Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Essays in Media Economics

A Thesis Submitted to Trinity College Dublin, the University of Dublin in Application for the Degree of Doctor of Philosophy in Economics

By

MATTEO POGRAXHA

Supervised By

DAVIDE ROMELLI AND NICOLA FONTANA

Dublin, 2026

Declaration, Online Access and the General Data Protection Regulation

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

Chapter 3 of this thesis is co-authored with Nicola Mastrorocco, Arianna Ornaghi and Stephane Wolton. Chapter 4 of this thesis is co-authored with Donato Masciandaro, Conor Parle and Davide Romelli.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

Matteo Pograxha
September 30, 2025

Summary

This thesis consists of three essays on how measurable signals of demand shape offline decisions across politics, news production, and financial markets. The unifying theme is gatekeeping in contemporary information systems, decision-makers increasingly rely on observable cues, some originating in digital discourse, others in audience composition and ratings, to decide who to amplify, what to cover, and when to act. Empirically, the chapters combine large-scale text and transcript data with quasi-experimental designs and simple theoretical frameworks, methodologically, they integrate modern NLP, classification and embeddings, with high-frequency outcomes to trace selection mechanisms across domains.

The first essay, Chapter 2, studies whether social-media virality serves as a booking signal for television producers. Using approximately 700,000 tweets from Italian MPs, a content-agnostic top-1 percent virality threshold, and a dynamic difference-in-differences design, it shows that viral politicians receive more next-day invitations and speaking time on TV. A simple model explains why, producers treat virality as a noisy proxy for audience interest, which in turn incentivises attention-maximising rhetoric online. The results reveal a cross-platform feedback loop whereby platform metrics reshape the supply of political speech by altering visibility incentives.

The second essay, Chapter 3, moves from "who to invite" to "what to cover," examining editorial choices in U.S. local television news using the weather as an empirical laboratory where underlying facts are observable. By linking near-universe newscast transcripts to market-day temperature deviations, the chapter documents a clear "man bites dog" pattern, rarer events receive more coverage, and shows systematic heterogeneity with audience partisanship, markets emphasise deviations consistent with local ideological priors. Ownership switches to a large conservative broadcaster leave these gradients largely unchanged, pointing to audience-driven rather than owner-driven selection. A simple demand model with taste for surprise and belief-consistent content rationalises the findings.

The third essay, Chapter 4, turns to financial markets and timing, testing whether pre-announcement social-media discourse around central-bank meetings predicts announcement-day asset-price moves. Using supervised classification to isolate policy-relevant tweets and embedding-based measures of disagreement, the chapter shows, focusing on the ECB for institutional consistency, that higher pre-meeting dispersion of narratives foreshadows larger price reactions when policy news arrives. The evidence highlights an anticipation channel, markets partly react to the resolution of publicly visible disagreement, not only to the information in the announcement itself.

Together, the essays map complementary facets of gatekeeping, who is amplified, what is covered, and when reactions occur, and show that observable signals such as engagement metrics, audience composition, and narrative dispersion do not merely mirror demand, they feed back into editorial and policy choices.

Acknowledgements

And now it’s time for the acknowledgments. But how can one possibly thank, in such little space, an experience that has changed my life? I arrived in Dublin weighed down by fatigue and with great disappointments behind me, yet I felt it was finally time to change, to give myself permission to try again. Well, what can I say, here we are at the end, and I would never have imagined being able to write these words. It has been by far the most important experience of my life, where I have learned so much both professionally and personally, and where I discovered a new sky filled with countless bright stars. Perhaps it is precisely to these stars that I want to address my infinite gratitude, because without their light, finding my way again, when sometimes lost, would have been almost impossible. I could write endless pages about each of them, but for reasons of space I will limit myself to those with whom I have shared the most.

The first thanks goes to the “boy crew,” Enrico, Francesco, and Simone. Thank you, Enri, for being the best emotional support I could ever have hoped for, and at the same time a great study and fun companion. Thank you, Francesco, for always being able to make my days lighter and for putting up with my endless daily complaints. Thank you, Simone, for being the beating heart of this long journey, every day, all day, through countless coffees (an infinite list). The kind of person everyone would want in their office, everyone would want by their side.

Thank you, Iris, for allowing me to share and discover so much about you, and for never giving up on me despite all my whims.

Thank you, Gio, for giving me strength and companionship during one of the most difficult moments of the PhD, when it was just you and me. It was wonderful to discover a new side to our friendship, a side I will cherish carefully. Eugenia, what can I say, thank you for the beautiful (though too little) time we spent together, for all the advice, and for the teasing you always knew how to give me. Each one, in its own way, was useful (yes, the teasing included).

Of course, this list would not be complete without the great Alexis. Thank you simply for existing, and for existing exactly as you are. You are a person full of resources who has given me so much, a friend I know I can always count on, and the person with the biggest heart in the world.

A special thank you also goes to Johann and Pippo, who, even though they were in Dublin for only a short time, managed to fill my life with a million new things. Perhaps the two most talented people I know (after Iris ofc :D). Thank you for teaching me so much and for being people who played a special and fundamental role in my life. I would always want you here with me!

Of course, a huge thank you to the entire Triss family, old and new, like Max, Prachi, Vince, Flavia, Bhavya, Andrej, Mike, and Lorenzo. Thank you for making Triss, over time, the kind of place where everyone would want to work.

A special thank you goes to my three supervisors, Nicola, Nicola, and Davide. How patient you have been with me. Thank you for all the effort, dedication, and example you provided me with in every circumstance. Without you, none of my achievements would have been possible. I owe you so much.

I also want to briefly thank you, Mayra, you were a fundamental part of my life during my first real experience outside of my comfort zone. You are beautiful and so intelligent. I hope to see you soon.

Another special thanks goes to my favorite couple: Tobia and Ludo. Guys, I don't know if you realize how important you are to me. You are the concrete example of why I don't "hate" all the couples I see around, and beyond that, you are amazing people with huge hearts. Thank you for everything you do for me.

Now comes the slightly more complicated part. Another enormous thank you goes to you, Giulia. You left halfway through, and yet you walked me by the hand for almost a third of my life. I will never stop thanking you. You are a wonderful person who made me a better one. I will be grateful to you forever, for life!

Now I would also like to thank Megan, my favorite American. Thank you for helping me discover a new side of myself, just as I was facing a Dublin full of painful memories. Your support, your company, your cheerfulness had a huge impact at a time when I no longer even knew who I was or what I wanted.

And of course, how could I forget them, Dad and Mom, my lenders of last resort. Thank you for all the daily love and support (including financial) that only two such special people could have given me. Day and night, every single day, 365 days a year. I love you.

Then, I want to thank Dublin itself, so often criticized by many, yet it never asks for anything in return, only gives. It gives career prospects, friendships, life, everything. Dublin is home. This experience would not have been the same without that constant wind blowing against me, as if to remind me each time that everything changes, everything moves.

And one last thank you I want to give to myself, for being courageous, for giving myself permission to stay, to explore, to take risks, and I couldn't be happier I did.

I feel extremely fortunate to have met such special people, not only because they helped me when I needed it, but because through them I saw life itself. All of us on the same journey, some facing one thing, some another; some breaking up, some starting new relationships; some just receiving a rejection, others finding a new job. A constantly evolving laboratory from which I was able to learn something from each person. And perhaps the most important thing this journey and these wonderful people taught me is that no matter how afraid you are to face something, you face it anyway, hand in hand with fear itself, because there is always a world beyond waiting for you.

On that note, I'd like to conclude with one of the phrases that best captures the meaning this journey has had for me:

"What world lies beyond this sea I do not know, but every sea has another shore, and I shall reach it." And so shall we, sooner or later.

Contents

List of Figures

List of Tables

Chapter 1

General Introduction

This dissertation studies how gatekeepers translate observable indicators of demand into offline choices across political media, news production, and financial markets. The unifying theme is that decision-makers respond to measurable signals, some originating in digital discourse, others in audience composition and ratings. The three chapters examine complementary facets of this process. Chapter 2 shows how politicians deploy extreme rhetoric on social media to increase virality, and how that virality functions as a booking signal for Italian television producers, creating a feedback loop between online traction and TV invitations. Chapter 3, operating outside the social-media realm, analyzes editorial decision-making in U.S. local television news through the lens of weather coverage, revealing that partisan audience composition pressure shape what is deemed newsworthy even for ostensibly neutral events. Conceptually, Chapters 2 and 3 address two sides of editorial selection: *who is worthy to invite* (politicians) versus *what is worthy to cover* (events and their salience). Chapter 4 turns to financial markets, showing that pre-announcement disagreement in public discourse around central bank decisions serves as an anticipation signal, predicting larger asset-price volatility on the press conference announcement day. Together, the essays map *who*, *what*, and *when* in modern information environments.

The first essay (Chapter 2) *From Virality to Visibility: How Social Media Shapes Television Coverage of Politicians* examines whether politicians can strategically leverage social media virality to gain access to traditional media platforms. It begins with an analysis of approximately 700,000 tweets from Italian MPs during the XVIII Legislature (2018–2022), classifying each tweet according to three theoretically grounded extreme-rhetoric strategies: institutional delegitimization, group-based discrimination, and zero-sum framing. I show that extreme rhetoric increases social media engagement by 39–54% and the probability of achieving top 1% virality by 64–118%. *Crucially, in the second step the treatment is defined purely mechanically as attaining the top 1% of engagement across all MPs, regardless of content, so viral tweets include both extreme and non-extreme messages.* Leveraging a dynamic difference-in-differences approach where treatment is the day when a politician achieves this top-1% virality threshold, I find that viral tweets increase the probability of receiving television invitations by 23.7 percentage points the following day, more than doubling the baseline rate. To understand the incentives underlying these patterns, the chapter develops a theoretical model of political communication in which politicians trade off reputational costs against expected media visibility gains. The model shows that even when citizens rationally discount rhetorical exaggeration, expressive engagement with extreme content can raise virality, and producers use virality as a noisy signal of audience interest. The results indicate that social media creates a new pathway to traditional media visibility. This essay demonstrates that the hybrid media environment generates incentives for extreme political communication, as politicians can use low-barrier digital platforms as bridges to high-barrier traditional media gatekeepers. The main contribution of this study consists of uncovering a previously unexplored cross-platform mechanism: when *politicians* go viral, television gatekeepers treat virality as a booking signal and allocate visibility accordingly. This paper is the first to provide evidence that content-agnostic top-1% tweets are followed by more next-day invitations and speaking time, showing that media respond not only to viral *stories* but to viral *people*. In addition, it introduces a theoretically grounded taxonomy of extreme rhetoric (delegitimization, group-based discrimination, zero-sum framing) and documents ideological heterogeneity in its use and its returns in engagement. A simple model clarifies the conditions under which producers rely on virality as a demand signal, highlighting the incentive it creates for rhetorical exaggeration.

The second essay (Chapter 3) *Man Bites Dog? Editorial Choices and Biases in the Reporting of Weather Events* asks what becomes newsworthy when the underlying facts are transparent. We study U.S. local television and exploit an ideal setting—the weather—where basic facts (temperatures) are observable to researchers independently of coverage and where the rarity of events varies across markets and seasons. By linking (near-universe) local TV newscast transcripts to market-day realizations of maximum temperature deviations from the historical mean, we estimate an empirical newsworthiness pattern. We document a clear "man bites dog" pattern: infrequent, large deviations receive substantially more coverage than days near the norm, and even moderate deviations exceed "normal" days in coverage and placement, showing that big dog's bites can be news too. We then move beyond averages to study heterogeneity: coverage varies systematically with audience partisanship. Relative to the norm, unseasonably cold *summer* days are covered more in Republican-leaning DMAs, while unseasonably hot days are emphasized more in Democratic-leaning DMAs. Salience (story placement) goes in the same direction. Yet explicit climate framing is scarce: references to "climate change" or "global warming" are rare across temperatures. To probe mechanisms, we distinguish demand from supply. Using the staggered acquisition of local stations by a conservative broadcast group (Sinclair) as quasi-experimental variation, we find that coverage patterns remain largely unchanged after ownership switches, pointing to audience-driven news generation rather than top-down ideological directives. We also show that our results are not explained by other correlates of markets or events. Finally, a simple model with audiences who value surprise and prefer belief-consistent content accounts for the smooth, partisan gradients we observe: a profit-maximizing outlet scales coverage with event rarity and varies emphasis on hot versus cold deviations according to market ideology, whereas supply-driven bias would produce suppression or saturation that we do not see. This paper contributes to the media economics literature in three ways. First, we establish an empirical benchmark for unbiased reporting by comparing coverage across politically neutral markets, demonstrating that publication bias extends to ostensibly neutral content like weather. Second, using Sinclair acquisitions as quasi-experimental variation, we show that audience preferences, not ownership ideology, drive coverage patterns. Third, we develop a model where confirmation bias about climate change shapes news demand, explaining why profit-maximizing outlets produce smooth partisan coverage gradients rather than stark all-or-nothing patterns.

The third essay (Chapter 4) *Anticipation Effects in Central Bank Communication: Evidence from Social Media* examines whether pre-announcement social media discourse influences market reactions to monetary policy decisions. While existing literature focuses on how central bank communications move markets through information content or policy surprises, we investigate an overlooked channel: anticipation effects arising from pre-meeting narratives. We construct a high-frequency Twitter/X dataset around Bank of England, European Central Bank, and Federal Reserve policy meetings, focusing on tweets posted within seven-day windows before and after decisions. Using supervised classification to identify monetary-policy-relevant tweets and embedding-based metrics to quantify disagreement (dispersion) and content tone/stance (hawkish/dovish/neutral stances), we test whether pre-meeting social media patterns predict announcement-day market moves. Focusing on the ECB for institutional consistency, we find that higher pre-meeting users' disagreement is associated with stronger stock market reactions at announcement, consistent with decisions resolving competing narratives visible in advance. Post-meeting disagreement also correlates with larger moves, aligning with standard policy-surprise channels. These patterns suggest volatility partly builds before decisions arrive, not just from their content. Conceptually, if Chapter 2 examines how social media activity generates television invitations for politicians, this chapter shows how social media discourse creates market anticipation effects around policy events. Substantively, it reveals that central banks face not just the challenge of crafting optimal messages but also managing pre-existing narrative environments, implying that between meeting communication could reduce noise from competing expectations and support orderly market adjustment. This paper contributes to the monetary policy communication literature by examining a previously unexplored channel: anticipation effects via social media. While existing work focuses on decomposing monetary policy shocks into policy versus information components using high-frequency financial data or professional forecasts, we introduce social media discourse as a source of pre-meeting monetary policy announcements meeting signals. Our approach complements recent advances using AI and NLP techniques to analyze central bank communications (audio tone, facial expressions), but shifts focus from the content of announcements to the pre-existing narrative environment. The key innovation is showing that dispersed social media opinions before meetings predict announcement day market volatility, revealing that anticipation effects and not just policy surprises drive market reactions.

This dissertation traces how observable signals shape decisions across politics, news, and markets. Taken together, the essays map complementary facets of gatekeeping, that is, who is amplified, what is covered, and when reactions occur, and show how offline choices are increasingly conditioned by measurable cues in contemporary information systems.

Chapter 2

Online Virality, Broadcast Visibility, and the Incentive for Extreme Political Rhetoric

2.1 Introduction

In contemporary democracies, information is abundant while public attention has become an increasingly scarce resource. For politicians, who depend on visibility to succeed, capturing attention is essential. Yet not all attention is equally valuable. Despite the rise of digital platforms, television remains central to politics: it shapes electoral outcomes (Durante and Knight, 2012; Caprini, 2023) and confers political credibility (Van Erkel and Van Aelst, 2021). For many voters, it is the main stage on which they observe national politicians and form judgments about their viability (Newman et al., 2025), making airtime a strategically valuable resource for political actors.

Access to this airtime, however, lies beyond politicians' direct control, as editorial decisions about whom to feature rest with television producers. Unlike the past, when traditional media were the main channel for political communication, politicians today can readily seek visibility on social media platforms that offer low-cost, lightly regulated access with minimal entry barriers. Crucially, social media do more than provide politicians with a direct communication channel: they also generate observable signals of public interest that television producers can use when making programming decisions. In an environment where audiences are fragmented (Wu, 2017; Allcott and Gentzkow, 2017) and competition for attention is fierce, viral success online may signal what viewers want to see, potentially influencing whom producers invite on air.

This paper examines whether social media provide a pathway into television and, if so, what communication strategies allow politicians to exploit it. Existing research has extensively documented how social media directly benefits politicians through increased campaign contributions and vote shares (Petrova, Sen, and Yildirim, 2021; Boken et al., 2023; Kowal, 2023; Bär, Pröllochs, and Feuerriegel, 2025) and has shown that online popularity influences which stories traditional media choose to cover (Hatte, Madinier, and Zhuravskaya, 2021; Cagé, Hervé, and Mazoyer, 2022). Yet we know surprisingly little about how these dynamics affect traditional media's editorial strategies toward politicians themselves. In particular, it remains unclear whether politicians' online prominence translates into increased access to airtime. This distinction is crucial: whereas content spillovers represent passive information diffusion, editorial decisions about whom to feature represent active strategic choices that create powerful incentives for individual political actors.

The analysis therefore proceeds by addressing two research questions. First, does online virality translate into increased television access? Second, what communication strategies drive virality, specifically, does extreme rhetoric increase the likelihood of online viral success? Crucially, if extreme rhetoric drives virality and some politicians can deploy it more readily than others, this feedback mechanism does more than reward individual success: it structurally advantages particular ideological positions.

To address this, I construct a dataset of approximately 700,000 tweets posted by Italian Members of Parliament (MPs) with an active X (previously Twitter) account during the XVIII Legislature (2018 to 2022). I define virality at the MP-day level: for each MP and day, I mark the MP-day as viral if at least one tweet posted that day lies in the top 1 percent of the tweet-level engagement[1] distribution computed across all MPs over the entire sample period. I link each MP's social media activity to comprehensive daily television appearance data from major talk shows and newscasts on the nine most influential and viewed channels in Italy. I define a television invitation as an appearance in which an MP speaks in voice for at least 60 seconds. Television data come from the Italian national communications authority (AGCOM) and, crucially for the analysis, include episodes from all key political formats that broadcast political information across the country: daily news programs and political talk shows (Mangani and Tarrini, 2018; Newman et al., 2025).

[1] Engagement is defined as the total sum of likes, replies, retweets, and quote retweets that a tweet generates.

To investigate whether online virality translates into traditional media visibility, I leverage an event-study design around the timing of MPs' first viral day. The design exploits within-MP variation in when MPs first go viral, estimating dynamic treatment effects in the days before and after the event. The specification includes day fixed effects, MP-by-week fixed effects, and controls for daily tweet volume, accounting for aggregate shocks, MP-specific trends, and strategic tweeting behavior. The identifying assumption is that, conditional on these controls, MPs who experience viral tweets would have exhibited parallel trends in television invitations absent the viral event.

I document a substantial rise in television visibility following viral success. The probability of receiving a television invitation on the next day increases by 23.7 percentage points, more than doubling from a baseline of 11 percent. The same pattern emerges when requiring longer speaking time at the two and three minute thresholds, indicating that viral MPs secure substantial airtime rather than brief appearances. Effects concentrate in the day immediately following virality and fade by the second or third day, consistent with the rapid production cycles of daily news programs.

To understand the mechanism driving this relationship, I examine whether effects differ between public and private television channels. Under a market-driven gatekeeping mechanism, channels that face competitive pressures and are more reliant on ratings should be most responsive to viral signals, whereas public broadcasters with guaranteed funding should be less responsive. The heterogeneity analysis reveals a clear divergence: private channels exhibit large and statistically significant responses to viral MPs, while public channels show no systematic change in invitations. This pattern is difficult to reconcile with pure newsworthiness. If viral MPs were simply covering more important stories, both public and private channels should respond similarly. Instead, the divergence points to market incentives as the main driver: private channels, which depend on ratings, use virality as a signal of audience interest when deciding whom to invite on air.

I then examine what communication strategies drive virality and test whether extreme rhetoric increases online engagement. To do this, I develop a theoretically grounded classification that moves beyond broad measures of negativity or toxicity (Gervais, 2015; Theocharis et al., 2016). Drawing from democratic theory, I identify three specific rhetorical strategies that undermine pluralistic norms: institutional delegitimization (Levitsky and Ziblatt, 2024), group-based discrimination (Grainger, 2018; Demir, 2025), and zero-sum framing (Mudde and Kaltwasser, 2017; Chinoy et al., 2025). Each tweet is classified as extreme if at least one of these strategies is detected using an open-source large language model (*Mistral 24B*).

Comparing tweets from the same MP on the same topic and day, I document that extreme rhetoric substantially increases online engagement. Tweets employing these strategies receive between 39 and 54 percent more engagement than conventional tweets, with group-based discrimination generating the strongest effects. These increases translate directly into virality: extreme rhetoric raises the probability of achieving top 1 percent viral status by up to 1.3 percentage points relative to a baseline of 1.1 percent. A potential concern is that MPs strategically time extreme tweets for moments when engagement would be high regardless of content. To test this, I estimate dynamic specifications examining engagement patterns around extreme tweets. The results reveal no anticipatory spikes in engagement, while contemporaneous effects are substantial and concentrated. The absence of pre-trends suggests that extreme rhetoric generates engagement rather than merely coinciding with moments of pre-existing attention.

As anticipated, the use of extreme rhetoric is asymmetrically distributed across the political spectrum. Extreme right MPs use group-based discriminatory rhetoric at more than double the rate of centrists. Zero-sum framing displays a similar asymmetry, concentrated among extreme right and right MPs, while delegitimizing rhetoric does not vary systematically by ideology. This creates systematic advantages for MPs whose ideological positions make attention-maximizing strategies more readily available, enabling them to more effectively convert online virality into traditional media visibility.

Lastly, to understand the incentives underlying these patterns, I develop a supply-side theoretical model of political communication in a hybrid media environment. A politician strategically chooses the level of rhetorical exaggeration, trading off reputational costs from citizens who verify the truth against the expected benefits of increased media visibility. The key insight is that exaggeration can be optimal even when citizens are not deceived: while rational audiences punish distortion, their expressive engagement with extreme content (Vosoughi, Roy, and Aral, 2018; Buchanan et al., 2024) generates algorithmic signals that television producers use when making programming decisions. This supply-side framework provides a lens for interpreting cross-ideological differences in rhetoric: if certain politicians face lower audience discipline (lower reputational costs) or higher marginal returns to visibility (higher television exposure value), they will rationally choose more exaggeration in equilibrium.

Related literature. This paper contributes to three main strands of the literature. First, I provide evidence on feedback mechanisms between social media and traditional media. Prior work shows that online popularity influences which stories traditional media cover (Hatte, Madinier, and Zhuravskaya, 2021; Cagé, Hervé, and Mazoyer, 2022; Sen and Yildirim, 2015), but focuses on editorial decisions about news content. I examine a different gatekeeping decision, which political actors who receive airtime as guests, and show that social media virality is linked to next day television invitations. The pattern is consistent with market incentives: responses concentrate on private channels that face competitive pressures, while public broadcasters show no systematic response. This establishes a pathway through which online engagement translates into traditional media access, with implications for how visibility is allocated in democratic competition.

Second, I contribute to understanding strategic political communication in hybrid media environments. Prior work documents direct returns to social media activity: opening accounts increases contributions (Petrova, Sen, and Yildirim, 2021), viral attention boosts donations (Boken et al., 2023), and conspiracy theories raise fundraising (Hilden, Kistner, et al., 2025). I identify an indirect pathway through which virality generates returns: by signaling audience interest to television producers, online activity translates into broadcast invitations. This two-stage mechanism, where extreme rhetoric raises engagement and engagement increases TV access, creates incentives that favor attention-maximizing communication. Crucially, because the propensity to use extreme rhetoric varies systematically by ideology, this pathway generates asymmetric advantages in who gains broadcast visibility and voice in democratic competition.

Third, I contribute to the literature on extreme content and political discourse. Prior work documents that social media amplifies emotional, negative, or extreme content (Brady et al., 2017; Bail et al., 2018; Rathje, Van Bavel, and Van Der Linden, 2021; Boken et al., 2023), but typically relies on broad measures of negativity, incivility, or sentiment (Gervais, 2015; Theocharis et al., 2016). I develop a theoretically grounded classification that distinguishes three distinct rhetorical strategies, namely, institutional delegitimization (Levitsky and Ziblatt, 2024), group-based discrimination, and zero-sum framing (Mudde and Kaltwasser, 2017), and show that they have differential effects on engagement, with group-based discrimination generating the strongest response.

The rest of the paper is organized as follows. Section 2.2 presents Italy's media landscape and institutional context. Section 4.2 presents the data sources, descriptive statistics, and measurement strategy for extreme rhetoric. Section 2.4

documents the link between virality and subsequent television visibility. Section 2.5 presents the main empirical results on extreme rhetoric's effects on online engagement. Section 2.6 discusses robustness checks. Sections 2.7 and 2.8 present the theoretical model and the conclusions.

2.2 Institutional Setting

Italian Politics (2018-2022 Legislature). The 2018-2022 legislature (XVIII) was marked by an exceptional political instability, witnessing three different government coalitions during its tenure. Following the March 2018 elections, no party secured a majority, resulting in a prolonged government formation process. The first government, formed in June 2018, was a populist coalition between the populist party (5 Star Movement) and the extreme right party (League). This coalition collapsed in August 2019 when the League leader withdrew support, anticipating new elections would strengthen his position.

A second government emerged in September 2019, this time uniting the ideologically opposed 5 Star Movement with the center-left Democratic Party, alongside smaller parties. This alliance survived until January 2021, when former Prime Minister Matteo Renzi's party withdrew support, triggering another crisis. In February 2021, the former European Central Bank president Mario Draghi was appointed to lead a national unity government, drawing support from across the political spectrum, from the extreme right wing League Party to the Democratic Party, with only the party, Brothers of Italy, remaining in opposition.

Despite these shifting coalitions, parties had stable positions along the ideological spectrum. I classify MPs according to their party's ideological position, as ideological differences shape the rhetorical strategies available to politicians and their ability to capture attention in a competitive media environment.[2]

Italian TV, Daily Talk Shows, and Organizations. Italy's television landscape comprises numerous national channels, though political coverage remains concentrated among a handful of major broadcasters. RAI (the public broadcaster), Mediaset, and La7, along with newer entrants like TV8 and NOVE, dominate

[2]Ideological leanings differ among parties, which can be grouped into distinct categories based on their values and political agendas: Left (Partito Democratico and Liberi e Uguali); Centrist (Italia Viva, Azione, +Europa, and allied formations such as Unione di Centro, Noi con l'Italia, MAIE, Centristi); Right (Forza Italia); Extreme-right (Lega, Fratelli d'Italia, Partito Sardo d'Azione, and Italexit); and Populist (Movimento 5 Stelle). Independents and unaffiliated MPs are grouped with Centrists if their parliamentary behavior aligned with liberal-centrist blocs. I group parties using the Chapel Hill Expert Survey (2019) for left–right and EU-position anchors, cross-checked with European Parliament group affiliations for party families (PD to S&D; Forza Italia to EPP; FdI to ECR; Lega to ID; +Europa/Azione/Italia Viva to Renew). *Populist.* I classify M5S as populist (syncretic) following the comparative politics literature. *Extreme Right.* I code Lega, Fratelli d'Italia, Partito Sardo d'Azione (PSdA), and Italexit as extreme right. PSdA is grouped with Lega because it ran on a joint list with Lega in 2018 and its deputies sat in Lega's parliamentary group in the XVIII Legislature. Italexit is classified as extreme right given its hard Eurosceptic and sovereignist profile and 2022 reports of neo-fascist (CasaPound-linked) candidates that led Alternativa to break an electoral pact. Independents/unaffiliated MPs are assigned to the nearest liberal-centrist bloc only when their parliamentary/coalitional behavior clearly aligns with that bloc.

political programming through their newscasts and political talk shows, attracting millions of viewers (Auditel, 2024).[3] The communications authority, AGCOM, monitors political pluralism by tracking each politician's *speaking time*, the seconds they are directly heard on air, across programs and channels.

Political guests are typically invited to television programs through direct contact between program press offices and party press offices.[4] Unlike journalists or commentators who may receive fees, politicians appear unpaid, with the media exposure itself serving as compensation. Invitation timelines vary by program format: daily newscasts and political talk shows typically extend invitations with one day's or same-day notice, while weekly programs generally plan four to five days ahead. This production timeline is crucial for the analysis. Next day invitation windows allow producers to incorporate real time signals of audience interest, including politicians' recent viral success online.

The Role of Political TV in Italy. Television remains at the center of political communication in Italy. Legacy broadcasters dominate news consumption and enjoy higher levels of public trust than digital platforms (Newman, 2025).[5] Television airtime carries clear electoral value: exposure on television has been shown to directly influence vote shares and electoral outcomes (Durante and Knight, 2012; Caprini, 2023), making media appearances a strategically important resource for politicians. Italian television audiences, as in many advanced democracies, are older on average, highly politicised, and broadly representative of the national electorate (Gambaro et al., 2025; Newman et al., 2025).

Italy's television centric news environment is not an outlier but part of a wider pattern across western democracies. In the United States, television remains a primary source of election news and is perceived as more reliable than social media (Pew Research, 2024).[6] In the United Kingdom, audiences consistently rate traditional television above online platforms on trust, accuracy, and impartiality (Ofcom, 2024),[7] while in France, television remains the dominant source of news, especially during election periods (European Parliament, 2024).[8] Taken together, these patterns make Italy a particularly informative setting for examining how television and social media interact in allocating political visibility in a hybrid media system.

[3] `https://www.auditel.it/dati/`.
[4] Personal interview with television industry professionals and media analyst in April 2024.
[5] Reuters Institute *Digital News Report 2025* overview here.
[6] Full article here.
[7] Top-line findings here.
[8] Press release here.

2.3 Data and Descriptives

This paper combines two main data sources.

2.3.1 AGCOM

Data on MPs' television appearances come from AGCOM (Autorità per le Garanzie nelle Comunicazioni), Italy's national communications authority. AGCOM monitors political pluralism on television by collecting standardized measures of airtime dedicated to political actors. It distinguishes between (i) *news time*, the time a journalist devotes to reporting on a political figure or institution; and (ii) *speaking time*, the time a politician speaks directly on air. I focus on speaking time from national daily political talk shows and newscasts aired on channels 1–9 as a direct and comparable measure of televised visibility across MPs (full program list in Appendix 2.E.1).[9]

Importantly, AGCOM's speaking-time measure does not label appearance formats. Speaking time covers both live guest segments and pre-recorded material (e.g., clips or archival footage) within programs. Because invitations would be the ideal indicator of broadcaster intent but are not separately coded, I use a transparent duration rule to separate short voice clips from sustained on-air presence. Specifically, I label an appearance as an *invitation* when an MP's *total* speaking time *within that program on that day* is at least 60 seconds. Speaking time is cumulative, not necessarily continuous, e.g., "20 seconds" can be several short voice cuts that add up within the same program, so the 60-second rule is applied at the program level. Because clips are typically brief and split across multiple voices, it is uncommon for purely clip-based coverage to sum to ≥ 60 seconds for a single MP; reaching this duration therefore *increases the likelihood* that the exposure reflects an editorial choice to allocate a seat and sustained on-air time.

