

Louvain School of Management

Do Donald Trump's tweets reduce the market relevance of official U.S. economic announcements?

Author(s): Duthoit Jean, Mathieu Sacha
Supervisor(s): Shrestha Prabal
Academic year 2025-2026
Dissertation for the master of management
Finance
Daytime schedule

Supervisor's name : Shrestha Prabal

DECLARATION

The declaration below is **mandatory and must appear on the first page (after the front page) of your master's thesis.**

STATEMENTS ON YOUR USE OF GENERATIVE AI

Confirm the following statements:

Statements on responsible use of generative AI	Your answer
The content of our master's thesis is our original ~~output~~ , even if we used generative AI to help prepare it.	☒ Yes ☐ No
We master the theories, tools and methods that we use for the thesis.	☒ Yes ☐ No
We verified that the output presented in the thesis is accurate and valid, and that it contains no hallucination, no false reference, etc.	☒ Yes ☐ No
We acknowledge that we master the final output and that we are able to explain it and answer questions about it.	☒ Yes ☐ No
We made sure not to provide generative AI tools with access to confidential data or information	☒ Yes ☐ No

SIGNATURE

Sign your declaration:

Author last name, first name(s)	Date	Signature
1 Mathieu, Sacha	26/05/2026	
2 Duthoit, Jean	26/05/2026	

Acknowledgements:

We would like to thank all the people who helped us during the writing of this thesis.

We would first like to thank our thesis supervisor, Mr. Shrestha, for his exceptional availability, patience and above all his wise advice, which helped to enrich our thesis.

We would also like to thank our parents, Marie and Fabrice, Aline and Jean-Yves, as well as Elise and Aurore for their support and encouragement which were of great help to us throughout our work.

Table of contents:

Introduction	- 1 -
Literature review.....	- 3 -
I. The Role of Political Leaders on Financial Markets: From General Influence to Trump's Disruption.....	- 3 -
<i>1.1. Political Leaders as Market Actors: A General Framework.....</i>	<i>- 3 -</i>
<i>1.2. Historical Evolution of American Presidential Communication: From Institutionalization to Disintermediation</i>	<i>- 4 -</i>
<i>1.2.1. The Institutional Era and the Genesis of Economic Influence (1933-1980)</i>	<i>- 4 -</i>
<i>1.2.2. The Strategy of 'Going Public' and Bypassing Intermediaries (1980-2008).....</i>	<i>- 5 -</i>
<i>1.2.3. The Era of Radical Disintermediation and the 'Liberal-Individualist' Style (2008-Present)</i>	<i>- 5 -</i>
II. Behavioral Mechanisms and the Amplification of Presidential Noise	- 5 -
<i>2.1. Noise Traders and the Attention Economy</i>	<i>- 6 -</i>
<i>2.2. Trump as the Ultimate Noise-Maker: The Volfefe Index</i>	<i>- 7 -</i>
<i>2.3. Cryptocurrency Markets as a Structural Amplifier of Political Noise</i>	<i>- 8 -</i>
<i>2.4. Sentiment Asymmetry and the Negativity Channel.....</i>	<i>- 9 -</i>
<i>H3: The Sentiment Asymmetry Effect</i>	<i>- 10 -</i>
<i>2.5. Digital Speed, Information Asymmetry, and the Crowding-Out of Official Signals....</i>	<i>- 11 -</i>
<i>H1: The Direct Volatility Effect.....</i>	<i>- 13 -</i>
III. The Institutional Pillar: Monetary Policy and Market Expectations.....	- 14 -
<i>3.1. Official Communication as the Basis for Rational Expectations</i>	<i>- 14 -</i>
<i>3.2. The Breakdown of 'Neutrality' Through Political Interference</i>	<i>- 15 -</i>
<i>H2: The Crowding-Out Effect.....</i>	<i>- 16 -</i>
Methodology.....	- 17 -
I. Data Description	- 17 -
II. Variable Description	- 19 -
<i>2.1 Dependent Variables</i>	<i>- 19 -</i>
<i>2.2 Independent Variables</i>	<i>- 20 -</i>
<i>2.3 Control Variables</i>	<i>- 21 -</i>
III. H1: Tweet Volume and Conditional Volatility: GJR-GARCH(1,1).....	- 22 -
IV. H2: The Crowding-Out Effect: Event Study and OLS	- 25 -
V. H3: Negative Tweets and Asymmetric Volatility: GJR-GARCH(1,1).....	- 27 -
VI. Robustness Tests	- 29 -
Results	- 33 -
I. Preliminary Statistics and Diagnostics	- 33 -

II.	H1: Tweet Volume and Conditional Volatility	- 33 -
III.	H2: The Crowding-Out Effect.....	- 34 -
	3.1 <i>Timing Check</i>	<i>- 34 -</i>
	3.2 <i>Main Results</i>	<i>- 35 -</i>
IV.	H3: Sentiment Asymmetry and Conditional Volatility	- 35 -
V.	Robustness Checks	- 36 -
	R1: <i>COVID-19 and Structural Stability (see Appendix, Table 7)</i>	<i>- 36 -</i>
	R2: <i>Subsample by Presidential Term (see Appendix, Table 7)</i>	<i>- 36 -</i>
	R3: <i>Alternative Pre-Announcement Windows (see Appendix, Table 7)</i>	<i>- 37 -</i>
	R4: <i>FIGARCH Alternative Model (see Appendix, Table 7)</i>	<i>- 37 -</i>
	R5: <i>Financially Filtered Tweets (see Appendix, Table 8-10)</i>	<i>- 37 -</i>
Discussion		- 39 -
I.	The Volume Channel	- 39 -
II.	The Crowding-Out Effect	- 40 -
III.	Why Sentiment Does Not Explain Variance	- 41 -
IV.	Putting It Together	- 43 -
V.	An Open Question	- 44 -
Conclusion		- 46 -
Bibliography.....		- 51 -
Appendix.....		- 55 -

Introduction

On a Sunday morning in May 2019, Donald Trump sent some tweets from his personal account. He announced that tariffs on 200 billion dollars of Chinese goods would rise from 10% to 25% within the week. There had been no warning. Negotiations with China were in progress, the month before, the Secretary of the Treasury said “both sides have a desire to reach an agreement”. Markets were closed. And yet, just after the publications, Dow Jones futures fell by more than 500 points. When trading opened on Monday, the Dow lost 450 points at the bell, the S&P 500 dropped 1.55%, and Chinese markets fell 6% (Bayly, 2019). No earnings release, no Federal Reserve decision, no macroeconomic data had moved. A tweet had.

That episode is not an isolated anecdote. It is a symptom of something more structural. Over the course of his two presidencies, Donald Trump used social media at a scale and with a directness that had no precedent in the history of executive communication. From 2016 to 2025, he posted tens of thousands of original messages on Twitter/X, and Truth Social. Many of his messages touch directly on markets, the Federal Reserve, trade policy, and the economy in general. Financial markets are designed to process information efficiently and rapidly, but have no established framework for dealing with this kind of communication channel. The existing literature documented that Trump’s tweets moved prices. However, we do not know if they did so in ways that systematically disrupted the institutional architecture through which markets normally form expectations.

It was a different kind of source that first raised this question for us. In an interview conducted by M. Lewis in 2018, Steve Bannon, one of Trump's former chief strategists, described the administration's communication approach. The strategy was to deliberately “flood the zone”. The idea is to saturate the informational space so thoroughly that no coherent counter-narrative could form (Lewis, 2018). Presidential communication may have been used as a deliberate instrument of informational disruption. If so, its effects on financial markets should be both measurable and systematic. This thesis is an attempt to measure them.

One question runs through the entire analysis. Can the social media communication of a political leader alter the informational architecture through which financial markets form expectations? And if so, through what mechanism? The central empirical expression of that

question is more specific: do Donald Trump's tweets reduce the market relevance of official U.S. economic announcements? That is the question this thesis is designed to answer.

Three hypotheses structure the empirical analysis. First of all, H1 asks whether Trump's tweet volume generates an independent and measurable increase in conditional market volatility. And this, above and beyond what the standard GARCH dynamics and macroeconomic controls already explain. Then, H2 asks whether elevated tweet activity in the days before a scheduled macroeconomic announcement reduces the market-moving capacity of that announcement. We are looking for a crowding-out effect consistent with the attention-driven buying framework of Barber and Odean (2008). And finally, H3 asks whether negatively framed tweets generate larger volatility responses than positive ones. It is what Prospect Theory would predict given the well-documented asymmetry in how investors process potential losses relative to equivalent gains (Kahneman & Tversky, 1979).

To test these three hypotheses, the analysis draws on a dataset covering the full 2016 to 2025 period, taking both Trump presidencies and the inter-term gap during which he was largely banned from major social media platforms (DataDrivenDecisionLab, 2026). The primary financial variables are daily log returns on the S&P 500 and daily changes in two-year and ten-year U.S. Treasury yields. Two hypotheses, H1 and H3, are estimated using GJR-GARCH(1,1) models on announcement-free trading days. These models allow the tweet variable to be identified against a clean volatility background. The remaining hypothesis, H2, is estimated through an event study framework covering 252 scheduled macroeconomic releases across three announcement types, FOMC decisions, Non-Farm Payroll reports, and CPI releases, using an OLS regression with HC3 robust standard errors. Five sets of robustness checks were done. A COVID structural break test, subsample analysis by presidential term, alternative pre-announcement windows, a FIGARCH model comparison, and a re-estimation using a keyword-filtered tweet variable restricted to financially relevant posts.

To conclude this introduction, the thesis is organized as follows. Next section develops the theoretical and empirical literature underpinning the three hypotheses, covering noise trading, the attention economy, sentiment asymmetry, and the institutional framework of central bank communication. Section “Methodology” presents the data and variable construction, details the empirical specifications for H1, H2, and H3 respectively and describes the robustness tests. After this, the Results, the Discussion and the Conclusion follow.

Literature review

I. The Role of Political Leaders on Financial Markets: From General Influence to Trump's Disruption

1.1. Political Leaders as Market Actors: A General Framework

In the academic literature, the relationship between political leadership and financial market behavior is well documented. It has long been established that presidential communication, electoral outcomes, and executive decisions form independent pricing factors. For instance, markets price a political risk premium generated by the policy environment itself (Pástor and Veronesi, 2010, 2013). Also, Montone (2022) shows that presidential approval ratings predict subsequent monthly stock returns. He demonstrates that political evaluations affect asset prices throughout presidential terms rather than only around election dates. For their part, Ficcadenti and Cerqueti (2024) show that the linguistic content of presidential speeches systematically influences trading volume and market activity. These findings suggest that political leadership is a structural input into the expectations through which asset prices are formed. Ficcadenti and Cerqueti's (2024) findings suggest that speech content and communication style are not merely rhetorical, they shape the informational environment in which investors operate. Thus, the perceived reliability and institutional posture of an executive become pricing inputs in their own right.

The Donald Trump election has put this relationship in the spotlight as never before. Studies show that Trump's candidacy, perceived as ideologically radical and behaviorally unpredictable, generated significant policy uncertainty and sector-specific market reactions around the election. Positive abnormal returns were observed in the days following the vote for specific sectors (Bouoiyour and Selmi, 2016). Beyond the electoral cycle, political announcements have an impact on the market, showing that political uncertainty can deeply affect market segments. Moreover, Kandi et al. (2021) study cases where presidential tweets are followed by sharp market reactions, before a gradual return to more stable price levels. Meanwhile, Born et al. (2017) find evidence that company-specific tweets produce sharp price movements accompanied by surges in Google search volume for the mentioned firms, a behavioral signature consistent with retail attention-driven trading. Salem et al. (2019) introduced J.P. Morgan's Volfefe Index¹. They analyzed it and discovered that a robust relationship exists, at a fixed income level, between Trump's posting activity and implied

¹ See section 2.2 of the literature review

volatility in two- and five-year Treasury instruments. This is a market segment conventionally governed by macroeconomic fundamentals, not social media. Furthermore, Klaus and Koser (2021) demonstrate that this effect extends to major European equity indices including the DAX and CAC 40. More recently, Krause (2025) documents that Trump's advocacy for cryptocurrency generates significant excess returns in Bitcoin markets even when devoid of concrete policy content. While, on their part, Ben Jabeur et al. (2025) confirm that political signals predict cryptocurrency returns. Together, this evidence suggests that digital asset markets are among the most sensitive segments of the financial system to presidential communication.

What the existing literature has established, therefore, is that Trump's communication moves markets. What it has not directly addressed is whether that communication systematically attenuates the capacity of formal institutional signals to anchor market expectations. Furthermore, it remains unknown whether the emotional valence of his messages produces asymmetric effects across different market conditions. These are the questions this thesis is designed to test.

1.2. Historical Evolution of American Presidential Communication: From Institutionalization to Disintermediation

A presidential tweet can shift the market capitalization within minutes. This is the result of a long transformation in how executive power communicates with the public and, by extension, with financial markets. For our study, a brief historical arc matters to understand Trump's communication style. And why his communication approach is not merely a quantitative intensification of what came before, but rather a qualitative break. Three distinct phases characterize this evolution.

1.2.1. The Institutional Era and the Genesis of Economic Influence (1933-1980)

The modern relationship between presidential communication and economic expectations took shape during the Roosevelt administration. Confronted with a collapse in consumer confidence during the Great Depression, Roosevelt deployed radio broadcasts, the 'Fireside Chats'. This was a direct channel to households, bypassing the institutional intermediaries that had previously filtered executive communication. In his study, Pedemonte (2019) demonstrates that these broadcasts stimulated consumer spending. They shape household expectations in a direction consistent with economic recovery. This historical precedent establishes that presidential communication alone, without any immediate policy implementation, can produce

measurable macroeconomic effects by shaping how economic agents form their expectations. Indeed, the executive branch of this era was constrained by a stable, procedure-bound bureaucratic apparatus. Presidential speech supported institutions, it did not seek to bypass them.

1.2.2. The Strategy of 'Going Public' and Bypassing Intermediaries (1980-2008)

The first significant rupture in this institutional logic occurred during the Reagan presidency. Kernell (2006) theorized this shift under the rubric of 'Going Public'. The idea is that rather than resolving legislative disputes through traditional bargaining with Congress, presidents began appealing directly to popular opinion as a source of political leverage. Traditional media continued to serve as gatekeepers during this period. Journalists still filtered, contextualized, and sometimes contested presidential statements before they reached mass audiences. The presidential narrative had gained partial autonomy but remained embedded in the broader institutional framework of journalism and regulated broadcasting.

1.2.3. The Era of Radical Disintermediation and the 'Liberal-Individualist' Style (2008-Present)

Social media eliminated the final institutional filter. Obama institutionalized the use of digital platforms as part of a coherent communication strategy, but Trump transformed them into something qualitatively different. Scacco et al. (2020) capture this distinction with precision. Traditional presidential communication is 'deliberative', measured, and institutionally embedded through the work of communication professionals. By contrast, Trump's style is characterized as 'liberal-individualistic', which is spontaneous, frequent, emotionally raw, and deliberately personal. Applied to financial markets, the structural consequence is a saturation of the media and informational space. This concept is a continuous stream of presidential statements that markets can neither ignore nor fully process. Its speed and volume raise the question of whether they are capable of displacing the institutional signals through which markets have historically formed stable expectations.

II. Behavioral Mechanisms and the Amplification of Presidential Noise

Standard financial theory provides a useful but ultimately insufficient benchmark for understanding why markets respond systematically to presidential statements that carry no verifiable policy content. Fama's (1970) Efficient Market Hypothesis holds that prices fully reflect all available information, such that no trading strategy based on publicly available

information should yield systematic abnormal returns. Rational investors should therefore adjust portfolios only when announcements generate genuinely new information about fundamentals. Similarly, Pástor and Veronesi (2010, 2013) establish that uncertainty about government policy is priced into equity markets. However, their framework does not account for the extreme speed and scale at which direct social media communication moves markets, independent of any formal policy decision. Accounting for these anomalies requires the conceptual apparatus of behavioral finance. Shiller (2003) and Kahneman and Tversky (1979) have documented that financial agents are systematically subject to cognitive biases that distort their information processing. It makes them susceptible to signals that a fully rational agent would discard as noise. The disruption operates through three compounding channels: the behavioral profile of noise traders and retail investors, the asymmetric weight that markets assign to negative sentiment, and the digital speed advantage that allows presidential communication to pre-empt institutional information channels. These three channels converge on a common outcome which is the crowding-out of official economic signals.

