic expressions—e.g., some tweets may convey sarcasm or humour, which traditional sentiment analysis models may misinterpret. Third, GAI it has been trained on extensive datasets and as a result can handle large, noisy and unstructured text, typical of social media data.

"LATEST Trump told his team to stop talks on fiscal stimulus just hours after Fed Chair Jerome Powells strongest call yet for greater spending. Here's what you need to know. Reply stop to opt out of notifications".

[7]We corroborated this with a classification exercise where two of the authors manually coded 4k tweets into sentiment buckets and we compared the performance of 6 lexicon-based models, 11 supervised machine-learning models, and both the web and API-versions of ChatGPT. We found that the ChatGPT models significantly outperformed the rest, with the lexicon-based models performing especially poorly. See Appendix A.3 for details.

2.2.1 Our procedure

Our goal is to assign each of the tweets in our dataset to one (and only one) of four categories: {'positive sentiment', 'neutral sentiment', 'negative sentiment', 'unrelated'}. Prior to deploying GAI for this purpose, we first pre-allocated a small number of tweets to these categories based on whether they include certain words or phrases that we judge to have unambiguous meaning—e.g., any tweet containing the hashtag #endthefed is pre-allocated to the negative sentiment category. We also deal with certain cases of remaining homonymy in this way. Table A.2 in the appendix lists the specific phrases used for this purpose.

For the remaining tweets in our sample, we used ChatGPT to assess the sentiment being expressed towards the Fed.[8] The prompt we use has the following structure: "Please classify each of the following tweets according to whether they have a positive, neutral, or negative sentiment toward the Fed, the Federal Reserve System, or its leadership (FOMC members). If a tweet is unrelated to the Fed, please classify it as 'unrelated'. Consider each tweet independently."[9] We also include information in the prompt about the format required of the answer—see appendix Appendix A.2 for details. We set the "temperature" parameter—a model property that controls response determinacy—to 0, ensuring that the same prompts consistently produce identical answers.[10]

Our dataset contains a large number of duplicate tweets—either in the form of retweets, likes or repeated tweets. In particular, there are a total of 1,519,077 duplicate tweets in our sample. Among these, 526 unique tweets are repeated more than 100

[8]We use a combination of ChatGPT 4o-mini and ChatGPT 4 to produce our index. We also experimented with ChatGPT 3.5 (gpt-3.5-turbo) and NotebookLM. However, the ChatGPT 4 models outperform these alternatives and are more commonly referenced in recent literature (Li et al., 2023).

[9]We explored several variants of this prompt. In one case, we replaced the words "positive, neutral or negative sentiment" with "express support for, or criticism of, or a neutral stance toward". In another, we included at the beginning of the prompt a set of six example tweets whose sentiment was decided by the authors. Neither variant had any impact on our results.

[10]Processing this number of tweets via ChatGPT is an extremely time-consuming task. It takes the API version of ChatGPT 4o-mini around 1 hour to process 5,000 tweets. In our case, that equates to around 760 hours of processing time.

times, while 413,139 unique tweets are repeated fewer than 100 times. To reduce processing time and ensure consistency in our results, we assessed sentiment in just one instance of these duplicate tweets and assigned the same value to all retweets/repeats. For tweets that are repeated or retweeted a large number of times (>100 times), we used the full ChatGPT 4 interface for this purpose. In our experience, this interface gives consistently outstanding results and using this method minimises the chance of costly error. For tweets that reoccur less often than this, we used the API version of ChatGPT (ChatGPT 4o-mini).

In total, our procedure assigns 383,267 tweets to the positive category (10.1% of the sample), 1,496,969 to the neutral category (39.4% of the sample), 1,670,352 to the negative category (44.0% of the sample), and 247,646 to the unrelated category (6.5% of the sample).

3 The index of trust and its covariates

We first describe how we use the categorised tweets to construct our index of public trust in the Federal Reserve. We then present the index, and analyse its behaviour and determinants.

3.1 An index of trust in the Fed

Let N_t be the number of relevant tweets at time t, N_t^+ be the number of positive tweets, N_t^0 be the number of neutral tweets, and N_t^- be the number of negative tweets—that is, $N_t = N_t^+ + N_t^0 + N_t^-$. We then define the trust index as the net share of positive tweets:

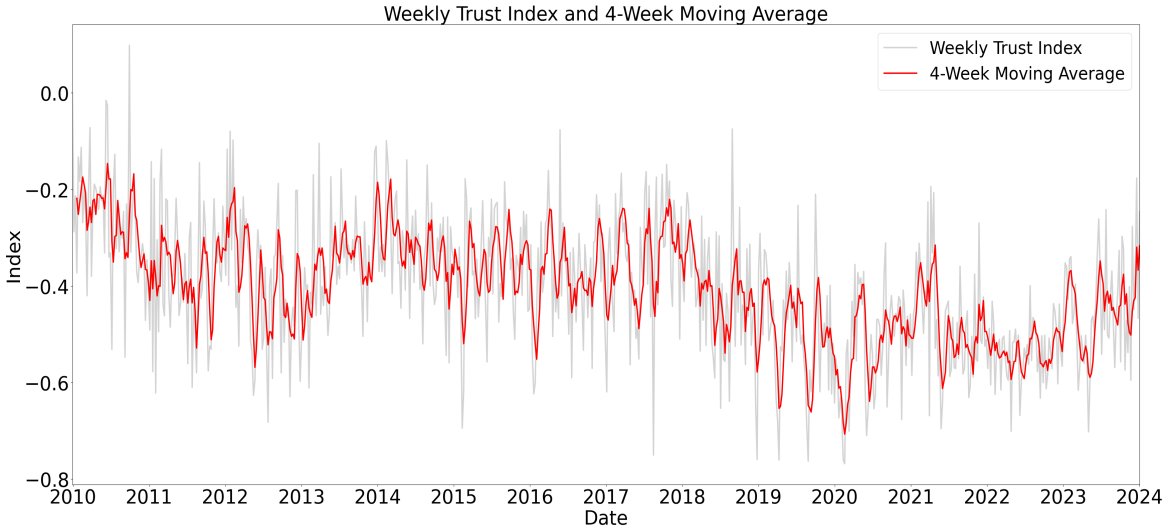
$$T_t = \frac{N_t^+ - N_t^-}{N_t} \tag{1}$$

Under this specification, $T_t \in [-1, 1]$. Note that neutral tweets only affect the

denominator of our index, while unrelated tweets are removed altogether. In our baseline approach, we include both likes and retweets within N_t; in effect, this means that a tweet that is liked or retweeted n times will have n times the weight in our index of a tweet that is neither liked nor retweeted.

Figure 1 presents our public trust index in the Federal Reserve. To aid visualisation, the index is presented at weekly frequency, along with its (trailing) four-week moving average. The mean of the index is clearly negative (-0.41) and there appears to be a mean-shift in 2018 when the index declines significantly. The index is highly volatile, with a weekly standard deviation of 0.14. It reaches its lowest point in February 2020, a period that coincides with the early stages of the Covid-19 pandemic. There are other troughs in August 2019, April 2019, and December 2018. From April 2020, there is a partial recovery, albeit this was halted for a period by the surge in inflation.

Figure 1: Index of public trust in the Fed



Notes: This figure presents our index of public trust in the U.S. Federal Reserve. The lightly-shaded grey line shows weekly values of the index, obtained by summing daily values Monday-Sunday. The bold red line shows a 4-week moving average of the index. The y-axis shows the difference between the positive and negative share of tweets, measured in decimals.

We can also decompose our index into positive and negative components (i.e., $\frac{N_t^+}{N_t}$ and $-\frac{N_t^-}{N_t}$). Our index is overwhelmingly driven by the movements in the negative component—the correlation between the index and its negative component is 0.95.

Appendix B.3 presents further analysis at the positive and negative component level.

3.2 What drives trust in Federal Reserve?

The academic literature on institutional trust (e.g, Mayer et al. (1995)) has identified three factors as important—summarised as ability, benevolence and integrity. Ability relates to the perception that policymakers have the required skills and competencies to do their job. Benevolence captures the idea that policymakers are pursuing policies in the interest of society at large rather than some specific group. Integrity captures policymakers' adherence to acceptable behavioural principles.

We would therefore expect our trust index to react to macro-financial news relating to the Fed's dual mandate of price stability and maximum employment, as well as to news about financial stability risks. We might also expect decisions over certain Fed policies that have large perceived distributional effects, such as Quantitative Easing and emergency lending programmes, to influence the index. Finally, the index should react to news of perceived conflicts of interest relating to the personal financial dealings of senior Federal Reserve officials, or to more general questions over their probity, as well as to positive information about policymakers' impacts on society beyond their policy actions. An example of this is the #popyourcollar tribute to Janet Yellen hailing her as a role model for women in STEM, launched on Twitter in the final days of her tenure as Fed Chair.

With this in mind, we examine the following set of potential determinants of our public trust index in the Federal Reserve.

Macroeconomic and financial indicators We include: the 5-year break-even inflation rate, a measure of medium-term (MT) inflation expectations; the 5-year 5-yr forward inflation rate, a measure of long-term (LT) inflation expectations;the 30-Day Fed Funds Futures to capture market expectations of monetary policy actions in the near term; the 10-year Treasury yield, a key indicator of long-term interest rates and

economic sentiment; and an index of the strength of the U.S. dollar vis-à-vis other currencies. We also include the VIX index, which captures market expectations for near-term volatility in the S&P 500; the Chicago Fed's National Financial Conditions Index, which provides a comprehensive measure of U.S. financial conditions; daily returns on the S&P 500 Index and the NASDAQ index; and the West Texas Intermediate price of crude oil, as a salient measure of future inflation pressures. Lastly, we include the Aruoba-Diebold-Scotti Business Conditions Index, which provides a daily real-time measure of business conditions in the U.S. economy (S. Borağan Aruoba and Scotti, 2009);

Sentiment and policy uncertainty We consider the relationship between our trust index and the Shapiro-Sudhof-Wilson Daily News Sentiment metric, derived from daily news reports to account for the influence of media sentiment on public perception (Shapiro et al., 2022). And the impact of the Baker-Bloom-Davis Economic Policy Uncertainty Index, which captures the level of uncertainty in economic and policy-related news (Baker et al., 2016).

Monetary policy actions and public appearances We introduce dummy variables to capture the influence of key monetary policy events, including the publication of the semi-annual Monetary Policy Report and the Fed Chair's speech at the Jackson Hole Symposium. We also exploit existing measures of monetary policy shocks to capture the effects of unexpected shifts in policy on our public trust index. In particular we consider the measure provided by Jarociński and Karadi (2020), which decomposes high frequency unexpected shifts in interest rates in a window around the monetary policy press conference into a pure monetary policy shock series and an information effect relating to the central bank's assessment of the economic outlook. We consider both series in our regression analysis.

During our sample period, the Fed experienced two transitions in chairmanship—from Ben Bernanke to Janet Yellen in 2014, and from Janet Yellen to Jerome Powell

in 2018. Each transition involves a set of events that could potentially influence public trust in the Fed. These include the announcement of the nominee, his or her Senate confirmation, and his or her first day in office. These milestones are significant as they signal shifts in leadership and policy direction, drawing substantial attention from the public and markets. The announcement reflects the formal nomination of the new chair, often accompanied by public statements outlining their qualifications and vision. The Senate confirmation solidifies the nominee's position, affirming their mandate to lead the central bank. Finally, the first day in office marks the official commencement of their responsibilities. We consider the impact of each of these events in our empirical analysis.

President Trump tweets During his first term in office (January 2017 to January 2021), President Trump frequently posted tweets criticising the Federal Reserve and its leadership. It is plausible that these tweets shaped public perceptions of the institution, until his banishment from Twitter in January 2021 during the final days of his term. Following the classification in Bianchi et al. (2023), we distinguish tweets made by President Trump that explicitly pose a threat to central bank independence from his other tweets that relate to the Fed.