Figure 2.1 shows the distribution of speaking time across all appearances during the XVIII Legislature. Television access is highly concentrated: 12.7% of MPs with active X accounts (94 of 742) never appear on television over the entire legislature. Among those who do appear, the distribution is sharply right-skewed: about 70% of appearances fall below 60 seconds, and only 29.4% meet the ≥ 60-second invitation criterion.

Figure 2.2(a) shows the distribution of unique MPs, appearances, and airtime by ideological bloc for invitations (≥ 60 seconds). The three metrics capture different dimensions of television access: coverage (how many distinct MPs

[9] Channels 1–9 refer to the national logical channel numbering: RAI 1, RAI 2, RAI 3, Rete 4, Canale 5, Italia 1, La7, TV8, and NOVE.

Figure 2.1: Distribution of Speaking Time

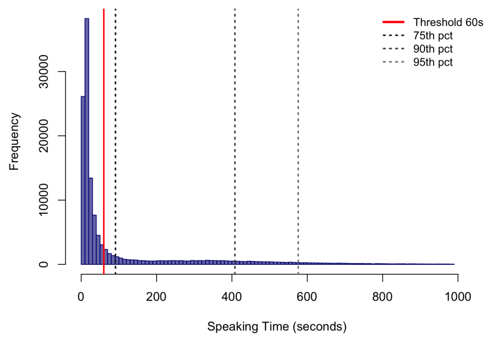

Notes: Distribution across MP–program–date observations with positive speaking time, XVIII Legislature. The red line marks the 60-second threshold used to proxy invitations (around the 71st percentile of the speaking-time distribution); dashed lines mark the 75th, 90th, and 95th percentiles. Shares of MPs who never appear (12.7%) are computed over MPs with active X accounts (N = 742). For visualization, the top 1% of speaking-time observations (99th percentile) is excluded.

appear), frequency (how often they appear), and prominence (how much they speak). Populist MPs represent the largest share of unique MPs invited to television (29.6%) yet receive only 14.4% of appearances and 12.4% of total airtime. This reflects fragmented visibility: many Populist MPs appear occasionally, but few achieve repeated invitations. In contrast, Extreme Right MPs constitute 23.5% of unique MPs invited but secure 35.8% of all appearances and 35.2% of airtime, indicating that Extreme Right politicians who access television are invited substantially more frequently than their counterparts from other blocs. Centrist MPs exhibit the sharpest concentration: just 7.2% of unique MPs invited account for 13.7% of appearances and 14.6% of airtime, nearly double their share of unique MPs. This concentration indicates that a small number of prominent centrist figures dominate television access within their bloc.[10]

These patterns in television access diverge sharply from electoral mandates. Figure 2.2(b) plots each bloc's share of 2018 parliamentary seats against its share of total airtime. The diagonal line represents proportional representation relative to electoral mandate. Centrist MPs exhibit the sharpest over-representation (5.8× their seat share): despite holding only 2.5% of seats, they capture 14.6% of airtime, a concentration driven by the small number of prominent centrists dominating access. Extreme Right MPs similarly over-perform (1.4×), while Populist MPs substantially under-perform (0.3×), their exposure spread across

[10]This pattern aligns with the structure of centrist parties during this period, where prominent figures such as Matteo Renzi (former PM, leader of Italia Viva) and Carlo Calenda (former Minister, leader of Azione) were party leaders commanding substantial media attention.

a larger parliamentary group. Left and Right MPs cluster near the diagonal, approximating proportional representation.

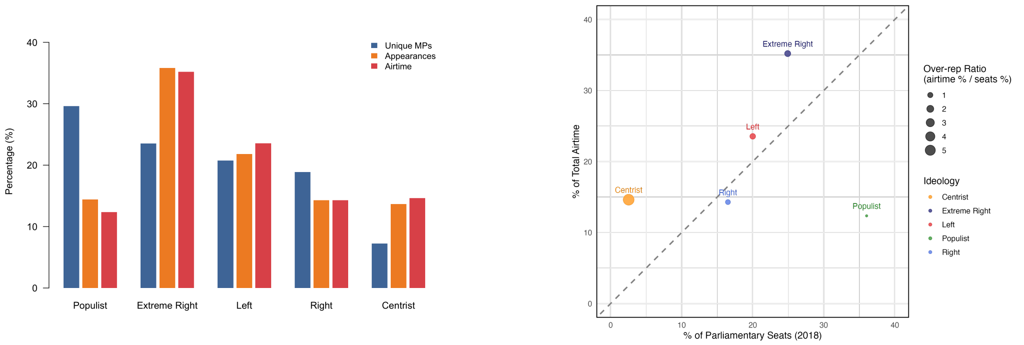

((a)) TV Access Shares by Ideology

((b)) Airtime vs. Electoral Mandate by Ideology

Figure 2.2: TV Access by Political Ideology (Invitations ≥ 60 seconds)

Notes: Panel (a) shows the percentage distribution of unique MPs, total appearances, and total airtime by ideological bloc, restricting to invitations with speaking time ≥ 60 seconds. Panel (b) plots each bloc's share of total airtime against its share of 2018 parliamentary seats. The diagonal line represents proportional representation, where airtime share equals seat share. Points above the diagonal indicate over-representation on television relative to parliamentary strength; points below indicate under-representation. Point size reflects the concentration ratio (airtime share divided by seat share). See Section 2.2 for ideology definitions.

2.3.2 X

Data on the online activity comes from the Twitter Academic API, which provided researchers (free) large-scale access to the full archive of public tweets until its closure in April 2023. I collected all available tweets from Italian MPs who maintained an active X account during the XVIII Legislature (2018–2022), defined as posting at least once over the legislative term. The final dataset comprises 705,695 tweets from 742 MPs, representing 78.5% of all parliamentarians who served during this period. X activity varies systematically across party ideology, gender, parliamentary chamber, and age (full set of descriptive statistics available in Appendix 2.C). Beyond posting frequency, a crucial dimension is engagement, measured as the total of likes, retweets, quote tweets, and replies. Engagement influences both algorithmic visibility within the platform and may signal relevance to external audiences such as TV producers. Table 2.1 presents engagement patterns by ideology. Extreme Right MPs average 402 engagements per tweet, four times the rate of Populist MPs (111) and Right MPs (98), and capture 52% of total engagement across the full sample despite posting only 33.4% of tweets. Left and Centrist MPs fall between these extremes, with average engagement rates of 256 and 364 per tweet respectively. In contrast, Populist MPs account for 24.1% of tweets but receive only 10.6% of engagement, reflecting fragmented online presence across many individual accounts.

Table 2.1: Social Media Engagement by Political Ideology

Political Ideology	MPs	% Tweets	% Engagement	Mean Engagement
Extreme Right	193	33.4	52.0	402
Centrist	40	9.4	13.5	364
Left	130	17.7	17.9	256
Populist	261	24.1	10.6	111
Right	118	15.4	5.9	98

Notes: This table reports the distribution of social media activity and engagement across ideological groups among Italian MPs during the XVIII Legislature (2018–2022). Column 1 shows the number of MPs in each group with active X accounts. Columns 2 and 3 report each group's share of all tweets and all engagement across the full sample, respectively. Column 4 shows mean engagement per tweet, defined as the average total of likes, retweets, quote tweets, and replies.

How to Measure Extreme Rhetoric. Political rhetoric becomes "extreme" not merely when it is emotionally charged, but when it adopts discursive forms that erode the foundations of democratic pluralism. I focus on three theoretically grounded rhetorical strategies: (i) *delegitimization*, where institutions or political actors are depicted as rigged, corrupt, or fundamentally untrustworthy (Levitsky and Ziblatt, 2024; Carey et al., 2019); (ii) *group-based discrimination*, targeting ethnic, religious, or cultural outgroups with exclusionary language (Mendelberg, 2001; Tajfel et al., 2001); and (iii) *zero-sum framing*, which presents democratic politics as a binary struggle without compromise (Iyengar and Westwood, 2015; Mason, 2018). Each strategy has been linked to declining institutional trust, increased affective polarization, and growing tolerance for political violence (Mason, 2018; Norris, 2019; Iyengar, Lelkes, et al., 2019).

This framework differs from sentiment-based or toxicity-based measures. What matters is not only whether a message is merely "toxic" or "angry," but how it frames political conflict and democratic legitimacy. Sentiment and toxicity classifiers, such as Google Perspective API, capture emotional tone (Brady et al., 2017; Rathje, Van Bavel, and Van Der Linden, 2021), but they measure a different dimension than rhetorical form. They can conflate legitimate disagreement with antidemocratic speech and may not capture indirect or coded strategies (Mamakos and Finkel, 2023; Hagen and Venturini, 2024; Bell et al., 2024). In Appendix 2.B, I examine these measurement challenges and compare the precision of categorical measures.

To operationalize this framework, I use *Mistral 24B*, an open-source instruction-tuned model. Tweets are classified using a structured prompt that defines each rhetorical strategy. The classification employs conservative coding rules: *delegitimization* is reserved for direct attacks on institutions (e.g., "rigged elections"); *group discrimination* requires negative generalizations based on identity (e.g., "criminal immigrants"); and *zero-sum framing* captures explicit win-lose constructions (e.g., "if they win, we lose"). Ambiguous cases are coded as 0.[11]

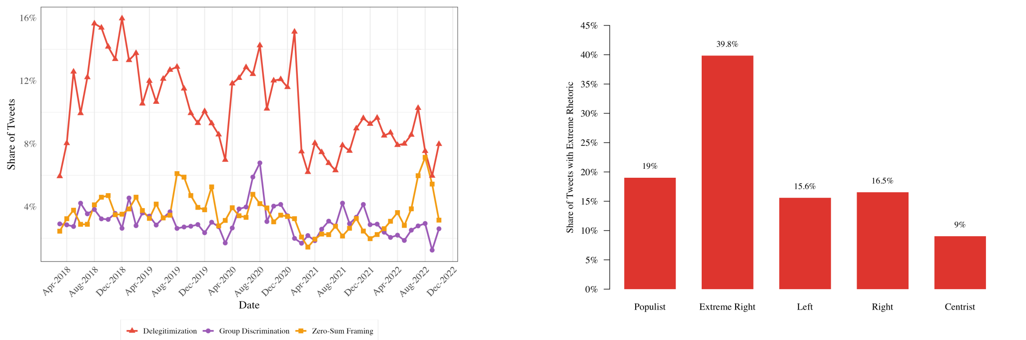

((a)) Monthly shares of extreme-rhetoric frames

((b)) Extreme rhetoric vs. tweet volume by ideology

Figure 2.3: Extreme rhetoric on X over time and by ideology

Notes: Panel (a) shows monthly shares of MPs' tweets coded as delegitimization, group discrimination, or zero-sum framing. Panel (b) shows the share of each ideological bloc's tweets that contain at least one definition of extreme rhetoric (delegitimization, group discrimination, or zero-sum framing). See Section 2.2 for ideology definitions.

Figure 2.3 illustrates the distribution and dynamics of extreme rhetoric online. Panel 2.3(a) plots monthly shares of MPs' tweets using the three extreme-rhetoric frames. The temporal patterns track salient shocks across the XVIII Legislature. *Delegitimization* surges in August 2018 amid the aftermath of the Genoa Morandi Bridge collapse, when public debate centered on accountability for concessions oversight and infrastructure safety,[12] and again in December 2018 during the EU–Italy budget dispute, which culminated in a negotiated deficit target of 2.04% (down from 2.4%).[13] *Group discrimination* rises in July and August 2020 alongside intensified coverage of landings in Sicily/Lampedusa and exceptional policy

[11] The full prompt can be found in Appendix 2.B.1.

[12] Reuters, "Anger at Italy bridge operator as hunt for survivors goes on," Aug. 15, 2018, https://www.reuters.com/article/world/anger-at-italy-bridge-operator-as-hunt-for-survivors-goes-on-idUSKBN1L00H6.

[13] The Washington Post, "Italy backs down in its budget standoff with the E.U.," Dec. 18, 2018, https://www.washingtonpost.com/world/europe/italy-backs-down-in-its-budget-standoff-with-the-eu/2018/12/18/3cb996f4-f8af-11e8-8642-c9718a256cbd_story.html; see also Reuters, "Italy lowers deficit target, EU sees progress on budget row," Dec. 11, 2018, https://www.reuters.com/article/business/italy-lowers-deficit-target-eu-sees-progress-on-budget-row-idUSKBN1OB00S.

signals, including a decree to close reception centres.[14] *Zero-sum framing* peaks in August 2019 during the government crisis,[15] and in August 2022 as campaigning unfolded under an energy-price shock that dominated the policy agenda.[16] These patterns further validate the LLM-based classification.

Panel 2.3(b) shows the share of tweets containing extreme rhetoric by ideological bloc. The pattern reveals systematic variation in rhetorical strategies: Extreme Right MPs produce extreme rhetoric in 39.8% of their tweets, more than double the rate of any other bloc. Populist MPs use extreme rhetoric in 19.0% of their tweets, while Left (15.6%) and Right (16.5%) MPs employ it at similar moderate rates. Centrist MPs exhibit the lowest incidence, with only 9.0% of their tweets containing extreme rhetoric.

I now turn to the empirical analysis. Section 2.4 examines whether viral success on X translates into television invitations. Section 2.5 analyzes what drives virality, testing whether extreme rhetoric increases online engagement and which ideological blocs are more likely to employ it.

2.4 From Online Virality to Television Visibility

If virality signals public interest to traditional media gatekeepers, then spikes in online engagement should translate into increased broadcast visibility. To estimate the relationship between online virality and television visibility, I employ an event-study design that exploits variation in the timing of viral tweets across MPs. I estimate the following equation:

$$Y_{it} = \theta_{iw} + \delta_t + \gamma V_{it} + \sum_{k=-6,k\neq-1}^{6} \beta_k D_{it}^k + \epsilon_{it} \tag{2.1}$$

[14] InfoMigrants, "Sicily's governor orders closure of all migrant hotspots and welcome centers," Aug. 24, 2020, https://www.infomigrants.net/en/post/26815/sicilys-governor-orders-closure-of-all-migrant-hotspots-and-welcome-centers; ANSA (English), "TAR suspends Musumeci migrant centre closure ordinance," Aug. 28, 2020, https://www.ansa.it/english/news/2020/08/28/tar-suspends-musumeci-migrant-centre-closure-ordinance_46e33c09-b357-4cc4-a507-b25925c2a1e7.html.

[15] Reuters, "Italian PM resigns, denounces Salvini for sinking government," Aug. 20, 2019, https://www.reuters.com/article/world/italian-pm-resigns-denounces-salvini-for-sinking-government-idUSKCN1VA1GU/.

[16] Reuters, "Italy unveils new $17.4 billion package against inflation," Aug. 4, 2022, https://www.reuters.com/markets/europe/italy-approve-145-billion-package-against-inflation-2022-08-04/; Reuters, "No time to waste, worried Italian business leaders warn politicians," Sept. 4, 2022, https://www.reuters.com/world/europe/no-time-to-waste-worried-italian-business-leaders-warn-politicians-2022-09-04/.

where Y_{it} is an indicator equal to one if MP i received a television invitation on day t (defined as speaking for at least 60 seconds in any program), and zero otherwise, across all daily political talk shows and newscasts.

The specification includes day fixed effects δ_t, which control for aggregate time-varying shocks common to all MPs, and MP-by-week fixed effects θ_{iw}, which account for all time-invariant MP-specific characteristics (such as baseline popularity, party affiliation, gender, or pre-existing media connections) as well as time-varying MP-specific trends at the weekly level, such as evolving personal campaigns, local controversies, or seasonal factors unique to each MP (e.g., an MP's weekly push on a policy issue). This flexible fixed-effects structure addresses potential differential pre-trends driven by MP-specific dynamics. The term V_{it} is the number of tweets posted by MP i on day t, controlling for daily tweeting intensity; this addresses potential confounding from strategic or effort-based tweeting behavior, as MPs who tweet more frequently may increase their chances of virality while also signaling newsworthiness to media outlets independently of engagement levels (e.g., high-volume posting on a controversial topic). D_{it}^k is an indicator variable equal to one if MP i experienced a viral tweet exactly k days before day t. Virality is defined at the MP-day level: a day is considered viral if the maximum engagement (sum of likes, retweets, replies, and quote retweets) of any tweet posted by the MP on that day reaches at least the 99th percentile of the distribution of daily maximum engagements across all MP-days with tweets in the sample. For MPs with multiple viral days, I define the first viral day τ_i as the event time and set $k = t - \tau_i$. Subsequent viral days that occur within the post window are not treated as separate events; any influence they have on Y_{it} is part of the treatment effect of experiencing an initial viral event at $k = 0$.

The event study coefficients β_k measure the change in the probability of receiving a television invitation at event time k relative to the baseline period $k = -1$. To assess whether the effect extends to more substantial appearances, I also estimate specifications that require at least 2 minutes and 3 minutes of speaking time. Standard errors are clustered at the MP level. Summary statistics on treatment timing are reported in Appendix Section 2.E.2.

Identification. The interpretation of the event-time coefficients as capturing the effect of virality on television access relies on a parallel trends assumption. Conditional on day fixed effects, MP-by-week fixed effects, and daily tweet volume, MPs who experience a first viral tweet at different points in time would have followed similar paths in their probability of receiving television invitations in the absence of the viral event. In practice, this assumption may fail for several reasons.

First, selection into the timing of virality could arise if MPs tend to go viral precisely when unobserved factors independently raise their television newsworthiness. For example, an MP involved in an emerging scandal might both generate viral tweets and attract producers' attention, which would make virality correlate with invitations even if it has no independent effect. MP-by-week fixed effects θ_{iw} absorb MP-specific factors that evolve at the weekly level, such as ongoing campaigns or local controversies, while day fixed effects δ_t control for aggregate shocks that affect all MPs, such as national political events. MP-specific shocks at the daily level (for instance, an MP being named in a breaking story on a particular day) remain a potential concern, as they are not absorbed by the fixed-effects structure.

Second, two related threats involve the timing and direction of the relationship. Anticipation could arise if MPs change their online behaviour in expectation of upcoming television appearances, for example by tweeting more frequently or adopting more extreme rhetoric when a TV invitation is scheduled, mechanically raising the probability of viral tweets just before they appear on air. Reverse causality could occur if television appearances themselves make subsequent virality more likely, for instance because increased visibility expands the potential online audience. The control for daily tweet volume V_{it} partially mitigates the anticipation concern by holding constant overall tweeting intensity. More generally, both threats imply that MPs should experience elevated invitation rates in the days prior to virality ($k < 0$). The event-study specification is therefore informative about these mechanisms through the pattern of pre-event coefficients, which allows me to assess the presence of systematic pre-trends.

2.4.1 Results

Figure 2.4 presents the dynamic effects of online virality on television appereance from equation 2.1. The figure shows effects on the probability of receiving TV invitations across three threshold definitions: speaking for at least 1 minute, 2 minutes, and 3 minutes in a program. Full results are presented in Table E.3.1 in Appendix 2.E.

My results indicate substantial increases in TV visibility following virality across all threshold specifications. Pre-treatment coefficients are jointly insignificant in all specifications (F-test, p-values > 0.10), supporting the parallel trends

Figure 2.4: Event-Study Estimates of Virality on TV Visibility

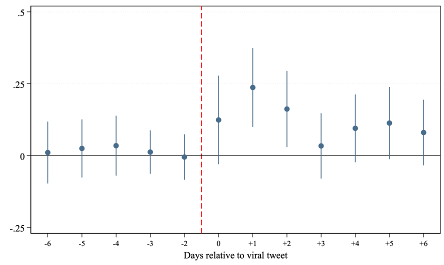

((a)) 1 Minute Threshold

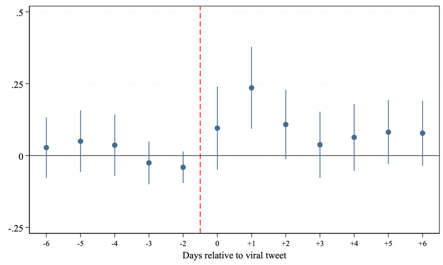

((b)) 2 Minutes Threshold

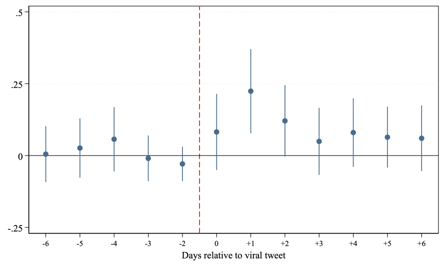

((c)) 3 Minutes Threshold

Notes: Figure shows event-study coefficients from Equation 2.1 using the most demanding specification including tweet volume control, MP, date, and MP×week fixed effects. Panels show effects on the probability of receiving TV invitations across three threshold definitions: Panel A uses a 60-second threshold, Panel B uses a 2-minute threshold, and Panel C uses a 3-minute threshold. Speaking time thresholds measure the total duration an MP speaks directly on air in their own voice; MPs meeting these thresholds are typically present in programs for substantially longer. 95% confidence intervals shown. Reference period is $t = -1$ (omitted).

assumption underlying my identification strategy. Effects emerge the following day after virality and concentrate in the immediate post-viral period.

Using the 1-minute threshold, the probability of receiving TV invitations increases by 23.7 percentage points the following day after virality ($p < 0.01$). This represents more than a doubling relative to the overall sample mean of 11.1%, with MPs transitioning from receiving invitations on roughly one in nine days to nearly one in three days. The effect remains significant at 16.2 percentage points on the second day after virality ($p < 0.05$), before fading thereafter. This demonstrates that viral MPs transition from rarely receiving television invitations to being actively chased by producers who recognize their enhanced newsworthiness.

Critically, these effects persist even when requiring longer speaking time. Using the 2-minute threshold, the next day invitation probability increases by 23.6 percentage points (from 8.1 to 31.7 percent; the increase is $23.6/8.1 \approx 2.91\times$ the baseline; $p < 0.01$). With the most stringent 3-minute threshold, the probability increases by 22.4 percentage points (from 6.8 to 29.2 percent; the increase is

$22.4/6.8 \approx 3.29\times$ the baseline; $p < 0.01$). The consistency of results across increasingly demanding thresholds demonstrates that viral tweets generate substantive television appearances rather than brief mentions.

The concentration of significant effects on the following day and second day after virality reflects the rapid response of TV producers to trending content within 24-48 hours, consistent with the production cycles of daily news programs and political talk shows. Effects fade by the third day in all specifications, supporting the identification strategy and highlighting the concentrated but ephemeral nature of virality-driven media attention.

The substantial relative increases observed across all thresholds reflect the inherently low baseline rates of television presence among MPs. With only 11.1% of MP-days resulting in invitations under the 60-second threshold, even the 23.7 percentage point increase represents a large relative effect. This pattern reflects how media attention operates in practice: producers make discrete decisions about whether someone is newsworthy enough to invite.

These results demonstrate that online virality serves as a powerful gateway to traditional media coverage, generating substantial increases in television visibility that persist even when requiring MPs to secure 2-3 minutes of speaking time. The findings provide evidence for the gatekeeping mechanism whereby social media engagement signals newsworthiness to traditional media producers.

2.4.1.1 Market Incentives: Private vs Public Channels

The main results demonstrate that online virality increases television visibility, but what drives this relationship? I propose that the mechanism operates through market-driven gatekeeping: television producers use viral engagement as a signal of audience interest and newsworthiness when making booking decisions. Under this mechanism, channels facing competitive pressures and reliant on ratings should be most responsive to viral signals, as they seek to attract viewers by featuring MPs who are already generating attention online. Channels with weaker market incentives — such as public broadcasters with guaranteed funding and public service mandates — should be less responsive to viral engagement metrics.

To test this mechanism, I examine whether the effects differ between public and private television channels by estimating Equation 2.1 separately for each channel type[17]. This distinction is particularly relevant in the Italian context, where public

[17] Channels 1–3 are public broadcasters (RAI 1, RAI 2, RAI 3). Channels 4–6 are private networks owned by Mediaset (Rete 4, Canale 5, Italia 1). Channel 7 is La7, a private network owned by Cairo Communication. Channels 8 and 9 are commercial broadcasters: TV8 (Sky Italia) and Nove (Warner Bros. Discovery).

and private broadcasters operate under different institutional constraints. Public broadcasters receive guaranteed government funding and public service mandates that may prioritize balanced political coverage and reduce responsiveness to audience engagement metrics (Gambaro et al., 2025). Private channels, in contrast, rely on advertising revenue tied directly to viewership and thus face stronger incentives to book viral MPs who can attract audiences.

The heterogeneity analysis reveals a clear divergence between public and private broadcasters that supports the market-driven gatekeeping mechanism. As shown in Panel (a) of Figure 2.5, private channels display substantial and statistically significant responses to viral MPs. The probability of receiving invitations on private channels increases by 20.1 percentage points the day after virality ($p < 0.01$), with the clearest and most systematic increases appearing on the day after virality ($t = 1$). In contrast, Panel (b) shows that public channels show virtually no systematic response. Coefficients are generally small, statistically insignificant, and scattered across time, with no clear post-viral pattern (full results available in Table E.3.2).

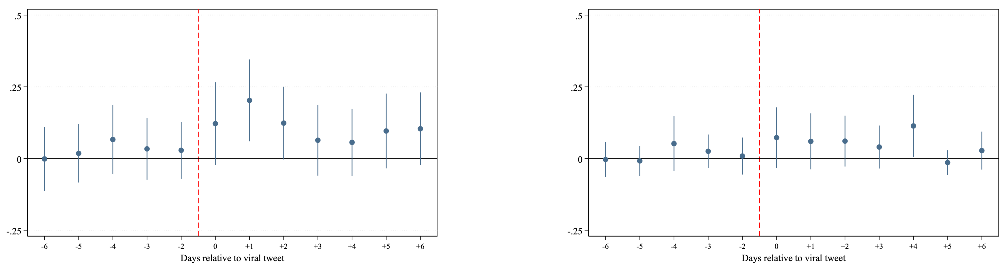

((a)) Private Channels **((b))** Public Channels

Figure 2.5: Effect of Viral Tweets on Television Invitation by Channel Type

Notes: Figure shows event-study coefficients from Equation 2.1 estimated separately for private channels (Panel a) and public channels (Panel b) using the most demanding specification including tweet volume control, MP, date, and MP×week fixed effects. The outcome is the probability of receiving TV invitations defined as speaking for at least 60 seconds in a daily political talk show or newscast program. 95% confidence intervals shown. Reference period is $t = -1$ (omitted).

This differential response is consistent with the market-driven gatekeeping mechanism. Private channels, operating under competitive pressures and reliant on advertising revenue, systematically convert online virality into television visibility. Public channels show little reaction to viral signals, consistent with institutional constraints — including guaranteed funding and public service mandates — that reduce the salience of market-based considerations in booking decisions. The clear contrast between channel types indicates that market incentives, rather than general newsworthiness or political importance alone, drive the translation of viral social media presence into broadcast exposure.

2.5 What Drives Online Virality?

Having established that online virality increases television visibility, I now examine what communication strategies drive viral success. Specifically, I test whether extreme rhetoric generates higher social media engagement than conventional communication, and analyze how the use of these strategies varies across political ideologies. Understanding what drives virality reveals the strategic incentives politicians face in this hybrid media environment.

2.5.1 Extreme Rhetoric and Social Media Engagement

To estimate the engagement response to extreme rhetoric, I estimate the following baseline OLS regression model at the tweet level:

$$Y_{ijkt} = \sum_{\rho} \beta_\rho \mathbb{I}\{\text{Rhetoric}_\rho\}_{ijkt} + \alpha_i + \gamma_k + \delta_t + \epsilon_{ijkt} \tag{2.2}$$

where Y_{ijkt} denotes, alternatively, the total engagement — measured as the sum of likes, retweets, replies, and quote tweets — for tweet j by politician i on topic k at time t, or a binary indicator equal to 1 if the tweet falls in the top 1% of the overall engagement distribution. The index j enumerates tweets within each MP-day-topic combination, as MPs can post multiple tweets on the same topic on the same day; the indices i, k, and t make explicit the structure of the fixed effects.

The key independent variables, $\mathbb{I}\{\text{Rhetoric}_\rho\}$, are indicators for three rhetorical strategies: *Delegitimization*, *Group Discrimination*, and *Zero-Sum Framing*, with conventional, non-extreme rhetoric as the omitted category. The model includes MP fixed effects (α_i) to control for all time-invariant politician characteristics, topic fixed effects (γ_k) to control for differences across topics, and day fixed effects (δ_t) to account for common daily shocks. The coefficients of interest, β_ρ, capture the average change in engagement — either in raw counts or in the probability of being among the top 1% of tweets — associated with each rhetorical strategy, relative to the baseline of non-extreme rhetoric. Standard errors are clustered at the MP level.

The key identifying assumption underpinning Equation (2.2) is that, conditional on MP, topic, and day fixed effects, the timing of extreme rhetorical strategies is exogenous with respect to other unobserved determinants of engagement. A potential violation occurs if politicians strategically deploy extreme rhetoric during high-attention periods (e.g., breaking news or trending events), creating confounding dynamics or pre-existing trends. To address this threat, I employ two robustness strategies. First, I use a more restrictive specification with MP-by-day fixed effects ($\theta_{i\times t}$), absorbing all MP-specific daily-level variation in engagement, thus accounting for any MP-specific short-term shocks or daily attention patterns. Second, I estimate dynamic models incorporating leads and lags of the rhetoric indicators, explicitly testing whether engagement systematically precedes rhetorical choices, which would suggest violations of the exogeneity assumption.

2.5.1.1 Results

Table 2.2 presents the main findings from the regression analysis. The results indicate that all three forms of extreme rhetoric are associated with a large and statistically significant increase in online engagement, a finding that is robust across both total engagement counts and the probability of achieving top-tier engagement. Focusing on the OLS specification with MP-by-day fixed effects for total engagement (Column 2), the coefficients for all three rhetoric types are large and highly significant. The coefficient on *Delegitimization* corresponds to approximately 111 additional engagements relative to a conventional tweet, representing a 39.4% increase relative to the mean of 281.4. Similarly, *Group Discrimination* and *Zero-Sum Framing* are associated with significant increases of approximately 151 and 140 additional engagements, respectively, representing 53.7% and 49.9% increases relative to the mean. The binary outcome specification examining the probability of being in the top 1% of engagement (Columns 3-4) confirms these findings. In the most restrictive specification (Column 4), *Delegitimization* increases the probability of achieving top-tier engagement by 0.7 percentage points, representing a 64% increase relative to the baseline probability of 1.1%. *Group Discrimination* and *Zero-Sum Framing* show even larger effects, increasing the probability by 1.3 percentage points each, representing increases of 118% relative to the mean. Results are robust to different counting-variable specifications, including log+1 transformation, inverse hyperbolic sine (IHS), and PPML models, as shown in Appendix Table D.0.1.