2.1. Noise Traders and the Attention Economy

The theoretical distinction between investor types was described by De Long et al. (1990). They distinguished two categories of market participants. Rational arbitrage traders identify and exploit mispricings on the basis of fundamental value. By contrast, noise traders treat signals without fundamental relevance as though they carry pricing information. And they act on them in ways that generate cumulative price distortions (Black, 1986). The key insight of De Long et al. (1990) is that noise traders are not simply irrational agents whose behavior is immediately corrected by the market. When sufficiently numerous and subject to a common sentiment shock, they impose genuine risk on rational counterparties. This makes arbitrage costly and therefore incomplete. This behavioral profile is precisely embodied by investors who adjust their portfolios in response to a presidential tweet that contains no concrete policy commitment.

In a complementary psychological account, Barber and Odean (2008) explain why retail investors are disproportionately susceptible to this kind of noise. Institutional actors survey the market systematically and across a broad asset universe. But unlike them, individual investors operate under severe attentional constraints. They focus on assets and events that achieve high media salience. They do not check whether that salience reflects informational significance. A presidential tweet that targets a company, sector, or asset class by name functions as an attention

shock. This tweet suddenly concentrates collective retail focus on a specific investment object. This generates a temporary price pressure, an effect highly consistent with Barber and Odean's (2008) attention-driven buying hypothesis. This mechanism is amplified by the architecture of social media platforms. It algorithmically favors emotionally charged and controversy-generating content. Thus, the presidential noise reaches retail investors faster and with greater force than technically measured institutional communication.

2.2. Trump as the Ultimate Noise-Maker: The Volfefe Index

The Trump presidency prompted the financial industry to develop new measurement tools, precisely because existing frameworks proved inadequate to capture the scale and nature of the disruption. In 2019, J.P. Morgan (Salem et al., 2019) constructed the 'Volfefe Index', designed specifically to measure the effect of Trump's tweet activity on fixed income market volatility. The choice of fixed income instruments is theoretically significant. Treasury bonds are conventionally governed by macroeconomic fundamentals and monetary policy expectations, not social media activity. Yet this index revealed a robust link. Indeed, a relationship between Trump's posting behavior and implied volatility in two- and five-year maturity instruments may be observed. Salem et al., creators of this index, analyzed it in detail and further identified a lexical trigger set. These are terms like 'China', 'Billion', 'Products', and they produce measurably stronger market responses. This empirical observation is highly consistent with the hypothesis that algorithmic trading systems have effectively encoded the presidential vocabulary as a pricing input.

Klaus and Koser (2021) extended the analysis geographically. They demonstrated that the Volfefe Index carries predictive power for returns on major European equity indices including the DAX and CAC 40. Therefore, the implication is stark. Presidential communication generates market disruptions that transmit beyond U.S. borders. It is a global phenomenon transmitted through the interconnected architecture of international capital markets. The presidential noise is powerful enough to displace local economic signals in jurisdictions where Trump has no direct political authority. For their part, Born et al. (2017) provide a complementary empirical test of the semi-strong form of the EMH, by using Trump's company-specific tweets. Following these tweets, the price movements coincide with sharp increases in Google search volume for the mentioned companies. This constitutes a behavioral signature consistent with retail noise trading. Furthermore, the price effects prove to be transitory. This effect dissipates within approximately five trading days, the time it takes for

the market to correct toward fundamental value. These disruptions are short, intense, and self-correcting. However, their systematic character, their predictability as a function of tweet content and volume, constitutes the empirical basis for the hypotheses developed later in this review.

2.3. Cryptocurrency Markets as a Structural Amplifier of Political Noise

The structural contrast between equity markets and cryptocurrency markets illuminates the range of noise vulnerability across different asset classes. Traditional equity markets benefit from powerful institutional moderating forces, despite their susceptibility to attention-driven distortions. These stabilization mechanisms include mandatory disclosure requirements, earnings-based valuation anchors, the disciplining presence of long-horizon institutional investors, and arbitrage mechanisms capable of correcting short-term mispricings. Cryptocurrency markets possess none of these structural moderators. In their study, Glaser et al. (2014) show that Bitcoin markets attract predominantly retail users who behave speculatively rather than as payment-system participants. And according to them, no standard fundamental pricing methodology exists for this asset class. Building on this, the structural properties of cryptocurrency markets amplify these dynamics further. There is no earnings calendar, no dividend anchor, and no fundamental valuation framework to which prices must eventually converge. Regulatory oversight remains jurisdiction-dependent, and markets operate continuously without the circuit breakers that equity exchanges deploy during periods of acute stress. These characteristics, documented in part by Glaser et al. and extended by the broader market structure literature, create conditions in which noise trading is not merely possible but structurally unconstrained.

If there are no economic fundamentals, then what drives the price? Maybe we can view cryptocurrencies as fundamentally narrative-driven assets. The collective belief, shaped and reshaped by influential voices, becomes the primary determinant of value. This structural reliance on narrative makes the asset class disproportionately sensitive to political communication. Specifically, it responds to the kind of emotionally charged, high-frequency signals that characterize Trump's social media presence. Moreover, Baur et al. (2018) classify cryptocurrencies as highly speculative assets that lack the conventional anchors of equity valuation. Consequently, Shiller's (2017) framework of narrative economics extends naturally to cryptocurrencies with particular force. He developed it to explain how popular stories drive economic fluctuations broadly. Indeed, the value of cryptocurrencies relies structurally on

collective belief as its primary driver. This dynamic finds a concrete application in Krause's work (2025). He shows that several of Trump's expressions of support for the cryptocurrency sector generated statistically significant excess returns in Bitcoin. At a broader level, Ben Jabeur et al. (2025) confirm that cryptocurrency markets react to political narrative signals. They found that the quality of these signals actually predicts returns for certain assets and horizons better than standard macroeconomic uncertainty proxies. These findings suggest that digital asset markets are among the most responsive segments of the financial system to presidential communication. This vulnerability is compounded by the algorithmic responsiveness of trading platforms and the 24-hour liquidity structure of crypto markets. As a result, all of this ensures that political signals are priced without delay regardless of the hour at which they arrive.

2.4. Sentiment Asymmetry and the Negativity Channel

If noise trading establishes that presidential communication moves markets, the sentiment channel addresses a further and more subtle dimension. The question here is whether the emotional valence of that communication generates systematically different outcomes depending on whether the framing is negative, positive, or neutral. The behavioral literature provides a clear theoretical prediction grounded in loss aversion.

In 2007, Tetlock established the foundational empirical link between linguistic tone and market behavior. He demonstrates that high negative sentiment in published media is associated with subsequent downward pressure on equity prices and increased trading volume. The causal interpretation is consistent with the behavioral framework established above. Indeed, negative language activates risk aversion in investors, prompting them to sell or reduce exposure before fully processing the underlying informational content. The linguistic tone is not merely reflective of conditions, it is partly constitutive of them.

Kahneman and Tversky's (1979) Prospect Theory provides the core theoretical mechanism. The subjective experience of a potential loss is approximately twice as intense as the experience of an equivalent gain. Investors do not evaluate outcomes on a symmetric utility scale. They are loss averse. They assign a disproportionate decision weight to messages that can be framed in terms of potential damage or threat. This asymmetry is formally reinforced, a few years later, by Tversky and Kahneman (1992). They extended the Prospect Theory and demonstrated that the value function is systematically steeper in the loss domain than in the gain domain. While Prospect Theory outlines this cognitive bias in general decision-making, Benartzi and Thaler (1995) provide direct empirical evidence of its impact on financial markets:

the myopic loss aversion² leads investors to demand historically anomalous return premia. Negativity bias and negativity dominance (Baumeister et al., 2001; Rozin & Royzman, 2001) further confirm that negative stimuli systematically command greater attentional resources. They also produce larger behavioral responses across cognitive domains.

Applied to presidential communication, this literature generates a clear prediction. The negatively framed tweets should activate the loss-aversion mechanism more powerfully than tweets of equivalent prominence framed positively or neutrally. Today's trading environments rely heavily on sentiment-detection algorithms calibrated to respond to linguistic cues within milliseconds. In such a high-frequency setting, this cognitive asymmetry becomes structurally amplified. Bollen et al. (2011) demonstrate that aggregate mood expressed on Twitter is a statistically significant predictor of subsequent Dow Jones movements. Of the mood dimensions tested, public calmness proves the strongest predictor of DJIA movements, with effects materializing three to four days after the initial signal. Trump's use of Twitter as his primary communication channel directly inserts his statements into this feedback loop between political rhetoric, collective mood, and asset prices. The ultimate impact of this loop depends critically on the speed at which his messages arrive.

H3: The Sentiment Asymmetry Effect

H3 introduces the qualitative dimension that Section 2.4 prepared. If H1 establishes that the volume of presidential communication matters, H3 asks whether the emotional valence of that volume matters too, and grounds this question in a precise theoretical prediction. As previously noted, Kahneman and Tversky's (1979) Prospect Theory establishes that investors are loss averse. This framework predicts that threatening, critical, or catastrophic tweets activate the loss-aversion mechanism more powerfully than positive or neutral ones. This holds true whether the messages are attacking the Fed, warning of economic disaster, or characterizing policy outcomes as failures.

This theoretical prediction is reinforced by the psychological literature on negativity bias and negativity dominance (Baumeister et al., 2001; Rozin & Royzman, 2001). This study shows that negative stimuli systematically command greater attentional resources. These negative stimuli produce larger behavioral responses across a wide range of cognitive domains. In high-

² *The tendency to evaluate outcomes over short horizons in ways that amplify the subjective weight of negative returns (Benartzi and Thaler, 1995)*

frequency trading environments, algorithmic systems are calibrated to detect and respond to linguistic sentiment. Consequently, this cognitive asymmetry translates into larger and faster order flows whenever negative signals appear. In line with this observation, Tetlock (2007) demonstrated the predictive power of media pessimism for equity prices. Trump's emotionally raw communication channel represents an amplified version of the same mechanism, with the filtering effects of journalistic mediation removed. The sentiment asymmetry framework developed in Section 2.4, encompassing Prospect Theory, myopic loss aversion (Benartzi & Thaler, 1995), and the negativity dominance literature, provides the full theoretical grounding for this hypothesis.

H3: Presidential tweets classified as negative in tone generate a larger and more persistent increase in market volatility than tweets classified as positive or neutral in tone, reflecting the asymmetric behavioral weight that investors assign to potential losses relative to equivalent gains.

2.5. Digital Speed, Information Asymmetry, and the Crowding-Out of Official Signals

The noise trading and sentiment channels described above operate through behavioral mechanisms. This final analytical dimension adds a structural complement. We will explore the concept of information asymmetry, and specifically the speed-based asymmetry introduced by the Trump communication model. The argument is that digital speed is not merely a feature of presidential communication, it is the mechanism through which behavioral disruption translates into institutional displacement.

The theoretical foundations of information asymmetry in economic analysis trace to Akerlof's (1970) canonical account of adverse selection. This concept has been extended to financial markets by Stiglitz and Weiss (1981). They demonstrated that even competitive markets fail to achieve efficient equilibria when information is asymmetrically distributed. Within this framework, official economic institutions occupy a specific and normatively significant role. The Federal Reserve and comparable agencies are designed to function as authoritative, symmetric information providers. They achieve this by releasing standardized, audited, and predictable data that all market participants access simultaneously. All these mechanisms minimize the informational advantage of any single actor and enable coordinated

expectation formation (Goldstein & Yang, 2017; Healy & Palepu, 2001). Their institutional design is, at its core, a structural response to the market failure that asymmetric information produces.

However, the emergence of social media has introduced a form of information asymmetry that this framework was not designed to address: digital speed asymmetry. Institutional communication is inherently deliberative. The preparation, internal review, and timed release of an FOMC statement or a labor market report follow protocols engineered to maximize signal quality and predictability. All this makes them time-consuming by design. By contrast, a presidential tweet is instantaneous, unreviewed, and arrives without any institutional filter. Modern market environments are dominated by high-frequency trading systems calibrated to respond to linguistic signals within milliseconds. In this setting, the speed differential creates a *de facto* informational advantage for the first signal to arrive, regardless of its quality. Kitvanitphasu et al. (2026) characterize this dynamic as ‘order flow toxicity’. The arrival of asymmetric, unverified information triggers a cascade of leveraged trading. This mechanism exploits the slower response times of more deliberative, rational agents. Similarly, Krause (2025) documents that Bitcoin generates statistically significant excess returns around key presidential advocacy events. This pattern is consistent with a mechanism in which political signals are incorporated into prices before standard disclosure-based information can be processed. This effectively widens the informational gap between participants capable of acting on political signals in real time and those who cannot.

Digital speed asymmetry does not merely add a competing signal to the information environment. It also actively degrades the informational value of official announcements. Presidential statements often move prices and reposition portfolios well before a scheduled data release. Consequently, when the formal announcement finally arrives, the market has already partially priced in the relevant news. Its capacity to surprise, to generate the abnormal return that reflects the informational content of the announcement, is diminished. The signal-to-noise ratio of official communication deteriorates. This happens not because the announcement itself is less informative, but because the informational space has already been partially occupied by presidential noise. This is the crowding-out mechanism at the heart of H2. The noise trading creates the behavioral vulnerability, sentiment asymmetry amplifies the magnitude of the disruption, and digital speed converts presidential communication from an additional source of

uncertainty into an active competitor that displaces the institutional signals designed precisely to reduce it.

H1: The Direct Volatility Effect

The theoretical foundation for H1 rests on Black's (1986) noise trading model and the theoretical framework provided by De Long et al. (1990). They demonstrate that investors who react to non-fundamental signals can, when sufficiently numerous and sharing common misperceptions, generate systematic and measurable price deviations. Multiple empirical studies have directly documented this effect in the context of Trump's communication. They provide the scientific grounding for H1. The Volfefe analysis of Salem et al. in 2019 constitutes the most direct empirical precedent. Thanks to their study, they show that specific lexical categories in Trump's tweets predict movements in implied volatility beyond what standard macroeconomic control variables explain. For their part, Kandi et al. (2021) document sharp short-term market reactions to presidential tweets, followed by a gradual correction toward prior price levels. Born et al. (2017) confirm that price movements following company-specific tweets are accompanied by sharp increases in retail investor attention, consistent with a noise trading mechanism. Collectively, this body of evidence establishes that presidential communication generates measurable volatility responses, but without cleanly resolving the endogeneity problem that H1 is designed to address.

H1: Donald Trump's tweet volume generates an immediate and independent increase in conditional volatility in U.S. equity and bond markets, holding constant the contemporaneous release of scheduled macroeconomic announcements.

The focus on volume rather than content is deliberate. Here, H1 tests whether the sheer quantity of presidential communication generates market uncertainty, independently of its emotion.

If confirmed, this finding would establish presidential communication as an autonomous source of conditional variance, operating independently of any concurrent policy developments. Crucially, it would provide the baseline identification needed for the displacement mechanism tested in H2.

III. The Institutional Pillar: Monetary Policy and Market Expectations

Having mapped the behavioral dimensions through which presidential noise is generated and amplified, we now turn to the institutional pillar and its features. A foundational principle of modern macroeconomics holds that monetary authority and political authority must be institutionally separated. The independence of the central bank from the executive branch is not a bureaucratic convention. It is a structural safeguard whose justification has been systematically documented.

The reference framework for analyzing the political economy of monetary policy is the one from Central Bank Independence (CBI). Its empirical foundations were established by Alesina and Summers (1993). Their cross-national analysis reveals a robust inverse relationship between a central bank's institutional independence and inflation. Specifically, higher independence correlates with a significant reduction in both the level and volatility of inflation. In conclusion, the more autonomous the institution is, the more stable is the macroeconomic environment it governs.

This result can be explained by two mechanisms. The first is what Alesina and Summers (1993) describe as the isolation of the political cycle. Indeed, an independent central bank is insulated from the electoral incentives. Otherwise, governments would be tempted to deploy monetary expansion before elections, at the cost of subsequent inflation. The second mechanism is credibility. An institution whose decisions are perceived as technically grounded and politically neutral communicates from a position of unique authority. For their part, elected officials are constantly subject to partisan pressures and can simply not replicate this credibility. For financial markets, this credibility is more than a desirable quality. Indeed, it is a pricing input. The predictability of monetary policy communication reduces uncertainty and compresses risk premia in ways that are directly reflected in asset valuations.