Exogenous events affecting trust/narrative proxy Trust in an institution and its leadership also depends on the public's view of the officials' integrity. For example, trust can be eroded if officials are perceived to be in breach of ethical guidelines, such as if they engaged in insider trading or leaked confidential information. Similarly, trust could be enhanced if a policymaker is deemed to have had a particularly positive impact on society (unrelated to policy actions). An example of this is the #popyourcollar tribute to Janet Yellen hailing her as a role model for women in STEM, launched on Twitter ahead of her last days in office.

In the latter part of our sample, there were several reports of high-profile Federal Reserve officials having engaged in investment trades ahead of important policy an-

nouncements. These news plausibly raised public scepticism about the integrity of the Fed decision-making process and could therefore have dented public trust in the Federal Reserve. Most of the exogenous events we consider are negative (scandals), but we also include in this set the #popyourcollar tribute to Janet Yellen discussed above. Table 3 lists the exogenous events we consider in our empirical analysis. We show in the appendix that each of these events is followed by a marked increase in Twitter/X activity relating the individuals concerned in the immediate aftermath of the news breaking.

These events are of particular interest for us because they can be used as narrative sign restrictions on the trust shock, as we discuss further in Section 4. We show below that these events are highly correlated with movements in trust. While the alleged ethical breaches could plausibly be endogenous to the economy (e.g., with trades happening ahead of policy moves), the variable and random delay of their reporting ensures that these events are exogenous.[11] We will use the shorthand "narrative proxy" for these events, as we will use them in the next section to construct a narrative proxy based on the shock sign restrictions they provide.

Table 3: Exogenous events affecting trust

Date	**Sign**	**Event Details**
04/04/2017	−	Richmond Fed President Jeffrey Lacker resigns abruptly after admitting to leaking information about confidential Fed deliberations in 2012. Published online in the Washington Post at 1.44pm ET and widely reported and shared thereafter by NYT, NPR, CNBC, CNN, WaPo, WSJ, Marketwatch and other outlets.
31/01/2018	+	#popyourcollar tribute to Janet Yellen launched by the New York Fed ahead of her last day as Fed Chair (3 February 2018).

Continued on next page

[11] In addition, in Appendix B.2 we document a material increase in the number of negative tweets about individuals named in these ethical scandals as the news of these events is reported.

Table 3: Exogenous events affecting trust (Continued)

07/09/2021	–	The Wall Street Journal reported that Kaplan was active buyer and seller of stocks in the course of 2020.
08/09/2021	–	Bloomberg reports that Rosengren made multiple purchases and sales in REITs and other securities in 2020. The Boston Fed said the trades were not done via a blind trust and were "consistent" with ethics rules.
09/09/2021	–	In response to reports that Rosengren and Kaplan had traded REITs and other stocks released the day before by the WSJ, but not picked up by other media nor the Twitterverse, the Federal Reserve Banks of Boston and Dallas simultaneously issue statements on 9 September noting that their financial transactions complied with the Fed's ethics rules. While the previous two articles are not picked up by the media nor Twitter users, there is a big reaction after the statement. Published online in Dallas News at 2.01 pm ET and widely reported and shared thereafter by WSJ, NYT, CNBC, Reuters, etc.
27/09/2021	–	Before the market opens, Boston Fed President Rosengren announces he will pull forward his retirement by nine months. After the market closes, Dallas Fed President Kaplan says he will step down on October 8, 2021, mentioning his financial dealings. News appears on CNBC at 8.30am and it is widely reported and shared by PBS, NPR, Bloomberg, WaPo, Dallas News, etc.
02/10/2021	–	Bloomberg reports on 1 October 2021 at 8.41pm ET that Fed Vice Chairman Clarida traded between $1 million and $5 million out of a Pimco bond fund and into two stock funds. The shift happened on February 27, 2020, a day before Powell issued a rare, impromptu statement reiterating the Fed's support of the economy. Most subsequent coverage by usual media outlets and tweets are mainly from 02/10/2021.
10/01/2022	–	Warren sends a letter to Jerome Powell, asking for more details on trades made by senior Fed officials. Later in the afternoon, Clarida sends a letter to President Joe Biden announcing that he will be stepping down on Jan. 14, instead of the Jan. 31 day that his term was originally set to expire. Press release at 4pm ET and is subequently reported that afternoon in all major outlets – CNBC, WaPo, NYT, NYPo.

Continued on next page

Table 3: Exogenous events affecting trust (Continued)

14/10/2022	–	Raphael Bostic reported that he had failed to previously disclose financial transactions on his official central bank forms. Release at 3pm ET by the Washington Post, followed by CNBC, Rueters, Bloomberg and others.
20/10/2022	–	The New York Times revealed that James Bullard, president of the Federal Reserve Bank of St. Louis, had spoken previous week on the 14th at an off-the-record, invitation-only forum held by Citigroup, and open to clients. Bloomberg picks up the news at 2.47pm ET and other media outlets follow.
15/06/2023	–	Bostic revealed transactions involving 19 exchange-traded funds made on May 2 of last year. At that time, officials were in a so-called blackout period during which public communications were limited and trading prohibited. News available on Bloomberg at 1.27pm ET and echo by Reuters, WSJ, etc.

3.3 Regression results

We use the daily version of our index of public trust in the Federal Reserve and regress it on the set of variables discussed above, plus its lags:

$$T_t = \alpha + \beta X_{t-1} + \gamma W_t + \sum_{i=1}^{5} \delta_i T_{t-i} + \varepsilon_t \tag{2}$$

where T_t is the trust index, X_{t-1} includes the set of macroeconomic and financial variables, and W_t incorporates dummy variables for the exogenous events listed in Table 3, plus other discrete event dates linked to monetary policy such as Jackson Hole speeches. We include 5 lagged values of the index to account for possible persistence in trust dynamics. The macro-financial variables are included with a one-day lag because the data we use are end-of-day values. The results from this regression are reported in Table 4.

We first regress the index on a set of macro-financial variables available at daily frequency. Increases in short-term inflation expectations have a modest negative impact on the trust index: a 1 standard deviation rise in expected inflation reduces the

index by around 2% points. In contrast, longer-term inflation expectations are not statistically significant. The Fed Funds rate has a consistently negative and economically large coefficient in the regression, with a 1 standard deviation reducing the trust index by 11-13% points. There is also some evidence that, holding constant the current Fed Funds rate and inflation expectations, expected hikes in the Fed Funds rate tend to boost the index. One interpretation of this is that policy rate hikes tend to happen in positive economic environments. The yield on the 10-year Treasury operates in a similar manner, with higher yields boosting the trust index—consistent with higher long-term interest rates reflecting stronger growth prospects for the economy. NASDAQ stock market returns have a modest positive impact on Fed trust, and the VIX "fear" index has a modest negative impact. An appreciation of the U.S. dollar exchange rate also tends to reduce our trust index, perhaps reflecting flight to safety effects into dollar assets during periods of financial turmoil. Finally, neither the oil price nor the Arouba-Diebold-Scotti business conditions index has a statistically significant effect.[12]

[12]We also considered the prices of gold and bitcoin in this regression. These indicators might be considered barometers of investor beliefs about the Fed's ability to maintain price stability and a reflection of distrust in traditional financial and economic markets. Neither indicator has a statistically significant impact on our trust index.

Table 4: Regression results, daily index of public trust in the Federal Reserve

Macro-financial variables:					
MT inflation expectations$_{t-1}$	-0.0180***	-	-	-	-0.0202***
LT inflation expectations$_{t-1}$	-0.0068	-	-	-	-0.0046
Fed Funds rate$_{t-1}$	-0.1350**	-	-	-	-0.1148**
30-day Fed Funds Futures rate$_{t-1}$	0.1008*	-	-	-	0.0891
10-year Treasury yield$_{t-1}$	0.0416***	-	-	-	0.0343***
Financial conditions$_{t-1}$	0.0019	-	-	-	0.0030
S&P500$_{t-1}$	-0.0122	-	-	-	-0.0118
NASDAQ$_{t-1}$	0.0146*	-	-	-	0.0143*
VIX$_{t-1}$	-0.0157***	-	-	-	-0.0126**
US dollar$_{t-1}$	-0.0182**	-	-	-	-0.0191**
ADS Business Conditions Index$_{t-1}$	-0.0063	-	-	-	-0.0057
Oil price$_{t-1}$	0.0058	-	-	-	0.0057
Sentiment and policy uncertainty:					
SSW News Sentiment Index$_{t-1}$	-	0.0183***	-	-	0.0085
Economic Policy Uncertainty$_{t-1}$	-	-0.0043	-	-	0.0021
Monetary policy:					
JK monetary policy shock$_{t-1}$	-	-	-0.0023	-	-0.0015
JK monetary policy information$_{t-1}$	-	-	0.0063***	-	0.0067***
Monetary policy report publication$_t$	-	-	0.1370***	-	0.1364***
Jackson Hole speech$_t$	-	-	0.0278	-	0.1033**
Leadership change$_{t-1}$	-	-	0.2136**	-	0.1937**
Exogenous drivers of trust:					
Ethical scandals$_t$	-	-	-	-0.2377***	-0.1952***
Trump tweets Bianchi et al.$_t$	-	-	-	-0.1661***	-0.1514***
Other Trump tweets about Fed$_t$	-	-	-	-0.0706*	-0.0667*
Lags	Yes	Yes	Yes	Yes	Yes
α	-0.2689***	-0.2237***	-0.2089***	-0.2090***	-0.2740***
Ad. R^2	0.1405	0.1206	0.1165	0.1206	0.1488
N	4,378	4,378	4,378	4,378	4,378

Notes: The dependent variable is the daily index of trust in the Fed. Continuous right-hand side variables have been standardised by subtracting the mean and dividing by the standard deviation. The regression includes 5 lags of the dependent variable. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels. Standard errors are heteroscedasticity robust (HC1).

Turning to the next category of indicators, perhaps surprisingly we find that newspaper sentiment has only a small impact on our trust index—while this variable is statistically significant in some regression specifications, the coefficient is very small. Evidently our trust index is capturing something distinct from newspaper sentiment, consistent with the discourse on Twitter/X evolving distinctly to that in major newspapers. We also find no link between the Baker-Bloom-Davis Economic Policy Uncertainty Index and our trust measure.

As might be expected, our trust index is systematically related to variables that are closely linked to Fed monetary policy. In particular, the publication of the semi-annual Monetary Policy Report, the announcement, confirmation and swearing in of recent Fed chairs, and Jackson Hole speeches by the Fed Chair all have large and statistically significant effects—albeit, with only the MPR having a robust effect across all specifications. We also find a positive impact from the Jarocinski-Karadi central bank information variable, although the economic magnitude of the effect is very small.

Finally, we note that the exogenous events listed in Table 3, defined as dummies on the date of the news release, have a large impact on the trust index. Our narrative proxy, which combines news of ethical breeches involving Fed officials and the #popyourcollar hashtag launched as a tribute to Chair Yellen (with reversed sign), has the largest measured impact on the index: taken together, these events reduce our index by 20-24% points on the same day. We also find a very material impact from the set of President Trump's tweets associated with threats to the Fed's independence: these tweets generate a 15-16% point drop in the index. Overall, both events appear to generate a substantial erosion of confidence in the Federal Reserve. Other tweets by President Trump that reference the Fed have a significantly smaller—albeit still material—impact.

In summary, our measure of public trust in the Federal Reserve has an intuitive relationship with a wide variety of plausible determinants. Macro-financial factors such as the yield curve, inflation expectations, and market volatility all have statistically

significant effects, with the largest impact coming from the Fed Funds rate. The index has only a weak relationship with newspaper sentiment. Monetary policy-related indicators and events also matter, with the semi-annual Monetary Policy Report, the Jackson Hole speech by the Fed Chair, and key dates in the leadership transition all entering the regression with large positive coefficients. Events unrelated to monetary and financial policy also matter. Notably, news of ethical scandals affecting senior FOMC members has the largest measured effect on the trust index, while tweets by President Trump perceived as threatening the Fed's independence also have large negative impacts.