To isolate the contemporaneous effect of rhetoric from spillovers around the focal tweet and to test for confounding pre-trends, I estimate a dynamic model that includes indicators for rhetoric in the two preceding tweets and the two following tweets by the same MP. Index each MP's tweets in chronological order by j. Let R_j denote whether the focal tweet uses extreme rhetoric, R_{j-1} and R_{j-2} the two immediately preceding tweets, and R_{j+1} and R_{j+2} the next two tweets.

Table 2.2: Effect of Extreme Rhetoric on Social Media Engagement

	Engagement Effects			
	Total Engagement		Top 1% Engagement	
	(1) Total Engagement	(2) Total Engagement	(3) Top 1% Engagement	(4) Top 1% Engagement
Group Discrimination	129.394***	150.970***	0.011**	0.013**
	(45.590)	(53.005)	(0.006)	(0.006)
Zero-Sum Framing	123.797***	140.364***	0.011**	0.013**
	(28.908)	(38.475)	(0.004)	(0.006)
Delegitimization	100.110***	110.913***	0.005***	0.007***
	(17.912)	(24.554)	(0.002)	(0.002)
Observations	690773	511074	690773	511074
Mean Dep. Variable	254.691	281.424	0.010	0.011
MP FEs	✓	✓	✓	✓
Topic FEs	✓	✓	✓	✓
Day FEs	✓	✓	✓	✓
MP × Day FEs		✓		✓

Notes: This table presents estimates showing the effect of extreme rhetorical strategies on social media engagement. Columns (1)-(2) use OLS with total engagement count (replies + likes + retweets + quote retweets) as the dependent variable. Columns (3)-(4) use OLS with a binary indicator for being in the top 1% of engagement as the dependent variable. Columns (2) and (4) include MP × day fixed effects in addition to the baseline specification. Standard errors clustered at the MP level are reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 2.6 shows a clear pattern: coefficients on the preceding tweets $(j-2, j-1)$ are small and statistically insignificant across rhetoric types, the contemporaneous effect at j is large and significant, and effects dissipate for the following tweets $(j+1, j+2)$. These results indicate that contemporaneous effects dominate any pre-existing trends, which is inconsistent with strategic timing as the main driver of the findings. Full regression results are reported in Appendix Table D.0.2.

Figure 2.6: Dynamic Effects of Extreme Rhetoric on Engagement

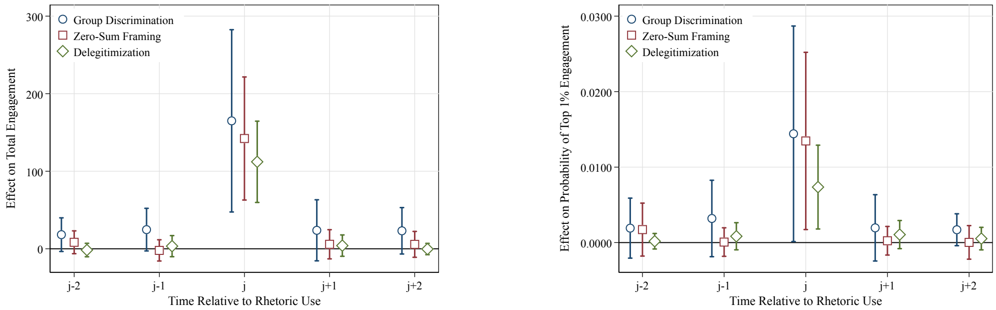

((a)) Total Engagement　　**((b))** Top 1% Engagement

Notes: Tweets for each MP are ordered chronologically by index j. The specification includes rhetoric on the focal tweet R_j, the two preceding tweets R_{j-1}, R_{j-2}, and the two following tweets R_{j+1}, R_{j+2}. Engagement is the sum of likes, retweets, replies, and quote retweets. With 75% of consecutive tweets posted within 24 hours, $j-1$ and $j-2$ serve as pre-trend checks and $j+1$ and $j+2$ capture short-run persistence. Panel (a) reports effects on total engagement. Panel (b) reports effects on the probability of achieving top 1% engagement. All specifications include MP fixed effects, topic fixed effects, day fixed effects, and MP×day fixed effects. 95% confidence intervals shown. Standard errors clustered at the MP level.

Overall, the evidence in Table 2.2 and Figure 2.6 indicates that the social media environment systematically rewards MPs for using extreme rhetoric. The findings are robust across multiple outcome specifications and estimators (including log+1, IHS, and PPML), and the dynamic specification shows contemporaneous effects that dominate small and insignificant pre-trends, consistent with genuine audience responses rather than strategic timing.

2.5.2 Ideological Patterns in Extreme Rhetoric Usage

The previous section shows that extreme rhetoric increases online engagement. This section asks which political actors most frequently employ these strategies. If certain ideological groups use extreme rhetoric more often, they may gain greater visibility through social media algorithms. I therefore analyze how the use of each rhetorical strategy varies across the political spectrum.

To uncover which ideological groups are more likely to employ extreme rhetoric, I estimate the following OLS regression model at the tweet level:

$$\mathbb{I}\{\text{Rhetoric}_\rho\}_{ijkt} = \sum_b \gamma_b \mathbb{I}\{\text{Ideology}_b\}_i + \gamma_k + \delta_t + \epsilon_{ijkt} \tag{2.3}$$

where $\mathbb{I}\{\text{Rhetoric}_\rho\}_{ijkt}$ is an indicator equal to one if tweet j by politician i on topic k at time t contains rhetorical strategy $\rho \in \{\text{Delegitimization}, \text{Group Discrimination}, \text{Zero-Sum Framing}\}$. $\mathbb{I}\{\text{Ideology}_{\text{ideology}}\}_i$ are indicators for four ideological categories: *Extreme-right*,

Right, *Populist*, and *Left*, with *Centrist* as the omitted category. The model includes topic fixed effects (γ_k), to control for systematic differences in rhetorical choices across policy domains and day fixed effects (δ_t) to account for common temporal factors that might influence all MPs' communication strategies simultaneously, such as major political events or platform algorithm changes. The coefficients of interest, γ_{b}, measure the average difference in the probability of using each rhetorical strategy relative to Centrist MPs. Standard errors are clustered at the MP level. I estimate this specification separately for each of the three rhetorical strategies to understand the distinct patterns of usage across the political spectrum.

2.5.2.1 Results

Figure 2.7 presents the differential usage of extreme rhetorical strategies across ideological groups, relative to centrist MPs. The analysis reveals striking asymmetries in the strategic deployment of these communication tactics. Extreme-right MPs employ group-based discriminatory rhetoric at rates 3.8 percentage points higher than centrists ($p < 0.01$), representing by far the largest ideological gap across all rhetorical strategies. This difference is both statistically and economically significant. Right MPs also show 1.1 percentage points higher usage compared to centrists ($p < 0.01$). The same does not hold for left MPs, whose engagement is only 0.2 percentage points higher than centrists, a difference that is not statistically significant. Populist MPs similarly show no meaningful difference from centrists, with an estimated increase of just 0.1 percentage points.

Delegitimizing rhetoric is more evenly distributed across the political spectrum. Right MPs display the highest levels, at 1.5 percentage points above Centrists, while Extreme-right MPs show a smaller increase of 0.7 points, though neither difference is statistically significant. Interestingly, this suggests that conventional right MPs, rather than their more extreme counterparts, are somewhat more inclined to question institutional legitimacy. Left MPs and populist MPs both use this rhetoric slightly less than Centrists (–0.4 and –1.2 points, respectively), but these estimates are also not significant. One possible explanation for this weak ideological gradient is that delegitimizing rhetoric is not tied to a single political ideology but serves as a general oppositional tool, used across the spectrum depending on government status, political crises, or corruption scandals.

Zero-sum framing reveals a notable ideological asymmetry. Extreme-right MPs lead with usage rates 1.2 percentage points above Centrists ($p < 0.05$), while Right MPs show a marginal, non-significant increase of 0.1 points. In contrast, Left and Populist MPs use this strategy less than Centrists (–0.5 and –0.9 points, respectively), though neither difference is statistically significant.

Taken together, these patterns point to systematic ideological asymmetries in the use of extreme rhetorical strategies (detailed regression results in Appendix 2.D, Table D.0.3). Right leaning politicians, especially those on the extreme right, more frequently employ attention oriented rhetoric than their centrist or left leaning counterparts. The clearest gap is in group based discrimination, and zero sum framing is also higher for the extreme right (the mainstream Right coefficient is small and not statistically significant). By contrast, delegitimizing rhetoric shows no clear ideological gradient. Combined with the engagement differentials in Table 2.2, these asymmetries suggest that some ideological groups are better positioned to exploit online attention, potentially producing unequal patterns of visibility across the political spectrum.

Figure 2.7: Extreme Rhetoric Usage by Political Ideology

((a)) Group Discrimination

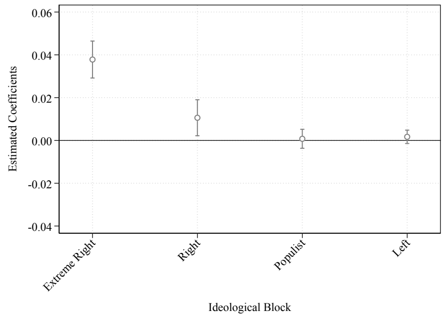

((b)) Delegitimization

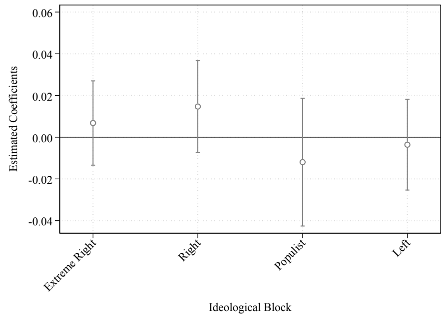

((c)) Zero-Sum Framing

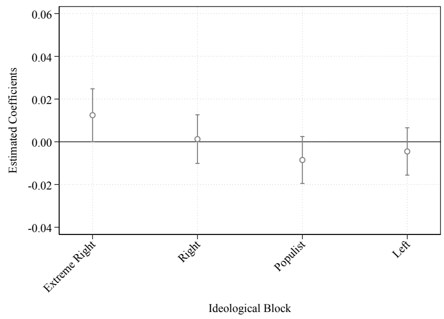

Notes: This figure shows the differential usage of extreme rhetorical strategies across ideological groups, relative to centrist MPs. Each panel presents coefficients γ_b from separate regressions following equation (2.3), controlling for topic and day fixed effects. Standard errors are clustered at the MP level. Extreme-right MPs show the highest usage rates for group discrimination and zero-sum framing, with statistically significant effects at 5% in both categories. For delegitimization, usage patterns are more distributed across ideologies with no statistically significant differences.

2.6 Robustness

Alternative Outcomes. To examine whether virality effects extend across different dimensions of television access, I estimate the event-study specification using three alternative outcomes: the number of distinct programs MPs appear in, total speaking time (IHS transformation), and MPs' share of speaking time within their ideological bloc. The first two outcomes test whether viral MPs gain access to multiple programs and receive substantial airtime, moving beyond the binary extensive margin. The third outcome tests for zero-sum competition under attention constraints: if television slots are limited, viral MPs should increase their share of total airtime within their ideology, crowding out non-viral same-ideology competitors. Results show a consistent pattern across all three outcomes and they mirror the main results: effects concentrate in the day immediately following virality and dissipate by the third day, consistent with the rapid production cycles of daily news programs. Results are presented in Appendix figure F.1.1.

Placebo test. I implement a permutation test that randomly reassigns each MP's viral dates $B = 500$ times, preserving the number of viral days per MP and enforcing the no overlap rule used in the main design. The distribution of placebo day after virality coefficients is centered near zero (mean -0.002, s.d. 0.049). The actual estimate, 0.237, exceeds all 500 placebo draws. Fewer than 0.2% of random assignments generate effects as large as the observed one. This supports the interpretation that television responses follow genuine viral shocks rather than statistical artifacts (Appendix Figure F.2.1).

Staggered Treatment Timing. The staggered timing of virality in the event-study design raises potential concerns about biases in conventional two-way fixed effects estimators, particularly under treatment effect heterogeneity (Goodman-Bacon, 2021; Sun and Abraham, 2021). To address this, I implement the estimator developed by Sun and Abraham (2021), which corrects for the problematic comparisons and negative weighting inherent in TWFE. Re-estimating the primary specification with this method yields nearly identical point estimates and standard errors for the three outcome of interests: 1-, 2-, and 3-minutes speaking time threshold (see Appendix Figure F.3.1).

2.7 A Model of Media Feedback and Political Exaggeration

I develop a theoretical framework to explain why politicians strategically exaggerate rhetoric to generate online virality and secure television exposure. The model captures a career concerns logic (Holmström, 1999), where exaggeration acts as a signal device to boost media rewards. This yields an equilibrium in which politicians exaggerate not to fool citizens, who verify the true state and punish the distortion, but to manipulate algorithmic signals for career visibility, mirroring how agents exert costly effort to influence future opportunities despite rational observers. The model thus formalizes the politician's *strategic supply* of rhetoric, showing how this choice endogenously responds to the conflicting incentives of the media environment.

2.7.1 Setup

The model features a single strategic **Politician (P)**, a continuum of **Citizens**, and a **Television (TV)** station.

The State and the Message. Nature draws a latent state $\omega \sim \mathcal{N}(0, \sigma_\omega^2)$ truncated to $[-\Omega, \Omega]$, where $|\omega|$ denotes event severity. The politician commits ex ante to a linear communication strategy $x = \rho\omega$ with $\rho \geq 0$, reflecting a fixed communication style set by the party; ρ is not publicly observed and can be read as a party guideline that the politician follows on average.

Citizen Engagement: Rational vs. Expressive. The politician's message generates two distinct outcomes:
Rational Audience (A): Citizens face quadratic loss $U_i(a_i, \omega) = -(a_i - \omega)^2$ from mismatching their action a_i to the state. With heterogeneous attention costs $C_i \sim U[0, \bar{C}]$, a share of citizens will pay the cost to become informed. Citizens observe the message x but are uncertain about the specific exaggeration factor ρ, which makes x an imperfect signal of ω; in addition, the public virality index does not reveal ω (see below). To micro-found the attention decision, I assume citizens decide whether to pay C_i after observing the event's severity ($s = \omega^2$) but before choosing their action; the uninformed default is $a_i = 0$. The benefit from becoming informed is therefore the reduction in quadratic loss, ω^2. Hence all citizens with $C_i \leq \omega^2$ pay the cost, and the attentive audience size is $A(\omega) = \omega^2/\bar{C}$.
Expressive Virality (v): Separately, expressive reactions (likes, replies, reposts) are driven by the message's extremity $|x|$. The platform produces a publicly

visible virality index

$$v(\omega, \rho) = \gamma(\rho\omega)^2 + \varepsilon, \qquad \mathbb{E}[\varepsilon] = 0, \ \varepsilon \ \perp \ (\omega, \rho),$$

where $\gamma > 0$ is an amplification parameter. The quadratic form is a tractable way to capture empirically consistent convex engagement, where more extreme messages gain disproportionately larger reactions and algorithmic amplification.[18]. Results are unchanged if $v = \gamma|\rho\omega|^{\eta}$ with $\eta > 1$, or with asymmetric curvature that is steeper for $\rho > 1$.

The Television Station. TV maximizes ratings, observing the public virality index v as a signal of audience buzz. It faces a daily newsworthiness shock $K \sim U[0, \bar{K}]$, independent of v, representing competing stories/people. TV invites the politician if $\beta v \geq K$, where $\beta > 0$ is a ratings conversion parameter. This implies an invitation probability $p(v) = \beta v/\bar{K}$ on its support. The assumption $\beta\, v_{\max} < \bar{K}$ ensures $p(v) < 1$.

2.7.2 The Politician's Problem

The politician chooses ρ to maximize expected utility, balancing the benefits of audience and TV exposure against the reputational cost of lying. The per-state payoff is:

$$U_P = A(\omega) + B \cdot \mathbf{1}\{\beta v \geq K\} - [\lambda_o + \lambda_T \cdot \mathbf{1}\{\beta v \geq K\}](x - \omega)^2,$$

where $B > 0$ is the TV bonus, $\lambda_o > 0$ is the baseline reputational cost for the statement made online, and $\lambda_T > 0$ is the additional penalty incurred from the greater scrutiny and larger audience of a television appearance. The rational audience $A(\omega)$ provides a baseline reach term, while reputational costs are imposed by that same audience for the observed distortion $(x - \omega)^2$. Averaging over K, ε, and ω (with $\mathbb{E}[\varepsilon] = 0$) yields:

$$V(\rho) = \frac{\mathbb{E}[\omega^2]}{\bar{C}} + \frac{B\beta\gamma}{\bar{K}}\rho^2\mathbb{E}[\omega^2] - \lambda_o(\rho - 1)^2\mathbb{E}[\omega^2] - \frac{\lambda_T\beta\gamma}{\bar{K}}\rho^2(\rho - 1)^2\mathbb{E}[\omega^4].$$

This framework makes a precise distinction between the *supply* and *demand* of rhetoric. The model is supply-side, as the politician's strategic supply is the endogenous choice of ρ. The demand side is treated as exogenous parameters that shape incentives in the short run: (1) audience discipline, captured by reputational

[18] I model virality as a function of message extremity x^2 rather than the magnitude of the lie $(x - \omega)^2$. The focus is on exaggeration. For simplicity I restrict attention to $\rho \geq 1$; allowing $\rho \in [0, 1)$ ("downplaying") cannot raise virality (since $v \propto (\rho\omega)^2$) and still increases reputational costs relative to honesty ($\rho = 1$)

costs (λ_o, λ_T); and (2) the environmental conversion of extremity into visibility, determined by platform algorithms and TV gatekeeping ($\gamma, \beta, B, \bar{K}$). Differences across parties or ideologies are interpreted as differences in these parameters.

This supply-side framework rationalizes empirical asymmetries in extreme rhetoric. If extreme-right politicians supply more extreme rhetoric, the model treats this as a rational response to their incentive structure through two channels. First, they may face lower audience discipline (lower λ_o or λ_T), consistent with evidence on asymmetric tolerance for in-group deception (Roets et al., 2019; DeVerna et al., 2024). Second, they may face a higher environmental reward for visibility (higher B), especially for insurgents or smaller parties (e.g., Brothers of Italy, 4.33% in the 2018 elections and now the leading party), where the marginal value of a single TV appearance is greater. The choice of exaggeration is therefore a rational supply response to these trade-offs.

2.7.3 Equilibrium and Comparative Statics

The optimal ρ^* solves the first-order condition $V'(\rho^*) = 0$. The second-order condition, derived in the Appendix, ensures the solution is a unique maximum, reflecting a trade-off where disciplinary forces (audience discipline and media competitiveness) outweigh temptations (rewards from virality and TV exposure).

Proposition 1 (Incentive for Exaggeration) *The unique equilibrium involves a strictly positive level of exaggeration ($\rho^* > 1$). This is because the marginal benefit of the first infinitesimal amount of exaggeration is positive, while its marginal cost is zero. Honesty ($\rho = 1$) is therefore never an optimal strategy.*

Proposition 2 (Media and Platform Incentives (Environmental Conversion)) *The level of exaggeration (ρ^*) increases with the value of the TV prize (B), the effectiveness of platform algorithms (γ), and the conversion of buzz into ratings (β), but falls with the competitiveness of the TV news environment ($\bar{K}$):*

$$\frac{\partial \rho^*}{\partial B} > 0, \quad \frac{\partial \rho^*}{\partial \gamma} > 0, \quad \frac{\partial \rho^*}{\partial \beta} > 0, \quad \frac{\partial \rho^*}{\partial \bar{K}} < 0.$$

Proposition 3 (Reputational Costs (Audience Discipline)) *The level of exaggeration (ρ^*) falls as the penalties for dishonesty (audience discipline) increase:*

$$\frac{\partial \rho^*}{\partial \lambda_o} < 0, \quad \frac{\partial \rho^*}{\partial \lambda_T} < 0.$$

A result of the model is that the cost of rational online attention, $\bar{C}$, does not affect the politician's exaggeration strategy ($\partial\rho^*/\partial\bar{C} = 0$), because the incentive to exaggerate is driven entirely by the TV pathway.

2.8 Conclusion

This paper examines whether viral success on social media serves as a bridge to traditional media access, and what communication strategies enable politicians to cross that bridge. Using comprehensive data on Italian MPs' Twitter activity and television appearances during 2018–2022, I document a media feedback loop operating through online virality.

I find evidence of two interconnected mechanisms. First, online virality translates into increased television visibility: MPs experiencing viral tweets see substantial increases in their probability of receiving TV invitations the following day, with effects persisting even when requiring substantial speaking time. These effects concentrate in private channels facing more competitive pressures, while public broadcasters show no systematic response, consistent with market-driven rather than newsworthiness based gatekeeping. Second, extreme rhetoric drives virality: tweets employing institutional delegitimization, group-based discrimination, or zero-sum framing generate markedly more engagement than conventional communication. Third, this pathway operates asymmetrically: extreme-right MPs employ discriminatory rhetoric at substantially higher rates than other ideological groups, revealing systematic variation in who can exploit online attention dynamics. My theoretical model clarifies why this feedback loop persists even when citizens are not deceived. Television producers observe viral engagement as a noisy but publicly visible signal of audience interest. A single politician choose rhetorical strategies trading off reputational costs against increased visibility benefits. Citizens engage expressively with extreme content due to arousal or entertainment value while rationally discounting exaggeration when making decisions. This creates documented incentives for rhetorical intensification: extreme rhetoric increases the probability of achieving viral status, and viral status systematically increases television access, particularly on market-driven channels. The implications extend beyond individual communication choices. The systematic reward structure creates differential access to traditional media based on willingness and capacity to employ extreme rhetoric. Certain ideological groups demonstrate higher propensity to exploit this pathway, potentially reshaping who gains visibility in hybrid media environments where online engagement serves as gateway to broadcast exposure.

Future research should examine whether politicians adapt their communication strategies when transitioning from social media to television — are online and offline communication strategies aligned? — and whether television functions primarily as a communication channel or as an entertainment-driven political arena even when it comes to political information.

Appendix

Appendix 2.A Appendix: Full Model Derivations

1. Information-Seeking Audience (A): A citizen i with attention cost C_i decides whether to pay the cost to learn the true state ω. As argued in the main text, citizens observe x but are uncertain about the specific ρ, so x is an imperfect signal of ω. To micro-found this decision, I assume citizens first observe the event's *severity* ($s = \omega^2$) but not its *direction* (ω). Their default (uninformed) action is $a_i = \mathbb{E}[\omega|s] = 0$. The loss from this default action is $-(0-\omega)^2 = -\omega^2$. Paying C_i reveals the true state ω, allowing an optimal action $a_i = \omega$ (loss of 0). The benefit of paying is thus the reduction in quadratic loss, ω^2. Therefore, a citizen attends if $C_i \leq \omega^2$. Given $C_i \sim U[0, \bar{C}]$, the attentive share is

$$A(\omega) = \frac{\omega^2}{\bar{C}}.$$

This shows that the size of the rational audience in equilibrium depends on the true state, not on the level of exaggeration.

2. Expressive Virality (v): Separately, the observable virality score is a convex function of message extremity, amplified by platform algorithms:

$$v = \gamma(\rho\omega)^2 + \varepsilon, \qquad \mathbb{E}[\varepsilon] = 0,\ \varepsilon \perp (\omega, \rho).$$

This parsimoniously captures that more extreme messages trigger disproportionately larger engagement and amplification, while the public index v is an imperfect signal of the underlying state.

2.A.1 The Model's Core Tension: Rational Audience vs. Expressive Virality

A key feature of the model is the distinction between two public responses to the message: the formation of a rational, informed audience (A) and the generation of an expressive virality score (v). These components are not mechanically linked; they arise from different motivations and affect incentives through separate channels. The rational audience A consists of citizens who learn ω and act as truth auditors, imposing reputational costs via $(x-\omega)^2$ while providing a baseline reach term $A(\omega)$. In contrast, virality $v = \gamma(\rho\omega)^2 + \varepsilon$ signals newsworthiness to TV through expressive engagement (likes, shares) driven by extremity; it raises the invitation probability and expected TV bonus B but does not reveal ω exactly. Exaggeration is therefore costly through audience discipline yet beneficial through TV rewards, and the choice of ρ^* resolves this trade-off.

Step 2: Expectation over the State (ω)

Next, take expectations of U_P over K, ε, and ω. Substituting A, v, and $p(v)$ gives

$$V(\rho) = \mathbb{E}\left[\frac{\omega^2}{\bar{C}}\right] + \frac{B\beta}{\bar{K}}\,\mathbb{E}[v] - \lambda_o(\rho-1)^2\,\mathbb{E}[\omega^2] - \frac{\lambda_T\beta}{\bar{K}}\,\mathbb{E}[v\,\omega^2].$$

With $v = \gamma(\rho\omega)^2 + \varepsilon$, $\mathbb{E}[\varepsilon] = 0$, and $\varepsilon \perp (\omega, \rho)$,

$$\mathbb{E}[v] = \gamma\rho^2\,\mathbb{E}[\omega^2], \qquad \mathbb{E}[v\,\omega^2] = \gamma\rho^2\,\mathbb{E}[\omega^4].$$

Therefore,

$$V(\rho) = \mathbb{E}_\omega\left[\frac{\omega^2}{\bar{C}} + B\left(\frac{\beta\gamma(\rho\omega)^2}{\bar{K}}\right) - \left[\lambda_o + \lambda_T\left(\frac{\beta\gamma(\rho\omega)^2}{\bar{K}}\right)\right](\rho-1)^2\omega^2\right].$$

Taking $\mathbb{E}_\omega[\cdot]$ term by term yields:

- **Expected Online Audience:** $\mathbb{E}_\omega\left[\frac{\omega^2}{\bar{C}}\right] = \frac{1}{\bar{C}}\mathbb{E}[\omega^2]$.
- **Expected TV Bonus:** $\mathbb{E}_\omega\left[\frac{B\beta\gamma\rho^2\omega^2}{\bar{K}}\right] = \frac{B\beta\gamma\rho^2}{\bar{K}}\mathbb{E}[\omega^2]$.
- **Expected Baseline Lying Cost:** $\mathbb{E}_\omega\left[-\lambda_o(\rho-1)^2\omega^2\right] = -\lambda_o(\rho-1)^2\mathbb{E}[\omega^2]$.
- **Expected TV Lying Cost:** $\mathbb{E}_\omega\left[-\lambda_T\left(\frac{\beta\gamma(\rho\omega)^2}{\bar{K}}\right)(\rho-1)^2\omega^2\right] = -\frac{\lambda_T\beta\gamma\rho^2(\rho-1)^2}{\bar{K}}\mathbb{E}[\omega^4]$.

Resulting Expected Utility

Combining these four components gives the final expression for $V(\rho)$:

$$V(\rho) = \left(\frac{1}{\bar{C}}\mathbb{E}[\omega^2]\right) + \left(\frac{B\beta\gamma}{\bar{K}}\mathbb{E}[\omega^2]\right)\rho^2 - \left(\lambda_o\mathbb{E}[\omega^2]\right)(\rho-1)^2 - \left(\frac{\lambda_T\beta\gamma}{\bar{K}}\mathbb{E}[\omega^4]\right)\rho^2(\rho-1)^2.$$

2.A.2 First-Order Condition (FOC)

The FOC that defines the optimal ρ^* is $V'(\rho^*) = 0$, which can be expressed as $MB(\rho^*) = MC(\rho^*)$.

- **Marginal Benefit (MB):** $MB(\rho) = \left(\frac{2B\beta\gamma\mathbb{E}[\omega^2]}{\bar{K}}\right)\rho$.
- **Marginal Cost (MC):** $MC(\rho) = 2\lambda_o(\rho-1)\mathbb{E}[\omega^2] + \left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1)$.

At $\rho = 1$, $MB(1) > 0$ while $MC(1) = 0$. Since the marginal gain from the first unit of exaggeration exceeds its cost, the solution must involve $\rho^* > 1$.

2.A.3 Second-Order Condition (SOC)

To ensure the FOC identifies a maximum, require $V''(\rho) < 0$, equivalently $MC'(\rho) > MB'(\rho)$.

- **Slope of MB:** $MB'(\rho) = \dfrac{2B\beta\gamma\,\mathbb{E}[\omega^2]}{\bar{K}} > 0$ (a positive constant).
- **Slope of MC:** $MC'(\rho) = 2\lambda_o\,\mathbb{E}[\omega^2] + \left(\dfrac{2\lambda_T\beta\gamma\,\mathbb{E}[\omega^4]}{\bar{K}}\right)(6\rho^2 - 6\rho + 1) > 0$ (positive and accelerating in ρ).

Since $MC'(\rho)$ accelerates with ρ^2 while $MB'(\rho)$ is constant, the cost of each additional unit of exaggeration eventually outweighs its benefit, guaranteeing a finite solution ρ^*.

To ensure uniqueness and stability for all $\rho \geq 1$, impose the sufficient condition $\lambda_o\bar{K} > B\beta\gamma$. This is a deliberately conservative bound: it drops the positive λ_T term from the disciplinary side of the full $V'' < 0$ condition, which only strengthens the requirement while keeping the expression transparent. The economics are a trade-off between "discipline" and "temptation." The right-hand side, $B\beta\gamma$, summarizes the temptation to exaggerate (TV prize size B, algorithmic amplification γ, and conversion to ratings β). The left-hand side, $\lambda_o\bar{K}$, summarizes discipline (baseline aversion to lying λ_o and a more competitive TV environment $\bar{K}$ that raises the bar for coverage). The condition requires discipline to exceed temptation.

2.A.4 Comparative Statics Derivations

Apply the Implicit Function Theorem to the FOC, $F(\rho,\theta) = MB(\rho,\theta) - MC(\rho,\theta) = 0$. Since the SOC implies $\partial F/\partial\rho < 0$, the sign of the comparative static $\partial\rho^*/\partial\theta$ is determined by the sign of $\partial F/\partial\theta$. The sign of $\partial F/\partial\theta$ is evaluated for each parameter θ below.