3.1. Official Communication as the Basis for Rational Expectations

Credibility, however, must be continuously operationalized through communication. Neuenkirch (2012) and Mokhtarzadeh and Petersen (2020) demonstrate that central bank communication plays a critical role in shaping how agents form expectations. Their work shows that consistent, transparent, and data-driven communication is necessary to reduce the dispersion of market beliefs and coordinate expectations around a common forecast trajectory.

When the Federal Reserve releases projections, minutes, or forward guidance, it performs an act of informational coordination that anchors market expectations.

This communication infrastructure is the product of deliberate institutional design. Several mechanisms protect the signal quality of official announcements. These include the formal protocols governing FOMC minutes, the measured language of forward guidance, and the systematic deployment of economic projections. In a well-functioning institutional environment, these signals are sufficient to anchor market expectations even during periods of genuine economic turbulence. Communication is one of the primary instruments of monetary policy. (Vayid, 2013)

3.2. The Breakdown of 'Neutrality' Through Political Interference

The theoretical integrity of the CBI framework rests on a single critical assumption. The monetary and political spheres remain functionally separated. What happens when a sitting president publicly attacks the Federal Reserve's leadership, characterizes its decisions as politically motivated, or demands rate adjustments through social media posts? This is not a hypothetical scenario in the context of Trump's presidencies. It became a recurring pattern of institutional interference whose implications for market functioning are severe.

Baker et al. (2024) show that central bank independence is performatively sustained through an ongoing communicative strategy in which central banks cultivate a technical, measured, and depoliticized tone in their communications. The logical implication of their framework is that when that communicative performance breaks down, investors can no longer be certain that an interest rate decision reflects technical economic analysis rather than political accommodation. Consequently, the informational content of official announcements is severely degraded. The credibility identified by Alesina and Summers (1993) as the central asset of independent central banking does not disappear instantaneously, but it erodes. And as it erodes, attention migrates. Markets begin to monitor the presidential communication channel rather than the institutional one. This is explained because it is there that the genuinely market-moving surprises are being generated.

The credibility framework established by Alesina and Summers (1993) and further developed by Aguir (2018) and Cepeda et al. (2023) rests on a fundamental premise. This premise is the political neutrality of the informational environment and it was repeatedly put under strain during the Trump era. If sustained political interference erodes that neutrality,

institutional signals lose their effectiveness. The theoretical inference is that the market-moving capacity of these announcements should diminish in proportion to the intensity of the competing noise. The existing literature has not directly addressed this question.

H2: The Crowding-Out Effect

H2 shifts the focus from direct market disruption to its institutional consequences. The question here is not whether presidential communication moves markets, but whether, in doing so, it reduces the capacity of official announcements to move them. The theoretical mechanism builds upon Barber and Odean's (2008) foundational premise that attention is a scarce resource. The investor's attention is a finite and rivalrous cognitive resource. When the period surrounding a scheduled macroeconomic announcement is saturated with presidential tweets, the cognitive bandwidth becomes constrained. Consequently, the capacity to process the official signal is heavily compressed. This way, the announcement becomes less capable of surprising the market. We cannot explain this result because the signal is less informative in absolute terms. The explanation rather lies in the competition for attention with a continuous stream of politically charged content that has already repositioned investor expectations.

Prior empirical work has documented that presidential tweets generate market responses. What has not been examined is whether this activity systematically attenuates the market-moving power of the formal institutional signals through which investor expectations are anchored. That untested relationship is what H2 addresses.

H2: A higher volume of presidential tweets in the pre-announcement period is associated with a reduction in the abnormal return generated by scheduled macroeconomic announcements, reflecting a weakened capacity of official signals to surprise the market.

Methodology

The theoretical framework developed in the preceding sections makes three specific predictions about how presidential communication interacts with financial markets. Each prediction requires a different empirical approach to test. H1 is about volatility: does tweet volume move it, independently of everything else? H2 is about information displacement: does tweet activity before an announcement reduce its market impact? H3 is about tone: does it matter whether the tweets are negative? These are not the same question, and they cannot be answered by the same method. What comes in this section is organized as follows: Section I. describes the data and sample period; Section II. defines the variables; Sections III., IV., and V. detail the empirical specifications for H1, H2, and H3 respectively; Section VI. describes the robustness tests.

I. Data Description


The sample runs from January 2016 to December 2025. This choice is deliberate on both ends. January 2016 marks the point at which Trump's presidential campaign had become serious enough to generate consistent and systematic market attention, meaning the tweet-market relationship we want to study was already structurally active at the start of the period. December 2025 is the natural closing point: it captures nearly a full year of Trump's second presidency, which began in January 2025, and includes his return to active posting on X (formerly Twitter) following Elon Musk's acquisition of the platform in late 2022. We do not extend the tweets sample to the end of 2025 only for a matter of adding observations. It is also a structural choice allowing us to test whether the mechanisms identified across the first term, hold across the second. These are distinct political cycles. One that began with different market conditions, a different institutional context, and a communication platform that had itself changed ownership and its policies regarding moderation.

The ten-year window also gives us something that the shorter sample could not: a change in the political regime. Indeed, Trump had been banned from Twitter entirely between January 2021 and November 2022. He was confined to Truth Social, a platform with a fraction of the market reach. This event creates a natural comparison period in which the primary variable of interest, which is presidential tweet volume reaching a mass financial audience, is suppressed. Including this interval as an implicit control period strengthens the identification of the tweet effects estimated across the active phases. It also raises a modelling question we return to in Section VI: whether the relationship between tweet volume and market outcomes is stable

across the first term, the inter-term gap, and the second term, or whether it evolved as markets and algorithms adapted to the communication style.

Financial market data, macroeconomic indicators, and volatility measures are retrieved from Bloomberg. The data sample only covers U.S. financial markets, reflecting both the primary object of the analysis. Those are the impact of an American president's social media communication on financial markets, and the institutional literature underpinning H2. Which centers on FED communication and the monetary policy of the US. We used publicly available archival records of Trump's posts on Twitter/X and Truth Social to gather the tweet data over the sample period (DataDrivenDecisionLab, 2026).

Two features of this sample warrant some honesty before estimation of any model. First, 2020 remains an outlier of historical proportions. Baker et al. (2020) show that COVID-19 caused a spike in EPU unequaled in history, generating volatility dynamics that were primarily driven by epidemiological updates and emergency policy announcements rather than presidential tweets. This motivates our first robustness exercise in Section VI. Second, between January 2021 and late 2022 Trump was effectively banned from most major social media platforms. Tweet frequency is zero or near-zero over this period for our primary channel of interest. We view this as an informative characteristic of the data rather than a nuisance to be purged: if our models are well-specified, then we should see correspondingly less predicted volatility when the tweet channel goes silent. If it does not, that is itself a diagnostic signal worth examining.



II. Variable Description

2.1 *Dependent Variables*

The primary dependent variables are daily log returns on the S&P 500 and daily changes in two-year and ten-year U.S. Treasury yields, all retrieved from Bloomberg. We use both equity and bond markets because the literature gives us a clear reason to. Salem et al. (2019) showed that Trump's tweets moved interest rate markets specifically, the Volfefe Index was built on fixed income, not equities. Thus, an analysis restricted to the equity market with the S&P 500 would miss a documented transmission channel.

The two-year yield is chosen because it captures short to medium term policy expectations. Gürkaynak, Sack, and Swanson (2005) show, in their event-study of FOMC announcements, that monetary policy communication operates through at least two distinct factors. One tied to the current federal funds rate and one to the future path of policy, and that shorter maturities are particularly responsive to these factors. Swanson and Williams (2014) further show that yields at maturities of one year and above remained responsive to macroeconomic news even during the zero lower bound period, suggesting that the two-year yield retains meaningful informational content under constrained monetary policy. Changes in the ten-year yield, by contrast, reflect a broader set of influences, including long-horizon growth expectations, inflation dynamics, and term premia. Therefore, it serves to separately identify a complementary channel through which presidential communications may affect risk perception and long-run uncertainty.

Changes are used instead of levels as a standard stationarity transform. It is well-known that Treasury yields are measured in levels that include unit roots, Nelson and Plosser (1982) show this for a large set of macroeconomic and financial variables, and Dickey and Fuller (1981) popularized the framework we use to formally test and mitigate this issue for our financial series. Regression in levels would thus be meaningless. First-differencing avoids this issue and aligns with the event-driven narrative of our inquiry. Specifically, we are concerned with day-to-day movements in response to presidential words, rather than long-run equilibria.

For the event study used to test H2, the dependent variable is the abnormal return around each scheduled macroeconomic announcement. For each of the three announcement types in the sample, CPI, NFP, and FOMC decisions, the abnormal return is the difference between the realized return on announcement day and the return predicted by a market model estimated over

the pre-event window running from 250 to 20 trading days before the announcement. The macro surprise is calculated as the announced value minus the Bloomberg consensus, which is published right before each release. The logic is as usual, it is not the level of an economic report that moves markets, but how much it misses market expectations (MacKinlay, 1997). Regressing abnormal returns on announcement values without controlling for surprise conflates the fact that an event happened with what the event revealed. This makes the crowding-out test impossible to interpret. The Bloomberg consensus is the standard benchmark for this kind of work, not because it is a perfect forecast, but because it is the best available proxy for what the marginal investor actually expected before the release. Over the ten-year sample, the combined announcement dataset covers approximately 323 event observations (82 FOMC rate decisions, 121 NFP releases, and 120 CPI announcements). Of these 323 announcements, 252 are retained after filtering for data availability, requiring a minimum estimation window of 250 trading days prior to each event date and non-missing values for SPX returns, VIX, and EPU (FOMC = 72, NFP = 76, CPI = 104).

2.2 Independent Variables

The primary independent variable for H1 and H2 is the daily tweet volume. More precisely, the total number of original posts made by Donald Trump on Twitter/X and Truth Social on each trading day in the sample. For H1 and H3, this count is log-transformed. It became $\log(\text{TweetCount} + 1)$, where the +1 handles zero-tweet days. Then we lagged by one trading day to make the model use the previous day's tweet activity rather than the same-day count. The lag reflects the same economic rationale as the T+1 timing used in H2, markets absorb and respond to presidential communication with a one-day delay. For H2, tweet volume is measured as the total count in the 48 hours preceding each announcement and log-transformed identically. This tweet count is divided into two separate variables for H3. The daily tweets labelled as negative in tone, and the daily tweets labelled as neutral or positives in tone. The VADER sentiment classification procedure described in section V was used for this classification. For H1 and H2, it is the quantity of tweets that matters, not their emotional content. It is important to highlight a distinction. The volume captures the informational noise channel (Black, 1986), and the attention-constrained trading channel (Barber & Odean, 2008), while sentiment captures the loss-aversion channel (Kahneman & Tversky, 1979), which is the object of H3.

Regarding the deleted posts and retweets, we excluded them from this count for two reasons. Retweets represent the amplification of third-party content, rather than the original presidential communication we are looking for. Including them would confound Trump's own messaging with the views of others. Then, deleted posts are excluded on the grounds that their content was not persistently accessible to market participants and therefore cannot be assumed to have generated a sustained informational effect.

2.3 Control Variables

Two control variables appear across all three models, both retrieved from Bloomberg: the VIX and the Economic Policy Uncertainty index of Baker, et al. (2016). Their roles are distinct.

Firstly, the VIX, published daily by the Chicago Board Options Exchange (CBOE), measures the 30-day implied volatility of S&P 500 options and is the standard proxy for market-wide risk aversion in the empirical finance literature (Whaley, 2000). Without the VIX, a positive coefficient on tweet volume in a volatility model would have been unclear, as Trump tweeted more during stress periods, which made the market more volatile. Therefore, we needed the VIX to break that circularity and to absorb investor anxiety regardless of Trump's tweets.

Then, the EPU index, constructed by Baker, Bloom, and Davis (2016) by aggregating the frequency with which financial media references economic policy uncertainty, the dispersion of economic forecasters' predictions about fiscal variables, and the volume of expiring tax provisions, controls for the broader political uncertainty environment. What we are trying to capture here is the reaction to specific tweets, not the general uncertainty that already existed throughout the Trump period. This is why EPU matters: without it, it would be impossible to know whether a spike in volatility was caused by a tweet or simply by the political climate of the moment. In H1 and H3, both variables are lagged by one trading day. It changes into $VIX(t-1)$ and $EPU(t-1)$, to ensure they are predetermined relative to the conditional variance they control for. Using same-day values would create a circularity: VIX and EPU are themselves measures of uncertainty. Indeed, including them contemporaneously would mean the model partially explains conditional variance with a variable that captures the same thing at the same point in time. The lag breaks this circularity. In addition, both variables are standardized to z-scores prior to estimation in H1 and H3. This is a numerical necessity as the raw scales of VIX (sample mean ≈ 18) and EPU (sample mean ≈ 153) are several orders of magnitude larger than the conditional variance they enter alongside (typically on the order of

10^{-4}). Without standardization, the GARCH solver faces scaling problems that prevent reliable convergence. In H2, VIX and EPU are used in their raw contemporaneous form. They are not lagged and not standardized, because OLS estimation is invariant to variable rescaling, and contemporaneous values are appropriate for a regression on announcement-day abnormal returns.

III. H1: Tweet Volume and Conditional Volatility: GJR-GARCH(1,1)

H1 requires a model that properly accounts for the fact that volatility is not constant over time. Looking at daily returns will show us that calm days tend to be followed by calm days, and turbulent days tend to come in clusters. This is a well-known pattern in financial data called volatility clustering, and it is precisely why standard OLS is not suitable here. OLS assumes that variance stays the same throughout the sample, which is simply not how financial markets work in practice.

The GJR-GARCH(1,1) developed by Glosten et al. (1993), which itself builds on Bollerslev (1986) and Engle (1982), was built to handle this. The model expresses conditional variance as a function of four things: a baseline long-run level, the size of the previous day's return shock, the previous day's estimated variance, and an extra term that allows bad news to hit volatility harder than good news of the same size. This last feature is the leverage effect and is one of the most consistently observed patterns in equity return data.

Choosing GJR-GARCH over the simpler GARCH(1,1) for H1 is not arbitrary. A symmetric model like GARCH(1,1) forces positive and negative shocks to have the same impact by construction. If that assumption is wrong, and we can note that in equity markets it usually is, then part of the volatility the model cannot explain will be incorrectly picked up by the tweet variable. GJR-GARCH avoids that problem. It also keeps the specification consistent with H3, where the asymmetric response to negative tweets is precisely what we are testing.

Finally, in equity markets, the sum of the ARCH, asymmetry, and GARCH coefficients tends to be close to one, meaning volatility shocks die out slowly. Over a ten-year sample covering two presidencies and a global pandemic, this persistence is likely to be particularly strong, it is something we address directly in Section VI.

The model is specified as follows:

Mean equation:

$$r_t = \mu + \varepsilon_t, \varepsilon_t | \Omega_{t-1} \sim t_v(0, \sigma_t^2)$$

Variance equation:

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 \cdot 1_{t-1}^- + \beta \sigma_{t-1}^2 + \delta \cdot TweetVol_t - 1 + \lambda_1 \cdot VIX_{t-1} + \lambda_2 \cdot EPU_{t-1}$$

where r_t is the daily log return (or yield change) on day t ; ε_t is the innovation conditional on the information set Ω_{t-1} ; σ_t^2 is the conditional variance; $\omega > 0$ is the long-run variance baseline; $\alpha \geq 0$ is the ARCH coefficient capturing the impact of past squared innovations; $\beta \geq 0$ is the GARCH coefficient capturing the persistence of conditional variance; $1_{t-1}^- = 1$ if $\varepsilon_{t-1} < 0$ and 0 otherwise; γ is the asymmetry coefficient capturing the leverage effect, expected to be positive; $VIX_{z,t-1}$ and $EPU_{z,t-1}$ are the control variables lagged by one trading day and standardized to z-scores for the same reasons we discussed in Section 2.3; and the covariance stationarity condition requires $\alpha + \gamma/2 + \beta < 1$; $TweetVol_{t-1} = \log(TweetCount_{t-1} + 1)$ is the natural log of the previous trading day's tweet count, lagged by one day; and δ is the coefficient of interest, expected to be positive under H1.

The error distribution is assumed to follow a Student-t distribution with v degrees of freedom, estimated as a free parameter. Daily financial returns are not normally distributed. Even after filtering through a GARCH model, extreme events happen more often than a Gaussian distribution would suggest. The Student-t distribution handles this naturally, which is why we use it here rather than forcing a normality assumption that the data would simply reject.