4 The effects of shocks to public trust in the Fed

To explore the macroeconomic and financial market impacts of shocks to public trust in the Fed, we estimate a Structural Vector Autoregression (VAR) model using a narrative identification approach based on shock sign-restrictions (Antolín-Díaz and Rubio-Ramírez, 2018). The reported ethical scandals that affected the Fed in 2017 and in 2021-2023 along with the emergence of the #popyourcollar Twitter trend launched on 31 January 2018 to honour Chair Yellen ahead of her stepping down as Fed Chair on 3 February 2018, provide an ideal set of events that are relevant for the dynamics of the trust index but exogenous to the state of the economy.

4.1 Identification strategy

Our index of trust in the Fed is affected by movements in the macroeconomy and financial markets, as well as by policy actions and public appearances of the FOMC members, which are taken to reveal information about their views on the state of the economy and, possibly, future policy. It is also affected by the ethics scandals, which erode trust in the institution and its leadership. As we will argue, these events are unrelated to the state of the economy, but negatively affect public trust in the Fed.

Our primary identification strategy is based on information about alleged breaches of ethical scandals involving FOMC members. These scandals, detailed in Table 3, can be grouped into two categories: insider trading allegations concerning Vice Chair Clarida, and Presidents Bostic, Kaplan, and Rosengren; and allegations of the leaking of information concerning Presidents Bullard and Lacker. In all cases, there are lags between the alleged breaches of Fed ethics guidelines and the dates these events are reported in the news, varying from 5 years (President Lacker) to 6 days (President Bullard). So, while the alleged behaviour that caused the scandal is potentially endogenous to the state of the economy, the date when these events are reported in the media is not. We assume that there is a negative shock to our trust index (Sign column in Table 3) on the day in which the news of each scandal is reported. We augment this with a single positive shock to our trust index that occurred in 2018, with the #popyourcollar Twitter trend to honour the achievements and the legacy of Janet Yellen as she was stepping down as Fed chair. This event too is unrelated to macroeconomic conditions: it coincided with the end of Chair Yellen's mandate, but was launched a few days ahead of the end of her tenure.

Giacomini et al. (2022) and Plagborg-Møller and Wolf (2021) demonstrate that it is possible to use information about the sign of a particular structural shock in specific periods by recasting it as a discrete-valued proxy for the structural shock, which takes the value 1 when the shock is known to be positive, -1 when it is known to be negative, and zero in all other periods. This proxy is "relevant" (it is positively correlated with the target structural shock) and is exogenous (uncorrelated with all other shocks) and is therefore a valid proxy for the structural shock, which can be used to identify impulse responses to that shock in a proxy-SVAR as discussed in Mertens and Ravn (2013) and Stock and Watson (2018). We exploit the narrative identification strategy around the events collected in Table 3 and, based on this insight, construct a proxy for the trust shock based on the signs for these events.

We compare the effects of the trust shocks identified exploiting information about

scandals with two alternative identification strategies: the first imposes similar sign restrictions for trust on the days in which Trump tweets reference the Fed in a way identified by Bianchi et al. (2023) as threatening to central bank independence; the second uses a Cholesky decomposition.

4.2 The VAR model

We specify a VAR for the trust index and a set of daily variables that we found to covary with it in Section 3: a measure of medium-term inflation expectations (the 5-yr break-even rate $E\pi_t^{5y}$); the Aruoba-Diebold-Scotti business conditions index of S. Borağan Aruoba and Scotti (2009) (ADS); the log of the NASDAQ index ($NASDAQ$); the daily federal funds effective rate (Fed Funds rate); the 30-day Fed Funds Futures rate ($FFF30$), which reflects market expectations about the future course of the Federal Reserve's monetary policy; the 10-year T-Bill yield; the daily news sentiment indicator (SSW) developed by Shapiro et al. (2022); and perceived financial market uncertainty, as measured by the VIX Index (VIX).

The variables enter the model in log-levels, except those already defined in terms of rates, such as interest rates and inflation expectations, which are taken in levels. Our baseline specification includes 28 lags, but our results are robust to various lag and model specifications.[13]

We estimate this model:

$$A(L)Y_t = u_t = B\varepsilon_t \tag{3}$$

where A and B are coefficient matrices, Y_t is the vector of endogenous variables described above, u_t is a vector of white noise innovations with covariance matrix Σ, and ε_t is the vector of white noise structural shocks, with the identity as their covariance

[13]We experimented with 7, 14 and 21 lags and with the log-differences for the NASDAQ index; we explored using the Secured Overnight Financing Rate instead of the Fed Funds Rate; and we tried including the S&P500 index along with or instead of the NASDAQ index.

matrix. We can partition these into two subvectors, $\varepsilon_t = [\varepsilon_t^1 \ \varepsilon_t^{2\prime}]'$ where ε_t^1 is the shock we aim to identify, i.e. the trust shock.

We consider two instruments for ε_t^1:

i) **Narrative proxy**: An indicator variable which is equal to -1 on the days indicated in Table 3 as having negative sign, i.e. the days in which the alleged ethical breaches are reported, equal to $+1$ on the day the #popyourcollar tribute was launched, and zero elsewhere;

ii) **Trump tweets**: An indicator variable that has sign -1 on the dates of the Trump tweets identified by Bianchi et al. (2023) as threats to central bank independence and zero elsewhere.

The case for the validity of instrument i) is very strong. It is possible that the alleged actions that contravened ethics guidelines were endogenous to the economy—for example, the Clarida trades happened just ahead of an extraordinary statement by Chair Powell in the height of the COVID-19 crisis. But lags in the reports of these alleged ethics breaches are random, thus ensuring the exogeneity of this instrument. The first-stage F-statistics with robust standard errors for these instruments from the VAR discussed in the next section are 11.6 for (i) and 26.1 for (ii)—both exceed 10, the rule of thumb suggested by Staiger and Stock (1997) as a threshold for instruments to not to be weak.

We also show the impulse response functions obtained with a Cholesky identification in which the trust index is ordered first, i.e. responds contemporaneously only to its own shock. This is justified by the fact that the other variables in the VAR are measured as end-of-day values and so cannot be fully reflected in the trust index at time t. This shock picks up all the residual movement in the index that is not explained by lags of the model's variables.

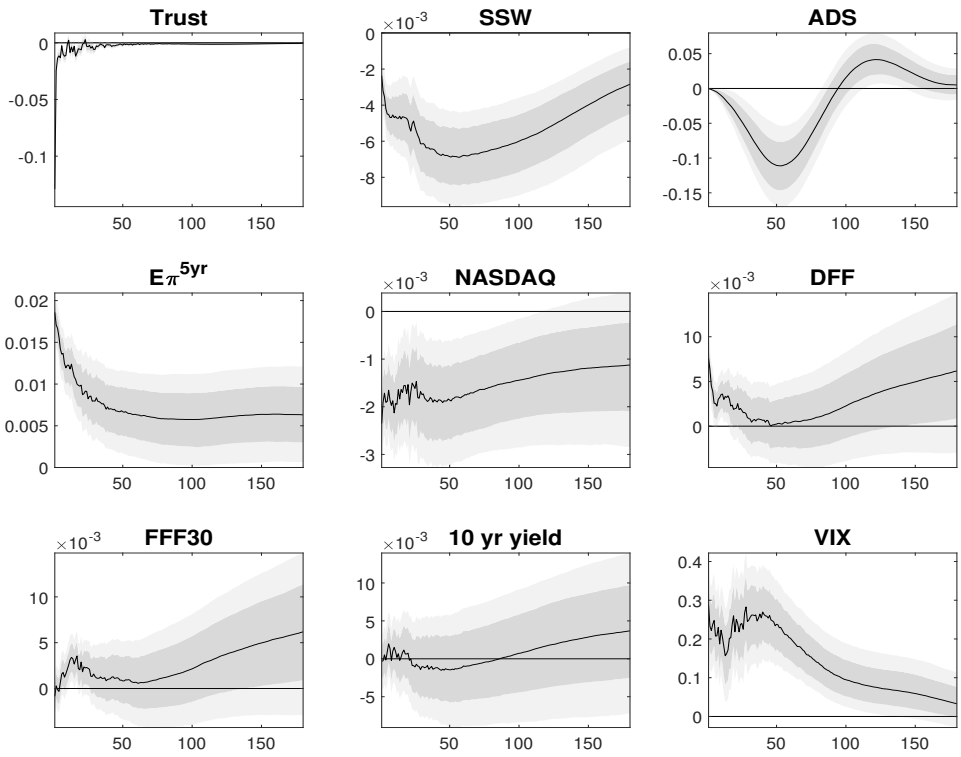


Figure 2: Narrative proxy: Impulse response functions to a trust shock identified using our narrative proxy (scandals + #popyourcollar) as instrument. The dark and light shaded areas correspond to the 68% and the 90% high probability density (HPD) sets, respectively. The x-axis measures days.

4.3 Responses to a trust shock

Figures 2, 3 and 4 present the dynamic effects of negative shocks to public trust in the Federal Reserve, using the various identification approaches described above. We begin with the narrative proxy approach, which is based on the narrative information about the effects of ethical scandals and the #popyourcollar trend. This, we argue, is the most credibly exogenous instrument for Fed trust.

As Figure 2 shows, a negative one standard deviation shock to public trust in the Fed causes a sharp contemporaneous drop in the trust index, which however recovers very quickly thereafter. The negative shock to trust lowers the daily news sentiment on

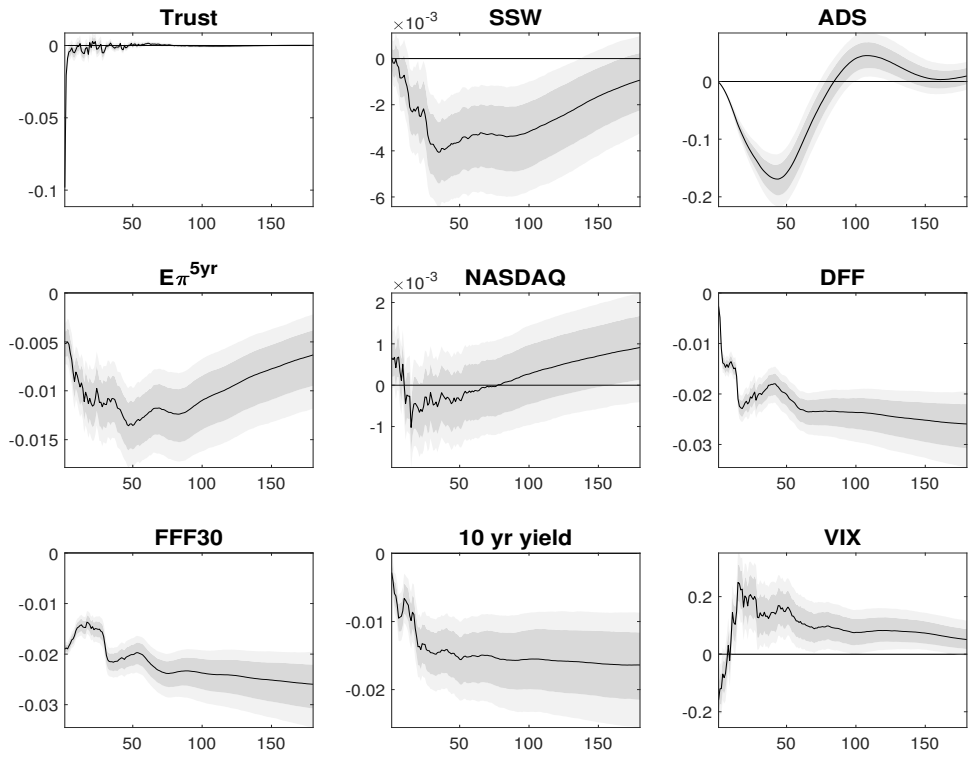


Figure 3: Trump tweets: Impulse response functions to a trust shock identified Trump tweets. The dark and light shaded areas correspond to the 68% and the 90% high probability density (HPD) sets, respectively. The x-axis measures days.

impact, while financial market uncertainty increases on impact and remains heightened for approximately half a year. The ADS index of business conditions and the NASDAQ both fall, the latter quite persistently. The Fed Funds Futures rate increases although the impact is not statistically significant. The break-even inflation rate over the next 5 years increases significantly and persistently before subsiding, taking approximately 2 years for the effect to be fully reabsorbed (see Figure B.11 in Appendix).