2.A.4.1 TV Bonus (B)

The parameter B only appears in the Marginal Benefit term.

$$\frac{\partial F}{\partial B} = \frac{\partial}{\partial B}MB(\rho) = \frac{\partial}{\partial B}\left[\left(\frac{2B\beta\gamma\mathbb{E}[\omega^2]}{\bar{K}}\right)\rho\right] = \left(\frac{2\beta\gamma\mathbb{E}[\omega^2]}{\bar{K}}\right)\rho.$$

Since all parameters are positive and $\rho^* > 1$, I have $\frac{\partial F}{\partial B} > 0$, which implies $\dfrac{\partial\rho^*}{\partial B} > 0$. *A larger prize for a TV appearance increases the marginal benefit of exaggeration, leading to more spin.*

2.A.4.2 Lying Penalties (λ_o, λ_T)

These parameters only appear in the Marginal Cost term.

$$\frac{\partial F}{\partial \lambda_o} = -\frac{\partial MC(\rho)}{\partial \lambda_o} = -\frac{\partial}{\partial \lambda_o}\left[2\lambda_o(\rho-1)\mathbb{E}[\omega^2] + \dots\right] = -2(\rho-1)\mathbb{E}[\omega^2].$$

Since $\rho^* > 1$, I have $\frac{\partial F}{\partial \lambda_o} < 0$, which implies $\dfrac{\partial \rho^*}{\partial \lambda_o} < 0$.

$$\frac{\partial F}{\partial \lambda_T} = -\frac{\partial MC(\rho)}{\partial \lambda_T} = -\frac{\partial}{\partial \lambda_T}\left[\dots + \left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1)\right] = -\left(\frac{2\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1).$$

Since $\rho^* > 1$, I have $\frac{\partial F}{\partial \lambda_T} < 0$, which implies $\dfrac{\partial \rho^*}{\partial \lambda_T} < 0$. *Higher baseline or TV-related reputational costs increase the marginal cost of exaggeration, forcing less spin.*

2.A.4.3 Algorithmic & Media Parameters ($\gamma, \beta, \bar{K}$)

These parameters appear in both the Marginal Benefit and Marginal Cost terms. I determine the sign of $\partial F/\partial\theta$ for each.

- **Virality Amplification (γ):** I first calculate the derivative of the FOC with respect to γ:

  $$\frac{\partial F}{\partial \gamma} = \frac{\partial MB(\rho)}{\partial \gamma} - \frac{\partial MC(\rho)}{\partial \gamma}$$

  Calculating each part separately:

  $$\frac{\partial MB}{\partial \gamma} = \frac{\partial}{\partial \gamma}\left[\left(\frac{2B\beta\gamma\mathbb{E}[\omega^2]}{\bar{K}}\right)\rho\right] = \left(\frac{2B\beta\mathbb{E}[\omega^2]}{\bar{K}}\right)\rho > 0$$

  $$\frac{\partial MC}{\partial \gamma} = \frac{\partial}{\partial \gamma}\left[\dots + \left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1)\right] = \left(\frac{2\lambda_T\beta\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1) > 0$$

  Since both terms are positive, the sign of $\frac{\partial F}{\partial \gamma}$ is ambiguous. To resolve this, I use the FOC, $MB(\rho^*) = MC(\rho^*)$, which I can write as:

  $$\left(\frac{2B\beta\gamma\mathbb{E}[\omega^2]}{\bar{K}}\right)\rho^* = 2\lambda_o(\rho^*-1)\mathbb{E}[\omega^2] + \left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho^*(\rho^*-1)(2\rho^*-1)$$

  Notice that $\frac{\partial MB}{\partial \gamma} = \frac{1}{\gamma}MB(\rho^*)$ (because the marginal benefit is proportional to γ). Substituting the full expression for $MB(\rho^*)$ from the FOC into the derivative gives:

  $$\frac{\partial F}{\partial \gamma} = \frac{1}{\gamma}\left(2\lambda_o(\rho-1)\mathbb{E}[\omega^2] + \left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1)\right) - \left(\frac{2\lambda_T\beta\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1)$$

The second and third terms cancel out (after distributing the $1/\gamma$ factor), leaving:

$$\frac{\partial F}{\partial \gamma} = \frac{1}{\gamma}\left[2\lambda_o(\rho-1)\mathbb{E}[\omega^2]\right] > 0. \quad \text{Thus, } \frac{\partial \rho^*}{\partial \gamma} > 0.$$

More effective platform algorithms make exaggeration a more potent tool, increasing spin.

- **Ratings Converter (β):** The derivation is mathematically identical to that for γ, yielding:

$$\frac{\partial F}{\partial \beta} > 0 \implies \frac{\partial \rho^*}{\partial \beta} > 0.$$

If online buzz converts more easily to TV ratings, the incentive to create buzz via exaggeration increases.

- **TV News Threshold ($\bar{K}$):** I calculate the derivative of the FOC with respect to $\bar{K}$:

$$\frac{\partial F}{\partial \bar{K}} = \frac{\partial MB(\rho)}{\partial \bar{K}} - \frac{\partial MC(\rho)}{\partial \bar{K}}$$

Calculating each part:

$$\frac{\partial MB}{\partial \bar{K}} = -\left(\frac{2B\beta\gamma\mathbb{E}[\omega^2]}{\bar{K}^2}\right)\rho < 0$$

$$\frac{\partial MC}{\partial \bar{K}} = -\left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}^2}\right)\rho(\rho-1)(2\rho-1) < 0$$

Again, the sign is ambiguous. I use the FOC by noting that $\frac{\partial MB}{\partial \bar{K}} = -\frac{1}{\bar{K}}MB(\rho^*)$. Substituting for $MB(\rho^*)$:

$$\frac{\partial F}{\partial \bar{K}} = -\frac{1}{\bar{K}}\left(2\lambda_o(\rho-1)\mathbb{E}[\omega^2] + \left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}}\right)\rho(\rho-1)(2\rho-1)\right) + \left(\frac{2\lambda_T\beta\gamma\mathbb{E}[\omega^4]}{\bar{K}^2}\right)\rho(\rho-1)(2\rho-1)$$

The second and third terms cancel out (after distributing the $-1/\bar{K}$ factor), leaving:

$$\frac{\partial F}{\partial \bar{K}} = -\frac{1}{\bar{K}}\left[2\lambda_o(\rho-1)\mathbb{E}[\omega^2]\right] < 0. \quad \text{Thus, } \frac{\partial \rho^*}{\partial \bar{K}} < 0.$$

A more competitive TV market (higher bar for entry) reduces the marginal value of exaggeration, leading to less spin.

2.A.4.4 Online Attention Cost ($\bar{C}$)

The parameter $\bar{C}$ determines the size of the rational audience, A, but does not appear in the Marginal Benefit or Marginal Cost functions that determine the choice of ρ^*.

$$\frac{\partial F}{\partial \bar{C}} = 0 \implies \frac{\partial \rho^*}{\partial \bar{C}} = 0.$$

The politician's exaggeration strategy is driven by the incentive to get on TV and is therefore independent of the difficulty of attracting a separate, rational online audience.

Appendix 2.B Measurement of Extreme Political Rhetoric

Researchers studying political communication increasingly rely on bounded measures like Google Perspective API's toxicity score to identify problematic content. While these tools offer scalability, this appendix demonstrates that they conflate emotional intensity with anti-democratic extreme rhetoric, limiting their utility for research on extreme political discourse. I show that categorical measures targeting specific rhetorical strategies provide greater precision for identifying the mechanisms through which political communication happens.

This appendix proceeds in three parts. First, I present the complete LLM prompt used to classify extreme rhetoric, detailing the conservative coding approach that prioritizes specificity over sensitivity. Second, I demonstrate empirically that toxicity measures conflate distinct rhetorical strategies, while my categorical measures capture theoretically distinct mechanisms of democratic norm violation. Third, I validate the LLM classification through manual coding of 10,000 tweets, confirming high precision and substantial inter-rater reliability across all three rhetorical dimensions.

2.B.1 Fine-tuned Prompt

This section provides the complete prompt used for the automated classification of extreme rhetoric in Italian MPs' tweets, as described in Section **??**. The classification employs *Mistral 24B* to identify three theoretically grounded rhetorical strategies: delegitimization, group-based discrimination, and zero-sum framing.

The prompt follows a conservative coding rule, flagging binary dimensions only when there is explicit evidence of the relevant rhetorical pattern. This design prioritizes specificity over sensitivity to minimize false positives and ensure that legitimate political disagreement is not conflated with anti-democratic discourse.

> You are an expert in Italian political discourse analysis.
>
> Your task is to analyze a political tweet written in Italian, detecting three specific types of extreme rhetorical elements. For each element, return 0 (absent) or 1 (clearly present). Choose 1 only if the language explicitly and unambiguously matches the definition.
>
> Required JSON schema:

```
{
  "tweet_id": integer,
  "group_discrimination": 0 or 1,
  "delegitimization": 0 or 1,
  "zero_sum_framing": 0 or 1
```

```
}
```

Return exactly this structure – no extra keys, no comments.

Dimension definitions with Italian examples:

group_discrimination – Negative generalizations, stereotyping, or dehumanizing language targeting people based on ethnicity, nationality, religion, gender, sexuality, or immigration status.

- Examples: "i Rom rubano sempre", "clandestini invasori", "islamici terroristi", "comunità LGBT corrompono", "stranieri delinquenti", "zingari sporchi"
- Include: Coded language like "risorsa" (ironic), "diversamente italiani", demographic scapegoating
- Exclude: Criticism of specific policies, individual politicians, or political parties

delegitimization – Direct attacks on the legitimacy, integrity, or democratic nature of institutions, electoral processes, judicial systems, or media.

- Examples: "elezioni truccate", "magistratura comunista", "giustizia a orologeria", "regime mediatico", "dittatura sanitaria", "golpe giudiziario", "stampa venduta"
- Include: Claims of systemic corruption, rigged systems, fake democracy, puppet institutions
- Exclude: Criticism of specific decisions, calls for reform, or disagreement with outcomes

zero_sum_framing – Explicit presentation of political competition as a winner-takes-all scenario where one group's success necessarily requires another's defeat or loss.

- Examples: "o noi o loro", "se vincono loro, l'Italia perde", "ci sostituiranno", "o si salva l'Italia o vincono i nemici", "prima gli italiani significa dopo tutti gli altri"
- Include: Replacement narratives, existential political framing, mutually exclusive victory language
- Exclude: General competitive language, standard electoral rhetoric, policy trade-offs

Coding guidelines:

- group_discrimination: Code 1 if negative traits attributed to entire demographic groups, dehumanizing language, ethnic/religious stereotypes. Code 0 if criticism focuses on policies, specific leaders, or actions rather than group identity.

- delegitimization: Code 1 if claims that institutions are fundamentally corrupt, rigged, or illegitimate. Code 0 if standard criticism, calls for improvement, or disagreement with decisions.
- zero_sum_framing: Code 1 if explicit "us vs them" where one side's victory means the other's destruction/replacement. Code 0 if normal political competition, policy debates, or electoral messaging.

Italian political context considerations:

- Regional rhetoric: North vs South generalizations count as group discrimination
- Anti-establishment language: Distinguish between populist criticism (code 0) and institutional delegitimization (code 1)
- Immigration discourse: "Invasione" narratives often combine group discrimination with zero-sum framing
- EU skepticism: Sovereignty arguments vs institutional conspiracy claims
- Sarcasm/irony: Code based on intended meaning, not literal interpretation

Decision rule: When uncertain, choose 0. Only code 1 when the rhetorical element is explicit and unambiguous.

2.B.2 Conflation vs. Precision

Table B.1: Correlation Matrix: Toxicity Conflation vs. Categorical Precision

Variable	Toxicity	Group Disc.	Delegit.	Zero-Sum
Toxicity	1.00	0.33	0.29	0.17
Group Discrimination	0.33	1.00	0.01	0.08
Delegitimization	0.29	0.01	1.00	0.09
Zero-Sum Framing	0.17	0.08	0.09	1.00

Notes: The correlation matrix reveals the fundamental difference between bounded and categorical approaches. My categorical measures exhibit low inter-correlations (average 0.062), indicating they capture theoretically distinct mechanisms. In contrast, toxicity shows much stronger correlations with each measure (average 0.263), demonstrating that it conflates different types of anti-democratic rhetoric into a single instrument.

Table B.1 provides evidence of the conflation problem. My three categorical measures, namely, group discrimination, delegitimization, and zero sum framing, exhibit very low inter-correlations (average correlation of 0.062), indicating they capture theoretically distinct mechanisms of democratic norm violation. In contrast, toxicity shows substantially higher correlations with each of my measures (average correlation of 0.263). Rather than isolating specific mechanisms, toxicity conflates different types of problematic content into a single measure. This conflation prevents from distinguishing between different rhetorical strategies that may have distinct effects on online engagement.

The limitations of toxicity measures become clear when examining specific content. In fact, 50.1% of tweets in the top decile of toxicity contain no extreme rhetoric according to my categorical measures, while 2.0% of below-median toxicity tweets do contain such rhetoric. This mismatch demonstrates that emotional intensity and anti-democratic content might not always overlap.

2.B.2.1 High Toxicity, No Extreme Rhetoric

Some high toxicity content consists of personal insults and emotional attacks without employing the specific extreme rhetorical strategies I measure. For example, one tweet with very high toxicity reads: "you are a fantastic idiot, it bothers you that we exposed you by sharing your video where you said that coronavirus was nonsense? so you take revenge.. but how stupid are you? shut yourself in the toilet and stay there." Another highly toxic tweet states: "monsters! you are monsters! you suck." A third example declares: "one word: imbecile, you have the virus in your brain!" These tweets contain personal insults and abusive language but do not employ the systematic rhetorical strategies I classify as delegitimization, group discrimination, or zero-sum framing. They attack specific individuals without delegitimizing institutions, targeting groups based on identity, or framing politics as zero-sum conflict.

2.B.2.2 Low Toxicity, Clear Extreme Rhetoric

Some of the rhetorically extreme tweets exhibit very low toxicity scores, yet still employ clear delegitimizing and discriminatory frames. One low-toxicity tweet classified as delegitimization states: "While everyone discusses the 'Russian funds' case, the Senate approves a constitutional reform that, disguised as a reduction in the number of parliamentarians, actually transforms the electoral system — opening for the first time in post-war history the path to a possible authoritarian drift." Another tweet, flagged for group-based discrimination, reads: "What's even more shocking is that the Moroccan behind the wheel, Farah Marouane, was walking free despite the fact that, last April, 225 kg of hashish were found in his home." A third example frames the political elite in derogatory, dehumanizing terms: "They took a series of people who resembled the characters of Ciccio, Franco, Albertone, Bombolo, and Er Munnezza — cleaned them up, shaved them, bought them suits, gave them jobs — and now they nominate them

as the ruling class of a nation."
These examples underscore how extreme rhetoric can be conveyed through a calm register devoid of personal insults or overt toxicity. The first tweet delegitimizes democratic processes by framing institutional reform as a covert authoritarian maneuver. The second invokes group-based discrimination by linking criminality and impunity to a specific ethnic identity. The third undermines the legitimacy of political representatives through ridicule and cultural stereotyping.

2.B.3 Classification Validation

To validate the LLM classification approach, I manually coded two random samples: 5000 tweets from the full dataset and 5000 tweets from those flagged by the LLM as containing extreme rhetoric. This validation strategy serves two purposes: assessing overall classification accuracy across the population and evaluating precision on positive cases where the model detected extreme rhetoric.

The manual coding followed identical definitions and guidelines as the LLM prompt, with binary coding for each of the three dimensions (group discrimination, delegitimization, zero-sum framing). Table B.2 presents the validation metrics. For the random sample from the full dataset, the LLM achieves high precision across all three dimensions (0.89-0.92), with recall ranging from 0.76-0.84. The lower recall indicates the conservative coding approach successfully minimizes false positives, though it may miss some borderline cases of extreme rhetoric. For the LLM-flagged sample, precision remains high (0.87-0.90), confirming that when the LLM identifies extreme rhetoric, it is typically correct.

The validation confirms that the categorical measures successfully isolate specific rhetorical strategies while maintaining high precision.

Table B.2: LLM Classification Validation Results

Random Sample from Full Dataset (n = 5000)				
Dimension	Precision	Recall	F1-Score	Accuracy
Group Discrimination	0.92	0.76	0.83	0.96
Delegitimization	0.89	0.84	0.87	0.95
Zero-Sum Framing	0.91	0.78	0.84	0.97
LLM (Mistral 24B) Flagged Sample (n = 5000)				
Group Discrimination	0.87	0.94	0.90	0.89
Delegitimization	0.90	0.92	0.91	0.91
Zero-Sum Framing	0.89	0.91	0.90	0.90

Notes: Validation based on manual coding of two samples: random tweets from full dataset and LLM flagged tweets containing potential extreme rhetoric. High precision across both samples confirms conservative coding approach effectively minimizes false positives while maintaining substantial accuracy across all three rhetorical dimensions.

Appendix 2.C Appendix: Descriptive Online Evidence

This section presents descriptive evidence on the communication strategies of Italian MPs during the XVIII Legislature (2018–2022) on X. The online analysis is based on a comprehensive dataset of 705,695 tweets from 742 MPs active on X during this period.

2.C.1 Social Media Adoption and Usage Patterns

X adoption among Italian MPs reveals substantial variation across the political spectrum. Of the 742 MPs who posted at least one tweet during the XVIII Legislature, usage intensity differs markedly by ideological positioning (see Table C.1.1). centrist MPs demonstrate the highest individual activity levels, averaging 1,659 tweets per MP over the five-year period[19], followed by extreme-right MPs at 1,222 tweets per MP. In contrast, populist MPs, despite representing the largest group with 261 active users, average only 650 tweets per MP, less than half the rate of Centrist MPs. Left MPs (960 tweets per MP) and right MPs (919 tweets per MP) fall in the middle range of activity levels. Usage also differs by age, gender, and branch of parliament (Table C.1.2). The distribution of activity reveals pronounced heterogeneity within each political camp, with centrist MPs showing the most concentrated distribution around higher activity levels (Figure C.1.1). Extreme-right MPs exhibit substantial variation, with some MPs posting very few tweets while others are highly active, creating a more dispersed distribution. Left MPs display a relatively concentrated distribution around their mean, indicating more uniform adoption of social media as a communication tool. Populist and right MPs show similar patterns, with most MPs clustering around lower activity levels but with some highly active outliers. Temporal patterns underscore these ideological differences in platform investment (Figure C.1.2). Extreme-right MPs maintained the steepest trajectory of cumulative activity throughout the legislature, ultimately producing the highest total volume (235,775 tweets), despite having a lower per-MP average than Centrists. This reflects their larger group size (193 MPs) combined with high overall engagement. Populist MPs followed a similar steep growth path, accumulating 169,685 total tweets. Centrist MPs, while having the highest per-capita activity, contributed a smaller total volume

[19] The elevated activity among Centrist MPs largely reflects the breakaway trajectories of figures from the Democratic Party during the XVIII legislature. Matteo Renzi (Italia Viva) and Matteo Richetti (Azione) both left the PD in 2019 to found new centrist parties, emphasizing direct communication and visibility outside traditional party structures. Lacking mass party infrastructure, these MPs relied heavily on social media to assert an independent profile, build recognition, and stake claims in the evolving political center. Their online engagement was thus both a necessity and a strategic tool for differentiation.

(66,368 tweets) due to their smaller group size (40 MPs). Left and right MPs maintained more moderate growth trajectories throughout the period. These usage patterns provide a first indication that digital communication decisions may be shaped by strategic incentives. The high activity levels among Centrist MPs, combined with the substantial aggregate volume from extreme-right MPs, suggest different approaches to social media strategy. Centrists appear to invest heavily on an individual basis, while the extreme-right achieves high visibility through broader participation across the ideological group. This suggests that social media activity is not merely a reflection of communication preferences, but may reflect calculated attempts to influence attention dynamics in a media environment where visibility is both scarce and politically valuable.

Table C.1.1: X Usage by Political Ideology Among MPs with Accounts (2018-2022)

Party ideology	Num. MPs	Total Tweets	Mean Tweets	Median Tweets	SD Tweets
Centrist	40	66,368	1,659.20	1,133.00	1,685.50
Extreme Right	193	235,775	1,221.60	392.00	2,569.90
Left	130	124,790	959.90	653.50	1,094.60
Populist	261	169,685	650.10	250.00	2,999.50
Right	118	108,480	919.30	401.50	1,500.00
Total	742	705,098	950.30	386.50	2,379.00

Notes: Sample includes 742 MPs who posted at least one tweet during the XVIII Legislature. Tweets per MP statistics refer to individual usage patterns within each ideological group.

Table C.1.2: Demographic Characteristics by Political Ideology

Party ideology	Num. MPs	Mean Age	% Female	% Deputies	% Senators
Centrist	40	54.60	37.50	55.00	45.00
Extreme Right	193	51.90	28.00	65.30	34.70
Left	130	56.30	31.50	66.20	33.80
Populist	261	46.70	42.90	69.70	30.30
Right	118	56.30	38.10	68.60	31.40
Total	742	51.70	36.00	67.00	33.00

Notes: Sample includes 742 MPs who posted at least one tweet during the XVIII Legislature. Age is measured at the start of the legislative term. Chamber composition reflects the bicameral structure of the Italian Parliament.

Figure C.1.1: Distribution of Tweet Volume by Ideology

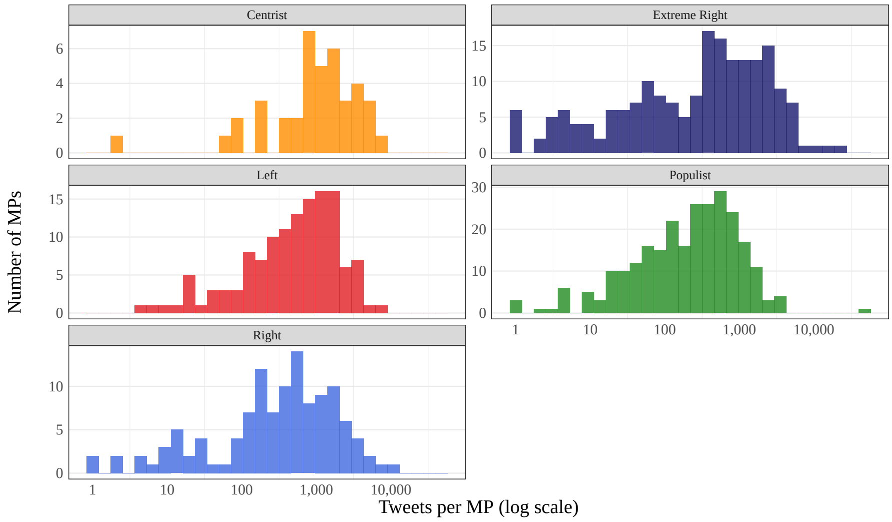

Notes: The figure shows the distribution of individual MP tweet counts across ideological groups using log-scale histograms. Each panel represents one ideological category, revealing substantial within-group heterogeneity in social media usage patterns.

Figure C.1.2: Cumulative Tweet Activity by Ideology Over Time

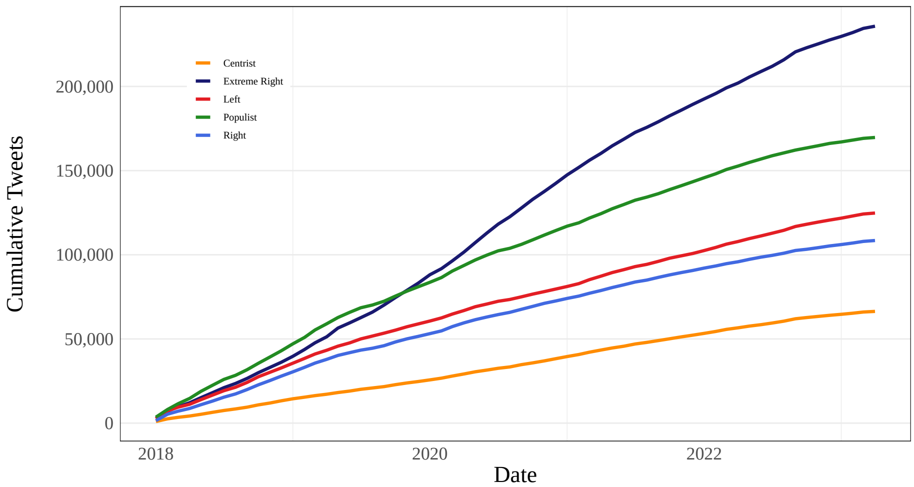

Notes: The figure tracks cumulative tweet volume for each ideological group throughout the XVIII Legislature (2018–2022). Lines represent monthly cumulative totals, showing divergent growth trajectories across political ideologies.

2.C.2 Topic Specialization Across the Political Spectrum

Parliamentary X discourse reveals systematic ideological specialization in topical emphasis, with patterns that may have important implications for online visibility. Overall, MPs dedicate 58.2% of their communication to "other" topics (ceremonial announcements, institutional commentary, and general political discourse), 18.4% to economic policy discussions, and 23.3% to cultural topics, defined as debates over social identity, values, or group belonging (the full list of topics' category is available in). However, this aggregate distribution masks meaningful ideological differences that reflect distinct strategic positioning. Extreme-right MPs demonstrate the highest focus on cultural issues at 25.2% of their tweets, followed closely by left MPs at 23.4% and centrist MPs at 22.7%. Populist MPs devote a similar share to cultural topics (22.4%), while right MPs show the lowest cultural focus at 21.1% (Figure C.2.2). The pattern for economic topics is reversed: right and populist MPs prioritize economic discussions more heavily (21.5% and 21.7% respectively), while extreme-right MPs devote the smallest share to economic topics (15.2%). These differences are documented in Table C.2.1. The specialization becomes more pronounced when examining which ideologies "own" specific topics, defined as contributing over 50% of all discussion on that subject. Extreme-right MPs dominate immigration debates, discussions linking immigration and crime, law and order themes, and a range of political mobilization topics centered on security and national identity.[20] Populist MPs control institutional reform discussions, including topics like "Parliament Privileges," "Banking Scandal Debate," and the "Superbonus Debate" (Figure C.2.1). This is unsurprising given their anti-establishment roots and sustained focus on institutional accountability, combined with the fact that key policies like the "Superbonus" originated from within the populist party itself.

[20]During the XVIII Legislature, Matteo Salvini (Lega) served as Interior Minister and implemented a hardline anti-immigration agenda, including the Security Decrees and port closures to migrant rescue ships, making immigration and public safety central to parliamentary and online discourse. Lega and Fratelli d'Italia MPs consistently linked crime to immigration and used social media to amplify these concerns. Giorgia Meloni (FdI), in opposition throughout the period, gained visibility by focusing on similar themes and framing her party as the sole defender of traditional values and national sovereignty.

Table C.2.1: Topic Distribution by Political Ideology (% of tweets within ideology)

Party ideology	Cultural	Economic	Other
Centrist	22.70	18.90	58.40
Extreme Right	25.20	15.20	59.70
Left	23.40	17.40	59.20
Populist	22.40	21.70	56.00
Right	21.10	21.50	57.40
Total	23.30	18.40	58.20

Notes: The table shows the percentage distribution of tweets by topic category (cultural, economic, and other) for each ideological group. Centrist and Extreme Right MPs exhibit the highest focus on cultural topics, while Populist and Right MPs prioritize economic discussions. Cultural topics include debates over social identity, values, immigration, civil rights, and family norms. Economic topics encompass fiscal policy, budget debates, and regulatory issues. Other topics include ceremonial content, institutional announcements, and general political commentary.

Figure C.2.1: Topics Dominated by Single Ideologies

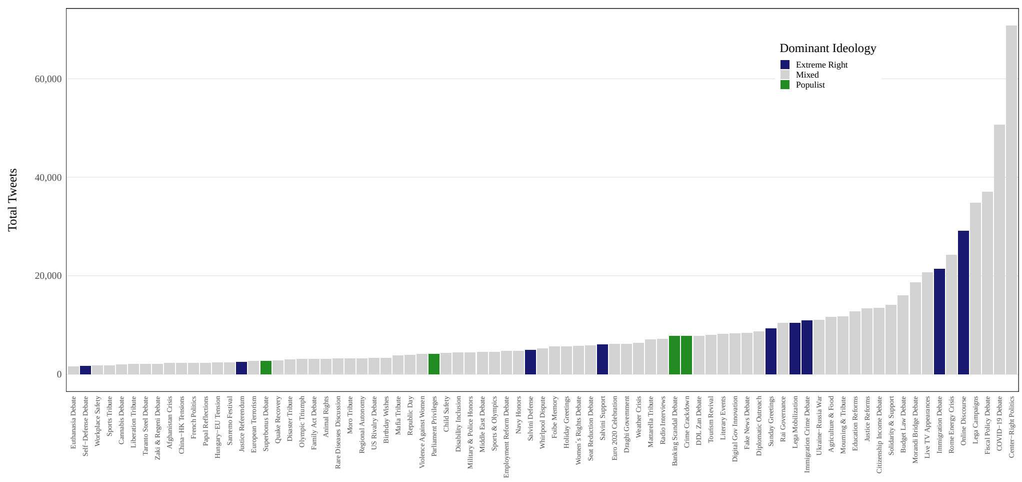

Notes: The figure shows topics where one ideology contributes > 50% of all tweets. Extreme-right dominates immigration and law-and-order discussions while populists control institutional reform topics and banking debates

Figure C.2.2: Cultural Topics Focus by Political Ideology

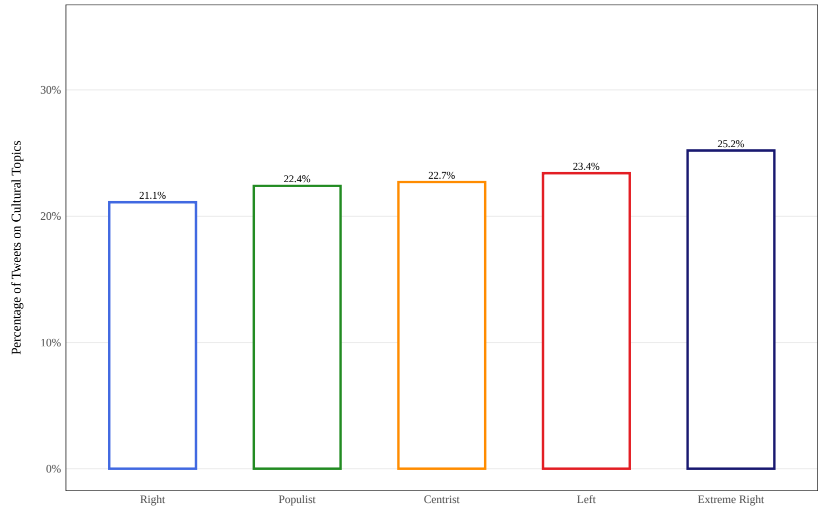

Notes: The figure shows the percentage share of each ideology's total tweets devoted to cultural topics. It shows clear ideological hierarchy in cultural emphasis.