We have to note that tweet volume enters as an exogenous variable in the variance equation, not in the mean equation. This distinction matters. We are not claiming that tweets predict the direction of market returns. We are claiming that they inflate the level of conditional uncertainty on the days they occur. That is the noise trading prediction: noise does not move markets in a systematic direction, it makes them more volatile (Black, 1986). If δ is positive and statistically significant, it means that markets are measurably more uncertain on high-tweet days than the GARCH dynamics alone would predict, after controlling for VIX and EPU. One limitation we have to acknowledge is that the parameters δ , α , γ , and β are assumed to stay constant across the entire sample period. In other words, we do not allow them to shift between

the first mandate, the gap between the two terms, and the second mandate. This is a simplifying assumption, not something we formally test. A rolling-window approach or a Markov-switching model could handle this more flexibly, but both go beyond what this thesis sets out to do. That being said, the subsample checks in Section VI, in which we re-estimate everything separately for each presidential term, give a reasonable indication of whether the relationships hold across periods that are structurally quite different from one another.

To make the identification credible, we remove from the sample all days on which a major macro announcement was scheduled. Specifically, we exclude FOMC rate decisions, Non-Farm Payroll releases, and CPI announcements, following the Bloomberg Economic Calendar and staying consistent with the announcement set used in H2. The logic is simple: on those days, volatility goes up because of the information in the release itself, and there is no clean way to separate that from whatever Trump may or may not have tweeted. Announcement days are high-volatility days, and Trump often tweeted around them. Including them creates a confound that no control variable fully resolves. By removing them, we ensure that the variance equation is estimated on days where the institutional information channel is quiet. Any residual relationship between tweet volume and conditional variance in that restricted sample cannot be attributed to the announcement itself. This is what makes H1's independence clause testable rather than aspirational.

H1 is confirmed if $\delta > 0$ and statistically significant on the announcement-free subsample.

One identification concern merits explicit acknowledgment: the relationship between tweet volume and conditional variance may not be unidirectional. Trump may have posted more actively during periods of market turbulence, reacting to adverse price movements rather than causing them. There is, however, an important limitation to keep in mind. Trump may well have tweeted in response to market movements rather than the other way around, and without an instrumental variable that reverse causality concern cannot be fully ruled out. The specifications estimated here are therefore associational rather than causal in the strict sense. That being said, the Ljung-Box tests on standardized residuals reported alongside each model in Section VI. help verify that the GARCH dynamics are properly specified and that no systematic pattern remains unexplained. And throughout this section, the language reflects that limitation: the tweet coefficient δ captures the association between elevated posting activity and conditional variance, not a structural causal effect.

IV. H2: The Crowding-Out Effect: Event Study and OLS

H2 is the core of the thesis. The question is not whether Trump moves markets, H1 already tests that. The question is whether, in doing so, he makes official institutional announcements less powerful. The right methodological framework for this is the event study. Event studies isolate the market impact of specific, dateable information events by comparing what actually happened to what would normally have happened in the absence of that event (MacKinlay, 1997). In our case, the events are CPI releases, NFP reports, and FOMC decisions. And what would normally have happened, the counterfactual, is estimated from a pre-event window running from 250 to 20 trading days before each announcement date. The 250-day lower bound is there for stability, it gives roughly one full trading year of data to build a reliable baseline from. On the other hand, the 20-day upper bound acts as a buffer. Without it, markets that start pricing in an upcoming announcement would contaminate the very baseline we are trying to estimate, and the whole logic of the event study would break down.

The abnormal return on announcement day is defined as:

$$AR_{i,t} = R_{m,t} - \text{mean}(R_m) \text{ over the estimation window } [-250, -20] \text{ trading days}$$

where $R_{m,t}$ is the S&P 500 log return on day t , and $\text{mean}(R_m)$ is its average over the estimation window running from 250 to 20 trading days before the announcement. Since the S&P 500 serves as both the instrument and the market benchmark in H2, the abnormal return reduces to a mean-adjusted return. Specifically, it equals the realized return minus the historical average over the pre-event window.

The abnormal return is measured on Day $T+1$ rather than Day T . This choice reflects the institutional release schedule. Indeed, FOMC rate decisions are published after market close, which means markets can only respond the following trading day. NFP and CPI releases, while published before the open, require time for institutional repositioning that extends into the following session. Day $T+1$ is therefore the more appropriate measurement window across all three announcement types. It is the preferred specification throughout H2.

Before running the regression, the macro surprise has been standardized to a z-score separately within each announcement type: FOMC, NFP, and CPI. This way, a one-unit surprise carries the same scale across release types that are otherwise measured in very different units.

VIX and EPU enter the regression in their raw contemporaneous form and are not standardized, as discussed in Section 2.3.

To test the crowding-out hypothesis, we then estimate the following OLS regression:

$$AR_{i,t+1} = \alpha + \beta_1 \cdot \text{Surprise}_{z,i,t} + \beta_2 \cdot \log(\text{TweetVol}_{i,t-48h} + 1) + \beta_3 \cdot \text{Surprise}_{z,i,t} \cdot \log(\text{TweetVol}_{i,t-48h} + 1) + \beta_4 \cdot \text{VIX}_t + \beta_5 \cdot \text{EPU}_t + FE_{type} + \varepsilon_{i,t}$$

where $\text{Surprise}_{z,i,t}$ is the standardized macro surprise: (Actual – Consensus) / SD, computed separately within each announcement type; $\log(\text{TweetVol}_{i,t-48h} + 1)$ is the natural log of total tweet count in the 48 hours preceding the announcement; and FE_{type} are announcement-type fixed effects (CPI, NFP, FOMC) that absorb the systematic differences in how each event type moves markets.

The 48-hour window is motivated by the documented evidence that market participants begin repositioning in the one to two trading days preceding scheduled institutional announcements, as both equity and fixed income markets appear to price in anticipated news over this horizon prior to formal releases (Lucca & Moench, 2013). A presidential communication saturation that occurs within this pre-positioning window is therefore most likely to compete directly with the attention investors would otherwise direct toward the incoming announcement signal. As a robustness check, the specification is re-estimated with 24-hour and 72-hour alternatives in Section VI. to verify that results are not sensitive to this particular window choice.

The interaction coefficient β_3 is what does the work. A negative value here tells a specific story. When tweet volume in the 48 hours before an announcement has been high, the market reacts less strongly to the surprise itself, a bigger-than-expected inflation reading, for instance, moves prices by less than it normally would. In other words, the relationship between the size of the surprise and the abnormal return weakens when Twitter activity has been elevated in the run-up to the release. That pattern is consistent with a crowding-out mechanism, not the disappearance of institutional information, but its dilution in the attention set of investors, extending the scarce attention principle of Barber & Odean (2008). H2 is confirmed if $\beta_3 < 0$ and statistically significant. The model is estimated with heteroskedasticity-robust standard errors. With roughly 320 announcement events across ten years, we have substantially more statistical power than prior event studies on this topic, and the larger announcement sample also

allows us to test whether the crowding-out effect differs across the first and second presidential terms.

V. H3: Negative Tweets and Asymmetric Volatility: GJR-GARCH(1,1)

H3 pushes the analysis one step further. It is not enough to show that tweets raise volatility, H1 already tests that. The claim here is that negatively framed tweets raise it more. Testing this requires a model built for asymmetry. The GARCH(1,1) is not: it treats positive and negative shocks to returns as having symmetric effects on future variance, which is a known empirical failure for equity markets. The leverage effect, the well-documented regularity that negative return shocks generate larger volatility responses than equivalent positive shocks, violates this assumption systematically (Glosten et al., 1993).

The GJR-GARCH(1,1) of Glosten et al. (1993) addresses this by adding a single asymmetric term to the variance equation. The model is specified as follows:

Mean equation:

$$r_t = \mu + \varepsilon_t, \varepsilon_t \mid \Omega_{t-1} \sim t_\nu(0, \sigma_t^2)$$

Variance equation:

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 \cdot 1_{t-1}^- + \beta \sigma_{t-1}^2 + \delta^- \cdot \text{NegTweet}_{t-1} + \delta^+ \cdot \text{PosTweet}_{t-1} + \lambda_1 \cdot \text{VIX}_{z,t-1} + \lambda_2 \cdot \text{EPU}_{z,t-1}$$

where $1_{t-1}^- = 1$ if $\varepsilon_{t-1} < 0$ and 0 otherwise (an indicator variable activating when the previous period's return shock was negative); γ is the asymmetry coefficient capturing the leverage effect, expected to be positive; δ^- is the coefficient on $\text{NegTweet}_{t-1} = \log(\text{NegCount}_{t-1} + 1)$, the natural log of the previous day's negatively classified tweet count; δ^+ is the coefficient on $\text{PosTweet}_{t-1} = \log(\text{PosCount}_{t-1} + 1)$, the natural log of the previous day's positively or neutrally classified tweet count; $\text{VIX}_{z,t-1}$ and $\text{EPU}_{z,t-1}$ are lagged one day and standardized to z-scores as in Section III.; and all other parameters are as defined in Section III.. The covariance stationarity condition for this model requires $\alpha + \gamma/2 + \beta < 1$.

As in the H1 specification, errors are assumed to follow a Student-t distribution with estimated degrees of freedom, for the same reasons of fat-tail accommodation discussed in Section III.

Before attributing asymmetric volatility effects to tweet sentiment, we need a model that already accounts for the baseline asymmetry in how the market responds to its own return history. The γ parameter handles this. If $\gamma > 0$ and significant, the model confirms the leverage effect as a structural feature of the data. What we are then testing, through δ^- and δ^+ , is whether negatively classified tweets impose an additional and distinct upward pressure on conditional variance beyond what this structural asymmetry already explains.

The sentiment classification of individual tweets relies on VADER (Valence Aware Dictionary and sentiment Reasoner), developed by Hutto and Gilbert (2014). VADER assigns a compound score between -1 and $+1$ to each text, derived from a lexicon of words with pre-scored emotional valence and a set of rules that account for features typical of social media, capitalization, repetition, exclamation marks, intensifiers. It was built for exactly this kind of text: a general-purpose classifier trained on news articles would struggle with the compressed, emotionally loaded language of Trump's tweets. That being said, VADER has a well-known limitation. It reads words, not context. A tweet calling the Fed 'incredible' might score as positive when Trump clearly meant the opposite, and that kind of misclassification is not rare in politically coded language. The practical consequence is attenuation bias: when sentiment labels are noisy, the estimated difference between positive and negative tweets gets pushed toward zero, and H3 becomes harder to confirm rather than easier. We therefore treat this as a conservative feature of the test. If anything, VADER's imprecision works against us, which means that any significant result we do find is unlikely to be an artefact of the classification.

To find out how well VADER actually works on this specific corpus, a random sample of 100 original tweets was chosen, excluding retweets and posts consisting solely of URLs. The agreement rate between VADER classifications and manual labels came out at 82%. The disagreements are not random: VADER tends to penalize assertive or rhetorically charged language even when the intended sentiment is clearly positive or neutral, because it reads words individually rather than in context. Irony and sarcasm are the clearest illustration of this, a sarcastic tweet built around a negative term will score negative regardless of what the author actually meant, and Trump's posting style makes heavy use of both. Some misclassification is therefore unavoidable, but that is not necessarily a problem. What matters for H3 is not whether VADER makes mistakes, but to know in which direction those mistakes go. Here, the direction is clear: tweets misclassified in this way push the estimated difference between δ^- and δ^+ toward zero, making H3 harder to confirm rather than easier. The 82% agreement rate, and the

conservative nature of the bias it produces, support the use of VADER as a reasonable, if imperfect, classifier for this analysis.

VADER is preferred over financial-specific dictionaries such as Loughran-McDonald (2011) because Trump's tweets are informal social media posts rather than financial disclosures. We thought a general-purpose social media sentiment tool would be more appropriate for this corpus.

The specification makes no symmetry assumption. Positive and negative tweets are free to have whatever effects the data actually shows, we do not constrain them to be equal and opposite, but estimate each independently. H3 is confirmed if $\delta^- > \delta^+$ and δ^- is statistically significant. This condition captures the differential impact of negative versus positive tweet tone on conditional variance, and follows directly from the Prospect Theory framework (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) and the myopic loss aversion evidence of Benartzi and Thaler (1995) developed in Section 2.4 of the literature review.

VI. Robustness Tests

Five sets of robustness checks test whether the main results depend on specific features of the sample, the model specification, or the definition of the tweet variable.

The first addresses the COVID-19 period, defined here as March 2020 through the end of 2021. While the most acute volatility spike was concentrated in early 2020, policy uncertainty remained structurally elevated throughout 2021 as emergency fiscal and monetary measures continued to dominate the macroeconomic environment (Baker et al., 2020). All models are re-estimated on a sample that excludes this window entirely. COVID-19 introduced a structural break with no historical precedent: the volatility of early 2020 was driven by epidemiological news and emergency policy decisions, not by presidential tweets. If the coefficients estimated in H1 and H3 collapse when 2020 is removed, the honest conclusion is that the identified effects were partly an artifact of that exceptional period. If the results hold, that gives us confidence that the relationship between tweets and volatility is real and not just coming from that particular period. The decision is driven by the data: a Chow test around March 2020 checks for a structural break, and if instability is detected, the pre-COVID subsample becomes the preferred specification.

Then, a related diagnostic concerns volatility persistence. In GARCH-type models, the sum of α , $\gamma/2$, and β tells us how quickly conditional variance reverts toward its long-run level, and when that sum gets close or reaches one, the implication is that volatility shocks are nearly permanent. We check this directly on the estimated coefficients. If the persistence sum stays strictly below 1, no further adjustment is needed. If it equals or exceeds 1, then it indicates near-integration. Under these circumstances, we acknowledge this explicitly and re-estimate the affected models under a FIGARCH(1,d,1) specification. This handles fractionally integrated persistence and is available in rugarch under model = 'fiGARCH'. The GJR-GARCH(1,1) baseline remains the preferred specification on grounds of parsimony and interpretability. This holds true provided the FIGARCH coefficients do not differ materially in sign or significance from the main results.

Whether the tweet-volatility relationship is stable across two structurally different mandates is a question the full-sample estimates would not be able to answer. That is what this second check is designed for. All models are re-estimated separately, firstly for the January 2016 through January 2021 term, then for the second, running from January to December 2025 with the inter-term gap seen as a reference period in which the primary channel of interest is suppressed. And the dimensions along which the two mandates differ are not trivial: the market microstructure has changed, algorithmic trading became far more advanced and used, the communication platform itself shifted from Twitter to X, and the broader institutional environment surrounding Trump's return was very different from his first term.

The prediction is directional. If investors and automated trading systems learned over the course of the first mandate to identify, anticipate, and partially discount Trump's communication patterns, the tweet coefficient δ should be smaller in magnitude during the second term than the first. Conversely, if the effect is driven by the content of political communication rather than its stylistic novelty, no systematic attenuation should be observed. Results are presented as separate coefficient estimates alongside the full-sample baseline.

H2 gets a separate treatment. The crowding-out regression is re-estimated with three alternative pre-announcement windows. We took 24 hours, 48 hours, and 72 hours to verify that the results are not an artefact of this particular measurement choice. If β_3 remains negative and significant across all three windows, the crowding-out mechanism is robust to reasonable variations in how we define the pre-announcement tweet saturation period.

A fifth robustness check tests whether the main results depend on how the tweet variable is defined. The baseline specifications count all original tweets posted by Donald Trump, regardless of their content. A tweet about a sporting event, a personal endorsement, or a political rally is unlikely to carry any information relevant to the conditional variance of the S&P 500 or to investor attention before a macro announcement. Including such tweets in the regressor introduces attenuation bias. Indeed, the financially irrelevant posts dilute the signal, pushing the estimated coefficients toward zero and making it harder to confirm the hypotheses even when a genuine relationship exists.

To address this, we construct a keyword-filtered tweet variable. This variable retains only posts containing at least one term from a lexicon of eleven semantic domains. Those are covering financial markets and indices, macroeconomic conditions, labor markets, trade policy and geopolitics, monetary policy and banking, currencies, fiscal policy, energy and commodities, the corporate sector, economic regulation, and cryptocurrency. The matching uses case-insensitive regular expressions with word-boundary anchors, so that partial matches on unrelated words are excluded. Out of 62,932 original tweets posted over the 2016–2025 sample period, 11,604 passed this filter, which is 18.4% of the corpus, with a daily mean of 3.2 tweets versus 13.2 for the unfiltered variable. Approximately 30% of trading days record zero financially relevant tweets, which motivates the same $\log(x+1)$ transformation used in the baseline. Sentiment classification using VADER is applied to the filtered tweets only, yielding daily counts of FinNegCount and FinPosCount for H3.