These results, whereby an erosion in trust resembles a "supply" shock, are in line with Bursian and Faia (2018). These authors propose a monetary model where trust emerges as an equilibrium in a game of strategic interaction between a continuum of betrayal-averse agents and the monetary authority whose incentives are stochastic

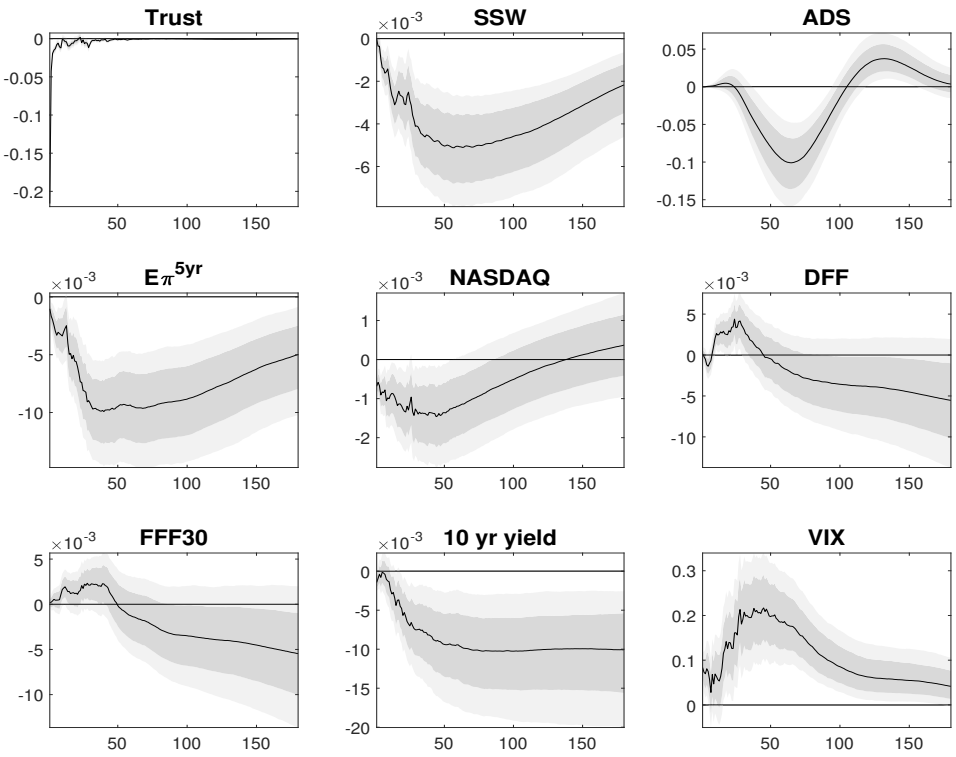


Figure 4: Cholesky identification: Impulse response functions to a trust shock identified with Cholesky decomposition where the trust index is ordered first (responds contemporaneously only to a trust shock). The dark and light shaded areas correspond to the 68% and the 90% high probability density (HPD) sets, respectively. The x-axis measures days.

ex post. In their model, as trust declines, consumers behave more cautiously and reduce demand. At the same time, the monetary transmission mechanism becomes less effective and the Phillips curve steepens. As agents perceive a worsening in the inflation-output trade-off, inflation increases in anticipation, despite the fall in aggregate demand.

In Figure 3, we focus on shocks to trust in the Fed identified using only President Trump's tweets. The dynamic responses in this case resemble an economic confidence shock. Consistent with the evidence reported by Bianchi et al. (2023), President Trump's tweets affected market expectations about future monetary policy: the ex-

pected Fed Funds rate falls on impact and is persistently lower, followed by the daily effective Fed Funds rate. The daily news sentiment remains negative for almost 6 months, while the ADS index of business conditions falls but recovers somewhat more quickly. After an initial decline, the VIX Index increases persistently. The impact on the NASDAQ is no longer significant.

Finally, we consider responses to a trust shock identified by orthogonalising the residuals of the VAR through a Cholesky decomposition (Figure 4). This approach uses all variation in the trust index that cannot be explained by macro-financial variables. The responses in this case are broadly comparable to those obtained with the other identification strategies. New sentiment declines, business conditions contract, the NASDAQ declines, and the VIX Index increases. The behaviour of the yield curve and inflation expectations is similar to that found following trust shocks identified using tweets by President Trump: inflation expectations decline, as do short and long-term yields (the former after a temporary increase).

How can we reconcile the persistence of the effects on the economy described in all three figures with the short-lived impact on our trust measure? One interpretation of these findings is that our trust index is not especially well suited to capturing persistence in the dynamics of trust itself. As a social media-based measure, attention will inevitably move on after a few days or at most weeks. As such, it might not capture latent erosions in trust, which would become apparent only when the next shock occurs. For this reason, we think our measure is perhaps best viewed as providing a fast-moving early read on shifts in trust.

5 Conclusion

Our research proposes a novel method to measure trust in central banking. By leveraging Generative AI (GAI) to assess the sentiment in millions of tweets about the Federal Reserve and its leadership, our measure provide a more dynamic and compre-

hensive perspective of the public's trust in the Fed than has hitherto been available. Our measure of trust, calculated as the daily difference between positive and negative sentiment tweets, fluctuates significantly over time, with a clear mean-shift downwards in 2018.

Macro-financial factors including the yield curve, inflation expectations, and market volatility, significantly influence our trust index, with the Fed Funds rate having the largest effect—a 1 standard deviation increase in the Fed Funds rate reduces net trust in the Fed by 11% points. Other monetary policy-related events also matter, with the publication of the Monetary Policy Report and the annual Jackson Hole speech by the Fed Chair boosting trust. Events unrelated to monetary and financial policy, such as financial scandals involving FOMC members and critical tweets from President Trump, erode it. Negative shocks to trust—captured using a narrative identification approach that maps the news of alleged financial scandals involving FOMC members into negative shocks to trust—reduce economic activity and increase inflation expectations, i.e., they operate akin to a supply shock. The effects are highly persistent, and continue to be felt in macro-financial variables 6 months after the shock.

What message should central bank policymakers take from our results? First, ensuring integrity and accountability of central bank officials is key to maintaining public trust. The events that cause the largest erosion in trust in our sample are scandals that bring into question the integrity of key officials. The number one priority for shoring up trust in the central bank must therefore be ensuring that central bank policymakers and staff have an appropriate formalised code of conduct, which includes restrictions on actively trading financial instruments that might be perceived as benefitting from inside information, while dealing swiftly with any incidents of unethical behaviour.

Second, our results reinforce the notion that pro-active and transparent communication are important tools for bolstering the public's trust in the central bank. We find large positive effects on trust from the publication of the Monetary Policy Report and the annual Jackson Hole speech by the Fed Chair. We also find positive effects

from the #popyourcollar trend that was initiated by the New York Fed as a tribute to Chair Yellen in her final days in the role.

Third, our analysis highlights the need for central banks to develop timely measures of the public's trust. Shocks to trust depress economic activity and may cause inflation expectations to increase. While these shocks have clearly not been of sufficient magnitude to influence U.S. business cycles materially over the recent past, there is no guarantee that this will continue to be the case in future. With the potential for greater political interference in U.S. monetary policy in future, understanding how central banks should act to retain public trust have become ever more urgent.

References

Angelico, C., J. Marcucci, M. Miccoli, and F. Quarta (2022): "Can we measure inflation expectations using Twitter?" *Journal of Econometrics*, 228, 259–277.

Antolín-Díaz, J. and J. F. Rubio-Ramírez (2018): "Narrative Sign Restrictions for SVARs," *American Economic Review*, 108, 2802–29.

Baker, S., N. Bloom, and S. Davis (2016): "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 131, 1593–1636.

Bianchi, F., R. Gómez-Cram, T. Kind, and H. Kung (2023): "Threats to central bank independence: High-frequency identification with twitter," *Journal of Monetary Economics*, 135, 37–54.

Blinder, A. S., M. Ehrmann, J. de Haan, and D.-J. Jansen (2024): "Central Bank Communication with the General Public: Promise or False Hope?" *Journal of Economic Literature*, 62, 425–57.

Brouwer, N. and J. de Haan (2022): "Trust in the ECB: Drivers and consequences," *European Journal of Political Economy*, 74, 102262.

Bursian, D. and E. Faia (2018): "Trust in the monetary authority," *Journal of Monetary Economics*, 98, 66–79.

Bursian, D. and S. Fürth (2015): "Trust Me! I am a European Central Banker," *Journal of Money, Credit and Banking*, 47, 1503–1530.

Carstens, A. (2023): "Trust and public policies," Tech. rep., Bank for International Settlements, Speech 19 May 2023.

Christelis, D., D. Georgarakos, T. Jappelli, and M. van Rooij (2020): "Trust in the central bank and inflation expectations," *International Journal of Central Banking*, 16, 1–37.

De Grauwe, P. and Y. Ji (2024): "Trust and monetary policy," *Journal of Forecasting*, 43, 903–931.

Ehrmann, M. (2024): "Trust in Central Banks," Tech. rep., European Central Bank mimeo.

Ehrmann, M., M. Soudan, and L. Stracca (2013): "Explaining European Union Citizens' Trust in the European Central Bank in Normal and Crisis Times," *The Scandinavian Journal of Economics*, 115, 781–807.

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., T. Kitagawa, and M. Read (2022): "Narrative Restrictions and Proxies," *Journal of Business & Economic Statistics*, 40, 1415–1425.

Jarociński, M. and P. Karadi (2020): "Deconstructing monetary policy surprises—the role of information shocks," *American Economic Journal: Macroeconomics*, 12, 1–43.

Li, K., F. Mai, R. Shen, C. Yang, and T. Zhang (2023): "Dissecting corporate culture using generative AI–Insights from analyst reports," *Available at SSRN 4558295*.

Mayer, R. C., J. H. Davis, and F. D. Schoorman (1995): "An Integrative Model of Organizational Trust," *The Academy of Management Review*, 20, 709–734.

Mertens, K. and M. O. Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–47.

Plagborg-Møller, M. and C. K. Wolf (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89, 955–980.

S. Boraǧan Aruoba, F. X. D. and C. Scotti (2009): "Real-Time Measurement of Business Conditions," *Journal of Business & Economic Statistics*, 27, 417–427.

Schnabel, I. (2020): "The importance of trust for the ECB's monetary policy," Tech. rep., European Central Bank, Speech 16 December 2020.

Shapiro, A. H., M. Sudhof, and D. J. Wilson (2022): "Measuring news sentiment," *Journal of Econometrics*, 228, 221–243.

Staiger, D. and J. H. Stock (1997): "Instrumental Variables Regression with Weak Instruments," *Econometrica*, 65, 557–586.

Stock, J. H. and M. W. Watson (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, 128, 917–948.

Tabassum, F., S. Mubarak, L. Liu, and J. Du (2023): "How Many Features Do We Need to Identify Bots on Twitter?" *International Conference on Information*, 312–327, themed Issue: Vector Autoregressions.

A Appendix A

In this appendix, we describe details of the procedure employed to compile a dataset of relevant tweets about the Federal Reserve, its leadership, or its policy framework and decisions, and to categorise the sentiment in these tweets.

A.1 Details of the query used for downloading relevant tweets

As described in the main text, the query we use in the Twitter API has three components. In particular, it returns a tweet if it:

1. $EITHER$ contains one element from the set {"the fed", "federal reserve", "thefed", "fedreserve", "federalreserve", "federalreservebank", "usfed", "fomc", "#Fed", "#FOMC"}

2. OR contains the name of a Fed Chair or key Fed official AND one element from {"market", "economy", "business"} OR references to specific regional Federal Reserve Banks OR to the individual's title, e.g., "chair" or "president".

3. AND contains one of the words listed in the lexicon in Table 1 in the main text.

We list the specific strings used in step 2 in Table A.1 below.