2.C.3 Online Engagement Patterns

The topic specialization patterns documented above raise a natural question: do these different communication strategies translate into measurable differences in audience response? Engagement performance varies across the political spectrum (Table C.3.1). Extreme-right MPs achieve the highest average engagement at 402 interactions per tweet, followed by centrist MPs at 364. Left MPs average 256 engagements per tweet, while populist and right MPs perform significantly worse at 111 and 98 respectively. These substantial gaps suggest that ideological groups differ not only in what and how often they tweet, but also in how effectively they capture audience attention. Critically, these differences are not driven by baseline popularity. While centrist MPs have the largest average follower count (136,186), followed by left MPs (59,366), this follower advantage does not translate into engagement success. Table C.3.2 reveals the opposite pattern when measuring engagement efficiency, engagements per 1,000 followers. Extreme-right MPs generate 14.9 engagements per 1,000 followers, making them over five times more efficient than centrist MPs (2.7) despite having far fewer followers on average (26,951 vs. 136,186). This stark reversal suggests that engagement success is driven by communication strategy rather than audience size. Topic choice also fails to explain ideological differences in engagement. To test whether certain groups simply choose more engaging topics, I examine 606 MPs who post across multiple content categories (Table C.3.3). When the same politicians tweet about different topics, engagement levels remain remarkably similar: cultural topics average 139 engagements, economic topics 125, and other topics 118. The near-identical performance across topics within the same MPs rules out topic selection as the primary driver of ideological engagement differences (Figure C.3.2). These findings point to systematic differences in rhetorical style, rather than popularity or topic choice, as the driver of engagement disparities. The data suggests that some MPs, particularly those on the extreme-right, have developed communication strategies that consistently generate higher audience response rates regardless of their follower count or chosen topics.

Table C.3.1: Engagement Performance by Political Ideology

Party Ideology	Num. Tweets	Num MPs	Mean Engagement	Median Engagement
Extreme Right	227852	193	401.80	25.00
Centrist	65364	40	364.10	71.00
Left	123237	130	256.10	38.00
Populist	168054	261	110.70	14.00
Right	106186	118	98.10	19.00

Notes: The table shows engagement performance metrics across ideological groups for 690,693 tweets from 742 MPs during the XVIII Legislature (2018-2022). Mean engagement represents the average number of total interactions (likes, retweets, replies, and quote tweets) per tweet. Median engagement shows the middle value to account for skewed distributions. Pct high engagement indicates the percentage of tweets receiving engagement above the 90th percentile threshold. Sample sizes vary by ideology reflecting different adoption rates and activity levels across political groups.

Table C.3.2: Engagement Efficiency by Political Ideology

Political Ideology	Mean Engagement	Mean Followers	Engagement × 1000 followers
Extreme Right	401.85	26951.00	14.91
Populist	110.71	11079.00	9.99
Right	98.09	18971.00	5.17
Left	256.10	59366.00	4.31
Centrist	364.12	136186.00	2.67

Notes: The table measures engagement efficiency by normalizing total interactions against follower counts to control for baseline popularity differences. Mean engagement represents average interactions per tweet, while mean followers shows average follower counts per MP within each ideological group. Engagement per 1,000 followers provides a standardized metric revealing which groups generate disproportionate audience response relative to their social media reach. Higher values indicate more efficient conversion of follower base into active engagement.

Figure C.3.1: Engagement Efficiency by Political Ideology

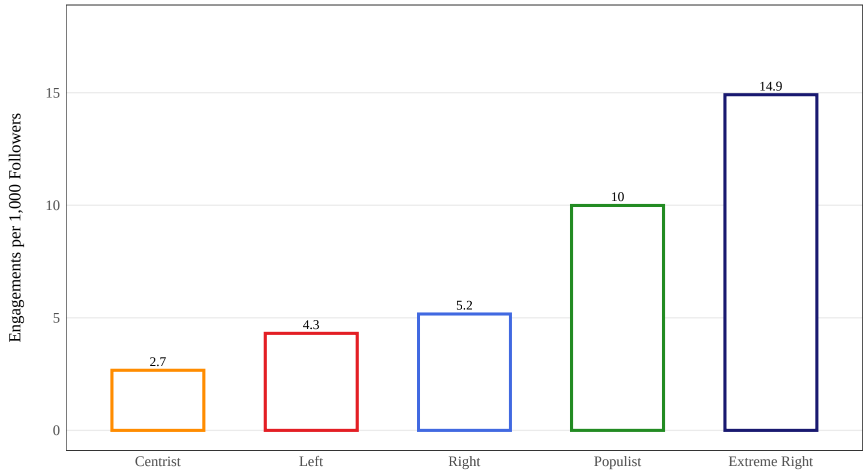

Notes: The figure shows engagements per 1,000 followers, controlling for baseline popularity differences across ideological groups. Extreme-right MPs achieve the highest efficiency at 14.9 engagements per 1,000 followers, over five times higher than Centrist MPs (2.7) despite having substantially fewer total followers. This metric reveals that engagement success is driven by communication strategy rather than audience size, as groups with larger follower bases perform worse when normalized for reach.

Figure C.3.2: Within-MP Topic Comparison: Same Politicians, Different Content Types

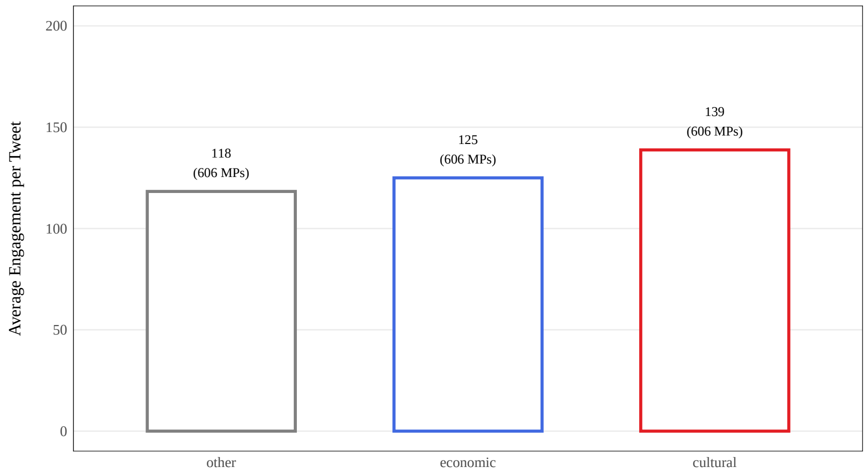

Notes: The figure compares engagement performance across topic categories for the same 606 MPs who posted content in multiple categories. Cultural topics generate slightly higher engagement (139) compared to economic (125) and other topics (118), but differences are minimal. This within-politician design controls for all MP-specific characteristics, demonstrating that topic choice alone cannot explain the substantial ideological differences in engagement documented in Table C.3.1.

Table C.3.3: Within-MP Topic Comparison: Same Politicians, Different Content Types

Category	Num. MPs	Mean Engagement	Median Engagement
Cultural	606	138.80	29.00
Economic	606	125.00	28.70
Other	606	118.30	28.70

Notes: The table presents within-MP topic comparison based on 606 MPs who posted content across multiple topic categories with sufficient volume (minimum 50 tweets total). This design controls for all politician-specific characteristics by comparing the same individuals' performance across different content types. Mean and median engagement show remarkably similar levels across cultural, economic, and other topics, indicating that topic choice alone does not drive engagement differences observed between ideological groups. The consistent sample size (606 MPs) across all categories ensures balanced comparison.

2.C.4 Deployment of Extreme Rhetoric

Having established that engagement differences cannot be explained by popularity or topic choice, this section examines the specific rhetorical strategies that drive these patterns. The analysis focuses on three forms of extreme rhetoric: delegitimization, group-based discrimination, and zero-sum framing. Extreme rhetoric is strategically deployed rather than uniformly distributed across the political spectrum (Table C.4.4). Extreme-right MPs exhibit the highest overall incidence of extreme rhetoric at 18.1% of their tweets, followed by right MPs (16.1%), centrist MPs (14.5%), left MPs (13.4%), and populist MPs (12.6%). The gap is especially pronounced for group-based discrimination, which appears in 5.7% of Extreme-right tweets, nearly three times the rate of left MPs (1.8%) and over four times that of populist MPs (1.6%). Delegitimization shows a more even distribution across ideologies, ranging from 9.6% (left) to 11.7% (right), while zero-sum framing peaks among extreme-right MPs at 4.6%. These rhetorical strategies deliver substantial engagement returns (Table C.4.5). Discriminatory tweets generate the highest average engagement at 563 interactions per tweet, more than double the baseline of 220 for non-extreme content. Zero-sum framing performs similarly well at 544 average interactions, while delegitimization achieves 396 interactions per tweet. All three strategies significantly outperform conventional rhetoric, providing clear incentives for their strategic deployment (Figure C.4.3). Extreme rhetoric concentrates in cultural discussions, amplifying the engagement advantages documented in previous sections. Cultural topics contain extreme rhetoric in 21.1% of tweets, compared to 17.3% for economic topics and just 12.3% for other content (Figure C.4.4). This pattern creates a compounding effect: ideological groups that both focus heavily on cultural topics (like the extreme-right) and deploy extreme rhetoric more frequently are positioned to maximize engagement returns through the combination of topic choice and rhetorical strategy.

Table C.4.4: Prevalence of Extreme Rhetoric by Political Ideology (% of tweets)

Political Ideology	Num. Tweets	% Delegitimization	% Discrimination	% Zero-Sum Game	% Any Extreme
Extreme Right	227852	10.17	5.74	4.58	18.09
Right	106186	11.70	2.29	3.53	16.12
Centrist	65364	10.89	1.27	3.24	14.51
Left	123237	9.57	1.82	2.99	13.43
Populist	168054	9.75	1.60	2.54	12.56

Notes: The table shows the percentage of tweets containing each rhetorical strategy by ideological group. Extreme Right MPs exhibit the highest overall usage of extreme rhetoric (18.1%), with particularly pronounced advantages in group-based discrimination (5.7% vs 1.6% for Populists). Delegitimization is the most prevalent strategy overall, while zero-sum framing and discrimination show stronger ideological concentration.

Table C.4.5: Engagement Performance by Rhetorical Strategy

Rhetoric Type	Num. Tweets	Mean Engagement	Median Engagement
Discrimination	18540	562.70	49.00
Zero-sum Game	16025	543.60	73.00
Delegitimization	70890	395.80	56.00
None	585238	220.00	21.00

Notes: The table compares average engagement across rhetorical strategies for 690,693 tweets. Discriminatory tweets achieve the highest engagement (563 interactions), followed by zero-sum framing (544) and delegitimization (396). All extreme rhetoric strategies substantially outperform conventional content (220 interactions), providing clear incentives for strategic deployment. Sample sizes indicate sufficient observations for reliable estimates.

Figure C.4.3: Average Engagement by Rhetorical Strategy

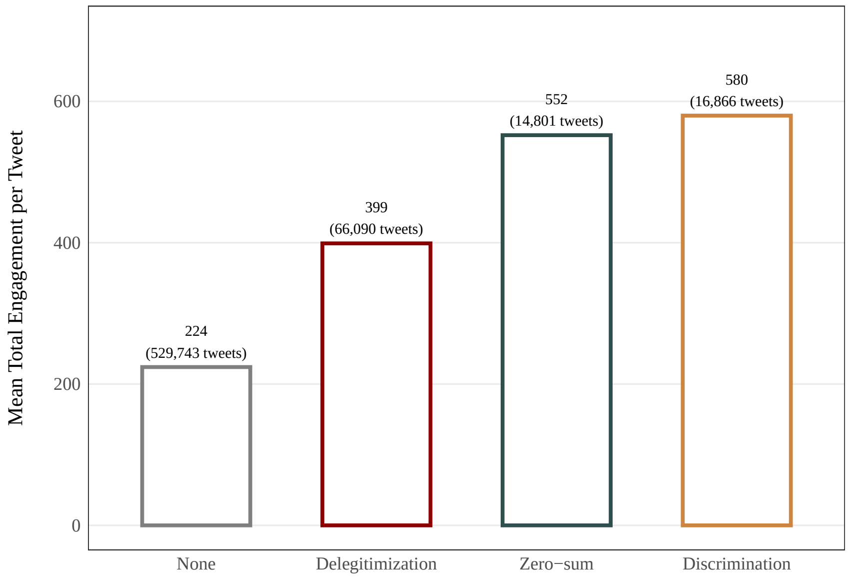

Notes: The figure visualizes engagement premiums for extreme rhetorical strategies. Discriminatory and zero-sum rhetoric generate over 2.5 times the engagement of baseline content, while delegitimization provides an 80% increase. The substantial differences demonstrate clear strategic incentives for adopting extreme communication approaches on social media platforms.

Figure C.4.4: Extreme Rhetoric Usage by Topic Category

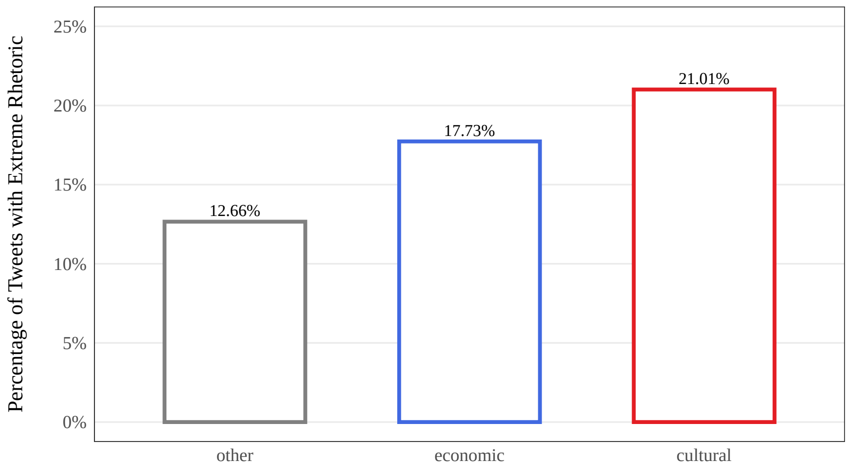

Notes: The figure shows the percentage of tweets containing any form of extreme rhetoric across topic categories. Cultural topics exhibit the highest concentration (21.1%), followed by economic topics (17.3%) and other content (12.3%). This pattern amplifies the engagement advantages for ideological groups that both focus on cultural issues and deploy extreme rhetoric more frequently.

Appendix 2.D Appendix: Online Engagement - Robustness and Additional Results

This appendix presents additional results that complement the main findings on extreme rhetoric and social media engagement presented in the results section. The tables below provide robustness checks using alternative outcome specifications and dynamic analysis with leads and lags to test for strategic timing and pre-trends, offering extensions to the core engagement analysis.

Table D.0.1: Robustness: Alternative Engagement Specifications

	Engagement Effects - Alternative Specifications					
	Log(Total Engagement+1)		IHS(Total Engagement)		PPML	
	(1)	(2)	(3)	(4)	(5)	(6)
	Log Total Engagement	Log Total Engagement	IHS Total Engagement	IHS Total Engagement	PPML Total Engagement	PPML Total Engagement
Group Discrimination	0.315***	0.316***	0.333***	0.331***	0.272***	0.257***
	(0.030)	(0.035)	(0.033)	(0.039)	(0.032)	(0.030)
Zero-Sum Framing	0.368***	0.348***	0.389***	0.366***	0.240***	0.231***
	(0.033)	(0.060)	(0.036)	(0.066)	(0.034)	(0.023)
Delegitimization	0.481***	0.470***	0.511***	0.500***	0.300***	0.290***
	(0.057)	(0.097)	(0.064)	(0.108)	(0.053)	(0.049)
Observations	691041	511181	691041	511181	691025	502069
Mean Dep. Variable	3.431	3.529	3.993	4.096	254.632	286.494
MP FEs	✓	✓	✓	✓	✓	✓
Topic FEs	✓	✓	✓	✓	✓	✓
Day FEs	✓	✓	✓	✓	✓	✓
MP × Day FEs		✓		✓		✓

Notes: This table presents regression results using alternative specifications for the engagement outcome variable to test robustness of the main findings. Total engagement is the sum of likes, replies, retweets, and quote retweets. Columns (1)-(2) use log(Total Engagement+1) as the dependent variable. Columns (3)-(4) use the inverse hyperbolic sine (IHS) transformation of total engagement. Columns (5)-(6) use Poisson Pseudo-Maximum Likelihood (PPML) estimation with raw total engagement counts. Even-numbered columns include MP × day fixed effects in addition to the baseline specification. All specifications include MP fixed effects, topic fixed effects, and day fixed effects. Standard errors clustered at the MP level are reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table D.0.2: Dynamic Effects of Extreme Rhetoric - Leads and Lags Analysis

	Dynamic Effects of Extreme Rhetoric on Engagement	
	(1)	(2)
	Total Engagement	Top 1% Engagement
Group Discrimination (j+2)	18.220*	0.002
	(11.040)	(0.002)
Zero-Sum Framing (j+2)	8.383	0.002
	(7.494)	(0.002)
Delegitimization (j+2)	-1.617	0.000
	(4.425)	(0.001)
Group Discrimination (j+1)	24.709*	0.003
	(13.993)	(0.003)
Zero-Sum Framing (j+1)	-2.074	0.000
	(6.973)	(0.001)
Delegitimization (j+1)	3.400	0.001
	(6.937)	(0.001)
Group Discrimination (j)	164.997***	0.014**
	(59.852)	(0.007)
Zero-Sum Framing (j)	142.215***	0.013**
	(40.414)	(0.006)
Delegitimization (j)	112.158***	0.007***
	(26.687)	(0.003)
Group Discrimination (j-1)	23.817	0.002
	(20.049)	(0.002)
Zero-Sum Framing (j-1)	5.795	0.000
	(9.578)	(0.001)
Delegitimization (j-1)	4.088	0.001
	(7.013)	(0.001)
Group Discrimination (j-2)	23.188	0.002
	(15.231)	(0.001)
Zero-Sum Framing (j-2)	5.714	0.000
	(8.498)	(0.001)
Delegitimization (j-2)	-0.330	0.001
	(3.730)	(0.001)
Observations	509824	509824
Mean Dep. Variable	281.646	0.011
MP FEs	✓	✓
Topic FEs	✓	✓
Day FEs	✓	✓
MP × Day FEs	✓	✓

Notes: This table presents regression results including leads and lags of extreme rhetorical strategies to test for strategic timing and pre-trends. Total engagement is the sum of likes, replies, retweets, and quote retweets. The table shows results across multiple outcome specifications: total engagement (columns 1-2), top 1% engagement indicator (columns 3-4), log(engagement+1) (columns 5-6), IHS transformation (columns 7-8), and PPML estimation (columns 9-10). Lead variables test whether politicians strategically time extreme rhetoric around predictable engagement spikes. Lag variables capture engagement persistence. Even-numbered columns include MP × day fixed effects in addition to the baseline specification. Standard errors clustered at the MP level are reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table D.0.3: Extreme Rhetoric Usage by Ideological Block - Regression Results

	Extreme Rhetoric Usage by Ideology		
	(1)	(2)	(3)
	Group Discrimination	Delegitimization	Zero-Sum Framing
Extreme Right	0.038***	0.007	0.012**
	(0.004)	(0.010)	(0.006)
Right	0.011**	0.015	0.001
	(0.004)	(0.011)	(0.006)
Populist	0.001	-0.012	-0.009
	(0.002)	(0.016)	(0.006)
Left	0.002	-0.004	-0.005
	(0.002)	(0.011)	(0.006)
Observations	705455	705455	705455
Mean Dep. Variable	0.031	0.103	0.035
Topic FEs	✓	✓	✓
Day FEs	✓	✓	✓

Notes: This table presents regression results examining the differential usage of extreme rhetorical strategies across ideological blocks, relative to centrist MPs (omitted category). Each column represents a separate regression with the indicated rhetorical strategy as the dependent variable. All specifications include topic fixed effects and day fixed effects. Standard errors clustered at the MP level are reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$ indicate significance at 10%, 5%, 1% levels respectively.

Appendix 2.E Online Virality and Television visibility: Additional Results and Statistics

2.E.1 Neswcasts and Political Talk Shows list

Rai 1 (Channel 1) includes TG1, CHE TEMPO CHE FA (2018-2019), IL CAVALLO E LA TORRE, IL CONFRONTO, MEZZ'ORA IN PIU', and PORTA A PORTA. **Rai 2** (Channel 2) includes TG2, MORNING NEWS, ONOREVOLI CONFESSIONI, POPOLO SOVRANO, POVERA PATRIA, RESTART, and SECONDA LINEA. **Rai 3** (Channel 3) includes TG3, #CARTABIANCA, CHE C'E' DI NUOVO, CHE TEMPO CHE FA (2021-2022), FILOROSSO, FRONTIERE, TITOLO QUINTO, and ZONA ROSSA. **Rete4** (Channel 4) includes TG4, CONTROCORRENTE, DALLA VOSTRA PARTE, DRITTO E ROVESCIO, FUORI DAL CORO, QUARTA REPUBBLICA, QUINTA COLONNA, STASERA ITALIA, and ZONA BIANCA. **Canale5** (Channel 5) includes TG5 and L'INTERVISTA. **Italia1** (Channel 6) includes STUDIO APERTO. **La7** (Channel 7) includes TG LA7, AGORA', BERSAGLIO MOBILE, COFFEE BREAK, DIMARTEDI', FACCIA A FACCIA, IN ONDA, L'ARIA CHE TIRA, NON E' L'ARENA, OMNIBUS LA7, OTTO E MEZZO, PIAZZAPULITA, PROPAGANDA LIVE, and TAGADA'. **TV8** (Channel 8) includes TG8 and SKY TG24. **NOVE** (Channel 9) includes ACCORDI & DISACCORDI, #FAKE - LA FABBRICA DELLE NOTIZIE, CHE TEMPO CHE FA (2019-2021), L'ASSEDIO, and SONO LE VENTI.

2.E.2 Treatment Assignment Statistics

The event study sample includes 742 MPs, of whom 32 experienced at least one viral tweet (99th percentile engagement) and 710 never achieved viral status during the observation period. Table E.2.1 shows how treatment rates vary systematically with online popularity. MPs in the bottom 60% of follower distributions never achieve top 1% virality, while treatment rates rise sharply in the upper percentiles: 2.7% for the 60-70th percentile, 5.4% for the 80-90th percentile, 18.9% for the 90-95th percentile, 41.4% for the 95-99th percentile, and 53.8% for the top 1% of MPs by followers. Notably, even MPs in the 70-80th percentile have zero treatment rate, highlighting how concentrated viral success is among the most followed politicians. This steep gradient reflects underlying network effects on social media platforms, where larger follower bases substantially increase the probability of viral success, but the concentration of treated units in the upper follower percentiles provides the identifying variation necessary for the event study design.

Table E.2.1: Treatment Assignment by Follower Percentiles (Top 1% Virality)

Follower Percentile	Total MPs	Treated MPs	Treatment Rate (%)
Bottom 1%	9	0	0.0
1-5%	28	0	0.0
5-10%	37	0	0.0
10-20%	76	0	0.0
20-30%	71	0	0.0
30-40%	74	0	0.0
40-50%	73	0	0.0
50-60%	74	0	0.0
60-70%	74	2	2.7
70-80%	73	0	0.0
80-90%	74	4	5.4
90-95%	37	7	18.9
95-99%	29	12	41.4
Top 1%	13	7	53.8
Total	742	32	100

Notes: Treatment is defined as having at least one day where the maximum engagement of any tweet reaches the 99th percentile of the distribution (top 1% virality). Follower percentile groups are based on MPs' follower counts at the beginning of the sample period.

2.E.3 Event Study: Virality and Television Invitation

Table E.3.1: Event Study: Private and Public TV Channels

	1 Minute		2 Minutes		3 Minutes	
	(1)	(2)	(3)	(4)	(5)	(6)
t-6	0.0083	0.0105	0.0254	0.0276	0.0030	0.0049
	(0.0528)	(0.0528)	(0.0515)	(0.0515)	(0.0477)	(0.0478)
t-5	0.0226	0.0247	0.0477	0.0499	0.0243	0.0262
	(0.0499)	(0.0496)	(0.0527)	(0.0526)	(0.0506)	(0.0506)
t-4	0.0321	0.0344	0.0339	0.0362	0.0551	0.0569
	(0.0513)	(0.0512)	(0.0523)	(0.0524)	(0.0547)	(0.0549)
t-3	0.0099	0.0122	-0.0277	-0.0254	-0.0115	-0.0096
	(0.0371)	(0.0370)	(0.0367)	(0.0365)	(0.0390)	(0.0390)
t-2	-0.0074	-0.0051	-0.0430	-0.0407	-0.0309	-0.0290
	(0.0391)	(0.0388)	(0.0272)	(0.0269)	(0.0291)	(0.0292)
t=0	0.1216	0.1239	0.0932	0.0955	0.0803	0.0822
	(0.0756)	(0.0756)	(0.0711)	(0.0712)	(0.0647)	(0.0650)
t+1	0.2347***	0.2369***	0.2334***	0.2357***	0.2222***	0.2241***
	(0.0666)	(0.0673)	(0.0692)	(0.0698)	(0.0714)	(0.0718)
t+2	0.1598**	0.1620**	0.1058*	0.1081*	0.1189*	0.1208*
	(0.0647)	(0.0650)	(0.0593)	(0.0595)	(0.0608)	(0.0610)
t+3	0.0313	0.0336	0.0353	0.0375	0.0475	0.0493
	(0.0560)	(0.0558)	(0.0566)	(0.0564)	(0.0573)	(0.0572)
t+4	0.0925	0.0948	0.0609	0.0632	0.0782	0.0801
	(0.0579)	(0.0579)	(0.0569)	(0.0571)	(0.0583)	(0.0585)
t+5	0.1107*	0.1131*	0.0792	0.0815	0.0621	0.0639
	(0.0611)	(0.0616)	(0.0543)	(0.0545)	(0.0516)	(0.0519)
t+6	0.0778	0.0802	0.0756	0.0779	0.0584	0.0602
	(0.0556)	(0.0559)	(0.0552)	(0.0555)	(0.0558)	(0.0561)
Observations	53760	53760	53760	53760	53760	53760
Mean Dep. Var.	0.111	0.111	0.081	0.081	0.068	0.068
MP FEs	✓	✓	✓	✓	✓	✓
Day FEs	✓	✓	✓	✓	✓	✓
MP × Week		✓		✓		✓

Notes: Table shows event-study coefficients from Equation 2.1. Columns (1)-(2) show effects on the probability of receiving TV invitations, defined as speaking for at least 60 seconds in a program. Columns (3)-(4) show effects on the probability of receiving TV invitations, defined as speaking for at least 60 seconds in a program. Columns (5)-(6) show effects on the probability of receiving TV invitations, defined as speaking for at least 60 seconds in a program. Even-numbered columns include MP×week fixed effects. All specifications include MP fixed effects, date fixed effects, and controls for the MP daily tweet volume. Standard errors clustered at the MP level. Reference period is $t = -1$ (omitted). $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$ indicate significance at 10%, 5%, 1% levels respectively.

2.E.3.1 Event Study: Heterogeneity by Channel Type

Table E.3.2: Event Study: Private and Public TV Channels

	Public Channels		Private Channels	
	(1)	(2)	(3)	(4)
t-6	-0.0036	-0.0034	-0.0028	-0.0014
	(0.0297)	(0.0298)	(0.0544)	(0.0545)
t-5	-0.0084	-0.0082	0.0168	0.0181
	(0.0253)	(0.0255)	(0.0499)	(0.0499)
t-4	0.0518	0.0520	0.0648	0.0663
	(0.0470)	(0.0470)	(0.0592)	(0.0592)
t-3	0.0250	0.0252	0.0322	0.0337
	(0.0286)	(0.0286)	(0.0526)	(0.0527)
t-2	0.0084	0.0087	0.0272	0.0287
	(0.0317)	(0.0316)	(0.0486)	(0.0486)
t=0	0.0726	0.0728	0.1202*	0.1216*
	(0.0519)	(0.0518)	(0.0703)	(0.0707)
t+1	0.0596	0.0598	0.2013***	0.2027***
	(0.0479)	(0.0478)	(0.0694)	(0.0700)
t+2	0.0606	0.0608	0.1222*	0.1235*
	(0.0433)	(0.0435)	(0.0618)	(0.0621)
t+3	0.0399	0.0401	0.0623	0.0638
	(0.0367)	(0.0367)	(0.0604)	(0.0606)
t+4	0.1134**	0.1136**	0.0548	0.0563
	(0.0533)	(0.0533)	(0.0570)	(0.0573)
t+5	-0.0147	-0.0141	0.0944	0.0961
	(0.0210)	(0.0211)	(0.0635)	(0.0639)
t+6	0.0271	0.0276	0.1020	0.1037
	(0.0324)	(0.0324)	(0.0618)	(0.0622)
Observations	53760	53760	53760	53760
Mean Dep. Var.	0.035	0.035	0.119	0.119
MP FEs	✓	✓	✓	✓
Day FEs	✓	✓	✓	✓
MP × Week		✓		✓

Notes: Table shows event-study coefficients from Equation 2.1 estimated separately for public and private TV channels. Public channels include RAI 1, RAI 2, and RAI 3 (Channels 1-3). Private channels include Rete 4, Canale 5, Italia 1, La7, TV8, and TV9 (Channels 4-9). Columns (1)-(2) show effects on the probability of receiving TV invitations on public channels, defined as speaking for at least 60 seconds in a program. Columns (3)-(4) show effects on the probability of receiving TV invitations on private channels. Even-numbered columns include MP×week fixed effects. All specifications include MP fixed effects, date fixed effects, and controls for the MP daily tweet volume. Standard errors clustered at the MP level. Reference period is $t = -1$ (omitted). $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$ indicate significance at 10%, 5%, 1% levels respectively.

Appendix 2.F Appendix: Robustness

2.F.1 Alternative Outcome

Figure F.1.1: Event study estimates for alternative television outcomes

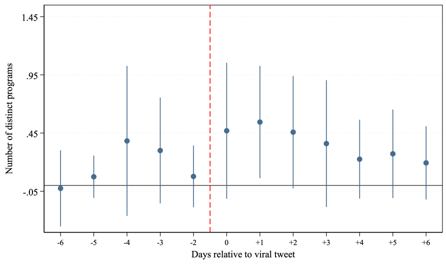

((a)) Number of distinct programs

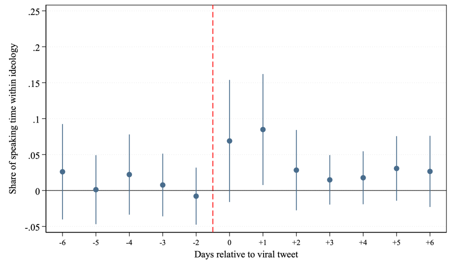

((b)) Share of speaking time within ideology

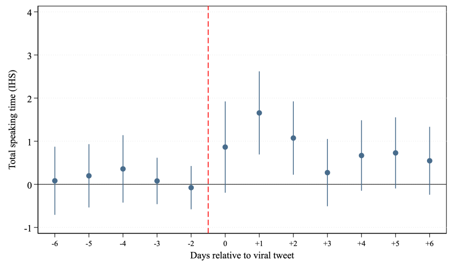

((c)) Total speaking time (IHS)

Notes: Figure reports event study coefficients from Equation 2.1 using the most demanding specification with tweet volume control, MP fixed effects, date fixed effects, and MP×week fixed effects. Panel A shows effects on the number of distinct programs an MP appears in. Panel B shows effects on the MP's share of total speaking time within the same ideological bloc. Panel C shows effects on total speaking time measured with the inverse hyperbolic sine transformation. Across outcomes, effects rise on the day after virality and dissipate by day three. 95% confidence intervals shown. The reference period is $t = -1$ (omitted).