All three hypotheses are re-estimated with these filtered variables under specifications otherwise identical to the baseline. If the results are qualitatively similar, the findings do not depend on the inclusion of tweets with no plausible financial content. If the filtered variable produces stronger results, larger coefficients or higher significance, this is evidence of attenuation bias in the baseline and of a more precisely identified signal in the financially relevant subset. This approach follows Baker et al. (2016), who construct the EPU index by filtering newspaper text for economic and policy terms. It is also consistent with the keyword-in-context methodology established by Bollen et al. (2011) for linking social media mood to equity markets. Furthermore, a similar approach was applied specifically to Trump's tweets by Colonescu (2018). It is also directly in the spirit of the Volfefe Index constructed by J.P. Morgan (Salem et al., 2019). They restrict their tweet measure to posts with plausible market-moving

content before linking them to fixed income volatility. It is the same filtering rationale we apply here.

Results

I. Preliminary Statistics and Diagnostics

Concerning the preliminary statistics, Table 2 reports those for all variables over the full 2016 to 2025 sample (see Appendix, Table 2). It shows the S&P 500 daily log returns average 0.048% with a standard deviation of 1.145%, and the excess kurtosis of 16.53 confirms what the data makes immediately visible. Financial returns are fat-tailed, and forcing a normal distribution on them would thus be wrong. Regarding the two Treasury yield series, they tell a similar story. They have negative skewness driven by the sharp downward moves that accompany flight-to-quality episodes. Trump's average daily tweet count stands at 13.24, but that figure is a bit misleading. The distribution is strongly right-skewed, with occasional bursts exceeding 100 posts on a single day, and a standard deviation that is almost as large as the mean itself.

Before turning to the estimations, the correlation matrix in Table 3 points to two things worth keeping in mind (see Appendix, Table 3). Firstly, the raw correlations between tweet variables and market returns are negligible across the board, ranging between 0.009 and 0.012 on SPX. This is not surprising. If tweets were systematically pushing prices in one direction, they would be informational signals rather than noise. Secondly, VIX and EPU are both positively correlated with tweet volume, at 0.087 and 0.141 respectively. Trump posted more actively when markets were already stressed. This is precisely the reverse causality concern raised in the methodology section, and it is the reason why both controls are included in every specification that follows.

II. H1: Tweet Volume and Conditional Volatility

Table 4 reports the GJR-GARCH(1,1) estimates for H1 across all three markets (see Appendix, Table 4).

Starting with the diagnostics, the model performs well on two of the three series. The persistence parameter on SPX is 0.857, and 0.979 for the 10-year yield. Both are below unity which confirms covariance stationarity. For the 2-year yield, persistence reaches 1.012. Near-integration in short-term interest rate series is a known empirical regularity for long-horizon samples that covers multiple monetary policy cycles. We address it directly through the FIGARCH robustness check below. The Ljung-Box tests pass for all three series, with p-values

of 0.725, 0.878, and 0.930. There is no evidence of residual autocorrelation in the standardized residuals.

The leverage effect captured by the γ coefficient is large and highly significant on SPX, at 0.293. This confirms the well-documented asymmetry in equity markets in which negative return shocks generate disproportionately larger volatility responses than equivalent positive shocks. For Treasury yields, this effect is a lot smaller, which is consistent with how the fixed income markets work.

Now, for the coefficient of interest. On SPX, δ is positive and statistically significant under QML standard errors, at 7.9×10^{-7} with $p = 0.021$. Under the Bollerslev-Wooldridge robust correction, the result no longer holds. For both yield series, the coefficient is essentially zero, with p-values of 0.999 and 0.981. H1 is therefore marginally supported on the S&P 500 under QML inference only, and not confirmed for Treasury markets.

Lastly, concerning H1, the sensitivity of the SPX result to the choice of standard errors deserves some attention. With a persistence at 0.857, the GARCH model already accounts for a large share of the daily variance through its own internal structure. What remains unexplained is only a narrow margin, and identifying a reliable external regressor within that margin is quite difficult. The tweet coefficient is positive in the expected direction. However, the signal is not strong enough to survive both inference approaches, and we treat H1 as only marginally supported at best.

III. H2: The Crowding-Out Effect

3.1 Timing Check

In the case of H2, Table 5 presents OLS estimates for both Day T and Day T+1 side by side for a quick comparison (see Appendix, Table 5). On Day T, the surprise coefficient β_1 is positive but entirely insignificant at $p = 0.832$. Markets do not react meaningfully to the announced figure on the release day itself. On Day T+1, everything changes. The same coefficient turns significant at $p = 0.015$. This is consistent with the institutional calendar as FOMC decisions are published after the market close, which means that the following trading day is the first real opportunity for markets to respond. Consequently, all subsequent analysis uses T+1 as the preferred specification.

3.2 Main Results

The interaction coefficient β_3 is the variable of interest for H2. On Day T, it is positive and statistically indistinguishable from zero at $p = 0.963$. On Day T+1, it becomes negative and significant at $p = 0.014$. Thus, H2 is confirmed.

The economic interpretation is very clear. Each additional unit of $\log(\text{TweetVol}_{48h} + 1)$ attenuates the abnormal return on Day T+1 by approximately 0.14 percentage points, after controlling for the size of the macro surprise, VIX, and EPU. When Trump has been posting in large volume in the 48 hours before a scheduled announcement, the market reacts less strongly to the news itself the following day. This phenomenon holds even when the announced figure deviates materially from what analysts had expected. The overall model explains 17.7% of the variance in announcement-day abnormal returns, which is a meaningful level for a cross-sectional event study of this type.

For the VIX control, it is marginally significant on Day T+1 at $p = 0.076$. This reflects the well-documented tendency for markets to arrive at scheduled releases already partially repositioned during periods of elevated stress. The NFP fixed effect is also marginally significant at $p = 0.076$, consistent with the historically larger market-moving content of labor market releases relative to the FOMC baseline.

IV. H3: Sentiment Asymmetry and Conditional Volatility

Table 6 reports the GJR-GARCH(1,1) estimates for H3. The model diagnostics are closely the same to those of H1 (see Appendix, Table 6). Persistence is 0.853 on SPX and 0.979 on the 10-year yield. The 2-year yield shows near-integration again at 1.012, and the Ljung-Box tests pass across all three series.

Turning to the empirical findings, the result on the sentiment coefficients leaves no room for doubt. Both δ^- and δ^+ are statistically indistinguishable from zero across all three markets. On SPX, δ^- equals 2.2×10^{-7} with $p = 0.215$, and δ^+ equals 4.1×10^{-7} with $p = 0.337$. The condition $\delta^- > \delta^+$ is not satisfied. For Treasury yields, both coefficients approach zero with p-values near unity. Hence, H3 is not confirmed on any of the three dependent variables.

Two factors can explain this result, and the first is structural. As in H1, near-integrated GARCH dynamics absorb the majority of daily variance variation, it leaves nearly no room for any external regressor to capture. The second is a measurement issue specific to H3. VADER assigns sentiment scores on the basis of individual word polarity, without having access to

context. As Trump's communication style relies heavily on the use of capital letters, rhetorical exaggeration, and politically coded language, it generates systematic misclassification noise. This explains the 82% manual agreement rate. While still being acceptable, this percentage implies that more or less one in five tweets receives an incorrect sentiment label. That level of noise is sufficient to push the estimated difference between δ^- and δ^+ toward zero, regardless of what the true underlying relationship might be.

V. Robustness Checks

Table 7 to 10 report the five robustness tests described in the methodology (see Appendix, Table 7-10).

R1: COVID-19 and Structural Stability (see Appendix, Table 7)

The Chow test around March 2020 returns $F = 0.114$ with $p = 0.893$. There is no evidence of a structural break in the model parameters at the outbreak of the pandemic. This result is not surprising once the mechanics of the GARCH model are considered. The Chow test evaluates whether the estimated parameters, α , β , γ , and δ , shift significantly across the break point. It does not test whether the level of conditional variance changes. A GARCH model absorbs volatility spikes endogenously. Therefore, when a large shock arrives, the lagged squared residual term ε_{t-1}^2 rises sharply, driving σ_t^2 upward through the model's own internal dynamics, without requiring any change in the underlying parameters. Thus, the COVID episode is captured as an unusually large innovation, not as a structural shift in how volatility propagates. Parameter stability and variance stability are not the same thing, and the Chow test only speaks to the former.

When the COVID-19 period is excluded entirely and H1 re-estimated on the remaining 1,828 observations, the tweet coefficient remains essentially zero at $p = 0.999$, and persistence drops to 0.772. Therefore, the main results are not being driven by the exceptional volatility conditions of 2020.

R2: Subsample by Presidential Term (see Appendix, Table 7)

Re-estimating H1 and H3 separately by presidential term produces one notable result. During Term 1, covering 2016 to 2021 with 1,135 observations, the tweet coefficient is not significant under either hypothesis. During Term 2, covering 2025 only with 223 observations, the coefficient becomes large and significant at $p < 0.001$. This finding should be treated with

caution. The Term 2 subsample covers only one calendar year, and short estimation windows mechanically amplify any persistent signal in the data. We therefore do not draw strong conclusions from it. The inter-term period from 2021 to 2024 could not be estimated due to a hessian inversion failure in the GARCH solver, attributable to near-zero tweet activity during Trump's platform suspension.

R3: Alternative Pre-Announcement Windows (see Appendix, Table 7)

The H2 result holds across all three window specifications. Re-estimating the crowding-out regression with 24 hour, 48 hour, and 72 hour pre-announcement windows yields $\beta_3 = -0.001516$ at $p = 0.028$, -0.001419 at $p = 0.014$, and -0.001225 at $p = 0.020$ respectively. All three confirm H2 at conventional significance levels, and coefficient magnitudes are stable across the three specifications. Accordingly, the baseline result is not due to the specific 48-hour measurement choice.

R4: FIGARCH Alternative Model (see Appendix, Table 7)

The re-estimation of the H1 specification on SPX, using a FIGARCH(1,d,1) model, yields a fractional integration parameter $d = 0.402$, significant at 1%. This confirms the presence of long memory in S&P 500 conditional variance. More importantly for our purposes, the tweet volume coefficient remains essentially unchanged at $\delta \approx 1.0 \times 10^{-7}$ with $p = 0.296$. Near-integration does not affect the main conclusions, and the GJR-GARCH(1,1) remains the preferred specification on grounds of parsimony.

R5: Financially Filtered Tweets (see Appendix, Table 8-10)

Regarding tables 8 to 10, they contain the results when the tweet variable is restricted to posts containing at least one term from the financial keyword lexicon. For H1, the effect on SPX sharpens considerably (see Appendix, Table 8). The FinTweetVol coefficient rises to $\delta = 1.06 \times 10^{-6}$ and becomes significant at 1% with $p = 0.006$, compared to $p = 0.021$ under QML in the baseline. Persistence remains stable at 0.855. This pattern is consistent with an attenuation bias interpretation. The raw tweet count includes a large share of posts with no plausible financial content, and that noise dilutes the estimated coefficient. When the variable is restricted to economically relevant communication, the signal is both stronger and more precisely estimated.

For H2, the filtered specification gives $\beta_3 = -0.001813$ at $p = 0.093$ (see Appendix, Table 9). The direction is consistent with the baseline, and the coefficient is actually larger in

absolute value, suggesting that each financially relevant tweet carries a stronger per-unit saturation effect than the average post. That said, the sparsity of the filtered variable reduces statistical power, and the result only reaches the 10% threshold.

Finally, H3 remains not confirmed with the filtered tweets variable (see Appendix, Table 10). Nevertheless, one result is worth acknowledging. On SPX, the positive tweet coefficient δ^+ is significant at 1% while δ^- is not. Financially positive tweets appear associated with higher market volatility rather than negative ones, which is the opposite of what Prospect Theory predicts. A plausible interpretation is that Trump's positive financial tweets tend to work like announcements. Declarations of market records or trade breakthroughs create uncertainty about whether the claimed outcome will actually materialize, generating volatility through an anticipation channel rather than a loss-aversion mechanism. This observation falls outside the scope of H3 as formulated, but it opens an interesting direction for future research.

Discussion

I. The Volume Channel

The results for H1 point to a relationship that exists but is fragile in the raw tweet specification (see Appendix, Table 4). The tweet volume coefficient on SPX is positive and significant under QML standard errors, but it does not survive the Bollerslev-Wooldridge robust correction. But it would be a mistake to present it as anything more than marginal evidence. What makes it more interesting, however, is what happens when the tweet variable is filtered.

When the analysis is restricted to financially relevant posts only, the signal strengthens considerably (see Appendix, Table 8). The FinTweetVol coefficient on SPX reaches $\delta = 1.06 \times 10^{-6}$ and becomes significant at 1% with $p = 0.006$. Persistence remains stable at 0.855. This contrast between the raw and filtered specifications is informative in itself. It suggests that the result in the baseline is not evidence that tweets have no effect on volatility. It is more likely evidence of attenuation bias. The raw tweet count includes a large proportion of posts about rally crowds, political opponents, and personal grievances. Those posts carry no financial signal. They dilute the coefficient rather than contribute to it. Once they are removed, the mechanism becomes detectable.

This finding connects directly to the noise trading framework of Black (1986) and De Long, Shleifer, Summers and Waldmann (1990). Their theories predict that noise expands asset price variance. The near-zero raw correlations between tweet volume and market returns in Table 3 are consistent with an environment where noise predominantly inflates uncertainty rather than exerting a persistent directional trend. Trump's tweets do not push markets up or down in a reliable way. They inflate the level of uncertainty around them. And this effect is most clearly visible when the tweet variable is restricted to the subset of posts that investors could plausibly read as relevant to their portfolios.

It is also worth connecting this to the analysis of the Volfefe Index (Salem et al., 2019). This index was itself built on a filtered subset of Trump's tweets. Unlike our daily-frequency specifications, their analysis operates at the intraday level, but the filtering rationale remains the same. They found that specific lexical categories produce measurably stronger market responses. It is consistent with our result. The financially relevant subset carries a stronger signal than the full corpus. The difference between the two is precisely what the filtering exercise is designed to isolate.

That being said, the result should not be overstated. The SPX finding remains model-specific, it holds under one inference approach and not the other. Furthermore it does not extend to Treasury markets under either specification. The most defensible interpretation is that tweet volume constitutes a weak but directionally consistent source of conditional variance on equity markets. But this effect is concentrated in the financially relevant portion of the corpus.

II. The Crowding-Out Effect

H2 is our clearest result and also the one with the most direct theoretical grounding (see Appendix, Table 5). The interaction coefficient β_3 is negative, significant at 5%, and stable across all three pre-announcement window specifications tested in the robustness checks. That kind of consistency is not common in event study designs, and it deserves an attentive interpretation.

The economic mechanism extends the core principle of scarce attention outlined by Barber and Odean (2008). Investor attention is finite. Presidential communication flows continuously in the days leading up to a scheduled announcement. People simply do not have the mental resources left to focus on the official message. The announcement still arrives. The surprise is still real. But the market's capacity to respond to it is reduced, because attention has already been allocated elsewhere. The result is a measurable attenuation of the abnormal return on Day T+1, exactly what H2 predicts.

What is notable about this result is that this effect is strictly limited to the following trading day. On Day T itself, β_3 is essentially zero. The crowding-out does not prevent markets from eventually pricing the information in the announcement. It delays and compresses the response. This is consistent with the pre-FOMC announcement drift documented by Lucca and Moench (2013). Markets partially anticipate institutional signals in the hours before their release. When that anticipation process is disrupted by competing information, the subsequent correction on Day T+1 is smaller.

Furthermore, the finding survives the filter (see Appendix, Table 9). The filtered H2 specification gives $\beta_3 = -0.001813$ at $p = 0.093$. The coefficient is actually larger in absolute value than with the raw data base. It suggests that financially relevant tweets carry a stronger per-unit saturation effect than the average post. The result only reaches the 10% threshold rather than 5%. We can explain this by the sparsity of the filtered variable. The daily mean drops from 13.2 tweets per day to 3.2 financially relevant posts. With the filtered variable we reach roughly

30% zero-tweet days, the variance is lower, and that mechanical reduction in variance inflates the standard error. The direction is consistent, the magnitude is plausible, and the attenuation of significance is structurally explained.