Next, we describe the process followed to obtain the lexicon described in Table 1 in the main text. The terms in this table are used to differentiate tweets that express confidence in, support for, or criticism and distrust of the Fed, its officials, or its policy framework, from those that provide objective, neutral reporting of events related to the Fed or Fed policy decisions.

We first downloaded a total of 40k tweets featuring the names of key Fed officials including Janet Yellen, Jerome Powell, Raphael Bostic, James Bullard, and Eric Rosengren. For the subset of tweets relating to each individual, we provided ChatGPT with a series of relevant vs irrelevant example tweets and then asked: "Given the

Table A.1: Query for Identifying Fed Officials

Fed official	**Query**
Ben Bernanke	bernanke AND (fed OR ben OR chair OR chairman OR market OR business OR economy)
Janet Yellen	yellen
Jerome Powell	powell AND (fed OR jerome OR jay OR jaypow OR chair OR chairman OR market OR business OR economy)
Donald Kohn	kohn AND (fed OR donald OR don OR fomc OR market OR business OR economy)
Stanley Fischer	fischer AND (fed OR fomc OR stan OR stanley OR market OR business OR economy)
Richard Clarida	clarida AND (fed OR chair OR richard OR market OR business OR economy)
Lael Brainard	brainard AND (fed OR chair OR lael OR markets OR business OR economy)
Philip Jefferson	jefferson AND (fed OR chair OR philip OR markets OR business OR economy)
Mary Daly	daly AND (fed OR president OR pres OR mary OR san francisco OR markets OR business OR economy)
Eric Rosengren	rosengren AND (fed OR president pres OR OR eric OR boston OR markets OR business OR economy)
Robert Kaplan	kaplan AND (fed OR president pres OR OR dallas OR markets OR business OR economy)
Raphael Bostic	bostic AND (fed OR raphael OR atlanta OR market OR business OR economy)
James Bullard	bullard AND (fed OR president OR pres OR louis OR james OR jim OR market OR business OR economy)
Jeffrey Lacker	kaplan AND (fed OR president OR pres OR dallas OR markets OR business OR economy)

above examples, which of these tweets are critical of or supportive of the fed, the federal reserve system or its leadership (fomc members), or of its ability to meet inflation and financial stability objectives? Provide your answers in the format: 'the number. Result: 1 for yes, 0 for no'." The example tweets are listed below.

1. Janet Yellen. Tweet: 'Janet Yellen's rookie mistake?' Result: 1. Tweet: '#yellen #Obama she flipped flopped enough to run for president, she def got the speech style all figured out.' Result: 1. Tweet: '#yellen they could not achieve 2% inflation with QEs, and talking now about inflation growth w/o QE and hiking rates 2015? #AreTheyonDrugs?' Result: 1. Tweet: '#MinutePostBiz Yellen Pursuing Bernanke's Fed Policies With More Direct Style via Businessweek' Result: 1. Tweet: '#fomc #forex Yellen handling of press is impressive so far. She is much more relaxed than Dr. B was at this time' Result: 1. Tweet: '#Yellen At 6.7%, the national #unemployment rate is still higher than it ever was during the 2001 recession' Result: 0. Tweet: '#Yellen 3 takeaways from Janet Yellen's day on the Hill, Thomson Reuters Reuters Insider' Result: 0. Tweet: '#Bloomberg #News #Business Yellen Meets Reporters The One Moment That Matters - Bloomberg... #Video #Finance...' Result: 0. Tweet: '#career #money Daily Investor Briefing Yellen strongly defends easy Fed... #business #news' Result: 0.

2. Jerome Powell. Tweet: '$21 Trillion And Counting Fed Misdirection and Powell's Powder Puff Presser (Part 4)' Result: 1. Tweet: 'Investors may not like it, but the Fed is doing what it should do. And the BOJ and ECB ignore that winter is coming' Result: 1. Tweet: '@KennethLFisher on Fed Chair Jerome Powell I'm not a fan of Mr. Powell.' Result: 1. Tweet: 'Jerome Powell left more questions than answers after leading his first meeting of the Fed s policy committee @tombuerkle' Result: 1. Tweet: 'U.S. college loan debt totals more than $1.4 trillion and holds back an entire generation of Americans from pursuing their dreams. I asked Federal Reserve Chairman Jay Powell about how this affects the economy. #DebtFreeCollege' Result: 0. Tweet: 'Yesterday, Federal Reserve Chairman Jerome Powell upped the prime interest rate by 1/4%.' Result: 0. Tweet: 'US Fed goes for 25bps rate hike in Jerome Powell s 1st meeting as chairman' Result: 0.

3. Raphael Bostic. Tweet: 'Him and Bostic (a legit crook) are one of the biggest doves at the Federal Reserve. Which means their opinions should be taken with huge bags of salt (or discarded altogether)' Result: 1. Tweet: 'Don't worry, @AtlantaFed Bostic made boatloads of money in his 4 trading mistakes. F*** Americans make money is his motto.' Result: 1. Tweet: 'Is the U.S. monetary policy too tight or just right? Atlanta Federal Reserve Bank President Raphael Bostic presents a compelling case against further U.S. interest rate hikes, emphasizing that current policies are appropriately restrictive to combat inflation.' Result: 1. Tweet: 'Comments like this will keep real yields at elevated levels as the stomach of the FED to get inflation to 2% (not 2.6%) is questioned. Bostic is a dove and also in camp transitory' Result: 1. Tweet: 'Good morning. The #preciousmetals are mostly higher in early U.S. trading. U.S. calendar features nonfarm payrolls (+160k exp.), Manufacturing PMI & ISM, Construction Spending, Auto Sales. #FedSpeak due from Bostic and Mester. #gold #silver #platinum #palladium' Result: 0. Tweet: '345 PM Fed's Bostic Economic Outlook 430 PM Fed Balance Sheet 700 PM Fed's Bostic Economic Mobility' Result: 0. Tweet: 'Boston Federal Reserve President Susan Collins and Atlanta Fed chief Raphael Bostic speak today' Result: 0. Tweet: 'Tomorrow's #economic reports Thursday, Aug. 31 - Atlanta #Fed President Raphael Bostic speaks in South Africa - Initial #jobless claims - Personal

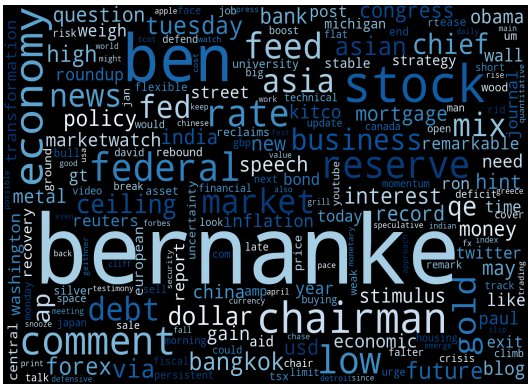
income (nominal) - Personal spending (nominal) - #PCE index - Core PCE index - PCE (year-over-year) - Core PCE (year-over-year)' Result: 0.

4. James Bullard. Tweet: 'Hahahahaha you crack me up You have to be joking. The Fed lost all their credibility a LONG time ago. Ask any working class American I mean, Bullard even laughed about it on national TV' Result: 1. Tweet: 'It's not even that he's leaving the same month, but the fact that both resigned before their term was up. Bullard was the outspoken one wanting extreme rate raises to kill inflation. Now they've over tightened and he doesn't want to be holding the grenade; be blamed' Result: 1. Tweet: 'The contamination in St. Louis has been significantly reduced by the resignation of James Bullard' Result: 1. Tweet: 'ICYMI Purdue has appointed Jim Bullard, president and CEO of , as the Dr. Samuel R. Allen Dean in #PurdueBusiness. Bullard is the longest-serving sitting president of the Federal Reserve Banks and a respected economist and scholar-leader' Result: 1. Tweet: 'Everyone made fun of Bullard, but he was right when he said rates will go to 5.75-6% as i think thats where we're headed' Result: 1. Tweet: 'James Bullard of the federal reserve bank of st Louis speaking on economic outlook' Result: 0. Tweet: 'St. Louis Fed President Bullard Offers Market Support Over Hot Jobs and Inflation' Result: 0. Tweet: 'Bullard looks back on eventful decade in Fed role' Result: 0.

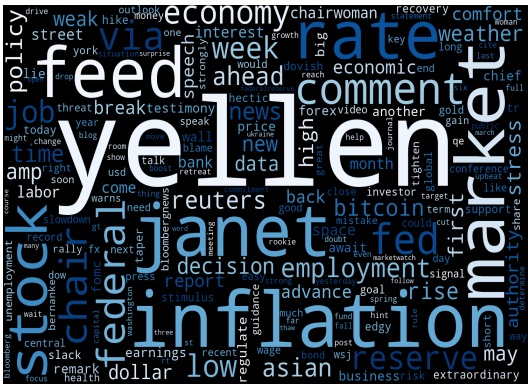
5. Eric Rosengren. Tweet: 'Want to clean up politics? Prohibit federal employees from trading securities. Not a panacea, but it would very much help.' Result: 1. Tweet: 'Fed s Kaplan, Rosengren to Sell All Stocks Amid Ethics Concerns' Result: 1. Tweet: 'BREAKING Powell Appearance of conflict of interest is unacceptable. Reminder - Powell owned $1.25-2.5M in bonds, while the fed bought $21.3B in bonds. - Rosengren owned 800K of REITs, while the Fed bought $649B - Barkin bought $3M in corporate bonds; Fed bought $46.5B' Result: 1. Tweet: 'Fed Leaders Eric Rosengren, Robert Kaplan to Resign Critics said the trading raised questions as to whether the men were setting monetary policy with the nation s well-being in mind, or whether they were acting for personal profit.' Result: 1. Tweet: 'In this @AP interview, the highly respected @BostonFed President Eric #Rosengren sets out the reasons why continued #Fed asset purchases (#QE) have few beneficial effects on the #economy and entail growing risks.' Result: 1. Tweet: '#Fed s Rosengren Sees Strong Growth Ahead for U.S. Economy - WSJ' Result: 0. Tweet: '#USA cannot afford #housing market boom and bust, warns #Fed official You don t want too much #exuberance, says Eric Rosengren amid fears over hot property sector #RealEstate #bubble' Result: 0. Tweet: 'Listen to President Rosengren s remarks from today s #RacismandtheEconomy event here' Result: 0.

Figure A.1 presents a wordcloud of the terms that feature in the tweets that received a '1' from the above process.

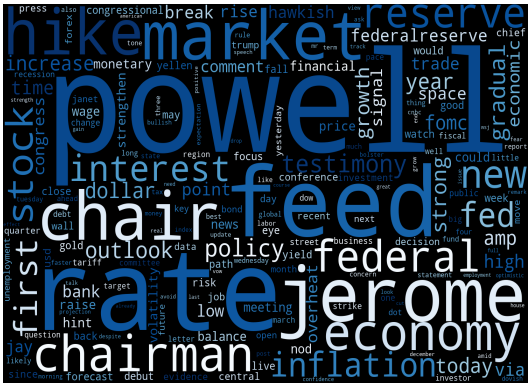
The result of this process was a set of subjective tweets about key Fed officials, expressing either confidence in or criticism of the individual in question. Finally, we produced a term histogram of these tweets, and retained those terms that we viewed as relating to the individual's perceived competence or their integrity, or terms relating to Fed policy.