2.F.2 Permutation Test

To assess whether the estimated next day effect could arise by chance, I conduct a permutation test that reassigns each MP's viral dates $B = 500$ times within MP, preserving the number of viral days per MP and the spacing restrictions used in the main design. The resulting placebo distribution of lag 1 coefficients is centered near zero (mean -0.002, s.d. 0.049). The actual estimate, 0.237, exceeds all 500 placebo draws. This supports the interpretation that television responses follow genuine viral shocks rather than statistical artifacts.

Figure F.2.1: Permutation placebo test for the lag 1 event study coefficient.

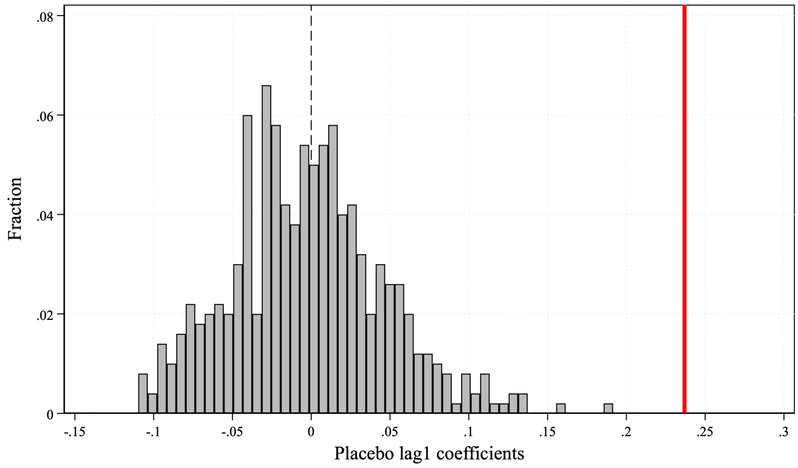

Notes: This figure shows the permutation (placebo) distribution of the *lag-1* event-study coefficient (next-day effect) under $B = 500$ random reassignments of viral dates *within MP*, preserving each MP's treatment frequency and spacing. Gray bars are placebo coefficients; the dashed line is the placebo mean (−0.002, s.d. 0.049). The red vertical line is the actual estimate (0.237), which exceeds all 500 placebo draws. The 95th and 99th percentiles of the placebo distribution are 0.078 and 0.129. Event time 0 is the day of the viral tweet/MP; the reference period is $t = -1$ (omitted).

2.F.3 Staggered Treatment Timing

The staggered timing of virality in the event-study design raises potential concerns about biases in conventional two-way fixed effects estimators, particularly under treatment effect heterogeneity (Goodman-Bacon, 2021; Sun and Abraham, 2021). To address this, I implement the estimator developed by Sun and Abraham (2021), which corrects for the problematic comparisons and negative weighting inherent in TWFE.

Figure F.3.1 compares the Sun-Abraham and standard TWFE estimates across the three speaking time thresholds. The Sun-Abraham estimator yields nearly identical point estimates and confidence intervals to the baseline TWFE specification. For the 1 minute threshold (Panel A), both estimators show a sharp increase in the probability of receiving invitations the day after virality, with effects concentrated in the immediate post-treatment period. The 2 minutes threshold results (Panel B) similarly demonstrate consistent patterns across estimators, with substantial increases the day following viral tweets that quickly dissipate. The 3 minutes threshold (Panel C) shows the same temporal pattern and magnitude of effects under both estimation approaches.

These results provide strong evidence that the main conclusions are robust to the choice of estimator and are not driven by biases from the staggered design.

The consistency across estimation methods reinforces the interpretation that viral social media content creates immediate but short-lived spillover effects to traditional media visibility.

Figure F.3.1: Robustness to Staggered Treatment Timing: Sun-Abraham vs. TWFE Estimators

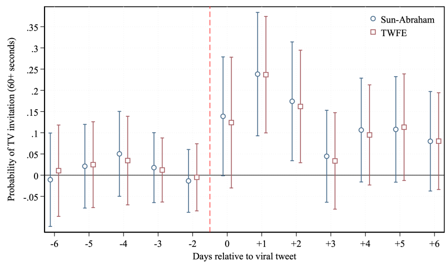

((a)) 1 Minute Threshold

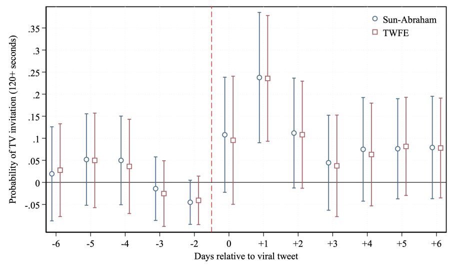

((b)) 2 Minute Threshold

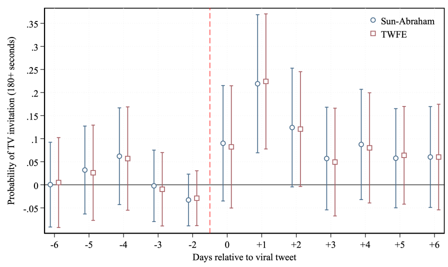

((c)) 3 Minute Threshold

Notes: Figure compares event-study coefficients estimated using the Sun-Abraham (2021) estimator and standard two-way fixed effects (TWFE) for three speaking time thresholds. Both estimators include tweet volume controls and MP, date, and MP×week fixed effects. Panel A shows effects on the probability of receiving TV invitations with at least 1 minute of speaking time. Panel B presents effects for the 2 minute threshold. Panel C shows effects for the 3 minute threshold. Navy circles represent Sun-Abraham estimates; maroon squares represent TWFE estimates. 95% confidence intervals shown. Reference period is $t = -1$ (omitted). Event time 0 corresponds to the day of viral tweet occurrence.

Chapter 3

Man Bites Dog? Editorial Choices and Biases in the Reporting of Weather Events

with Nicola Mastrorocco (Unibo), Arianna Ornaghi (Hertie School) and Stephane Wolton (LSE).

3.1 Introduction

"Dog bites man is no news; man bites dog is news." Every day, media outlets face the fundamental task of selecting which events to cover and how much prominence to assign to them. These editorial decisions—central to the news generation process—are often portrayed as objective judgements of newsworthiness. "All the news that's fit to print," the front page of the New York Times promises. "We deliver news, not noise," USA Today declares. But what makes an event newsworthy? And are editors just responding to the rarity or informational value of events, or do other strategic considerations shape coverage decisions?

In this paper, we unpack the production process of news by examining how local television stations in the United States cover one of the most ubiquitous categories of content: the weather. Weather is everywhere, it happens every day, and it affects everyone. Unlike scandals or natural disasters, the weather is not always newsworthy and the amount of coverage devoted to it is mostly an editorial choice. Unlike politics or crime, it lacks an obvious partisan connotation. Unlike most other news items, basic facts about the weather—temperatures, precipitations—are observable to researchers, independently of media attention. This, combined with variation in the frequency and severity of different weather events, allows us to quantify the informational content of different events depending on how rare or surprising a given day's weather truly is. All these reasons make weather an especially interesting setting to study how editors decide what counts as news.

We begin by estimating the newsworthiness of different weather events, defined as daily temperature deviations from the historical mean,[1] by linking coverage of local weather in TV newscasts and the occurrence of the events themselves, separately in the summer and in the winter. Exploiting close to the universe of newscast transcripts of local TV stations, we first document a clear "man bites dog" effect. In line with what one might expect, days that experience infrequent weather events, which happen to be large deviations from the historical mean, see substantially more coverage of local weather than days in which temperatures are in line with the norm. We also show that local TV news cover moderate deviations more than weather in line with the norm. Hence, "big dog's bites" are also sometimes news.

We then move beyond average coverage decisions to explore how editorial choices vary depending on audience characteristics and investigate the mechanisms behind variation in coverage. We focus in particular on heterogeneity depending on the political leaning of the market served by a station. While temperature deviations are ostensibly neutral, their link to climate change introduces an interpretive dimension that may interact with audience beliefs. We find that coverage of weather events in Democratic- and Republican-leaning media markets (also known as Designated Market Areas or DMAs) is markedly distinct. This difference is especially visible when it comes to uncommon events in summer. Relative to the baseline of temperatures in line with the historical norm, coverage of local weather in unseasonably cold days increases significantly more in Republican- than in Democratic-leaning markets, whereas the reverse holds true in unseasonably hot days. All the news that's fit to print has a clear partisan bend even for ubiquitous events like weather events.

While the first choice any editor makes is whether to broadcast (or publish) a story, editorial decisions do not stop there. Editors must choose how much salience to give to a story (i.e., where to place it in a newscast or a newspaper) and how to present it (i.e., the slant of the story). We find that the salience of weather stories presents the same direction of bias as the number of weather stories. We then look at the framing of these stories and, in particular, whether they are explicitly linked to climate change or global warming. Our results are

[1] Specifically, for each media market and day in the 2013-2019 period, we calculate the deviation of the maximum temperature from the mean temperature recorded on the same calendar day from 2000 to 2012 (the historical mean).

striking: the terms "climate change" and "global warming" are barely mentioned no matter the temperature of the day. We also train a topic model on our local weather stories and, again, find no topic that explicitly relates to climate change.

Next, we explore what mechanisms drive the systematic differences in coverage of weather events across media markets that we document. Specifically, we distinguish between demand- and supply-driven explanations. Using quasi-experimental variation from the staggered acquisition of local TV stations by a conservative broadcast group, Sinclair, we assess whether changes in ownership shift coverage. Our findings are more consistent with audience-driven news generation: coverage patterns remain stable after Sinclair acquisitions, suggesting that editorial strategies respond to anticipated audience interest rather than to top-down ideological directives. We also test for, and are able to exclude, that the patterns we document are explained by other correlates of media markets or weather events. Overall, our evidence suggests a demand-driven explanation for publication bias in the coverage of weather events.

In the final part of the paper, we present a formal model of coverage decisions and media consumption to make sense of the main empirical patterns we uncover. An outlet faces an audience which is constituted in part of Democratic citizens and for the rest of Republican citizens. All citizens like to learn about uncommon events as they value surprises. They also all suffer from a form of confirmation bias: their utility depends on what they learn about climate change from weather news.[2] While Democrats receive a payoff gain from holding higher beliefs in climate change following news, Republicans suffer a loss from increased beliefs. Citizens consider their expected utility from watching local news versus the utility from their outside option, which could include watching different, entertainment shows on TV. This expected utility is a function of the time allocated to weather news and other news by the local TV station times the consumption value of each type of news. This consumption value is fixed for non-weather news. It depends on the size of the weather events for weather news. Larger events are more unexpected and deviations above the mean also indicate that climate change is happening.

We compare the coverage decisions of a profit-maximizing outlet and of biased outlets which seek to minimize or maximize belief in climate change. The first type captures demand-driven coverage, the second corresponds to supply-driven coverage. We show that the strategy of a profit-maximizing outlet matches

[2] In the Online Appendix, we use individual-level survey data from the Cooperative Congressional Election Survey to show that, indeed, individuals' stated beliefs about climate change depend on weather events.

most of our empirical findings. Larger events receive more coverage. Events above (below) the norm, that increase (decrease) beliefs in climate change, see a greater increase in coverage when the TV station operates in a market with many Democratic (Republican) citizens. In contrast, supply-driven coverage fails to match any of our empirical results. In fact, even if a biased TV station can commit to a coverage strategy, it either hides events or spends as much time as possible covering them, generating either no or very stark variations in coverage as events become more extreme as opposed to the smooth increase we observe empirically.

Related Literature

With this paper, we contribute to several strands of literature within the economics of media. First and foremost, we contribute to an extensive theoretical and empirical literature that studies media bias. The theoretical literature (see, for example, Gentzkow, Shapiro, and Stone, 2015) broadly considers two non-mutually exclusive types of bias: filtering or publication bias (bias in *what* events are covered) and presentation bias (bias in *how* events are covered). In this paper, we primarily focus on the former—that is, on editors' decisions of what events to cover and what events not to cover. Recently, Armona et al., 2024 have noted that, to be able to identify publication bias, it is important to determine an unbiased reporting benchmark, to which to compare coverage against. In their paper, they focus on national media outlets and determine unbiased reporting based on a scoring rule (a statistical measure of the amount of "news" in the realization relative to the consumer's prior). Instead, thanks to our focus on local media outlets, we are able to determine the benchmark of unbiased reporting empirically, by looking at reporting of the closely related events in markets with a neutral political leaning. Our investigation also extends the empirical literature on publication bias in media coverage. Most of this literature has focused on explicitly political topics or events (e.g., scandals as in Puglisi and Snyder Jr, 2011 or economic news that are likely to impact voting behavior as in Larcinese, Puglisi, and Snyder Jr, 2011). Instead, we show that media bias is also present when studying a commonplace topic, the weather, which is not explicitly political, although it connects to the politically charged issue of climate change.

Our focus on an environmental issue also connects our work to a few recent advances in the literature on media strategies. Closest to us is Pianta and Sisco, 2020, who focus on coverage of extreme temperatures for a sample of European newspapers. Their analysis, however, does not address the determinants of coverage and does not examine how coverage varies with ideological leaning. Ash et al., 2023 and Djourelova et al., 2024 study the media coverage of climate change with a focus on presentation bias and its consequences for political preferences. Relative to these papers, we find only limited evidence of presentation bias being at play in the closely related domain we analyze, weather event. While this might appear to be surprising at first glance, recent studies have documented variation in coverage even across different natural disasters (see Fetzer and Garg, 2025).

After documenting media bias in the coverage of weather events, we conduct a comprehensive examination of its drivers. In particular, we contribute both theoretically and empirically to work in the economics of media literature highlighting supply-side explanations (see, among many others, Strömberg, 2004; Baron, 2006; Wolton, 2019) and demand-side explanations (see, among others, Gentzkow and Shapiro, 2006; Anand, Di Tella, and Galetovic, 2007; Mullainathan and Shleifer, 2005). Our evidence suggests that the empirical patterns are more likely driven by demand-side than supply-side factors, extending the findings of Gentzkow and Shapiro, 2010 to show that media bias also extends to everyday news items such as weather events. From a theoretical perspective, we enhance existing models of media bias by highlighting how learning and belief updating might play a fundamental role in determining the demand for news, thus adding to the experimental findings in Chopra, Haaland, and Roth, 2024.

In summary, our paper sheds light on the daily operations of news production, focusing on routine events like weather. We find that extreme weather deviations receive more coverage, and local TV stations' weather reporting is influenced by the political leanings of their audience, rather than supply-side forces. This has two broad implications. Newsworthiness cannot on its own explains the editorial choices we observe. Weather reporting may inadvertently shape climate change beliefs through daily coverage decisions.

3.2 Background

Each local TV station is licensed to operate in a specific media market (also referred to as Designated Market Area or DMA), which effectively represents the geographic reach of the station. More precisely, media markets are regions whose residents have access to the same television and radio broadcasts, and are defined by Nielsen based on household viewing patterns, which makes them non-overlapping geographies.[3] Within each market, our focus is on stations

[3]A county is assigned to a specific media market if the majority of its television audience watches stations from that market (Nielsen, 2019). Although counties can sometimes be divided between multiple media markets, this is relatively rare. According to Moskowitz (2021), only 16 out of 3,130 counties fall into this category. Additionally, while Nielsen updates media market boundaries annually, only 30 counties changed their market affiliation between 2008 and 2016.

affiliated with the four major broadcast networks (ABC, CBS, FOX, and NBC), as they tend to have the highest viewership and produce their own newscast (Papper, 2017).[4]

Local TV newscasts, which are locally produced by each station, feature a mix of national and local stories. Mastrorocco and Ornaghi (Forthcoming) show that the most common topics of local stories are crime, politics, weather, and sports. Despite a gradual decline in viewership, these TV newscasts remain a significant source of information for many Americans. A 2017 report by the Pew Research Center found that 50% of U.S. adults regularly consume news from television, surpassing the share who rely on online sources (43%), radio (25%), or print newspapers (18%) (Gottfried and Shearer, 2017). Among television news sources, local stations attract larger audiences than both cable networks and national broadcasters (Matsa, 2018). This makes studying local TV news particularly relevant.

3.3 Data & Measurement

This paper combines several data sources, which we detail below.

Local TV Stations. Our starting sample includes all local TV stations affiliated to one of the big-four networks (ABC, CBS, NBC, and FOX) in the continental United States. Information on the market served by each station and yearly affiliation is from BIA/Kelsey, an advisory firm focusing on the media industry. Additionally, we track historical call sign changes over our study period using data scraped from the Federal Communications Commission (FCC) website.

Newscast Transcripts. To measure coverage of local weather, we rely on comprehensive transcripts of (close to) the universe of all local TV newscasts 2013-2019. In particular, our sample includes 691 TV stations, active in 205 media markets. The transcripts are originally collected by TVEyes and archived by Harmony Labs. We segment each transcript into separate stories using chunks of 150 words.[5] A segment is classified to be about local weather if it meets two criteria: (i) it contains at least two terms from a weather-related lexicon compiled by Baylis et al. (2018); and (ii) it references at least one county or municipality within the

[4] Broadcast networks create and distribute content under a unified brand, while affiliated stations, which are independently owned, carry the network's programming alongside their own locally produced content. Typically, each media market has only one affiliate per network.

[5] We choose 150 words as it corresponds approximately to the number of words per minute spoken by a TV news anchor (e.g., Jensema, McCann, and Ramsey, 1996), but our results are robust to using segments of 300 words.

station's media market.[6] We aggregate this information at the station-by-day level and define our main outcome as the number of segments about local weather appearing in the station's local newscasts on a given day. To maintain consistency, we focus exclusively on weekday broadcasts, as weekend programming structures vary significantly across stations.

Weather Data. We define weather events using data from the AN81d dataset of the PRISM Climate Group. Our starting point are daily maximum temperatures measured in degree Celsius at the 4km-by-4km cell level from 2000 to 2019, which we aggregate at the media market level using population-weighted averages.[7] We construct weather events proceeding in three steps. First, we use data from 2000 to 2012 to calculate the historical mean for each calendar date and media market. Second, we compute the difference between the temperature recorded in each day and media market 2013-2019 and the historical mean for the same calendar date and media market. We call this difference the deviation from the historical mean (or deviation in short). Third, we look at how this deviation compares with the deviations in a given season in media market m over the 2013-2019 period and record its percentile in the within-media market distribution of season-specific deviations. We aggregate these percentiles in different bins, that become our weather events.

Several methodological choices warrant discussion. First, we focus on maximum temperatures, as they are more noticeable to the public and occur during waking hours, unlike minimum temperatures, which often materialize overnight.[8] Second, we limit our analysis to winter and summer, where temperature patterns exhibit greater stability, excluding the transitional and more volatile seasons of spring and fall. Finally, we define weather events looking at the within-media market distribution of deviations in each season. This means that our definition of weather event consists of a double comparison. First, we compare the temperature in a given date in a media market relative to its historical mean. Second, we look at whether the difference in temperature corresponds to a large deviation relative

[6]See Appendix C for further information and a detailed discussion of the classification procedure.

[7]To perform the weighting we use information on population for 1km by 1km cells from the 2000 population census, which is made available by the Socioeconomic Data and Applications Center (SEDAC).

[8]An additional reason that justifies the focus on maximum temperatures requires getting into the details of how a day is defined in PRISM. PRISM defines each day as covering the 1200 UTC-1200 UTC interval, corresponding for example to 7 AM-7AM in EST (note that we align our day definition in the content data for consistency). Considering that most local newscasts take place in the afternoon/evening, focusing on maximum temperatures allows us to study a day's news coverage of the day's weather event.

to all possible deviations from the mean during the same season within a specific media market.[9] As a result, a weather event is large if (i) the distance between the experienced temperature and its historical mean is big and (ii) this difference is substantial relative to the expected variation experienced during the season in the area. Our key assumption is that individuals experience and care about how daily temperatures compare to the usual local means, not some national means. Yet, our measure is such that weather events are comparable across media markets as they are all measured on the same scale: how uncommon they are relative to the usual within-media market temperature deviations (their percentile in the distribution of deviations).

DMAs' Political Leaning. We use county-level results for the 2008 Presidential elections from the MIT Election Lab, which we aggregate to the media market level, to define each media market's political leaning. Specifically, we define a media market to be Republican-leaning if the Republican vote share is in the top quartile (the average Republican vote share is 65% in these media markets); Democratic-leaning if the Republican vote share is in the bottom quartile (average of 38%); and "swing" otherwise (average of 52%).

Sinclair Ownership. We collect information on the date of each Sinclair acquisition 2013-2019, in addition to the name of stations controlled by the group at the beginning of the period, from the group's yearly company reports to shareholders and 10-K forms. Following Martin, Mastrorocco, et al. (2024), we consider a station to be under Sinclair control if it is directly owned and operated by Sinclair or has a local marketing agreement (LMA) with the group (which effectively gives Sinclair control over the station's programming). We use the same data source to identify starting dates of LMAs.

CCES. We proxy attitudes towards climate change using the Cooperative Congressional Election Study (CCES) waves from 2009 to 2012.[10] We use these data to study how weather events impact stated climate change beliefs (see Online Appendix 3.F), but also to check whether our estimation of publication bias is

[9] In theory, large deviations need not be rare events (e.g., if the distribution of deviations is double peaked at its extremes). In practice, the distribution of deviations assume a close approximation of a (truncated) Normal distribution. Hence, large deviations are also rare weather events.

[10] Specifically, we measure stated beliefs about climate change using the question "From what you know about global climate change or global warming, which of the following statements comes closest to your opinion?" We then define an indicator equal to one if the respondent answers "Global climate change has been established as a serious problem, and immediate action is necessary" or "There is enough evidence that climate change is taking place and some action should be taken."

Figure 3.1: Distribution of Deviations

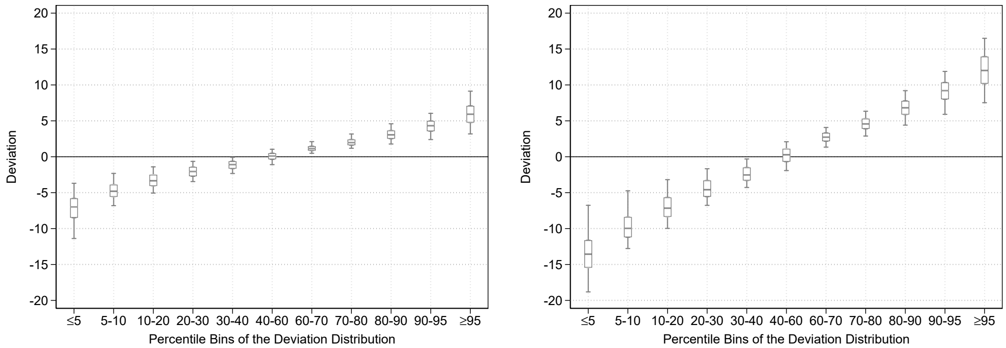

((a)) Summer **((b))** Winter

Notes: This figure shows the distribution of deviations for different weather events. In particular, we show the boxplot (5th percentile, 25th percentile, median, 75th percentile, 95th percentile) of the deviation from historical mean falling in each bin of the within-media market deviation distribution, separately for summer (panel (a)) and for winter (panel (b)).

robust to using heterogeneity in beliefs about climate change directly. To do so, we aggregate these responses at the media market level to create an alternative proxy of climate change scepticism. That is, we define a media market to be climate change sceptic if the share of respondents stating that climate change requires immediate action is in the bottom quartile (the average share of respondents reporting that climate change requires immediate action is 43% in these media markets); non-sceptic if it is in the top quartile (average of 61%); and neutral otherwise (average of 51%).

Additional Data Sources. In Appendix 3.D, we provide information on the data sources that we use to perform robustness checks and analyses only presented in the Online Appendix.

3.3.1 Descriptives

We begin by showing the variation in temperature deviations and local weather coverage that we use to identify editorial strategies. Figures 3.1 and 3.2 show the distribution of temperature deviations from the historical mean in each weather event, on average and by media market political leaning. The x-axis reports our weather events, that is, the different bins of the within-media market deviation distribution that we use throughout the analysis. Days that fall between the 40th and 60th percentile of the distribution are days in which temperatures are in line with the historical mean: deviations are close to zero. To the left, we have weather events that correspond to unseasonably cold days (days in which temperatures are below the historical mean), with increasing levels of severity. Instead, as we move to the right, we have weather events that correspond to unseasonably hot days (days in which temperatures are above the historical mean), again with increasing levels of severity.

Three notable patterns emerge. First, deviations from historical temperatures are more pronounced in winter than in summer, suggesting greater variation in cold-season weather anomalies. Second, extreme deviations—those in the bottom or top 5%—exhibit substantially larger temperature fluctuations than those in intermediate bins. Third, while there is some overlap in temperature distributions across media markets, deviations remain relatively well-demarcated. This ensures that extreme deviations in one media market generally correspond to similarly extreme deviations in others, allowing for meaningful cross-market comparisons. Importantly for our strategy, weather events do not systematically refer to different deviations in media markets with different political leanings as can be seen in Figure 3.2.

Appendix Figures 3.A.1 and 3.A.2 similarly show the average number of local weather stories in different weather events, again on average and by media market political leaning. The coverage–weather events relationship follows a U-shape in the raw data, a pattern that is more pronounced in the winter than in the summer. Also, stations in Democratic-leaning markets tend to have higher coverage of local weather, for every weather event, relative to both swing and Republican-leaning media markets. This hints to the importance of estimating proportional effects to take into account different base levels in coverage.

3.4 Publication Strategies

In this section, we investigate the average publication strategy of local TV stations. This allows us to understand what events are considered newsworthy on average across all stations.

3.4.1 Empirical Approach

To uncover the average publication strategy of local TV stations, we estimate the following regression:

$$Y_{st} = \sum_{\rho} \sum_{k \in \{-1,0,1\}} \beta_k^{\rho} \mathbb{I}\{\rho^{th} bin\}_{m(s)t+k} + \delta_s + \delta_t + \delta_{\eta(st)} + \epsilon_{st}, \tag{3.1}$$

where Y_{st} is the number of segments about local weather in the newscasts of station s on date t, $\mathbb{I}\{\rho^{th} bin\}_{m(s)t}$ are indicator variables equal to one if the deviation

Figure 3.2: Distribution of Deviations by DMA Political Leaning

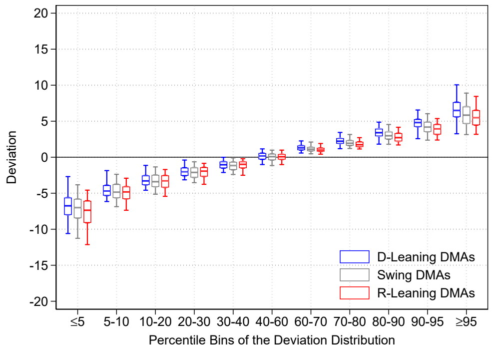

((a)) Summer

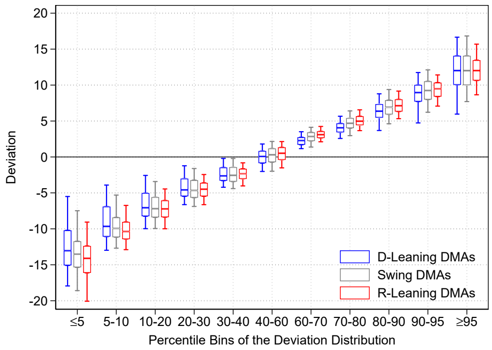

((b)) Winter

Notes: This figure shows the distribution of deviations for different weather events, by the political leaning of the media market. In particular, we show the boxplot (5th percentile, 25th percentile, median, 75th percentile, 95th percentile) of the deviation from historical mean falling in each bin of the within-media market deviation distribution for media markets that are Democratic-leaning, swing, and Republican-leaning, separately for summer (panel (a)) and for winter (panel (b)).

in temperatures from the historical mean of media market $m(s)$ on date t falls within the ρ_{th} bin of the within-market deviation distribution, δ_s are station fixed effects, δ_t are date fixed effects, and $\delta_{\eta(st)}$ are number of segments decile fixed effects. We estimate the specification separately for summer and winter, using a Poisson regression.[11] Standard errors are clustered at the media market level.

We estimate two versions of the regression that allow different degrees of flexibility in the relationship between coverage of local weather and weather events. We begin by estimating a fully-flexible specification in which we separate the within-media market deviation distribution into eleven different bins and allow each weather event (each bin) to have a different effect on coverage. Specifically, we split weather events as follows: bottom 5%, 5-10%, 10-20%, 20%-30%, 30%-40%, 40%-60%, 60%-70%, 70%-80%, 80-90%, 90-95%, and top 5%. The first five bins correspond to weather events with temperatures below the historical mean (unseasonably cold days), with decreasing degree of severity. The middle bin, 40%-60%, corresponds to days in which temperatures are approximately in line with the historical mean and is our omitted category. The last five bins capture events with temperatures higher than the historical mean (unseasonably hot days),

[11] Different stations have different baseline propensities to cover local weather, which means that results are best estimated as proportional effects. We chose a Poisson model over a log transformation since (i) our outcome variable is a count and (ii) our dataset includes a few days when the local weather is not talked about, so that using a log transformation would have required ad hoc adjustments (such as, using the log+1 transformation, or dropping these instances). We show that our results are robust to using OLS and outcomes in logs or shares in Online Appendix 3.E.

now with increasing degree of severity. The estimates from this specification are reported in our figures. In addition, we estimate a more parsimonious specification in which we split weather events according to five categories only: bottom 10%, 10%-40%, 40%-60%, 60%-90%, and top 10%. The estimates from this specification are reported in our tables.

In our analysis, we include the weather event of the same day as news coverage ($k = 0$ in Equation 3.1), the weather event of the day before news coverage ($k = -1$) and the weather event of the day after news coverage ($k = 1$). This is to avoid attributing daily coverage to past or future weather events, especially given the correlation between weather events over time. In other words, this allows us to isolate the effect we are interested in estimating: the contemporaneous effect of weather events on coverage of local weather. We also add date fixed effects (δ_t), which control for differences in deviations or reporting that affect all stations equally such as national events happening on the same date, station fixed effects (δ_s), which allow different stations to have a different baseline propensity to cover local weather, and fixed effects for different deciles of the number of segment distribution across all stations ($\delta_{\eta(st)}$), which control for the space allotted to local newscasts by station s in date t.

The β_0^ρ are our coefficient of interest. Given that we estimate a Poisson model and that the omitted category is always the 40-60% bin, the β_0^ρ coefficients correspond to the average percentage change in the coverage of local weather events when the deviation from the historical mean falls into the ρ^{th} bin (e.g., the bottom 5% bin) of the within-market deviation distribution relative to the coverage when temperatures are close to the norm (i.e., the deviation falls in the reference bin of 40-60%).

3.4.2 Results

Figure 3.3 reports the average publication strategy that we estimate using Equation 3.1, separately for summer (panel (a)) and winter (panel (b)). In turn, Table 3.1 displays the summary estimates (column (1) for summer and column (3) for winter).