Taken together, the main and filtered H2 results support the attention saturation interpretation over a pure noise hypothesis. By this expression, we mean the mechanism through which presidential tweets overwhelm investors' finite attentional resources, extending Barber and Odean's (2008) scarce attention principle to a context of competing information streams. It is not simply that Trump is producing noise. It is that his financially relevant communication specifically competes with official institutional signals for a limited pool of investor attention. This competition has measurable consequences for the market-moving power of scheduled macroeconomic releases.

From an institutional standpoint, this finding has implications that go beyond the Trump context. Neuenkirch (2012) demonstrates that the effectiveness of central bank communication depends on the assumption of an orderly informational environment. If that environment is continuously saturated by competing presidential signals, the transmission mechanism of monetary policy communication is structurally weakened. The credibility framework, which Alesina and Summers (1993) draw on to motivate their empirical analysis as the central asset of independent central banking requires not just institutional independence, but also informational space. Indeed, H2 suggests that social media can erode that space, and the erosion is empirically measurable.

III. Why Sentiment Does Not Explain Variance

Even if unconfirmed, H3 provides valuable insights (see Appendix, Table 6). The results are not the ones that we could expect, given the theoretical framework. The Prospect Theory argument for H3 is theoretically well-grounded. Loss aversion is a well-documented and robust cognitive asymmetry. Negative stimuli require greater attentional resources and produce larger behavioral responses across a wide range of domains. This is confirmed by several studies such as Baumeister et al. (2001) and Rozin and Royzman (2001). High-frequency trading environments are specifically calibrated to detect linguistic sentiment. Therefore, this asymmetry should lead to a detectable difference in how negative and positive presidential tweets affect volatility. And yet, it does not.

Two structural explanations account for this. The first is model-level. With GARCH persistence at 0.853 on SPX, the model already absorbs most of the daily variance variation through its own internal dynamics. The ARCH and GARCH terms capture the clustering and persistence of volatility shocks. Thus, it leaves a residual fraction that is too narrow for a sentiment-classified external regressor to explain with precision. This is not a failure of H3 specifically. It is a constraint that affects H1 and H3 equally. It reflects a genuine limitation of the GARCH external regressor approach when applied to long-horizon samples with high persistence.

The second explanation is measurement-specific. VADER classifies each tweet on the basis of individual word polarity, without any access to context or rhetorical intent. Trump's communication style creates exactly the conditions under which this approach breaks down. A tweet reading "The Fed is DESTROYING our economy" might score as negative when the intended effect is to create political pressure, not fear. A tweet about a "GREAT trade deal" might score as positive when the actual content introduces uncertainty about whether the deal will materialize. These misclassifications push the estimated difference between δ^- and δ^+ toward zero. And this, regardless of what the true underlying relationship might be. Furthermore, the 82% manual agreement rate seems reasonable, but in fact it implies one misclassification in every five tweets. That is sufficient to attenuate the signal materially.

However, the filtered specification reveals particularly interesting insights (see Appendix, Table 10). On SPX, the positive financial tweet coefficient δ^+ is significant at 1% in the H3 filtered model, while δ^- is not. This is the reverse of what Prospect Theory predicts. It deserves a careful interpretation rather than dismissal.

One plausible reading is the following. Trump's positive financial tweets tend to be announcement-like in character. Declarations of record market highs, trade breakthroughs, or manufacturing victories do not reassure investors in the way that neutral institutional communication does. Instead, they create uncertainty about whether the claimed outcome will actually materialize. It is because they arrive without the institutional verification infrastructure that gives official announcements their credibility. A presidential tweet claiming a historic trade deal raises immediate questions about the deal's terms, durability, and market implications. All of which are unresolved at the moment of posting. The result is that positive financial tweets generate volatility through an anticipation and verification channel. This is consistent with Shiller's (2017) narrative economics framework. According to him, the stories that economic

actors tell about the future shape expectations and uncertainty even when they carry no verifiable informational content.

IV. Putting It Together

Across our three hypotheses, the results converge on a coherent interpretation of how presidential social media communication interacts with financial markets. The mechanism is primarily volumetric rather than sentimental. What moves markets is not whether Trump's tweets are negative or positive. It is how many of them there are, and whether they are relevant enough to compete for investor attention.

H1 and H3 together establish that tweet volume has a weak but directionally consistent effect on conditional variance. This effect is concentrated in financially relevant communication. By contrast, sentiment adds no explanatory power beyond volume. This suggests that when applying the noise trading framework of De Long et al. (1990) to digital communication, the mechanism appears to operate through quantity rather than tone. This observation is only valid at the daily frequency and market level tested here.

H2 then provides the most concrete evidence of how this volume mechanism actually operates institutionally. High pre-announcement tweet activity does not merely add background noise. It specifically attenuates the market impact of official macroeconomic releases. The fact is that this crowding-out effect is robust to the choice of pre-announcement window. This robustness is consistent across both the raw and filtered tweet specifications. This strengthens confidence in the interpretation considerably.

There is also reason to believe that this informational saturation was not entirely accidental. Steve Bannon, one of Trump's former chief strategists, openly described the administration's communication approach as deliberately designed to flood the zone. This is a tactic aimed at purposefully overwhelming the opposition and the media so as to prevent any coherent counter-narrative from forming (Lewis, 2018). Therefore, H2 may empirically reflect a deliberate communication strategy. This suggests the effect is intentional, rather than an unintended byproduct of high posting frequency. This does not change the econometric interpretation of β_3 . But it situates the finding within a broader political economy of information. If the crowding-out of institutional macro signals is a measurable consequence of that strategy, then what this thesis documents is not simply an econometric finding. It is empirical evidence of a deliberate informational architecture at work.

Taken together, the three results describe a presidential communication channel that disrupts financial markets through volume and attention saturation. But not through the emotional valence of its content. This is perhaps the most counterintuitive finding of the thesis. The intuition that negative presidential tweets should matter more than positive ones is not confirmed by this analysis. What the data does confirm is that volume alone is sufficient to produce measurable market effect. The institutional signals, most central to expectation formation, are not immune to that disruption.

V. An Open Question

The robustness analysis by presidential term produces one result that deserves separate attention (see Appendix, Table 7). And this, even if it cannot be interpreted with full confidence. During Term 2, covering 2025 only, the H1 the tweet volume coefficient on SPX rises to $\delta = 5.5 \times 10^{-6}$. It is approximately seven times larger than the full-sample estimate, and becomes highly significant at $p < 0.001$. This stands in sharp contrast to Term 1, where no significant effect is detected across 1,135 trading days.

Several explanations for this divergence are plausible. The policy environment of 2025 was dominated by aggressive tariff announcements, executive orders with direct market implications, and an unusually direct relationship between presidential communication and asset prices. In this context, even a casual tweet referencing trade or markets could carry immediate informational weight. Furthermore, the communication platform had itself changed. The X platform under Elon Musk's ownership operates with different algorithmic amplification dynamics than the Twitter of 2016 to 2020. Posts on X potentially reach financial audiences more directly and with greater speed.

At the same time, these explanations must be weighed against a critical methodological limitation. The Term 2 subsample covers only 223 trading days. That is roughly eleven months of data, which is a genuinely short estimation window for a GARCH model. Short samples mechanically amplify any persistent signal that happens to be present in the data. Moreover the significance of the coefficient should not be treated as equivalent in reliability to the Term 1 result, which rests on more than five times as many observations. We therefore present this finding as a suggestive pattern rather than a confirmed result. It may be a direction for future research as the second term extends further into the data.

However, the Term 2 result does suggest that the tweet-volatility relationship may be time-varying in ways that a single full-sample coefficient cannot capture. Are markets adapting to presidential communication over time, or are they becoming more sensitive to it? Answering this question would require a longer second-term sample.

Conclusion

The core focus of this thesis was to examine whether and how Donald Trump's social media activity affects financial markets. Yet, presidential communication is noisy by design, financial markets are driven by dozens of intertwined forces. Indeed, isolating the marginal contribution of a single communication channel over a ten-year sample requires methodological choices that each carry their own assumptions and limitations. With that in mind, the results reported here should be read as careful empirical evidence rather than definitive causal claims.

Three hypotheses were tested across three markets: the S&P 500 and the 2-year and 10 year Treasury yield. This analysis was conducted using a combination of GJR-GARCH volatility models, an event study framework, and five sets of robustness checks to ensure the reliability of the findings. While being mixed, they are not uninformative, and taken together, they point in a consistent direction.

H1, which predicted that tweet volume increases conditional volatility, receives marginal support on the S&P 500 under QML standard errors. However, it does not survive the Bollerslev-Wooldridge robust correction. On its own, that result would be difficult to interpret with confidence. What changes the picture considerably is the filtered tweet specification. When the analysis is restricted to financially relevant posts only, the coefficient on SPX becomes significant at 1% and the signal strengthens meaningfully. This pattern is consistent with attenuation bias in the baseline, the raw tweet variable contains a large share of posts with no plausible financial content, and those posts dilute the estimated effect rather than contribute to it. The conclusion is not that tweet volume drives volatility in a clean and reliable way. It is that financially relevant presidential communication is associated with elevated market volatility on equity markets, and that this association is detectable once the measurement noise is removed.

H2 is the most robust finding of the three. The crowding-out of macro announcement abnormal returns by pre-announcement tweet volume is confirmed at 5% on Day T+1. It is stable across three pre-announcement window specifications and directionally consistent in the filtered tweet robustness check. It is worth to note that the coefficient is larger in absolute value but reaches only the 10% significance threshold due to the sparsity of the filtered variable. The mechanism applies the core insight of Barber and Odean (2008). Investor attention is finite, and when it is saturated by presidential communication in the days before a scheduled release, the market-moving power of that release is measurably compressed. The fact that the effect is

specific to Day T+1 and absent on Day T itself adds credibility to the interpretation of the findings. Markets are not ignoring the announcement, but they are too distracted to fully process it. And as the Bannon evidence suggests, this saturation may not have been accidental.

H3, which suggested that negative tweets generate more volatility than positive ones, is not confirmed under any specification. This result is not only disappointing, it is also instructive. Two structural factors account for it, the near-integrated GARCH dynamics that leave little residual variance for external regressors to explain, and the systematic misclassification noise introduced by applying VADER to Trump's rhetorically unusual communication style. What the filtered specification does reveal, however, is that positive financial tweets are associated with higher conditional variance on SPX, the reverse of what Prospect Theory predicts. This suggests that the mechanism at work in financial markets is not loss aversion activated by negative presidential language. It is uncertainty generated by announcement-like positive communication that lacks institutional verification. That finding falls outside the scope of H3, but it is one of the more interesting results to emerge from the analysis and points toward a genuinely different theoretical framework for future work.

Across the three hypotheses, the clearest conclusion is the following: Presidential social media communication affects financial markets mainly through volume and attention saturation, not through the emotional content of individual posts. The noise trading mechanism operates quantitatively rather than qualitatively. How much Trump tweets, and whether those tweets are financially relevant, matters more than whether they are negative or positive. H2 is the strongest empirical expression of this mechanism, and it is the finding on which the most confidence can be placed.

That said, this thesis has genuine limitations, and intellectual honesty requires stating them clearly.

The first one and maybe the most fundamental is the reverse causality. Trump may have tweeted more during periods of already elevated market stress. In other words, he was reacting to adverse price movements rather than causing them. Indeed, the positive correlation between tweet volume and VIX in the descriptive statistics is consistent with exactly this interpretation. We lack an instrumental variable that predicts tweet activity independently of market conditions. Consequently, the estimated coefficients cannot be given a structural causal interpretation. They are associational, and should be treated as such.

The second limitation concerns the temporal resolution of the analysis. All three hypotheses are estimated at the daily frequency. This means that intraday price dynamics are entirely invisible to the models. This is a genuine constraint. Trump's tweets are known to generate market reactions in minutes rather than hours. Several high-frequency studies, such as Salem et al. (2019) with the intraday Volfefe Index, show that daily data mechanically flattens or completely misses key volatility effects. By collapsing all tweet activity within a trading day into a single daily count, we inevitably smooth over the timing dimension of the communication shock. The estimated coefficients in H1 and H3 should therefore be understood as lower bounds on the true tweet-volatility relationship. If the mechanism operates primarily within the trading session rather than across sessions, daily data will always understate it. Moving to intraday data would require a fundamentally different estimation strategy, but it would also recover a layer of information that the current framework cannot access.

The third limitation concerns the near-integration of the 2-year Treasury yield variance process. Its persistence reaches 1.012. The FIGARCH robustness check confirms that this does not affect the conclusions on SPX, and the Ljung-Box diagnostics pass for all series. But the 2-year yield model is technically non-stationary. This means its conditional variance does not return to a long-run mean in the standard sense. The results for that series should be interpreted with more caution than those for SPX and the 10-year yield.

The fourth limitation is specific to H3 and concerns the accuracy of VADER sentiment classification on Trump's tweets. The 82% manual agreement rate is reasonable for a social media corpus of this size, but it implies that roughly one in five tweets receives an incorrect label. In this model the hypothesis depends on detecting a differential effect between two sentiment categories. This level of misclassification is sufficient to materially attenuate the signal. Maybe a context-aware language model trained on political discourse would produce a cleaner classification.

The fifth limitation relates to the Term 2 subsample result. The highly significant H1 coefficient during 2025 is based on only 223 trading days. Short samples inflate statistical power for any persistent pattern in the data. Hence, drawing strong conclusions from an eleven-month window would be methodologically premature. It is a suggestive finding, not a confirmed one. It will need a longer second-term sample to be properly evaluated.

The sixth and final limitation is one of external validity. This thesis focuses exclusively on Trump's Twitter and Truth Social activity. In this thesis, we identified mechanisms such as attention saturation and crowding-out of institutional signals. Can we generalize these mechanisms to other political leaders, other social media platforms, or other institutional contexts? It remains an open empirical question. The Bannon evidence suggests that the flooding strategy was deliberate and specific to this administration. That specificity is part of what makes the results interesting, but it also limits how far the conclusions can be extended.

These limitations point naturally toward several directions for future research. The most immediate is extending the second-term sample. As 2025 data accumulates, it will become possible to test whether the Term 2 volatility signal is structural or whether it reflects a short-window artifact. Beyond that, replacing VADER with a context-aware transformer model fine-tuned on political speech could substantially improve the sentiment classification. Perhaps it would resurrect an H3-type finding that our measurement approach cannot reliably detect. Finally, the event study could be extended to other central bank communication events, such as ECB press conferences or Bank of England inflation reports. This would test whether the crowding-out mechanism is specific to the Fed or operates at a broader institutional level.

Beyond the empirical findings, we believe this thesis raises a question that goes beyond the academic context. Indeed, if a single political actor can measurably compress the informational effectiveness of institutions designed to stabilize market expectations, then the resilience of those institutions is weaker than conventional frameworks assume. In our view, the crowding-out finding is not just an econometric result. It is a signal that the architecture of institutional communication is structurally vulnerable to the kind of deliberate saturation that social media now makes possible at zero cost. We do not think this vulnerability is specific to Trump, but he made it visible. It is a feature of any institutional system that competes for attention in an environment where a single actor with a large social media platform can flood the zone at will. The policy implication, in our view, is that central banks and regulatory bodies may need to think more carefully about how they protect the signal quality of their communications in an era of permanent informational noise.

Finally, the main contribution of our thesis is the following. Trump's presidential social media communication affects financial markets primarily through volume and attention saturation, not through the emotional content of individual posts. How much Trump tweets, and whether those tweets are financially relevant, matters more than whether they are negative or

positive. The noise trading mechanism, as documented here, operates quantitatively rather than qualitatively. H2 is the strongest empirical expression of this finding, and it is the result on which the most confidence can be placed. Overall, in this thesis, the results are honest about their limits. But they are consistent enough, across markets, specifications, and robustness checks, to support a clear and defensible conclusion. We live in an era where a political actor can flood the information environment at will. It shows that volume alone is sufficient to crowd out the market-moving power of official institutional signals is not a trivial finding.