(a) Ben Bernanke



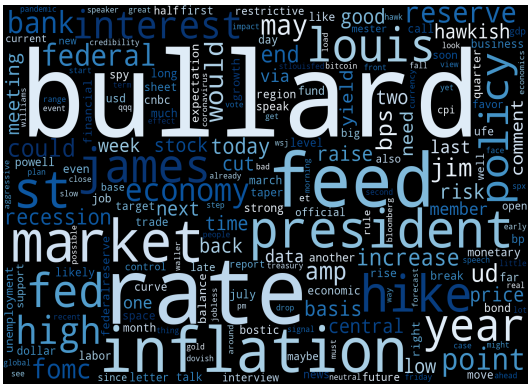
(b) Janet Yellen



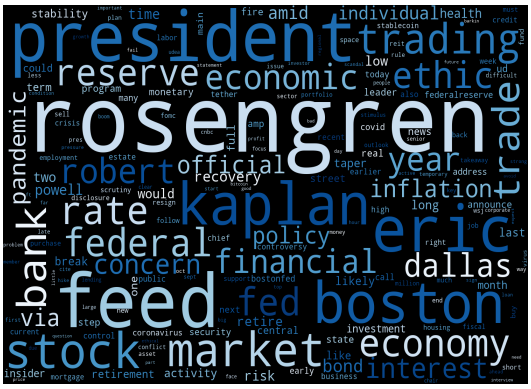
(c) Jerome Powell



(d) Raphael Bostic



(e) James Bullard



(f) Eric Rosengren

Figure A.1: Word Clouds of Identified Tweets

A.2 Approach to pre-allocating a subset of relevant tweets

As explained in the main text, to reduce processing time and aid accuracy we pre-allocate some tweets into one of our buckets {'positive','neutral','negative','unrelated'} if they contain certain words or phrases. These words and phrases are listed in Table A.2 below.

Table A.2: Words and phrases used to preallocate tweets

Category	**Phrases**	**No.**
unrelated	federal government, federal govt, fed government, fed govt, fed govnt, fed state government, fed state govt, federal supreme court, fed supreme court, federal investigation, fed investigation, federal budget deficit, fed budget deficit, federal budget, fed budget, federal election, fed election, federal jobs, jonathan h. yellen, jonathan h yellen, jonathan yellen, sheldon yellen, kayla yellen	13,782
negative	end the fed, endthefed, #endthefed, audit the fed, auditthefed, #auditthefed	42,770
positive	#popyourcollar, popyourcollar, pop your collar	93

A.3 Approaches to assessing tweet sentiment

In this section of the Appendix, we present our analysis of the performance of lexicon-based methods, supervised machine learning models, and GAI in categorising tweet sentiment.

For this exercise, we randomly downloaded 4k tweets, where the sampling was weighted towards periods of high Fed activity.[14] We also downloaded some tweets that are unrelated to the Federal Reserve—e.g., "Well he's honest! And only took the Fed 5 years to catch him. Doaps!"—to gauge the ability of the various approaches to appropriately categorise these.

Two of the authors then independently coded the 'true' sentiment in these tweets by hand, taking 2k each. This process was manually intensive and prone to errors and discrepancies and, despite the authors having almost 30 years' experience in central banking between them, there were often disagreements when we explored a cross-validation of the scores. In what follows, we used all 4k of the scored tweets, but we also considered methods' performance where we retained only 1k that received identical scores from both authors. Our main conclusions are unaffected by this, but we do find a modest deterioration in the performance of ML models given the smaller training sample.[15]

Table A.3 presents the results of this exercise, where "accuracy" is defined as the number of correctly classified tweets as a fraction of the overall number.

[14] We viewed active periods as: January 2008 to June 2009 (financial crisis), January 2013 to December 2013 (Taper Tantrum, Janet Yellen's nomination), February 2014 (Janet Yellen's term starts), November 2017 (Jerome Powell's nomination), February 2018 (Jerome Powell's term starts), March 2022 to April 2022 (COVID-19 crisis), November 2021 (Jerome Powell's reappointment), and May 2022 (Jerome Powell's second term starts). We chose 3k tweets from these periods, where each was weighted according to the number of tweets in that period relative to the total across all these active periods. We supplemented these by randomly choosing 1k tweets from non-active periods.

[15] We also considered using bootstrapping techniques to augment this dataset to 8k tweets. In particular, we adopted a back-translation approach for the retained tweets, whereby each tweet was first translated into Chinese using Google Translate and then translated back into English. This method allowed us to increase the sample size without altering the original meaning of the tweets. This approach increases the performance of all ML models, with some generating accuracy scores that exceed those of GAI. However, ... [TO ADD!]

Table A.3: Classification accuracy, using 4k manually coded tweets

Lexicon	Accuracy	Supervised ML	Accuracy	GAI	Accuracy
FS	56.07%	Logistic Regression	65.25%	GPT-4o-mini	70.05%
LM	53.63%	MLP Classifier	64.50%	GPT-4	73.74%
Opinion	44.70%	Bernoulli NB	63.88%		
AFINNLM	40.27%	Decision Tree	62.50%		
Harvard-IV	39.40%	Multinomial NB	62.25%		
VADER	38.10%	Linear SVC	61.00%		
		Ridge Classifier	60.50%		
		Random Forest	60.38%		
		Extra Tree	60.12%		
		Label Spreading	60.12%		
		Gaussian NB	50.00%		

Lexicon-based methods The lexicon-based models we considered included AFINN, Harvard-IV, Loughran-McDonald (LM), Financial Stability (FS), Opinion, and Valence Aware Dictionary for Sentiment Reasoning (VADER). These dictionaries did not generate satisfactory results in our context: as the first column of Table A.3 shows, the performance of these models is uniformly poor, with accuracy rates ranging from 56% (for the FS lexicon) to as low as 38% (for VADER). The primary limitation of these methods is that they assign scores to texts based on specific words or phrases identified as positive or negative, without considering each tweet as a whole. For example, a tweet like "The yen tumbled after the Fed's announcement" would be classified as negative by the lexicon-based model because of the word "tumbled". However, this tweet does not imply a negative sentiment toward the Fed.

Supervised machine learning models We considered the performance of various machine learning models, including the Ridge Classifier, Label Propagation, MLP Classifier, Linear SVC, Logistic Regression, Label Spreading, Multinomial NB, Extra Tree Classifier, Decision Tree Classifier, Bernoulli NB, and Gaussian NB. As Table A.3 shows, ML approaches have superior performance than lexicon-based methods in our context. But these models' accuracy remains rather poor, ranging from 65% for a

Logistic regression model to just 50% for the Gaussian NaiveBayes classifier.

One challenge for these models is identifying whether a tweet's sentiment is directed at the Fed or some other entity. We explored augmenting these models with similarity analysis (Cosine Similarity) and topic modeling (Latent Dirichlet Allocation) to identify tweets unrelated to the Fed. But these methods are time-intensive and produced poor results. Another problem in our context is that ML approaches have substantial limitations in interpreting punctuation and in recognising sarcasm. To address this, we further experiment with BERT models, but found that the results were less accurate than with GPT.

GAI The final two columns of Table A.3 report the accuracy scores of ChatGPT 4o-mini and ChatGPT 4 models. The results are higher than those of lexicon or ML models, despite these models having no access to the data used to train our ML models. On this basis, we felt that ChatGPT provided the most accurate and consistent approach to tweet classification.[16]

[16]In our experience, the web-based version of ChatGPT (ChatGPT 4) performed especially well—as a rule, cases of discrepancy with our manually scored tweets were open to interpretation. The version of ChatGPT we used via the API was ChatGPT 4o-mini, however. While this performed less well than ChatGPT 4, its performance was still very good and we used this as the cost implications of using ChatGPT 4 were prohibitive.

A.4 ChatGPT prompt

The full prompt we use is as follows:

"Please classify each of the following tweets according to whether they have a positive, neutral, or negative sentiment toward the Fed, the Federal Reserve System, or its leadership (FOMC members). If a tweet is unrelated to the Fed, please classify it as 'unrelated'. Consider each tweet independently.

Each tweet will be provided with a number. For each tweet, provide the sentiment classification as one of the following: [positive, neutral, negative, unrelated]. Format your response as follows: 'Tweet number', 'First 2 words of the tweet', 'Sentiment: [positive, neutral, negative, unrelated]'. Response format example: 'Tweet number 1' 'First 2 words' 'Sentiment: positive', 'Tweet number 2' 'First 2 words' 'Sentiment: neutral', 'Tweet number 3' 'First 2 words' 'Sentiment: negative', 'Tweet number 4' 'First 2 words' 'Sentiment: unrelated'.

Please start from tweet number X and end with tweet number X+20."

As explained in the main text, we set the "temperature" parameter—a model property that controls response determinacy—to 0, ensuring that the same prompts consistently produce identical answers.

B Appendix B

B.1 The impact of duplicate tweets on the trust index

To mitigate the impact of bots on our index, this figure reports a version of the index that excludes duplicate tweets. While this adjusted index exhibits slight differences vis-a-vis our baseline index reported in the main text, the overall trends are largely unchanged. This suggests that the influence of bots does not significantly distort the fundamental patterns captured by our measure.

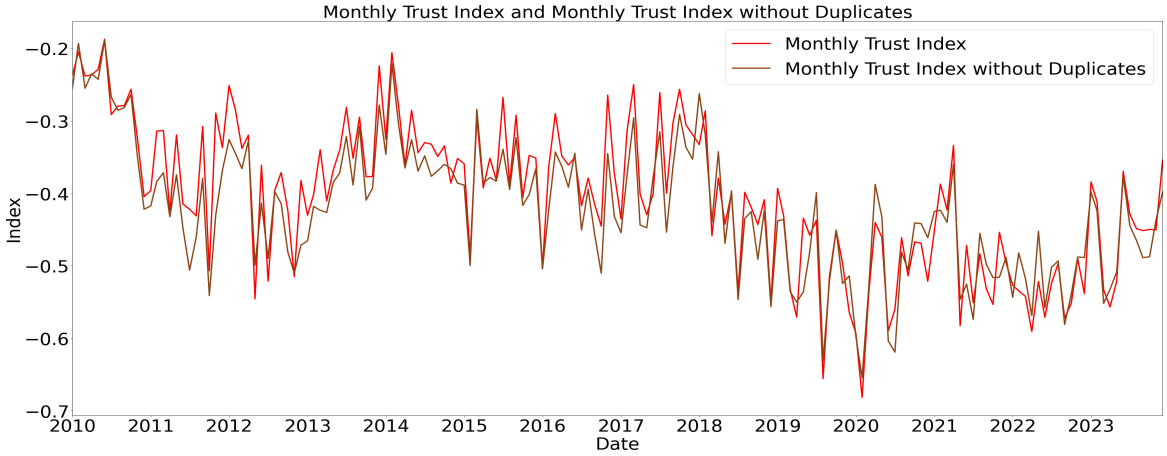


Figure B.1: 4-week moving average index with and without duplicates.

B.2 Tweet frequency before and after narrative proxy dates

Figure B.2 presents information about the number of tweets, retweets, and likes related to the tribute to Janet Yellen. The figure highlights the strong public interest and positive engagement with this Twitter trend.

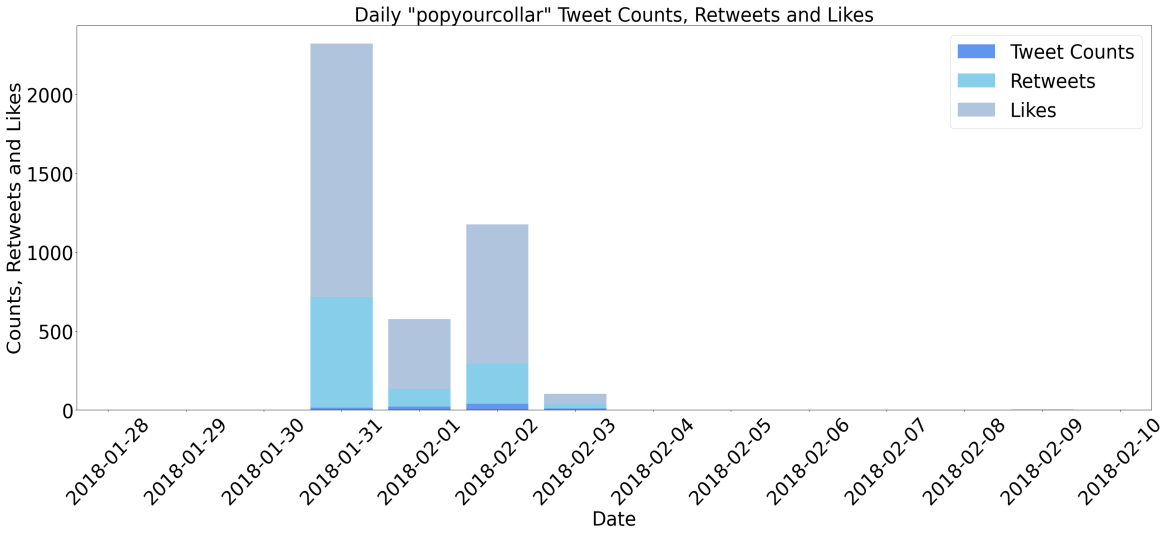


Figure B.2: Daily "#popyourcollar" tweet counts, retweets and likes

Figures B.3 to B.8 presents tweets, retweets, and likes relating to the Fed officials involved in ethical scandals. The light blue bars show the number of tweets related to Fed officials, while the dark blue bars highlight dates when the number of tweets ranked among the top 5 highest in our sample. Key dates corresponding to the exposure of Fed officials' scandals are also labeled for reference. While the number of tweets will depend on many factors including policy actions, speeches and so on, we generally find that the peak tweet count for the individuals concerned corresponds to news of their alleged involvement in an ethical scandal. There are two exceptions to this. First, while there is a pronounced spike in tweet activity surrounding Raphael Bostic after news of his first alleged scandal broke on 14 October 14 2022, his subsequent alleged scandal on 15 June 2023 did not garner the same level of attention. Second, the high number of tweets relating to James Bullard do not coincide with news breaking of his alleged scandal.