In summer, both unseasonably cold and unseasonably hot days see higher coverage of local weather relative to days with deviations in the 40% to 60% range of the deviation distribution. When temperatures are below the historical mean—that is, in unseasonably cold summer days—deviations have to be large to induce an increase in coverage of local weather. Instead, unseasonably hot days induce stations to increase their coverage of local weather also when the

Figure 3.3: Publication Strategies

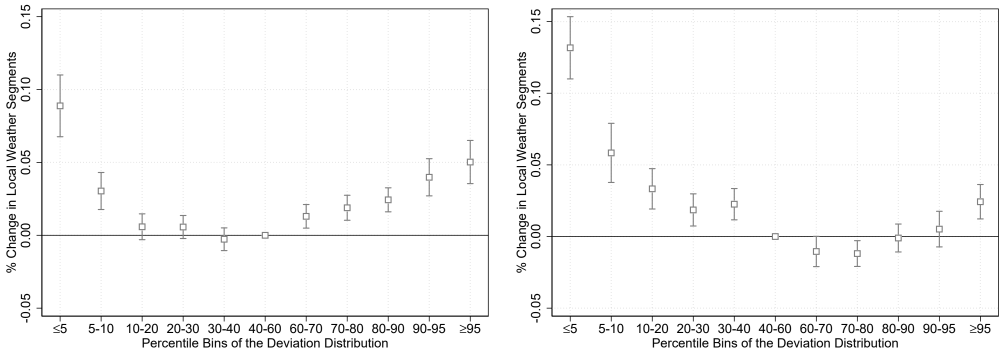

(a) Summer

(b) Winter

Notes: This figure shows the relationship between news coverage of local weather and weather events. We regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, one lead and one lag of the same indicators, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.1). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (panel (a)) and for winter (panel (b)), using a Poisson model. Standard errors are clustered at the media market level.

deviations are small. We again see that the larger the deviation the larger the increase, although the largest deviations above the mean induce smaller increases than the larger deviations below the mean. Specifically, days in the bottom 10% of the deviation distribution display 5.9% higher coverage than days in which temperatures are in line with the usual, while days in the top 10% of the deviation distribution increase coverage by 4.5% (Table 1 column (1)).

In winter, unseasonably cold days receive more attention relative to days in which temperatures are in line with the norm and, the larger the size of the deviation, the higher the increase in coverage. Table 1 column (3) shows that days in the top 10% of the deviation distribution see a 10.2% increase in coverage of local weather relative to days with temperatures in line with the historical mean. Smaller deviations also experience higher coverage, with an increase of 2.5%. Instead, coverage of local weather in unseasonably hot days only increases when temperatures are well above the historical mean (top 5% of the deviation distribution) and is otherwise constant or even lower for deviations of smaller magnitudes. The asymmetry in these patterns might be surprising at first glance. These weather events, however, may be less "remarkable." For example, Moore et al. (2019) use social media data for the United States and show that in cold periods, people tweet more (and more negatively) about the weather when temperatures are below the historical mean and less (and using more positive sentiment) when temperatures are above the mean.

Taken together, these results show that severe weather events are treated as fundamentally newsworthy—that is, "man bites do" is indeed news—, but intermediate temperature deviations are also sometimes covered—"big dog's bites" can also be news.

3.5 Publication Bias

In this section, we begin to entertain the possibility of factors other than newsworthiness influencing coverage of weather events. We ask whether coverage patterns vary with the political leaning of the media market in which stations operate. This is particularly compelling because weather—unlike elections, social issues, or economic policies—is at first glance a less explicitly political topic. Still, weather events carry a political connotation to the extent that they intersect with climate change, which suggests that media coverage of weather events could potentially vary based on political ideology. As above, before presenting our results, we describe our empirical strategy.

3.5.1 Empirical Approach

To test for the presence of publication bias, we estimate separate editorial strategies for stations that operate in Democratic-leaning, swing, or Republican-leaning media markets. We do so by including interactions between the indicators capturing the weather events and our measure of media market's political leaning. Our baseline specification to estimate publication bias is the following:

$$
\begin{aligned}
Y_{st} = & \sum_{\rho} \sum_{k \in \{-1,0,1\}} \beta_k^{\rho} \mathbb{I}\{\rho^{th} bin\}_{m(s),t+k} \\
& + \sum_{i \in \{D,R\}} \sum_{\rho} \sum_{k \in \{-1,0,1\}} \beta_{i,k}^{\rho} \mathbb{I}\{\rho^{th} bin\}_{m(s),t+k} \times \mathbb{I}\{ideology_{m(s)} = i\} \\
& + \delta_s + \delta_t + \delta_{\eta(st)} + \epsilon_{st}, \qquad (3.2)
\end{aligned}
$$

where $\mathbb{I}\{ideology_{m(s)} = i\}$ are indicator variables equal to 1 if media market $m(s)$ has political leaning i and all other variables are defined as before. Similarly, weather events can be either finely or coarsely defined, the specification is estimated through a Poisson regression, and standard errors are clustered at the media market level.

3.5.2 Results

Figure 3.4 reports the estimates we obtain from the fully flexible version of Equation 3.2. Specifically, the figure reports publication strategies for stations

Figure 3.4: Publication Bias

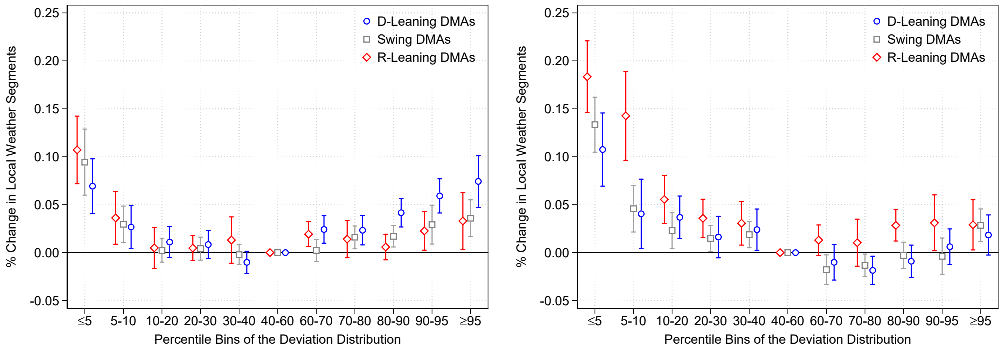

(a) Summer (b) Winter

Notes: This figure shows the relationship between news coverage of local weather and weather events by media market political leaning. We regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, the same indicators interacted with indicators for the market being either Democratic- or Republican-leaning, one lead and one lag of the same variables, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.2). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (panel (a)) and for winter (panel (b)), using a Poisson model. Standard errors are clustered at the media market level. Note that the figure reports the overall effect in each type of media market (that is, β_0^ρ for Swing media markets and $\beta_0^\rho + \beta_{i,0}^\rho$ for Democratic- and Republican-leaning DMAs).

that operate in different media markets. That is, for each weather event we show β_k^ρ for swing media markets, $\beta_k^\rho + \beta_{R,k}^\rho$ for Republican-leaning media markets, and $\beta_k^\rho + \beta_{D,k}^\rho$ for Democratic-leaning media markets. Instead, Table 3.1 columns (2) and (4) report estimates from the summary specification. The table reports the $\beta_{i,k}^\rho$ coefficients directly.

Figure 3.4 and Table 3.1 provide clear evidence of publication bias: stations that operate in Republican- and Democratic-leaning media markets follow drastically different publication strategies in their reporting of the same weather events. We begin by discussing the results for summer (reported in panel (a) in the figure and column (2) in the table).

No matter the media market, we observe that unseasonably cold and unseasonably hot days receive more coverage than days in which temperatures are in line with the historical mean. A more interesting pattern emerges when we compare by how much coverage increases. In days with temperatures below the historical mean, coverage increases more in stations in Republican-leaning markets than in stations in Democratic-leaning markets (with the difference just reaching statistical significance at the 5% level). Stations in swing media markets are somewhat in-between. The pattern is exactly reversed in days with temperatures above the historical mean. There, we see a greater increase in

Table 3.1: Publication Strategies and Bias

	Local Weather Segments			
	Summer		Winter	
	(1)	(2)	(3)	(4)
Bottom 10%	0.059***	0.061***	0.102***	0.096***
	(0.007)	(0.012)	(0.009)	(0.012)
Bottom 10% × D-Leaning		-0.015		-0.012
		(0.016)		(0.019)
Bottom 10% × R-Leaning		0.018		0.070***
		(0.018)		(0.020)
10%-40%	0.001	-0.001	0.025***	0.018***
	(0.003)	(0.004)	(0.005)	(0.006)
10%-40% × D-Leaning		0.004		0.009
		(0.007)		(0.010)
10%-40% × R-Leaning		0.008		0.021**
		(0.008)		(0.011)
60%-90%	0.018***	0.012***	-0.011***	-0.015***
	(0.003)	(0.004)	(0.004)	(0.006)
60%-90% × D-Leaning		0.017**		-0.001
		(0.007)		(0.009)
60%-90% × R-Leaning		0.000		0.030***
		(0.007)		(0.008)
Top 10%	0.045***	0.032***	0.007	0.003
	(0.006)	(0.008)	(0.005)	(0.007)
Top 10% × D-Leaning		0.035***		0.001
		(0.012)		(0.011)
Top 10% × R-Leaning		-0.005		0.022*
		(0.011)		(0.013)
Observations	297659	297659	288726	288726
Stations	698	698	699	699
DMAs (Clusters)	204	204	204	204
Mean Dep. Variable	27.905	27.905	27.915	27.915
Bottom 10% × D = Bottom 10% × R		0.050		0.000
10%-40% × D = 10%-40% × R		0.629		0.328
60%-90% × D = 60%-90% × R		0.023		0.001
Top 10% × D = Top 10% × R		0.001		0.155

Notes: This table shows the relationship between news coverage of local weather and weather events, on average and by media market political leaning. In columns (1) and (3), we regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, one lead and one lag of the same indicators, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.1). Columns (3) and (4) additionally include the weather events indicators interacted with interacted with dummies for the market being either Democratic- or Republican-leaning, with one lead and one lag of the same variables (see Equation 3.2). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (columns (1) and (2)) and for winter (columns (3) and (4)), using a Poisson model. Standard errors are clustered at the media market level.

coverage in stations operating in Democratic-leaning media market than in those with a Republican-leaning audience with, again, TV stations in swing media markets in-between. For those unseasonably warm days, the difference is statistically significant both for medium-sized deviations, falling within the 60th and 90th percentile, and large deviations, in the top 10% (p-value = 0.023 and 0.001 respectively).

Patterns are slightly different in winter. In particular, we observe larger increases in coverage of weather for both unseasonably cold and unseasonably hot days in stations located in Republican-leaning media markets than in stations operating in Democratic-leaning media markets. Again, the estimates for swing media markets are generally in-between. Below the mean, the difference in coverage increase is only statistically significant between Republican and Democratic media markets for very cold day (bottom 10% deviation). Above the mean, the difference is only statistically significant for intermediate shocks (deviations falling between the 60th and 90th percentile).

Overall, stations that operate in Republican- and Democratic-leaning media markets follow different publication strategies in their reporting of similar weather events. Editors do not just take into account the newsworthiness of an event (its rarity) in their coverage choices, political factors appear to also matter. In short, coverage decisions are biased.

3.6 Beyond Publication Bias

Editors do not just decide whether to cover an event, they also choose how much emphasis to give to a story and how to talk about it. In this section, we look at other aspects of outlets' editorial strategies: salience (that is, how prominently a weather event is featured) and framing (that is, how explicitly it is connected to the broader climate crisis).

Salience. To study the salience of weather news, we adopt the approach used for newspapers. Rather than the page in which the story appears (e.g., Wasow, 2020), we look at the rank of a segment in a newscast. More precisely, our dependent variable is the log of the average minimum rank in which a weather story appears in a given day. The specification is the same as Equation 3.2, but we now estimate it using OLS as the average rank can be non-integer.[12]

Figure 3.5 displays the results from this analysis.[13] Note that, for consistency with our previous exhibits, we have inverted the y-axis so that estimates visually above zero corresponds to higher salience of the news story. The patterns we observe match our findings for publication bias. Overall, the more extreme the weather event, the higher its salience in the newscast (with the exception of above

[12] Using a log transformation of the outcome allows us to be consistent with our approach in the rest of the paper and estimate proportional effects.

[13] Appendix Table 3.B.1 reports the corresponding estimates from our summary specification. We also report for specifications that control for local weather story decile fixed effects to avoid mechanical effects whereby the ranking of local weather story diminishes because there are simply more stories reported in the newscast. Our results remain the same.

Figure 3.5: Salience Bias

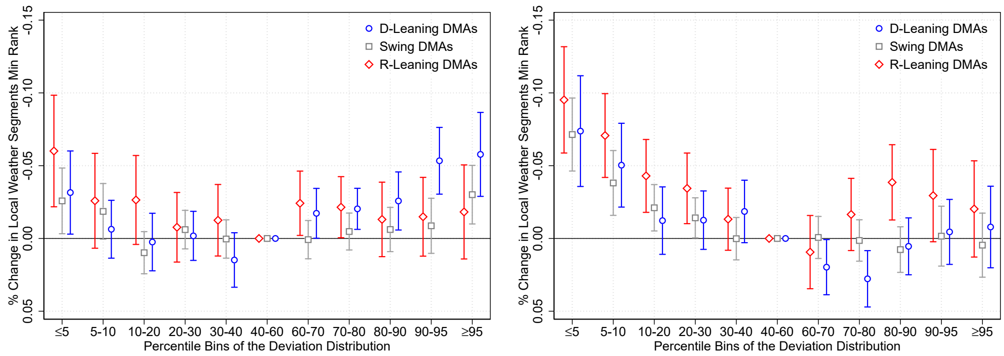

(a) Summer (b) Winter

Notes: This figure shows the relationship between the salience of news coverage of local weather and weather events by media market political leaning. We regress the log minimum rank of segments about local weather (averaged across the different newscasts) on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, the same indicators interacted with indicators for the market being either Democratic- or Republican-leaning, one lead and one lag of the same variables, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.2). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (panel (a)) and for winter (panel (b)), using OLS. Standard errors are clustered at the media market level. Note that the figure reports the overall effect in each type of media market (that is, β_0^ρ for Swing media markets and $\beta_0^\rho + \beta_{i,0}^\rho$ for Democratic- and Republican-leaning DMAs).

the mean events in winter). In summer, stations in Republican-leaning media markets increase the salience (that is, decrease the mean rank) of local weather news in unseasonably cold days more than stations in swing or Democratic-leaning media markets. Instead, unseasonably hot days in summer see a larger increase in the salience of local weather news in Democratic-leaning versus Republican-leaning markets. In winter, we see higher responsiveness of stations in Republican-leaning markets both in unseasonably cold and unseasonably hot days. Combining the results of Section 3.5 with those of this subsection, we observe that when weather events receive more coverage, they are also given more prominence. Salience bias and publication bias go hand-in-hand.

Framing. We also examine the framing of weather news by analyzing how local broadcasts present temperature deviations relative to historical norms. To do so, we take a straightforward approach and measure the frequency with which the terms "climate change" and "global warming" appear in stories about local weather. Figure 3.6 illustrates the extent to which these terms are mentioned in relation to contemporaneous weather events, separately by the political leaning of the media market. The most striking finding is the sheer scarcity of such references: only 0.1% to 0.3% of the local weather stories we identify mention climate change.

Figure 3.6: Climate Change in Local Weather Coverage, by DMA Political Leaning

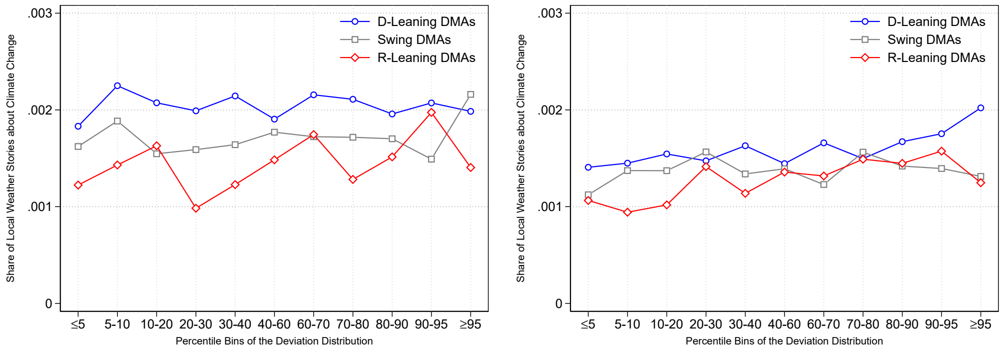

(a) Summer

(b) Winter

Notes: This figure shows climate change references in coverage of local weather in different weather events, by the political leaning of the media market. In particular, we show the average share of segments about local weather that mention climate change or global warming in each bin of the within-media market deviation distribution, separately for summer (panel (a)) and for winter (panel (b)).

To systematically investigate the thematic focus of weather coverage and the extent to which it engages with climate discourse, we also train a topic model on our local weather segments (see Appendix 3.G for details on the methodology and text inputs). The topic distributions across winter and summer, presented in Appendix 3.G, reveal that, regardless of the season, local weather segments overwhelmingly concentrate on immediate, short-term conditions, emphasizing terms such as "snow," "rain," "temperatures," "morning," "cold," and "heat." While extreme events such as hurricanes, floods, or tornadoes occasionally feature in these reports, broader contextualization is rare (see Appendix Tables 3.G.1 and 3.G.2).

These findings highlight a crucial pattern: day-to-day deviations from historical temperature means are rarely framed as evidence of climate change. Instead, local weather newscasts remain focused on short-term, practical concerns, with no discussions of climate largely.

3.7 Understanding Publication Bias

Our results so far point to the presence of publication bias in the coverage of weather events. Stations that operate in Republican-leaning and Democratic-leaning media markets follow different editorial strategies. In this section, we discuss several possible explanations for the patterns we uncover.

3.7.1 Is This Really Publication Bias?

To start with, we have to entertain the possibility that stations in Republican-leaning and Democratic-leaning media markets follow different publication strategies for reasons that have nothing to do with politics. In this subsection, we briefly discuss, and one by one present evidence against, several alternative explanations. We provide more details for each of the empirical tests we implement in Appendix 3.E.

Demand for Weather Coverage. First, it is possible that the patterns we uncover are simply explained by higher demand for reporting on rare weather events (deviations falling in the top or bottom of the distribution) in Republican-leaning markets. However, if it were so, the increase in coverage should be universally larger in Republican-leaning market than in Democratic-leaning markets, which is not what we observe for weather events above the norm in the summer. Also, the fact that on average local weather coverage is higher in Democratic- relative to Republican-leaning markets (see Appendix Figure 3.A.2) provides additional suggestive evidence against this explanation.

Differences in Weather Events. Second, it is possible that these patterns could be explained by Republican- and Democratic-leaning media markets experiencing substantively different weather events. Again, this is unlikely to be the case. Figure 3.2 already shows that the distribution of deviations within each weather event in the three types of media market is overlapping. And even if we were to take any differences at face value, they would be too small to explain our effects (see Appendix Table 3.E.1). In addition, our results are robust to only comparing media markets that are broadly exposed to similar weather patters, as we can show by estimating specifications with climate region and year fixed effects (see Appendix Table 3.E.2).

Confounding Market Characteristics. Third, we show that the patterns we documented above are not driven by media market's characteristics that are different than political leaning, but happen to correlate with it. Our estimates of publication bias survive controlling for weather shocks interacted with media market's area, population, the share of urban population, and industry shares (Appendix Table 3.E.3). In addition, they are also not driven by natural disasters or wildfires, than might both receive media attention and correlate with weather (Appendix Table 3.E.4).

Measurement and Model Choices. Finally, they are not driven by specific measurement and modeling choices that we made in the estimation, as they are robust to using longer segments, estimating an OLS specification with the outcome in logs or in share, and splitting markets by ideology using the 20th and 80th percentiles and the 30th and 70th percentiles (Appendix Table 3.E.5). No matter how we slice the data, publication bias persists.

3.7.2 Sinclair Acquisitions

If the political leaning of a media market is indeed behind the publication bias we document, as the previous subsection strongly suggests, what can explain these variations in coverage? We see two big classes of explanations here: the difference in coverage could be driven by supply factors (outlets attempting to persuade viewers) or demand factors (outlets responding to viewers' preferences). In this section, we put the supply-side theory to the test by examining what happens when a conservative media group, Sinclair, takes over local TV stations.

Over the 2013 to 2019 period, the Sinclair Broadcast Group (Sinclair) went from controlling 44 stations to almost 100 stations. Because Sinclair has a strong conservative leaning (Miho, 2024) and existing research has shown that Sinclair influences the content of the stations it acquires (Martin and McCrain, 2019), there are good reasons to anticipate that if coverage is meant to influence viewers, local TV stations acquired by Sinclair should change their publication strategies of weather events. In particular, weather events associated with climate change (i.e., temperature deviations above the historical mean) should receive less coverage after acquisition. At the same time, we would expect the opposite pattern for weather events possibly going against how climate change is generally understood (i.e., temperature deviations below the historical mean).

We test this hypothesis using a differences-in-differences design based on the staggered timing of Sinclair acquisitions. Identification rests on a parallel trends assumption: non-Sinclair stations provide the correct counterfactual for the editorial strategy that the stations that get acquired by Sinclair would have followed, had they had not been acquired. The main concern is that Sinclair acquisitions might be endogenous to station- or media market-level trends.[14] We perform two sets of analysis to alleviate these concerns. First, we estimate event study specifications in which we allow editorial strategies to vary in time since/to the acquisition, to provide suggestive evidence that pre-trends are not too much of an issue. Second, we estimate a triple differences specification in which we include media market by date fixed effects. This ensures that we are only using

[14]Importantly, Martin, Mastrorocco, et al. (2024) and Mastrorocco and Ornaghi (Forthcoming) provide evidence that Sinclair's acquisition strategy is mostly an expansion strategy. Sinclair buys stations that come to the market rather than targets specific stations or areas.

within-media market variation, and allows us therefore to non-parametrically control for media market specific shocks.

We begin by estimating the following differences-in-differences specification:

$$
\begin{aligned}
Y_{st} =& \gamma \mathbb{I}\{\text{Sinclair}\}_{st} + \sum_{\rho} \sum_{k \in \{-1,0,1\}} \beta_k^{\rho} \mathbb{I}\{\rho^{th} \text{ bin}\}_{m(s),t+k} \qquad (3.3)\\
&+ \sum_{\rho} \sum_{k \in \{-1,0,1\}} \lambda_k^{\rho} \mathbb{I}\{\rho^{th} \text{ bin}\}_{m(s),t+k} \times \mathbb{I}\{\text{Sinclair}\}_{st} \\
&+ \delta_{\eta(st)} + \delta_s + \delta_t + \epsilon_{st}, \qquad (3.4)
\end{aligned}
$$

where all variables are defined as before and $\mathbb{I}\{\text{Sinclair}\}_{st}$ is an indicator variable equal to 1 if the station is under Sinclair control. Our triple differences specification additionally includes $\delta_{m(s)t}$, that is, media market by date fixed effects. For the sake of the readability of the results, we only estimate the constrained version of our specification focusing on large and medium-sized weather events, above and below the mean. As before, we use a Poisson regression and we cluster standard errors at the media market level.

Table 3.2 presents the results with DMA fixed effect and day fixed effect (odd columns), with DMA-day fixed effect (even columns), and removing the always treated DMAs by restricting the sample to markets where Sinclair is not present in 2013 (columns (3) and (4) for winter and (7) and (8) for summer). In general, we observe that after Sinclair acquires a station, coverage of weather events in days in which temperatures are in line with the historical mean increases, perhaps because this type of coverage might be relatively cheap to produce. However, this increase in coverage is not differential by weather event with the exception of days with large deviations above the historical mean in winter, that experience even higher coverage. In no case do we observe lower coverage of weather deviations above the historical mean. Event studies in which we allow the effect of Sinclair to vary by time since/to acquisition (Appendix Figure 3.A.3 and 3.A.4) show no evidence of pre-trends and confirm these findings.

3.7.3 Discussion

The empirical findings presented in the preceding section challenge the notion that supply-side factors predominantly drive publication bias in weather coverage. Instead, by process of elimination, the results lend credence to the hypothesis that demand-side considerations play a pivotal role in shaping editorial strategies. Two additional pieces of evidence provide support for this conclusion.

Table 3.2: Sinclair Acquisitions and Weather Events Coverage

	Local Weather Segments							
	Summer				Winter			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sinclair	0.090*	0.115**	0.141***	0.158***	0.075	0.124**	0.054	0.079*
	(0.053)	(0.053)	(0.043)	(0.044)	(0.052)	(0.054)	(0.043)	(0.045)
Bottom 10%	0.061***	0.066***			0.104***	0.102***		
	(0.008)	(0.009)			(0.009)	(0.011)		
Bottom 10% × Sinclair	-0.012	-0.006	-0.002	0.007	-0.020	-0.025	-0.015	-0.013
	(0.015)	(0.018)	(0.013)	(0.016)	(0.016)	(0.023)	(0.014)	(0.024)
10%-40%	0.002	0.001			0.024***	0.021***		
	(0.003)	(0.003)			(0.005)	(0.006)		
10%-40% × Sinclair	-0.001	0.014	-0.005	0.000	0.004	-0.010	0.003	-0.006
	(0.007)	(0.010)	(0.008)	(0.010)	(0.009)	(0.011)	(0.008)	(0.014)
60%-90%	0.018***	0.016***			-0.012***	-0.015***		
	(0.003)	(0.004)			(0.004)	(0.004)		
60%-90% × Sinclair	0.004	0.012	0.006	0.016*	0.016*	0.004	0.006	-0.004
	(0.008)	(0.008)	(0.007)	(0.008)	(0.009)	(0.012)	(0.008)	(0.011)
Top 10%	0.043***	0.046***			0.004	0.003		
	(0.006)	(0.007)			(0.005)	(0.006)		
Top 10% × Sinclair	0.013	0.015	0.014	0.002	0.030***	0.019	0.028***	0.034**
	(0.010)	(0.013)	(0.010)	(0.015)	(0.010)	(0.013)	(0.011)	(0.015)
Date FEs	✓	✓			✓	✓		
DMA-By-Date FEs			✓	✓			✓	✓
Drops Always Treated		✓		✓		✓		✓
Observations	297659	228724	289165	220499	288726	222094	280292	213730
Stations	698	537	683	522	699	537	684	522
DMAs (Clusters)	204	164	189	149	204	164	189	149
Mean Dep. Variable	27.905	27.241	28.489	27.948	27.915	27.261	28.526	28.028

Notes: This table tests whether the relationship between news coverage of local weather and weather events changes after Sinclair acquires a local TV station. In columns (1) and (5), we regress the number of segments about local weather on an indicator for the station being under Sinclair control, indicator variables for the deviation from historical mean falling in a given bin of the within-media market deviation distribution, one lead and one lag of the same indicators, the interaction of the Sinclair and weather events indicators, one lead and one lag of these interactions, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.4). Columns (2) and (4) additionally include media market by date fixed effects. All even columns drop always treated media markets. The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (columns (1) to (4)) and for winter (columns (5) to (8)), using a Poisson model. Standard errors are clustered at the media market level.

First, we have assumed that the biased coverage of weather events may be due to the connection of such events with climate change. A large body of work investigates how temperatures affect climate change beliefs (see Howe et al., 2019 for a review). In Online Appendix 3.F we show using individual-level survey data from CCES and variation in interview dates, that, indeed, individuals do learn from weather events, and that this learning is heterogeneous depending on individuals' ideology.

Second, and in connection to this, the political leaning of media markets is proxying for attitudes towards climate change. Indeed, we can show using the same survey data that individuals who self-identify as liberal or belonging to the Democratic party are much more likely to believe that climate change requires immediate action than individuals who self-identify as conservative or belonging to the Republican party (see Appendix Figure 3.A.5). In line with this, we show in Appendix Table 3.B.2 patterns consistent with publication bias when we classify markets as climate change sceptic, non-sceptic, or neutral again using the CCES data.[15]

Taken together, these insights reinforce the view that editorial choices in local weather coverage are not merely the byproduct of journalistic discretion but are, to a considerable extent, shaped by audience demand and ideological predispositions. This interplay between demand-driven coverage and confirmation bias has profound implications for public discourse, influencing how individuals interpret and internalize information about climate change. In the subsequent section, we formalize these dynamics within a theoretical framework, offering a structured analysis of the determinants of media coverage.

3.8 A Model of Coverage Decision

We now provide a stylized theoretical framework to rationalize the empirical patterns we have documented above. Our main contribution here is to study how much an outlet devotes to the event in its report rather than whether to cover this event.

3.8.1 Model Setup

We consider a game with one media outlet M and a mass of citizens of size one divided into two groups, $J \in \{D, R\}$. To study demand-driven and supply-driven coverage, we will consider in the analysis two possible types of outlet: a profit-maximizing outlet and a biased outlet who seeks to persuade citizens (we will look at a climate-sceptic and a climate-believer outlet). The objective of a media outlet is known (ours is not a model of reputation building). The media outlet commits to a publication (or equivalently coverage) strategy, which consists in the amount of time devoted to weather news (denoted w) relative to other news (denoted n). When it comes to citizens, group D consists of Democratic citizens (which constitute a proportion α of the population) and group R consists of Republican citizens (a proportion $1-\alpha$). Each citizen decides whether to watch

[15] This raises the question of why not use climate change beliefs in the first place. The reason for this choice is data driven: CCES data, which to the best of our understanding are the most appropriate source of data for attitudes towards climate change in the early 2010s, are not representative at the media market level, which generates potential measurement issues.

the outlet based on the (anticipated) entertainment value of news and what they can learn from it. Given our interest in weather news, we suppose that only those are informative.

More precisely, we capture a weather event as a random variable $\tilde{c}$ drawn from the interval $[-1, 1]$. We assume that the distribution from which an event is drawn depends on an underlying state of the world $\omega \in \{0, 1\}$. State $\omega = 0$ captures the absence of climate change (or, perhaps more accurately these days, the idea that climate change is not man-made), while state $\omega = 1$ indicates that there is climate change (or climate change is man-made). The underlying state is unknown to all actors and we assume that the common prior is that the realisation of the state variable is 1 with probability π: $Pr(\omega = 1) = \pi \in (0, 1)$.[16] A weather event is drawn from the Cumulative Distribution Function (CDF) $F_\omega(\cdot)$, with continuous probability density function (pdf) $f_\omega(\cdot)$, itself differentiable and $f_\omega(\cdot)$ single-peaked around zero. We also assume that large weather events c are relatively more frequent when the realisation of the state is $\omega = 1$. We impose that the pdfs $f_\omega(\cdot)$, $\omega \in \{0, 1\}$ satisfy the strict monotone likelihood ratio property: $\frac{f_1(c)}{f_0(c)} > \frac{f_1(c')}{f_0(c')}$ for all $c > c'$.