Bibliography

- Aguir, A. (2018). Central Bank Credibility, Independence, and Monetary Policy. *Journal of Central Banking Theory and Practice*, 7, 91-110. doi:10.2478/jcbtp-2018-0025.
- Akerlof, G. A. (1970). The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500. <https://doi.org/10.2307/1879431>
- Alesina, A., & Summers, L. H. (1993). Central Bank Independence and Macroeconomic Performance: Some Comparative Evidence. *Journal of Money, Credit and Banking*, 25(2), 151-162. doi:10.2307/2077833
- Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593-1636. doi:10.1093/qje/qjw024
- Baker, S. R., Bloom, N., Davis, S. J., & Terry, S. J. (2020). COVID-induced economic uncertainty (NBER Working Paper No. 26983). *National Bureau of Economic Research*. doi:10.3386/w26983
- Baker, A., Hindmoor, A. and McDaniel, S. (2024), Performing central bank independence: The Bank of England's communicative financial stability strategy. *Regulation & Governance*, 18, 1000-1017. <https://doi.org/10.1111/rego.12564>
- Barber, B. M., & Odean, T. (2008). All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *The Review of Financial Studies*, 21(2), 785-818. doi:10.1093/rfs/hhm079
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323–370. <https://doi.org/10.1037/1089-2680.5.4.323>
- Baur, D. G., Hong, K., & Lee, A. D. (2018). Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets, Institutions and Money*, 54, 177–189. <https://doi.org/10.1016/j.intfin.2017.12.004>
- Bayly, L. (2019, May 6). Dow plunges 450 points after Trump tweets on China trigger massive global selloff. *NBC News*. <https://www.nbcnews.com/business/economy/trump-tweet-china-tariffs-throws-global-markets-tailspin-n1002261>
- Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics*, 110(1), 73–92. <https://doi.org/10.2307/2118511>
- Ben Jabeur, S., Dhifaoui, Z., Bakkar, Y., & Ballouk, H. (2025). 'Crypto president': Do narrative political signals drive cryptocurrency returns? *Finance Research Letters*, 78, Article 107194. <https://doi.org/10.1016/j.frl.2025.107194>
- Black, F. (1986). Noise. *The Journal of Finance*, 41(3), 529–543. <https://doi.org/10.1111/j.1540-6261.1986.tb04513.x>
- Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1), 1-8. doi:10.1016/j.jocs.2010.12.007
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
- Born, J. A., Myers, D. H., & Clark, W. J. (2017). Trump tweets and the efficient Market Hypothesis. *Algorithmic Finance*, 6(3-4), 103-109. doi:10.3233/AF-170211

- Bouoiyour, J., & Selmi, R. (2016). The Price of Political Uncertainty: Evidence from the 2016 U.S. Presidential Election and the U.S. Stock Markets. *Working Paper*. <https://doi.org/10.48550/arXiv.1612.06200>.
- Cepeda, V., Taboada-Arango, B., & Villamizar-Villegas, M. (2023). Can Central Bank Credibility Improve Monetary Policy? A Meta-Analysis. *International Finance*. <https://doi.org/10.32468/be.1239>.
- Colonescu, C. (2018). The effects of Donald Trump's tweets on US financial and foreign exchange markets. *Athens Journal of Business & Economics*, 4(4). doi:10.30958/ajbe.4-4-2
- DataDrivenDecisionLab. (2026). Trump Tweets (2009-2025). *Kaggle*. https://www.kaggle.com/datasets/datadrivendecision/trump-tweets-2009-2025?select=djt_posts_dec2025.csv
- De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy*, 98(4), 703-738. <http://dx.doi.org/10.1086/261703>.
- Dickey, D. A., & Fuller, W. A. (1981). Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root. *Econometrica*, 49(4), 1057-1072. doi:10.2307/1912517
- Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4), 987-1007. <https://doi.org/10.2307/1912773>
- Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2), 383-417. <https://doi.org/10.2307/2325486>
- Ficcadenti, V., & Cerqueti, R. (2024). Economic keywords in political communications and financial markets. *Annals of Operations Research*. doi:10.1007/s10479-024-05905-w
- Glaser, F., Zimmermann, K., Haferkorn, M., Weber, M. C., & Siering, M. (2014). Bitcoin - asset or currency? Revealing users' hidden intentions. *Proceedings of the European Conference on Information Systems (ECIS) 2014*. <https://ssrn.com/abstract=2425247>
- Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5), 1779-1801. <https://doi.org/10.1111/j.1540-6261.1993.tb05128.x>
- Goldstein, I., & Yang, L. (2017). Information Disclosure in Financial Markets. *Review of Financial Economics*, 9, 101-125. <https://doi.org/10.1146/annurev-financial-110716-032355>.
- Gürkaynak, R. S., Sack, B., & Swanson, E. (2005). Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. *International Journal of Central Banking*, 1(1), 55-93. <http://dx.doi.org/10.2139/ssrn.633281>.
- Healy, P., & Palepu, K. (2001). Information Asymmetry, Corporate Disclosure and the Capital Markets: A Review of the Empirical Disclosure Literature. *Journal of Accounting and Economics*, 31, 405-440. [https://doi.org/10.1016/s0165-4101\(01\)00018-0](https://doi.org/10.1016/s0165-4101(01)00018-0).
- Hutto, C., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, 8(1), 216-225. <https://doi.org/10.1609/icwsm.v8i1.14550>
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263-291. <https://doi.org/10.2307/1914185>

- Kandi, U., Gujjula, S., Buddha, V., & Bhagavan, V. S. (2021). Visualizing the Financial Impact of Presidential Tweets on Stock Markets. *International Journal of Advanced Science and Technology*, 29(5), 9941-9950. <https://doi.org/10.48550/arXiv.2101.03205>
- Kernell, S. (2006). *Going public: New strategies of presidential leadership*. CQ Press.
- Kitvanitphasu, A., Kyaw, K., Likitapiwat, T., & Treepongkaruna, S. (2026). Bitcoin wild moves: Evidence from order flow toxicity and price jumps. *Research in International Business and Finance*, 103163. <https://doi.org/10.1016/j.ribaf.2025.103163>
- Klaus, J., & Koser, C. (2021). Measuring Trump: The Volfefe Index and its impact on European financial markets. *Finance Research Letters*, 38, Article 101447. doi:10.1016/j.frl.2020.101447
- Krause, D. (2025). Trump, tokens, and tailwinds: An empirical study of digital asset reactions to presidential crypto advocacy. *Marquette University, Working Paper*. doi:10.13140/RG.2.2.34466.54725
- Lewis, M. (2018). Has anyone seen the president? *Bloomberg Businessweek*. <https://www.bloomberg.com/opinion/articles/2018-02-09/has-anyone-seen-the-president?embedded-checkout=true>
- Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, 66(1), 35-65. <https://doi.org/10.1111/j.1540-6261.2010.01625.x>
- Lucca, D. O., & Moench, E. (2013). The Pre-FOMC Announcement Drift (Revised ed.). *Federal Reserve Bank of New York Staff Reports*, No. 512. <http://dx.doi.org/10.2139/ssrn.1923197>
- MacKinlay, A. C. (1997). Event Studies in Economics and Finance. *Journal of Economic Literature*, 35(1), 13–39. <http://www.jstor.org/stable/2729691>
- Mokhtarzadeh, F., & Petersen, L. (2020). Coordinating expectations through central bank projections. *Experimental Economics*, 24, 883 - 918. <https://doi.org/10.1007/s10683-020-09684-6>.
- Montone, M. (2022). Does the U.S. president affect the stock market? *Journal of Financial Markets*, 61, Article 100704. <https://doi.org/10.1016/j.finmar.2021.100704>
- Nelson, C. R., & Plosser, C. R. (1982). Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications. *Journal of Monetary Economics*, 10(2), 139-162. <https://doi.org/10.2307/2527132>
- Neuenkirch, M. (2012). Managing financial market expectations: The role of central bank transparency and central bank communication. *European Journal of Political Economy*, 28, 1-13. <https://doi.org/10.2139/ssrn.1698187>.
- Pástor, L., & Veronesi, P. (2010). Uncertainty about Government Policy and Stock Prices. MFI Working Paper Series, No. 2010-08. *The Milton Friedman Institute for Research in Economics. The University of Chicago*. <https://doi.org/10.2139/ssrn.1625845>
- Pástor, L., & Veronesi, P. (2013). Political Uncertainty and Risk Premia. *Journal of Financial Economics*, 110(3), 520-545. doi:10.1016/j.jfineco.2013.08.007
- Pedemonte, M. (2019). Fireside Chats: Communication and Consumer Expectations in the Great Depression. *The Review of Economics and Statistics* 2024; 106 (4): 997–1011. https://doi.org/10.1162/rest_a_01200

- Rozin, P., & Royzman, E. B. (2001). Negativity Bias, Negativity Dominance, and Contagion. *Personality and Social Psychology Review*, 5(4), 296-320. https://doi.org/10.1207/S15327957PSPR0504_2
- Salem, M., Younger, J., & St John, H. (2019). Introducing the Volfefe Index: Quantifying the impact of presidential tweets on rates volatility. *J.P. Morgan Securities LLC; North America Fixed Income Strategy*. <https://fr.scribd.com/document/425120215/JPM-Introducing-the-Volf-2019-09-06-3116298>
- Scacco, J. M., Copeland, L., Becker, A. B., & Berger, J. (2020). When the President Tweets: Exploring the Normative Tensions of Contemporary Presidential Communication. *International Journal of Communication*, 14, 3119-3142. <https://ijoc.org/index.php/ijoc/article/view/13193>
- Shiller, R. J. (2003). From efficient markets theory to behavioral finance. *Journal of Economic Perspectives*, 17(1), 83–104. doi:10.1257/089533003321164967
- Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4), 967–1004. doi:10.1257/aer.107.4.967
- Stiglitz, J. E., & Weiss, A. (1981). Credit Rationing in Markets with Imperfect Information. *The American Economic Review*, 71(3), 393–410. <https://doi.org/10.7916/D8V12FT1>
- Swanson, E. T., & Williams, J. C. (2014). Measuring the effect of the zero lower bound on medium- and longer-term interest rates. *American Economic Review*, 104(10), 3154–3185. <https://doi.org/10.1257/aer.104.10.3154>
- Tetlock, P. C. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. *The Journal of Finance*, 62(3), 1139-1168. <https://doi.org/10.1111/j.1540-6261.2007.01232.x>
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323. <http://dx.doi.org/10.1007/BF00122574>
- Vayid, I. (2013). Central bank communications before, during and after the crisis: From open-market operations to open-mouth policy (Working Paper No. 2013-41). *Bank of Canada*. <https://doi.org/10.34989/swp-2013-41>
- Whaley, R. E. (2000). The Investor Fear Gauge. *The Journal of Portfolio Management*, 26(3), 12-17. doi:10.3905/jpm.2000.319728

Appendix

Trump Tweets and Financial Market Volatility : Empirical Results

Table 1. Variable Descriptions

Variable	Hypothesis	Definition
Panel A: Dependent Variables		
SPX Return	H1, H3	Daily log return of the S&P 500 index: $\ln(P_t/P_{t-1}) \times 100$. Source: Bloomberg.
$\Delta\text{Yield } 2Y$	H1, H3	Daily first difference of the 2-year U.S. Treasury yield, in basis points. Source: Bloomberg.
$\Delta\text{Yield } 10Y$	H1, H3	Daily first difference of the 10-year U.S. Treasury yield, in basis points. Source: Bloomberg.
AR (T+1)	H2	Abnormal return on the next trading day after a scheduled macro announcement: $AR_t = R_{m,t} - \text{mean}(R_m)$ over the estimation window $[-250, -20]$ trading days.
Panel B: Tweet Variables		
TweetVol	H1	$\log(\text{TweetCount}_{t-1} + 1)$. Lagged one trading day to reflect the one-day delay in market reaction.
NegTweet	H3	$\log(\text{NegCount}_{t-1} + 1)$. NegCount = tweets classified as negative by VADER (compound score < -0.05).
PosTweet	H3	$\log(\text{PosCount}_{t-1} + 1)$. PosCount = tweets classified as positive by VADER (compound score $> +0.05$).
TweetVol (H2)	H2	$\log(\text{tweet count in the 48-hour window before the announcement} + 1)$. Not lagged: captures pre-announcement informational saturation.
Panel C: Control Variables		
VIX	H1, H3	CBOE Volatility Index lagged one day and standardised (z-score): $(VIX_{t-1} - \text{mean}) / \text{SD}$. Standardisation ensures numerical convergence in GJR-GARCH estimation.
EPU	H1, H3	Baker, Bloom & Davis (2016) Economic Policy Uncertainty index, lagged one day and standardised (z-score). Captures macroeconomic policy risk.
VIX (H2)	H2	CBOE VIX on the announcement day, contemporaneous raw value. Controls for ambient market uncertainty at the time of the release.
EPU (H2)	H2	EPU index on the announcement day, contemporaneous raw value.
Surprise_z	H2	Standardised macro surprise = $(\text{Actual} - \text{Survey consensus}) / \text{SD}$, computed separately within each announcement type (FOMC, NFP, CPI) to ensure cross-type comparability.

Notes. All variables are defined as used in the final estimated models. H1 and H3 are estimated on announcement-free trading days only ($n = 2,243$ after one observation lost to lagging). VADER = Valence Aware Dictionary and sEntiment Reasoner (Hutto & Gilbert, 2014). EPU = Economic Policy Uncertainty index (Baker, Bloom & Davis, 2016).

Table 2. Descriptive Statistics

Variable	N	Mean	Median	SD	Min	Max	Skew	Kurt.
SPX Return (%)	2,514	0.048	0.076	1.145	-12.77	9.089	-0.684	16.53
Δ Yield 2Y (bps)	2,514	0.090	0.145	5.245	-60.98	29.09	-0.923	12.85
Δ Yield 10Y (bps)	2,514	0.070	0.010	5.437	-29.50	36.02	-0.098	2.629
TweetCount (daily)	2,514	13.24	10.00	13.02	0	167	2.722	16.09
NegCount (daily)	2,514	2.972	2.000	3.311	0	37	2.555	13.69
PosCount (daily)	2,514	5.310	4.000	4.824	0	59	1.914	9.208
VIX (raw)	2,514	18.50	16.69	7.331	9.14	82.69	2.570	12.57
EPU Index (raw)	2,514	152.8	110.8	123.4	5.39	1,026	2.130	5.565
Surprise_z (pooled)	322	0.000	-0.042	1.023	-3.91	4.12	0.112	3.24

Notes. Full sample: $N = 2,514$ trading days, January 2016 – December 2025. Surprise z computed separately by announcement type (FOMC=82, NFP=121, CPI=120). TweetCount, NegCount, and PosCount are raw daily counts (pre-transformation). Kurtosis = excess kurtosis (normal distribution = 3). SPX log returns in percentage form; yield changes in basis points.

Table 3. Correlation Matrix

Variable	SPX	2Y	10Y	TweetVol	NegTweet	PosTweet	VIX	EPU
SPX Return	1.000							
Δ Yield 2Y	-0.021	1.000						
Δ Yield 10Y	0.031	0.532***	1.000					
TweetVol	0.012	0.008	0.003	1.000				
NegTweet	0.011	0.007	0.002	0.842***	1.000			
PosTweet	0.009	0.006	0.001	0.876***	0.653***	1.000		
VIX	-0.243***	-0.071***	-0.042**	0.087***	0.076***	0.074***	1.000	
EPU	-0.098***	-0.011	-0.003	0.141***	0.126***	0.126***	0.382***	1.000

Notes. Pearson correlations, full sample ($N = 2,514$). Lower triangle. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. VIX and EPU in contemporaneous raw form for this table only.