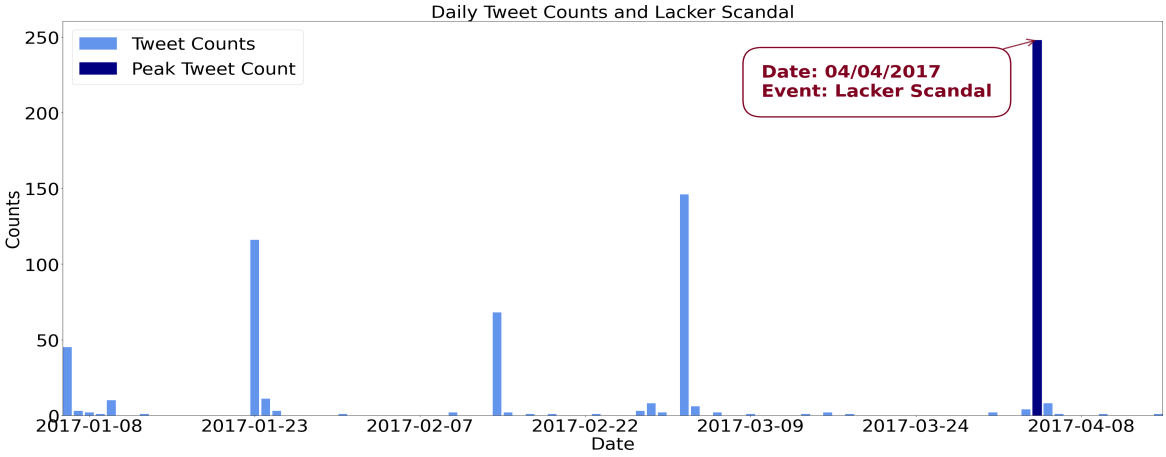


Figure B.3: Daily "Lacker" tweet counts and Lacker scandal

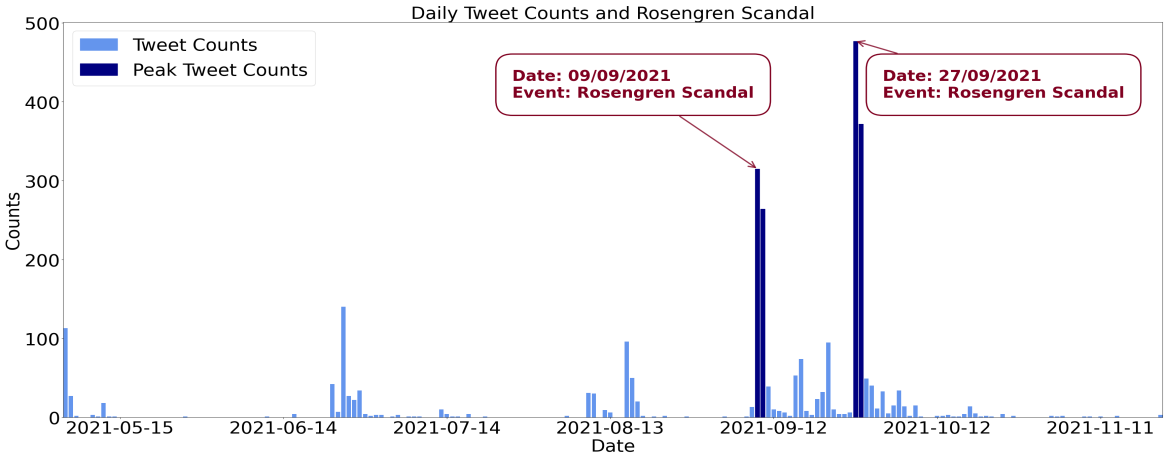


Figure B.4: Daily "Rosengren" tweet counts and Rosengren scandal

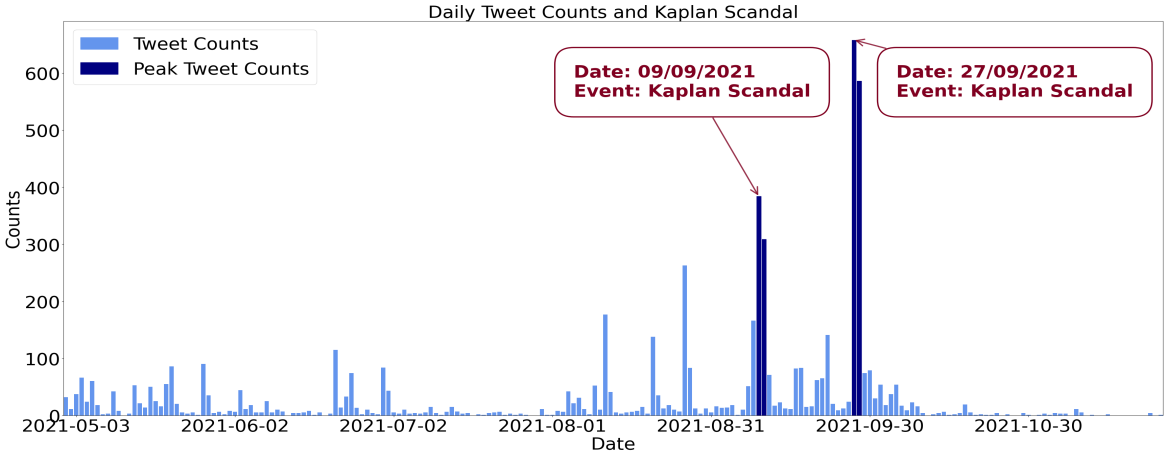


Figure B.5: Daily "Kaplan" tweet counts and Kaplan scandal

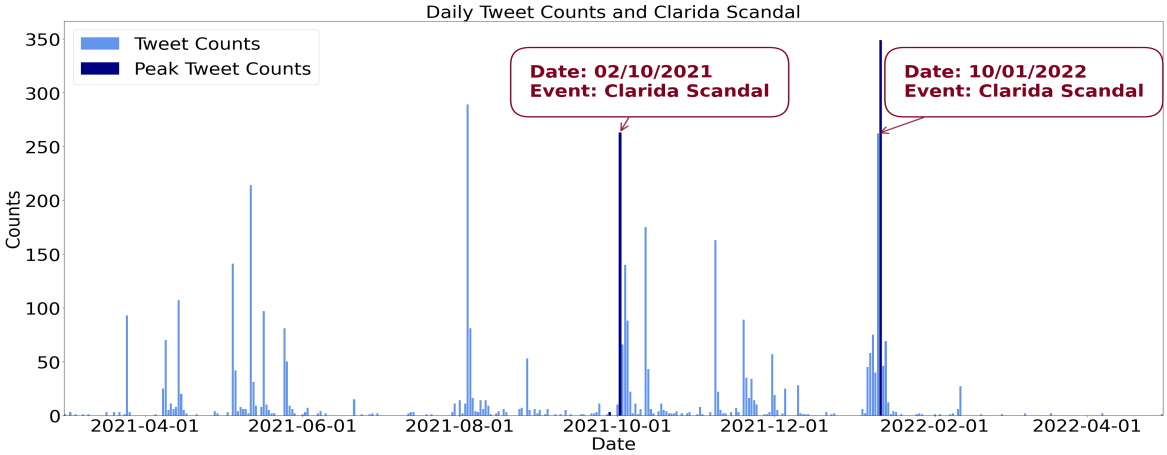


Figure B.6: Daily "Clarida" tweet counts and Clarida scandal

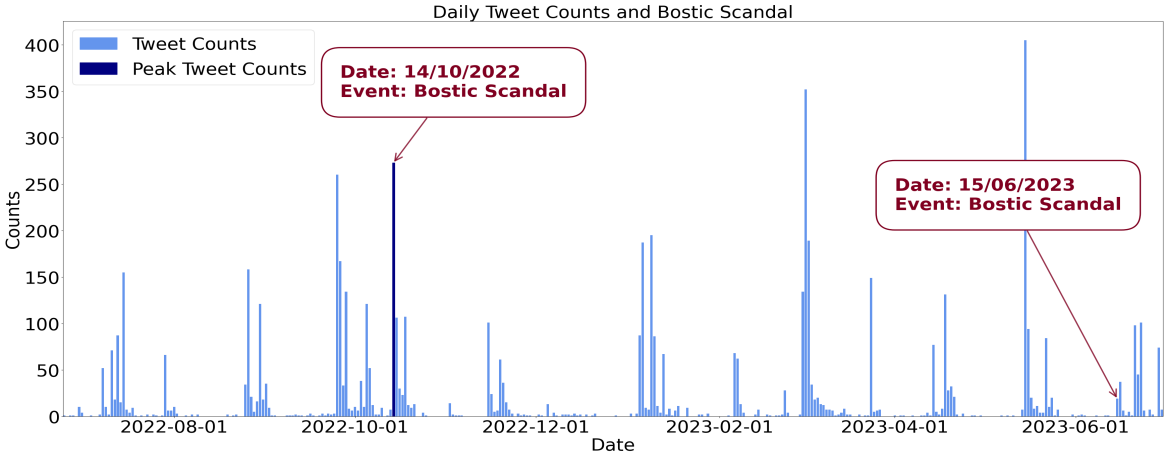


Figure B.7: Daily "Bostic" tweet counts and Bostic scandal

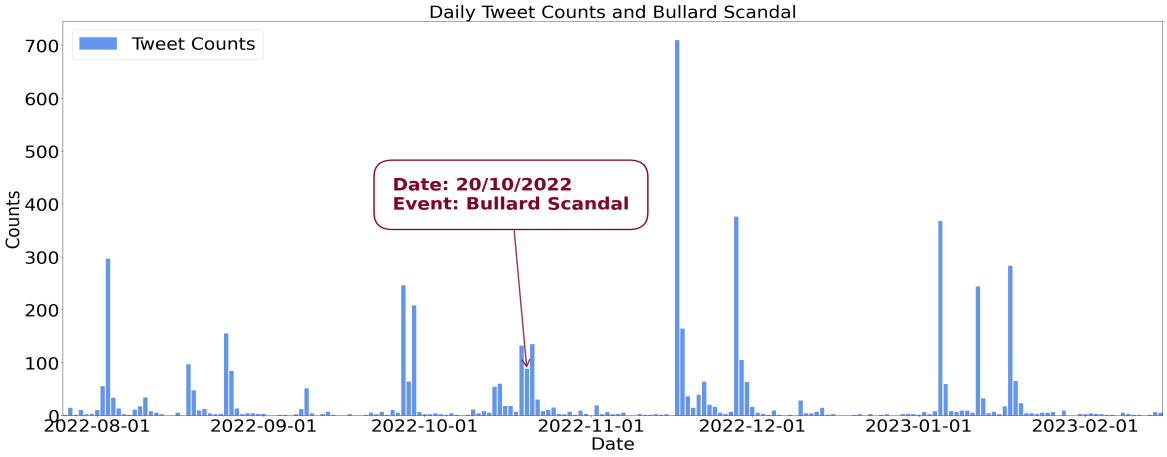


Figure B.8: Daily "Bullard" tweet counts and Bullard scandal

To better understand this phenomenon, we conduct further analysis. We hypothesise that the elevated tweet counts for influential figures such as Bostic and Bullard tend to be because of their frequent appearances in the media and their public comments on economic issues, rather than scandals. Even in the absence of controversies, their prominence and active participation in public discourse make them frequent topics of discussion. On the other hand, during periods of scandals, we expect tweets to focus more directly on the controversies.

We examine tweets during their scandal periods. The results, reported in Figure B.9, confirm our prior. During scandal periods, the tweets prominently feature scandal-related terms such as "blackout", "corruption", "culture", "disclosure", "ethics", "guidelines", "insider", "issues", "mockery", "scandal", "trouble", and "violation". However, during other periods, including those with high tweet volumes, the discussions focus on economic policies, commentary, and general market trends.



(a) Bostic scandal **(b)** Bullard scandal

Figure B.9: Word Clouds around Bostic and Bullard scandals

These findings demonstrate the significant impact of news, particularly scandals, on Twitter discourse about the Federal Reserve. News of alleged ethical scandals drew increased public attention, with the potential to alter perceptions of the Fed's credibility.

B.3 Positive and negative components of the trust index

In Figure B.10, we decompose the overall trust index into positive and negative components (i.e., $\frac{N_t^+}{N_t}$ and $-\frac{N_t^-}{N_t}$). It is clear that movements in our index by and large reflect developments in the negative component of the index. By contrast, the positive component tends to be rather stable over time, albeit there was a modest increase in the ratio of positive tweets between 2018 and 2021.

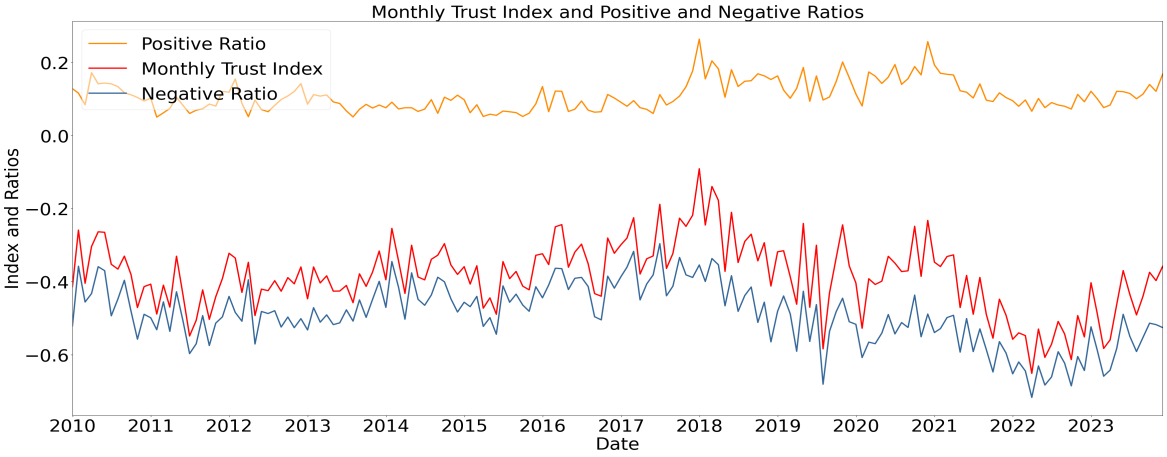


Figure B.10: Monthly Index and Percentages of Positive and Negative Tweets

In Tables B.1 and B.2, we present regression results for the drivers of the positive and negative subcomponents of the trust index. The results are largely in line with those reported in the Table **??** in the main text, but there are some interesting results. For instance, news of the release of ethical scandals results in a large increase in the negative tweet component (i.e., this component becomes more negative) with a much smaller fall in the number of positive tweets. We observe a similar phenomenon after President Trump tweets that threaten Fed independence.

Table B.1: Regression results, daily positive component of the public trust index in the Fed

Standard macro-financial variables:					
MT inflation expectations$_{t-1}$	-0.0048**	-	-	-	-0.0053**
LT inflation expectations$_{t-1}$	0.0091**	-	-	-	0.0098**
Fed Funds rate$_{t-1}$	-0.0268	-	-	-	-0.0213
30-day Fed Funds Futures rate$_{t-1}$	0.0343	-	-	-	0.0313
10-year Treasury yield$_{t-1}$	-0.0084***	-	-	-	-0.0107***
Financial conditions$_{t-1}$	-0.0022	-	-	-	-0.0014
S&P500$_{t-1}$	0.0006	-	-	-	0.0007
NASDAQ$_{t-1}$	-0.0001	-	-	-	-0.0002
VIX$_{t-1}$	0.0000	-	-	-	0.0011
US dollar$_{t-1}$	0.0026	-	-	-	0.0021
ADS Business Conditions Index$_{t-1}$	-0.0042**	-	-	-	-0.0038*
Oil price$_{t-1}$	0.0022	-	-	-	0.0023
Sentiment and policy uncertainty:					
SSW News Sentiment Index$_{t-1}$	-	0.0015	-	-	0.0033
Economic Policy Uncertainty$_{t-1}$	-	0.0033*	-	-	0.0008
Monetary policy:					
JK monetary policy shock$_{t-1}$	-	-	-0.0014	-	-0.0010
JK monetary policy information$_{t-1}$	-	-	0.0027***	-	0.0029***
Monetary policy report publication$_t$	-	-	0.0295*	-	0.0320**
Jackson Hole speech$_t$	-	-	0.0195	-	0.0249
Leadership change$_{t-1}$	-	-	0.0782**	-	0.0768**
Exogenous drivers of trust:					
Ethical scandals$_t$	-	-	-	-0.0698**	-0.0662**
Trump tweets Bianchi et al.$_t$	-	-	-	-0.0320***	-0.0389***
Other Trump tweets about Fed$_t$	-	-	-	0.0055	0.0016
Lags	Yes	Yes	Yes	Yes	Yes
α	0.1474***	0.1480***	0.1468***	0.1473***	0.1458***
Ad. R^2	0.0374	0.0340	0.0355	0.0354	0.0416
N	4,378	4,378	4,378	4,378	4,378

Notes: The dependent variable is the daily positive component of the public trust index. Continuous right-hand side variables have been standardised by subtracting the mean and dividing by the standard deviation. The regression includes 5 lags of the dependent variable. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels. Standard errors are heteroscedasticity robust (HC1).

Table B.2: Regression results, daily negative component of the public trust index in the Fed

Standard macro-financial variables:					
MT inflation expectations$_{t-1}$	-0.0140***	-	-	-	-0.0157***
LT inflation expectations$_{t-1}$	-0.0163**	-	-	-	-0.0150**
Fed Funds rate$_{t-1}$	-0.1226***	-	-	-	-0.1074***
30-day Fed Funds Futures rate$_{t-1}$	0.0800*	-	-	-	0.0702*
10-year Treasury yield$_{t-1}$	0.0500***	-	-	-	0.0454***
Financial conditions$_{t-1}$	0.0050	-	-	-	0.0046
S&P500$_{t-1}$	-0.0131**	-	-	-	-0.0130**
NASDAQ$_{t-1}$	0.0146**	-	-	-	0.0145**
VIX$_{t-1}$	-0.0179***	-	-	-	-0.0165***
US dollar$_{t-1}$	-0.0210***	-	-	-	-0.0210***
ADS Business Conditions Index$_{t-1}$	-0.0018	-	-	-	-0.0020
Oil price$_{t-1}$	0.0032	-	-	-	0.0030
Sentiment and policy uncertainty:					
SSW News Sentiment Index$_{t-1}$	-	0.0164***	-	-	0.0037
Economic Policy Uncertainty$_{t-1}$	-	-0.0086***	-	-	0.0007
Monetary policy:					
JK monetary policy shock$_{t-1}$	-	-	-0.0006	-	-0.0003
JK monetary policy information$_{t-1}$	-	-	0.0036**	-	0.0038**
Monetary policy report publication$_t$	-	-	0.1038***	-	0.1005***
Jackson Hole speech$_t$	-	-	0.0060	-	0.0777***
Leadership change$_{t-1}$	-	-	0.1447**	-	0.1277**
Exogenous drivers of trust:					
Ethical scandals$_t$	-	-	-	-0.1684***	-0.1284***
Trump tweets Bianchi et al.$_t$	-	-	-	-0.1446***	-0.1244***
Other Trump tweets about Fed$_t$	-	-	-	-0.0746**	-0.0681**
Lags	Yes	Yes	Yes	Yes	Yes
α	-0.4187***	-0.3720***	-0.3558***	-0.3565***	-0.4218***
Ad. R^2	0.1731	0.1347	0.1224	0.1288	0.1808
N	4,378	4,378	4,378	4,378	4,378

Notes: The dependent variable is the daily negative component of the public trust index. Continuous right-hand side variables have been standardised by subtracting the mean and dividing by the standard deviation. The regression includes 5 lags of the dependent variable. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels. Standard errors are heteroscedasticity robust (HC1).

B.4 Responses to a trust shock

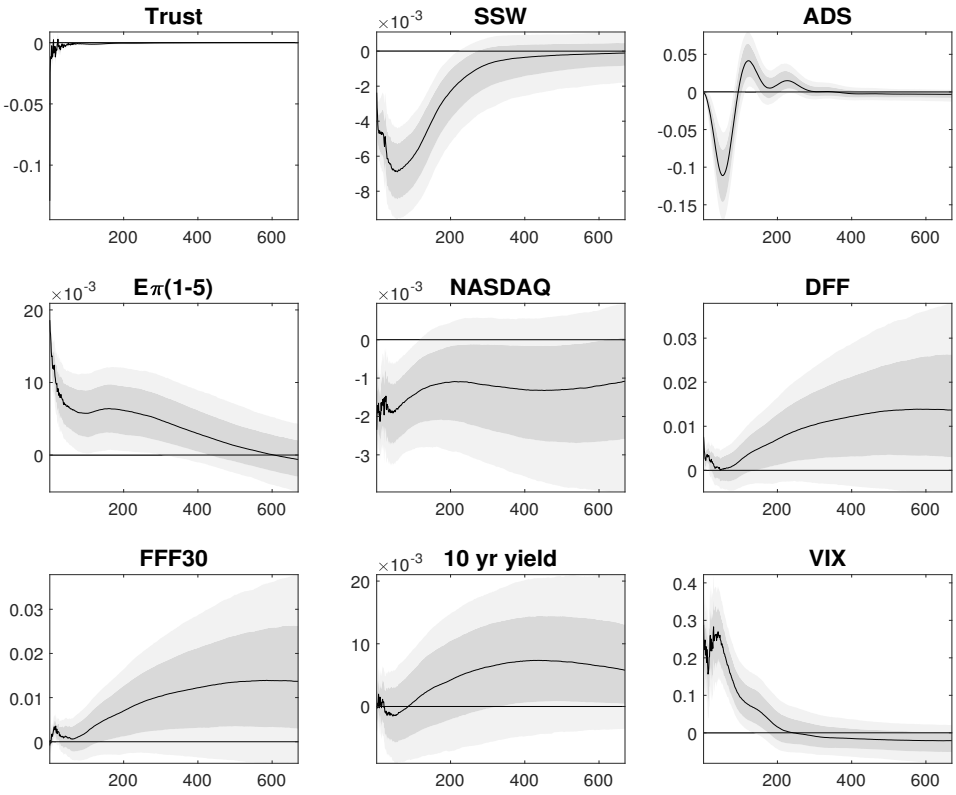


Figure B.11: Narrative proxy - 2 years: Impulse response functions to a trust shock identified using our narrative proxy (scandals + #popyourcollar) as instrument. The dark and light shaded areas correspond to the 68% and the 90% high probability density (HPD) sets, respectively.