Table 4. H1: Effect of Tweet Volume on Conditional Volatility: GJR-GARCH(1,1)

	SPX Return	Δ Yield 2Y	Δ Yield 10Y
TweetVol (δ)	$7.9 \times 10^{-7} \uparrow$ (0.021)	0.000000 (0.999)	0.000000 (0.981)
VIX (λ_1)	1.1×10^{-5} (0.000)	0.000000 (0.997)	0.000000*** (0.000)
EPU (λ_2)	≈ 0 (1.000)	0.000000 (1.000)	0.000000 (1.000)
μ	0.000674*** (0.000)	0.001315*** (0.000)	0.001575* (0.092)
ω	0.000013*** (0.000)	0.000001 (0.571)	0.000058** (0.019)
α (ARCH)	0.000300* (0.090)	0.110323*** (0.000)	0.083122*** (0.000)
β (GARCH)	0.710323*** (0.000)	0.910486*** (0.000)	0.896769*** (0.000)
γ (GJR)	0.292873*** (0.000)	-0.016850*** (0.000)	-0.001037 (0.788)
ν	6.488*** (0.000)	5.125*** (0.000)	12.361*** (0.000)
Persistence	0.857	1.012	0.979
Ljung-Box p	0.725	0.878	0.930
Log-Likelihood	7,611.1	4,381.2	3,599.3
AIC	-6.779	-3.899	-3.201
N	2,243	2,243	2,243

Notes. Estimated on announcement-free days ($N = 2,243$). Variance equation: $\sigma_t^2 = \omega + \alpha \cdot \varepsilon_{t-1}^2 + \gamma \cdot \varepsilon_{t-1}^2 \cdot I_{t-1} + \beta \cdot \sigma_{t-1}^2 + \delta \cdot \text{TweetVol}_t + \lambda_1 \cdot \text{VIX}_t + \lambda_2 \cdot \text{EPU}_t$. $\text{TweetVol} = \log(\text{TweetCount}_{t-1} + 1)$. VIX and EPU lagged one day and standardised (z-scores). Distribution: Student-t. p-values in parentheses (QML standard errors). $\uparrow$ On SPX, TweetVol QML $p = 0.021$ (**); Bollerslev-Wooldridge robust $p = 0.139$. Persistence = $\alpha + \gamma/2 + \beta$. Ljung-Box: H_0 = no autocorrelation, 10 lags. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5. H2: Crowding-Out Effect of Tweet Volume on Announcement Abnormal Returns:
OLS

	AR on Day T	AR on Day T+1	Preferred
		✓	T+1: FOMC after close
Surprise_z (β_1)	0.000389 (0.832)	0.001033 (0.426)	
TweetVol (β_2)	-0.000171 (0.754)	-0.000347 (0.620)	
Surprise_z $\times$ TweetVol (β_3)	0.000028 (0.963)	-0.001419** (0.014)	
VIX (β_4)	-0.000623*** (0.001)	-0.000624* (0.076)	
EPU (β_5)	0.000016* (0.051)	0.000005 (0.553)	
FE: NFP	-0.000648 (0.729)	0.003304* (0.076)	
FE: CPI	0.000106 (0.954)	0.000326 (0.876)	
R ²	0.109	0.177	
R ² adjusted	0.084	0.153	
N (events)	252	252	

Notes. Dependent variable: abnormal return $AR = R_{m,t} - \text{mean}(R_m)$ over the estimation window $[-250, -20]$ trading days. Specification: $AR = \alpha + \beta_1 \cdot \text{Surprise_z} + \beta_2 \cdot \text{TweetVol} + \beta_3 \cdot (\text{Surprise_z} \times \text{TweetVol}) + \beta_4 \cdot \text{VIX} + \beta_5 \cdot \text{EPU} + \text{FE_type} + \varepsilon$. $\text{TweetVol} = \log(\text{tweet count in 48h before announcement} + 1)$. Surprise_z pre-standardised within each announcement type (FOMC, NFP, CPI). VIX and EPU contemporaneous raw values. Fixed-effect baseline: FOMC. HC3 robust standard errors. Timing check (bivariate OLS, no controls): $\text{Surprise_z} \rightarrow \text{AR}_t$: $\beta = +0.000167$, $p = 0.880$; $\text{Surprise_z} \rightarrow \text{AR}_{t+1}$: $\beta = -0.002847$, $p = 0.015$ (**), confirming T+1 as the preferred specification. In the full interaction model above, β_1 represents the conditional marginal effect of Surprise_z when $\log(\text{TweetVol}_{48h} + 1) = 0$. When estimated without the interaction term but with all controls and fixed effects, the surprise coefficient on Day T+1 is $\beta_1 = -0.002491$ ($p = 0.052$), confirming the presence of an announcement effect. The attenuation of β_1 in the full interaction model reflects the conditional structure of the specification, not the absence of a macro announcement effect. $N = 252$ events (FOMC=72, NFP=76, CPI=104). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6. H3: Asymmetric Effect of Tweet Sentiment on Conditional Volatility: GJR-GARCH(1,1)

	SPX Return	Δ Yield 2Y	Δ Yield 10Y
NegTweet (δ^-)	2.2×10^{-7} (0.215)	0.000000 (0.999)	0.000000 (1.000)
PosTweet (δ^+)	4.1×10^{-7} (0.337)	0.000000 (0.998)	0.000000 (1.000)
VIX (λ_1)	1.1×10^{-5} (0.000)	0.000000 (0.999)	0.000000* (0.036)
EPU (λ_2)	≈ 0 (1.000)	0.000000 (1.000)	0.000000 (1.000)
μ	0.000677*** (0.000)	0.001315*** (0.000)	0.001575* (0.096)
ω	0.000014*** (0.000)	0.000001 (0.573)	0.000058 (0.224)
α (ARCH)	0.000091 (0.970)	0.110342*** (0.000)	0.083112*** (0.000)
β (GARCH)	0.707051*** (0.000)	0.910474*** (0.000)	0.896785*** (0.000)
γ (GJR)	0.291843*** (0.000)	-0.016855*** (0.000)	-0.001038 (0.949)
ν	6.493*** (0.000)	5.125*** (0.000)	12.363*** (0.000)
$\delta^- > \delta^+$	NO	NO	NO
Persistence	0.853	1.012	0.979
Ljung-Box p	0.719	0.878	0.930
Log-Likelihood	7,610.7	4,381.2	3,599.3
AIC	-6.777	-3.898	-3.201
N	2,243	2,243	2,243

Notes. Estimated on announcement-free trading days ($N = 2,243$). Variance equation: $\sigma_t^2 = \omega + \alpha \cdot \varepsilon_{t-1}^2 + \gamma \cdot \varepsilon_{t-1}^2 \cdot I_{t-1} + \beta \cdot \sigma_{t-1}^2 + \delta^- \cdot \text{NegTweet}_t + \delta^+ \cdot \text{PosTweet}_t + \lambda_1 \cdot \text{VIX}_t + \lambda_2 \cdot \text{EPU}_t$, where $\text{NegTweet} = \log(\text{NegCount}_{t-1} + 1)$ and $\text{PosTweet} = \log(\text{PosCount}_{t-1} + 1)$. H3 confirmed if $\delta^- > \delta^+$ and δ^- statistically significant. VADER manual validation: 82% agreement on 100-tweet holdout sample. VIX and EPU lagged one day and standardised. Distribution: Student-t. p-values in parentheses (robust standard errors). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7. Robustness Checks - R1, R2, R3, R4

#	Specification	Coefficient	p-value	Interpretation
Panel A: COVID-19 Exclusion and Structural Break Test (R1)				
R1	Chow Test - COVID structural break (March 2020)	$F = 0.114$	0.893	No structural break: GJR-GARCH specification is stable
R1	H1 GJR-GARCH - COVID excluded (n = 1,828)	$\delta \approx 0$	0.999	H1 not confirmed on COVID-free sample
Panel B: Subsample Analysis by Presidential Term (R2)				
R2	H1 - Term 1 (2016–2021, n = 1,135)	$\delta \approx 0$	0.999	H1 not confirmed in most active tweet period
R2	H1 - Term 2 (2025, n = 223)	$\delta = 0.0000055$	0.000***	Significant: short sample (1 year); treated with caution
R2	H3 - Term 1 (2016–2021, n = 1,135)	$\delta^- \approx \delta^+ \approx 0$	0.999	H3 not confirmed in most active tweet period
Panel C: Alternative Pre-Announcement Windows - H2 (R3)				
R3	H2 OLS - 24h window (AR on T+1)	$\beta_3 = -0.001516$	0.028**	H2 confirmed ✓
R3	H2 OLS - 48h window (baseline, AR on T+1)	$\beta_3 = -0.001419$	0.014**	H2 confirmed ✓
R3	H2 OLS - 72h window (AR on T+1)	$\beta_3 = -0.001225$	0.020**	H2 confirmed ✓
Panel D: Alternative Volatility Model - FIGARCH (R4)				
R4	FIGARCH(1,d,1) vs GJR-GARCH(1,1) - SPX	$d = 0.402***$; $\delta \approx 1.0 \times 10^{-7}$	$d: < 0.01$; $\delta: 0.296$	δ stable across specifications; GJR-GARCH preferred (lower AIC)

Notes. R1: Chow test for structural break at March 2020. When the COVID-19 period is excluded and H1 re-estimated on the remaining 1,828 observations, model persistence drops from 0.857 to 0.772, confirming that the near-integrated dynamics are partly driven by the exceptional volatility conditions of 2020. R2: GJR-GARCH re-estimated on presidential term subsamples. The Term 2 result ($p = 0.000$) should be interpreted with caution given the short sample ($n = 223$, year 2025 only). The Inter-term result (2021–2024) is excluded due to hessian inversion failure, attributable to near-zero tweet activity during the Twitter suspension period. R3: H2 OLS re-estimated with alternative pre-announcement windows; all specifications use AR on T+1. R4: FIGARCH(1,d,1) on SPX to address near-integration; $d = 0.402$ confirms long memory; tweet coefficient unchanged under FIGARCH; GJR-GARCH preferred on AIC. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Note: The following tables replicate the main analysis (Tables 4–6) using a keyword-filtered tweet variable. FinTweetVol counts only tweets containing at least one term from a 10-category financial lexicon (Fed, tariff, GDP, trade, market, inflation, interest rate, China, dollar, economy). Out of 62,932 original tweets (2016–2025), 11,604 pass this filter (18.4% of corpus, mean 3.2 tweets/day vs 13.2 for all tweets; approximately 30% of trading days record zero financially relevant tweets). Specifications are identical to the main analysis: tweet variable lagged one trading day, VIX and EPU standardised (z-scores) in H1 and H3, contemporaneous raw values in H2.

Table 8. R5 - H1 (Robustness): Effect of Financially Filtered Tweet Volume on Conditional Volatility: GJR-GARCH(1,1)

	SPX Return	Δ Yield 2Y	Δ Yield 10Y
FinTweetVol (δ)	0.000001*** (0.006)	0.000000 (0.999)	0.000000 (1.000)
VIX (λ_1)	0.000011*** (0.000)	0.000000 (0.997)	0.000000 (0.999)
EPU (λ_2)	0.000000 (1.000)	0.000000 (1.000)	0.000000 (1.000)
μ	0.000674*** (0.000)	0.001315*** (0.005)	0.001576* (0.096)
ω	0.000013*** (0.000)	0.000001 (0.618)	0.000058* (0.085)
α (ARCH)	0.000261 (0.939)	0.110333*** (0.000)	0.083110*** (0.000)
β (GARCH)	0.710930*** (0.000)	0.910485*** (0.000)	0.896791*** (0.000)
γ (GJR)	0.287946*** (0.000)	-0.016855*** (0.000)	-0.001044 (0.954)
ν	6.476*** (0.000)	5.125*** (0.000)	12.362*** (0.000)
Persistence	0.855	1.012	0.979
Ljung-Box p	0.715	0.878	0.930
N	2,243	2,243	2,243

Notes. Estimated on announcement-free trading days ($N = 2,243$). $\text{FinTweetVol} = \log(\text{FinTweetCount } t-1 + 1)$, where FinTweetCount counts keyword-matched tweets lagged one trading day. VIX and EPU lagged one day and standardised (z-scores). Distribution: Student-t. p-values in parentheses (QML standard errors). Persistence = $\alpha + \gamma/2 + \beta$. Ljung-Box: H_0 = no autocorrelation in standardised residuals, 10 lags. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 9. R5 - H2 (Robustness): Crowding-Out Effect of Financially Filtered Tweet Volume on Announcement Abnormal Returns: OLS

	AR on Day T	AR on Day T+1	Preferred
		✓	<i>T+1: FOMC after close</i>
Surprise_z (β_1)	-0.000147 (0.935)	0.000323 (0.820)	
FinTweetVol (β_2)	-0.000665 (0.376)	-0.000823 (0.412)	
Surprise_z $\times$ FinTweetVol (β_3)	0.000357 (0.655)	-0.001813* (0.093)	
VIX (β_4)	-0.000624*** (0.001)	-0.000632* (0.067)	
EPU (β_5)	0.000017** (0.036)	0.000006 (0.509)	
FE: NFP	-0.000765 (0.680)	0.003523* (0.066)	
FE: CPI	0.000225 (0.903)	0.000772 (0.706)	
R ²	0.113	0.174	
R ² adjusted	0.088	0.150	
N (events)	252	252	

Notes. $AR = R_{m,t} - \text{mean}(R_m)$ over estimation window $[-250, -20]$ trading days. $\text{FinTweetVol} = \log(\text{financial tweet count in 48h before announcement} + 1)$. Surprise_z pre-standardised within announcement type (FOMC, NFP, CPI). VIX and EPU contemporaneous raw values (not standardised — OLS specification). FE baseline: FOMC. HC3 robust standard errors. Timing check consistent with main analysis: $T+1$ is the preferred specification. $N = 252$ events (FOMC=72, NFP=76, CPI=104). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 10. R5 - H3 (Robustness): Asymmetric Effect of Financially Filtered Tweet Sentiment on Conditional Volatility: GJR-GARCH(1,1)

	SPX Return	Δ Yield 2Y	Δ Yield 10Y
FinNegTweet (δ^-)	0.000000 (0.421)	0.000000 (0.989)	0.000000 (1.000)
FinPosTweet (δ^+)	0.000001*** (0.000)	0.000000 (0.969)	0.000000 (1.000)
VIX (λ_1)	0.000011*** (0.000)	0.000000 (0.999)	0.000000 (0.999)
EPU (λ_2)	0.000000 (0.997)	0.000000 (1.000)	0.000000 (1.000)
μ	0.000670** (0.016)	0.001315*** (0.001)	0.001576* (0.096)
ω	0.000014*** (0.000)	0.000001 (0.578)	0.000058** (0.037)
α (ARCH)	0.000111 (0.575)	0.110374*** (0.000)	0.083111*** (0.000)
β (GARCH)	0.708582*** (0.000)	0.910432*** (0.000)	0.896788*** (0.000)
γ (GJR)	0.287697*** (0.000)	-0.016847*** (0.000)	-0.001043 (0.954)
ν	6.450*** (0.000)	5.125*** (0.000)	12.366*** (0.000)
$\delta^- > \delta^+$	NO	NO†	NO†
Persistence	0.853	1.012	0.979
Ljung-Box p	0.719	0.878	0.930
N	2,243	2,243	2,243

Notes. $FinNegTweet = \log(FinNegCount_{t-1} + 1)$; $FinPosTweet = \log(FinPosCount_{t-1} + 1)$. VADER sentiment classification applied to keyword-filtered tweets only. H3 confirmed if $\delta^- > \delta^+$ and δ^- statistically significant. On SPX, $FinPosTweet$ (δ^+) is significant at 1% while $FinNegTweet$ (δ^-) is not, indicating that positive financial tweets are associated with higher volatility — the reverse of H3 prediction. † For 2Y and 10Y yields, $\delta^- > \delta^+$ numerically but neither coefficient is statistically significant. VIX and EPU lagged one day and standardised. Distribution: Student-t. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Abstract :

This thesis examines whether and how Donald Trump's social media activity affects U.S. financial markets. We use a dataset covering the full 2016 to 2025 period across two presidential terms. Three hypotheses were tested on daily S&P 500 returns and changes in two-year and ten-year Treasury yields. Firstly, H1 asks whether tweet volume increases conditional market volatility. Secondly, H2 describes whether pre-announcement tweet activity reduces the market-moving capacity of scheduled macroeconomic releases. And thirdly, H3 asks whether negatively framed tweets generate larger volatility responses than positive ones.

H1 and H3 are estimated using GJR-GARCH(1,1) models on announcement-free trading days. H2 is tested through an event study framework covering 252 FOMC, NFP, and CPI announcements, estimated by OLS with HC3 robust standard errors. Five sets of robustness checks are reported.

The results are mixed but coherent. H1 receives marginal support on the S&P 500 under QML inference only. The signal strengthens to 1% significance when the analysis is restricted to financially relevant tweets. It is consistent with attenuation bias in the baseline. H2 is the most robust finding, confirmed at 5% and stable across all window specifications. H3 is not confirmed under any specification.

Overall, the evidence suggests that Trump's social media communication affects financial markets primarily through volume and attention saturation, not through the emotional content of individual posts. How much Trump tweets matters more than what he says.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
Louvain School of Management

Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve
Boulevard Emile Devreux 6, 6000 Charleroi, Belgique
Chaussée de Binche 151, 7000 Mons, Belgique
www.uclouvain.be/lsm