For each realization of weather event $c \in [-1, 1]$, a media outlet decides the amount of coverage $w(c)$. As newscasts are time constrained, and without loss of generality, we assume that the total time available for local weather news and other news is one so that $w(c) + n(c) = 1$. If a citizen watches the newscast, she can update about the likelihood of (human-made) climate change. We denote her posterior belief as a function of the coverage of weather news by $\rho(w(c))$. This posterior consists either of what a citizen learns from learning c if $w(c) > 0$ $(\rho(w(c)) = Pr(\omega = 1|c) =: \mu(c))$ or of what a citizen infers from not observing any news about c given the outlet's strategy if $w(c) = 0$ $(\rho(0) = Pr(\omega = 1|\emptyset) =: \mu(\emptyset))$. If a citizen does not watch the newscast, we assume that she learns nothing (a simplification without loss of generality) and her posterior is her prior, equal to π.

The utility of consuming the newscast for potential viewers depends on two of its features. First, the amount of reporting on each type of news captured by the functions $g(n)$ and $h(w)$, which are strictly increasing in their argument. Second, the utility depends on the value of each news item, which acts as a scale-up. For non-weather news, the value for all citizens is assumed to be constant and equal to a finite value $\underline{u} < 1$. In turn, the value of weather news for a viewer from group $J \in \{D, R\}$ is: $v(c) + z_J(\rho(w(c)))$. The function $v(c)$ captures the

[16] Notice that this is the prior of the population, the relevant one for our purpose, which can differ from the prior in the scientific community. We could add different priors for different groups without affecting our results.

entertainment benefit of watching news about weather events. We assume $v(c)$ is strictly decreasing in c for $c < 0$ and strictly increasing in c for $c > 0$ so that citizens prefer to watch news reports about extreme, rare weather events than common weather events. The second function $z_J(\rho(w(c)))$ captures the cost or benefit of learning. We assume that citizens suffer from a form of confirmation bias. We represent this by assuming that $z_D(\mu(c))$ is strictly increasing in its argument, whereas for Republican citizens $z_R(\mu(c))$ is strictly decreasing in its argument, with $z_D(0) = z_R(0) = 0$. We assume that $v(1) + z_D(1) < 1$ and $v(0) + z_R(\pi) > 0$, a normalization.

The overall value of consuming the newscast for a citizen from group $J \in \{D, R\}$ is then:

$$\overbrace{h(w)}^{\text{Value of time on weather news}} \times \underbrace{\big(v(c) + z_J(\rho(w(c)))\big)}_{\text{Comsumption value of weather news}} + \overbrace{g(n)}^{\text{Value of time on other news}} \times \underbrace{\underline{u}}_{\text{Consumption value of other}} \tag{3.5}$$

A citizen i can always decide not to watch the news in which case she receives her outside option payoff equal to an idiosyncratic event δ_i uniformly distributed (i.i.d.) over the interval $[0, \overline{\delta}]$, with $\overline{\delta} > 1$ (so that share of viewers is always less than the full population).

When it comes to media outlets, we first need to denote A as the proportion of citizens who watch the newscast. A profit-maximizing TV station only seeks to maximize its audience so its payoff is simply A (we assume that there is a monotonic relationship between audience size and revenues). In turn, a biased TV station seeks to persuade. A climate-believer outlet seeks to maximize the average belief in climate change. For each c, this is equivalent to maximizing $A\rho(w(\widetilde{c})) + (1 - A)\pi$ (viewers update about climate change from the newscast, non-viewers keep their prior). Hence, we represent a climate-believer station's payoff as:

$$\int_{\widetilde{c}} V_b(A\rho(w(\widetilde{c})) + (1 - A)\pi) d(\pi F_1(\widetilde{c}) + (1 - \pi)F_0(\widetilde{c})), \tag{3.6}$$

with $V_b(\cdot)$ a strictly increasing, continuous, twice differentiable, and strictly concave function. The outlet considers all possible weather events ($\int_{\widetilde{c}}$), since it commits to a strategy, as well as the expected distribution of weather events $(\pi F_1(\widetilde{c}) + (1 - \pi)F_0(\widetilde{c}))$ as it does not know ω when it chooses its coverage strategy.

In turn, a climate-sceptic outlet wants to minimize the average belief in climate change. As a result, we assume that its payoff assumes the following form

$$\int_{\widetilde{c}} V_s(A\rho(w(\widetilde{c}) + (1 - A)\pi) d(\pi F_1(\widetilde{c}) + (1 - \pi)F_0(\widetilde{c})), \tag{3.7}$$

with $V_s(\cdot)$ a strictly decreasing, continuous, twice differentiable, and strictly concave function.

The game, in turn, proceeds as follows:

0. Nature draws the state of the world $\omega \in \{0, 1\}$.
1. The outlet publicly commits to a publication strategy: $w : [0, 1] \to [0, 1]$.
2. Weather event c is drawn by Nature and coverage occurs according to the editorial strategy.
3. Citizens decide whether to watch the outlet. They observe what is reported if they watch the newscast and nothing if they don't.
4. Game ends and payoffs are realized.

The equilibrium concept is Perfect Bayesian Equilibrium. We focus on stationary pure strategies. A stationary pure strategy for the outlet is a mapping from the set of possible realizations of climate event to a coverage decision. (We are not looking for a coverage function, rather for a value of coverage for each c separately.) We also add a few assumptions to simplify the analysis. First, we impose that $\mu(0) = Pr(\omega = 1|c = 0) = \pi$. Second, we suppose that $h(w) = w$ (i.e., weather news acts as a numeraire), whereas $g(n)$ satisfies $g'(0) > 1$ and $g'(1) = 0$ and is $\mathscr{C}^\infty$ and strictly concave. Finally, throughout, we assume that the value of entertainment is greater than the value of learning: $v'(c) > -\mu'(c)z_R(\mu'(c))$ for all $c > 0$ and $v'(c) < -\mu'(c)z_D(\mu'(c))$ for all $c < 0$.

Before proceeding to the analysis, a few remarks are in order. First, building on previous works (e.g., DellaVigna and La Ferrara, 2015; Durante, Pinotti, and Tesei, 2019), our model assumes that citizens consume the outlet, and even news, primarily for entertainment value. For weather events, this corresponds to citizens finding more entertainment in large, unexpected weather events than small, common events. This can correspond to the idea that citizens like surprise as in Ely, Frankel, and Kamenica (2015), though we recognize we model the value of surprise very differently.

Yet, entertainment is not everything and we suppose that individuals learn from weather events. This assumption is crucial for our results below. Yet, we do not think it is unwarranted. As noted above, in Appendix 3.F, we use individual-level survey data from the CCES to show evidence of exactly the type of learning we assume in this model. While some individuals may observe the event directly, reducing the value of a newcast, media markets generally cover large area meaning that most possible watchers do not know about the weather in this relevant geographical area. On top of this, mentions of the weather event on the newscast could help a citizen interpret the event. A citizen can see that today is a hot day, but she may be unable to fully make sense of how hot it is relative to the historical mean if she doe not watch the newscast. In formal terms, her own experience of the weather can affect her belief about the distribution of $\tilde{c}$ that day, but she can only learn the realisation c of the random variable by watching the outlet.

Our set-up further assumes that citizens differ in their demand for weather news depending on the group they belong to. We model this as a form of confirmation bias as in Mullainathan and Shleifer (2005) and Anand, Di Tella, and Galetovic (2007) (though, beyond the differences in payoffs, all citizens are fully rational). Democrats experience a utility gain if they see evidence confirming (man-made) climate change is occurring. In contrast, Republicans, who tend to oppose call for actions on climate change following weather events above the norms (Appendix Table 3.F.1), suffer a loss when the evidence goes against their belief.[17] This approach captures the idea that "[d]enialism is motivated by conviction rather than evidence" (Kemp, Milne, and Reay, 2010) and the evidence that deniers simply reject evidence point to the existence of climate change (Washington, 2013).

Another payoff assumption worth commenting on is the complementarity between the amount of coverage received by a news item and the value of the news. More reporting on a news item provides higher utility, but the marginal value of increased coverage reduces with the amount of time spent on it. In turn, the value of an event acts as a scale-up. It increases the marginal value of an additional minute of coverage. These assumptions are meant to capture the choice of a media outlet of how to structure its newscast given the events that occurred that day. Fixing the value of other news to $\underline{u}$ is without loss of generality.

The attentive reader would have noted that we allow an outlet to commit to an editorial strategy. This is a strong assumption as such strategy is not a contract with viewers and there would be no institution to enforce a hypothetical contractual agreement. Why this assumption then? We look at the best case scenario for a biased outlet. Absent commitment, the biased station's problem becomes one of information disclosure. As it is well known in the literature, an unravelling argument would apply and a biased outlet's coverage strategy would

[17] Our approach to information is markedly different from Armona et al. (2024). They consider that viewers always want to learn about an underlying state of the world. In contrast, we assume an entertainment value of news and that the payoff from learning which depends on the event and individuals' ideological predisposition. As a result, we can dispense from the assumption that consumers do not learn if they observe nothing, which Armona et al. (2024) need for their results (otherwise, an unravelling argument would yield that outlets always cover an event no matter its informativeness in their model, contrary to what they seek to demonstrate).

be independent of the weather event (a previous version of the paper formally proves this point). By imposing commitment, we leave open the possibility that supply-driven coverage varies with the weather event (indeed, we will show it may). In other words, we make it possible that a supply-driven coverage can match our empirical patterns (though, we will show it does not).

Finally, we briefly describe how our theoretical parameters and choices map into our empirical quantities above. The theoretical weather event c can be understood as the percentile in the distribution of all possible deviations from the mean. It is the theoretical equivalent to our dependent variable in the regression above. We see 0 as the seasonal norm, any positive (negative) value as weather events above (below) the norm with values further away from zero indicating greater deviation. In turn, the choice variable $w(c)$ corresponds to the space devoted to the weather event in the newscast controlling for the number of segments in a broadcast (as per our normalization that total time available is equal to one).

3.8.2 Demand-driven Coverage

We start by analyzing demand-driven coverage in our model. In other words, we assume that the TV station is profit-maximizing. As the entertainment value of weather event is always valued more by possible viewers than learning (per assumptions), a TV station has never any interest in not covering weather event (under the assumption that $g'(1) = 0$ and $h(w) = w$). Hence, we can restrict attention to the case in which the outlet chooses $w(c) > 0$ for all c.

To then determine the editorial strategy, we first need to study what a viewer can learn about the state $\omega \in \{0, 1\}$ upon observing c. The viewer forms a posterior:

$$\mu(c) = \frac{1}{1 + \frac{1-\pi}{\pi}\frac{f_0(c)}{f_1(c)}} \tag{3.8}$$

Under the assumption of MLRP, the posterior is strictly increasing with c. Note that under our assumptions for all $c < 0$, $\mu(c) < \pi$, whereas $\mu(c) > \pi$ for all $c > 0$.

Given an editorial strategy $(w(c))_{\{c\in[-1,1]\}}$, citizens can compute their expected utility from consuming the outlet's news combining Equation 3.5 and Equation 3.8 for each weather event c. Citizens, however, do not know the value of the weather event that will be reported if they watch the outlet. They need to take into account the expected distribution of weather events, which we denote by $F^e(\tilde{c}) = \pi F_1(\tilde{c}) + (1 - \pi)F_0(\tilde{c})$ in what follows, to compute their expected utility from turning on the news. For a citizen from group $J \in \{D, R\}$, we obtain:

$$\int_0^1 w(\tilde{c})\big(v(\tilde{c}) + z_J(\mu(\tilde{c}))\big) + g(1 - w(\tilde{c}))\underline{u}\, dF^e(\tilde{c})$$

Given the outside option of citizen i ($\delta_i \sim \mathcal{U}[0, \overline{\delta}]$) and the assumption that this event is i.i.d. and we have a mass of citizens in each group, the *proportion* of individuals from group $J \in \{D, R\}$ who watch the outlet given its editorial strategy is: $A_J = \frac{1}{\delta} \int_0^1 w(\tilde{c})\big(v(\tilde{c}) + z_J(\mu(\tilde{c}))\big) + g(1 - w(\tilde{c}))\underline{u}\, dF^e(\tilde{c})$.

From this, we can easily define the total audience as a function of the outlet's editorial strategy. It is simply:

$$A = \frac{1}{\overline{\delta}} \int_0^1 w(\tilde{c})\big(v(\tilde{c}) + \alpha z_D(\mu(\tilde{c})) + (1 - \alpha) z_R(\mu(\tilde{c}))\big) + g(1 - w(\tilde{c}))\underline{u}\, dF^e(\tilde{c})$$

As the outlet seeks to maximize its audience size, it is immediate that its strategy corresponds to maximizing "point-by-point" the utility of the "average" citizen. We obtain that the equilibrium publication strategy satisfies (with superscript dem for demand-driven coverage):

Proposition 4 *A TV stations's editorial strategy is a function defined for all $c \in [-1, 1]$ by:*

$$g'(1 - w^{dem}(c; \alpha)) = \frac{v(c) + \alpha z_D(\mu(c)) + (1 - \alpha) z_R(\mu(c))}{\underline{u}} \tag{3.9}$$

Proof: The proof of Proposition 4 and all results can be found in Online Appendix 3.H.

The next result provides some simple comparative statics on the coverage of weather event.

Corollary 1 *The coverage of weather event $w^{dem}(c; \alpha)$:*

- *is strictly increasing with α,*
- *is strictly decreasing with c for $c < 0$ and strictly increasing with c for $c > 0$.*

This first comparative statics matches the basic descriptive for coverage collected in Appendix Figure 3.A.2. With this, we now turn to studying how the leaning of the media market affects the coverage of extreme weather events relative to events consistent with the seasonal norm ($c = 0$ in our model). To do so, we define the theoretical equivalent of our empirical estimate as: $\Delta^{dem}(c; \alpha) = w^{dem}(c; \alpha) - w^{dem}(0; \alpha)$. We obtain:

Proposition 5 *Suppose* $\alpha^D > \alpha^R$*. There exists* $\overline{g} > 0$ *such that if* $g'''(c) \leq \overline{g}$ *for all* $c \in [-1, 1]$*, then*

- $\Delta^{dem}(c; \alpha)$ *is strictly decreasing with* c *for* $c < 0$ *and strictly increasing in* c *for* $c > 0$*,*
- $\Delta^{dem}(c; \alpha^D) > (<) \Delta^{dem}(c; \alpha^R)$ *for all* $c > (<) 0$*,*
- $\dfrac{\partial \Delta^{dem}(c; \alpha^D) - \Delta^{dem}(c; \alpha^R)}{\partial c} > 0$ *for all* c*.*

Proposition 5 makes three points. The first states that the difference in coverage increases with the extremeness of weather events. The second highlights that the increase in coverage is greater for TV stations in Democratic-dominated media markets than in Republican markets for weather event above the seasonal norm, whereas the opposite holds true for events below the seasonal norm. Finally, the last result states that as we go from weather events below the norm to events below the norms, the difference-in-differences in coverage (i.e., the difference between stations in Democratic market versus Republican market on top of the difference between baseline events and other events) is increasing.

While the first result follows directly from Corollary 1, as the weather event becomes more extreme (away from zero), coverage increases relative to the baseline, the last two do not. Indeed, coverage increases in all types of markets, dominated by Republicans or Democrats. Hence, we need an additional condition so that the rate of increase is faster in Republican-dominated markets for weather events below the norm ($c < 0$) and in Democrat-dominated markets for weather events above the norm ($c > 0$). This is what the condition on the third derivative of $g(\cdot)$ in the proposition guarantees. Notice that the learning part in the citizens' payoff matters. Absent this, to recover the same pattern, one would need to assume that Republican citizens have a higher (lower) entertainment value for news about cold (hot) weather than Democratic citizens and this difference is increasing as events become more unlikely. A utility function that includes benefits and losses from learning reproduces the same type of variations in a simple and (we hope) elegant way.

The theoretical results from Proposition 5 fits our empirical findings on presentation bias in summer (Figure 3.4 panel (a)). They also correspond to the patterns we observe for weather events below the mean in winter (Figure 3.4 panel (b)). The main difference between our theoretical and empirical results regard above the norms weather event in winter. While we do not have a robust theoretical explanation for this, our best guess, as already noted above, is that the entertainment value of warmer days than the norms in winter is low since those days are cold nonetheless. Despite the failure to parallel all of our empirical results, a demand-driven model with learning performs relatively well. Can a supply-driven model of coverage outperform it? We turn to this question in the next subsection.

3.8.3 Supply Driven Coverage

We now consider how coverage would look like if the media outlet's objective was to move individuals' beliefs in its preferred direction. In general, it proves very difficult to define the optimal strategy of a biased outlet as there an infinite number of possible combinations. We can, however, provide a basic idea of what coverage looks like. As Proposition 6 highlights, a biased TV station either fully omits weather events or only covers them. Denoting $w_\tau^{sup}(\cdot)$ the equilibrium editorial strategy for a biased outlet of type $\tau \in \{b, s\}$ (with the superscript sup standing for supply-driven coverage), we obtain:

Proposition 6 *For all $c \in [-1, 1]$, the equilibrium coverage of weather event of a biased outlet of type $\tau \in \{b, s\}$ satisfies: $w_\tau^{sup}(c; \alpha) \in \{0, 1\}$ for all $\alpha \in [0, 1]$*

To understand the result in Proposition 6, recall that a biased media outlet has two considerations in mind when it determines its editorial strategy. The first is whether to reveal the weather events to manipulate beliefs (i.e., to choose $w(c) > 0$ or $w(c) = 0$). The second is to design coverage so as to manipulate its audience size, who observe the weather news. Because the objective function of the TV station is strictly concave, the outlet has more to lose from citizens observing bad news than from citizens observing good news (e.g., for a climate-believer outlet, the gain from moving beliefs up by a certain amount is lower than the loss from moving beliefs down by the same amount). As such, the media outlet never seeks to maximize its audience. If it reveals the weather event (if $w(c) > 0$), it chooses a suboptimal coverage and the least optimal coverage (given our assumption on the citizens' utility) is $w(c) = 1$.[18]

Building on Proposition 6, we obtain that as weather events become more extreme, coverage either remain constant or increase abruptly. These variations are starker than those we theoretically obtain for profit-maximizing outlets and they are a far cry from the smooth increase we empirically document. Denoting $\Delta_\tau^{sup}(c; \alpha) = w_\tau^{sup}(c; \alpha) - w_\tau^{sup}(0; \alpha)$, we get:

[18]The result requires that outlets have concave payoff function, that they are risk avoiding. If their utility is convex (they are risk seeking), then they can choose some interior amount of coverage. Indeed, they will choose the demand-driven level of coverage $w^{dem}(c; \alpha)$. This would imply that demand forces also shape the coverage of biased TV stations.

Corollary 2 *For all* $c \in [-1, 1]$*, all* $\alpha \in [0, 1]$*, and* $\tau \in \{b, s\}$*,* $\Delta_{\tau}^{sup}(c; \alpha) \in \{-1, 0, 1\}$*.*

Corollary 2 indicates that fixing all elements of our model, but the objective of the media outlet, supply-driven coverage fails to match our empirical patterns whereas demand-driven coverage works relatively well. A further implication of Corollary 2 is that if coverage is supply-driven (by TV stations who are biased), we should observe large swings in coverage when the ideology of the outlet changes (note that this is a reasonable conjecture, but we cannot prove this result as we cannot fully define an outlet's strategy). In turn, by definition, if the coverage is demand-driven, we should observe relatively little change in they weather events are covered. This last theoretical prediction is again relatively conform to our findings in Table 3.2 when we look at the effect of TV station acquisitions by Sinclair. Overall, our parsimonious set-up provides a way to rationalize (most of) our empirical findings and to justify our claim that those results come from demand-side forces rather than supply-side ones.

3.9 Conclusion

Our paper explores the daily decision of what to cover in the news by local TV stations in the United States. To do so, we look at weather events, which we define as temperature deviations from the historical mean. We document a clear "man bites dog" effect: Outlets cover significantly more severe (uncommon) weather events than typical ones. Interestingly, even moderate deviations can sometimes be deemed newsworthy, offering deeper insight into how news is produced. Beyond this, we document striking publication bias in the coverage of even seemingly ubiquitous events like the weather. In summer, Republican-leaning markets downplay above-average temperatures and expand coverage of below-average temperatures, whereas Democratic-leaning markets do the opposite, amplifying hot weather and minimizing cold spells. Crucially, our evidence suggests this bias is not driven by supply-side persuasion efforts but rather by audience demand. Using a stylized theoretical model, we demonstrate how media outlets tailor coverage to viewers' preferences. Additionally, we show that the salience of weather news follows the same pattern as publication bias. However, we find little evidence of systematic differences in how the these news are framed.

How are we to understand this last result? One possibility is that temperature deviations are easy to interpret in the context of the climate crisis. After all, climate change has been understood as global warming and unseasonably hot and cold days can have clear meaning in this context. Others, studying different media outlets, have robustly shown presentation bias in the coverage of disasters (Djourelova et al., 2024). If we combine the easiness of interpretation of temperature deviations and the complementarity with the coverage of disasters, one can maybe see why narratives supplied by the media (a la Eliaz and Spiegler, 2024) may well be unnecessary for the weather events we study. Our focus on the editors' choices of what to cover every day makes it impossible for us to test the consequences of the publication bias we uncover. Still, we believe that our findings paint a striking picture. Our paper suggests that publication bias is much more extensive than previously though, applying even to common place events like weather events. Doing so, little by little, day by day, without the need for sensational events like disasters, the media may shape divergent views on the existence and the causes of climate changes.

Appendix

Appendix 3.A Appendix Figures

Figure 3.A.1: Coverage of Local Weather

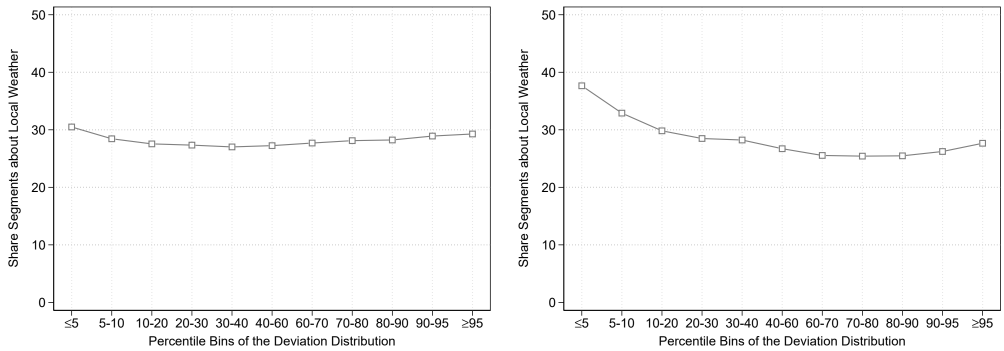

(a) Summer (b) Winter

Notes: This figure shows coverage of local weather in different weather events. In particular, we show the average number of segments about local weather in each bin of the within-media market deviation distribution, separately for summer (panel (a)) and for winter (panel (b)).

Figure 3.A.2: Coverage of Local Weather by DMA Political Leaning

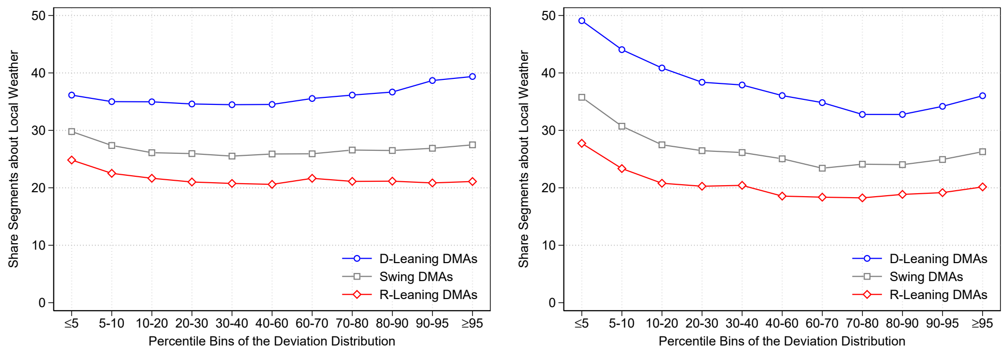

(a) Summer (b) Winter

Notes: This figure shows coverage of local weather in different weather events, by the political leaning of the media market. In particular, we show the average number of segments about local weather in each bin of the within-media market deviation distribution for media markets that are Democratic-leaning, swing, and Republican-leaning, separately for summer (panel (a)) and for winter (panel (b)).

Figure 3.A.3: Effect of Sinclair Acquisitions on Coverage of Weather Events, Event Studies for Differences-in-Differences Specification

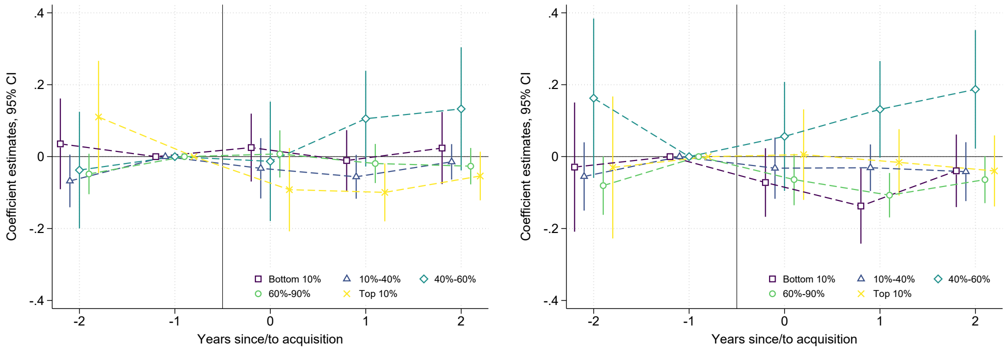

(a) Summer (b) Winter

Notes: This figure shows the dynamic effect of Sinclair acquisitions on the reporting of different weather events. We regress the number of segments about local weather on an indicators for years to/since the acquisition, indicator variables for the deviation from historical mean falling in a given bin of the within-media market deviation distribution, one lead and one lag of the same indicators, the interaction of the Sinclair and weather events indicators, one lead and one lag of these interactions, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.4). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (panel (a)) and for winter (panel (b)), using a Poisson model. Standard errors are clustered at the media market level.

Figure 3.A.4: Effect of Sinclair Acquisitions on Coverage of Weather Events, Event Studies for Triple Differences Specification

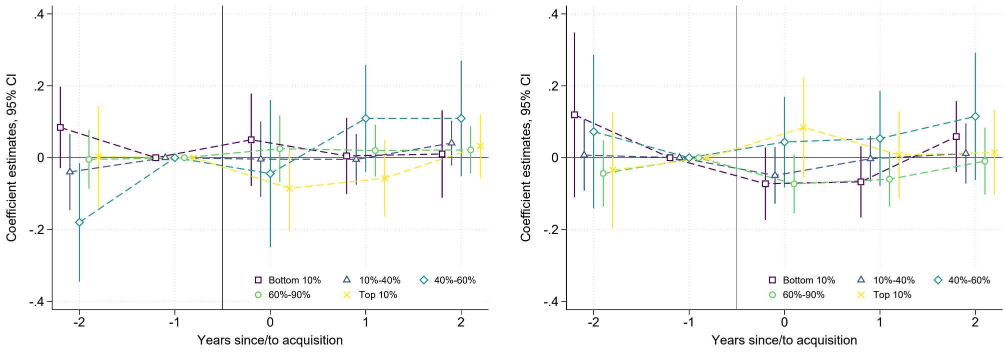

(a) Summer (b) Winter

Notes: This figure shows the dynamic effect of Sinclair acquisitions on the reporting of different weather events. We regress the number of segments about local weather on an indicators for years to/since the acquisition, indicator variables for the deviation from historical mean falling in a given bin of the within-media market deviation distribution, one lead and one lag of the same indicators, the interaction of the Sinclair and weather events indicators, one lead and one lag of these interactions, station fixed effects, media market by day fixed effects, and number of segments decile fixed effects (see Equation 3.4). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (panel (a)) and for winter (panel (b)), using a Poisson model. Standard errors are clustered at the media market level.

Figure 3.A.5: CCES Descriptive

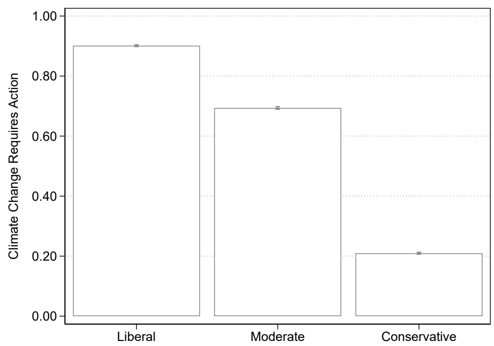

(a) By Ideology

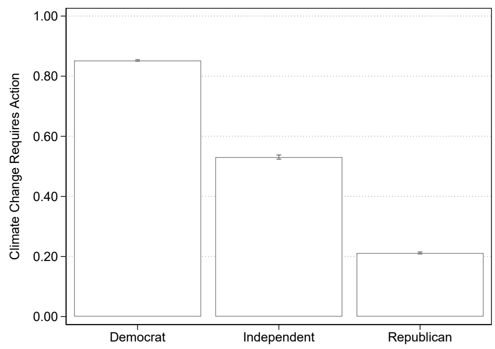

(b) By Party

Notes: This figure shows the mean share of respondents in CCES stating that climate change requires action by ideology (panel (a)) and by party identity (panel (b)).

Appendix 3.B Appendix Tables

Table 3.B.1: Salience Bias

	Local Weather Segments, Rank					
	Summer			Winter		
	(1)	(2)	(3)	(4)	(5)	(6)
Bottom 10%	-0.043***	-0.042***	-0.022***	-0.104***	-0.096***	-0.058***
	(0.007)	(0.009)	(0.008)	(0.009)	(0.012)	(0.010)
Bottom 10% × D-Leaning		0.012	0.008		-0.007	-0.017
		(0.015)	(0.013)		(0.019)	(0.017)
Bottom 10% × R-Leaning		-0.028	-0.016		-0.032	-0.018
		(0.017)	(0.016)		(0.019)	(0.017)
10%-40%	0.001	0.002	0.001	-0.025***	-0.020***	-0.013**
	(0.004)	(0.006)	(0.005)	(0.005)	(0.007)	(0.005)
10%-40% × D-Leaning		0.007	0.005		-0.002	-0.005
		(0.009)	(0.008)		(0.011)	(0.009)
10%-40% × R-Leaning		-0.014	-0.014		-0.017	-0.018
		(0.013)	(0.012)		(0.013)	(0.011)
60%-90%	-0.018***	-0.009*	-0.004	0.009**	0.010	0.006
	(0.004)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)
60%-90% × D-Leaning		-0.022**	-0.018**		0.012	0.013
		(0.009)	(0.008)		(0.010)	(0.009)
60%-90% × R-Leaning		-0.010	-0.013		-0.024**	-0.017
		(0.010)	(0.009)		(0.011)	(0.011)
Top 10%	-0.045***	-0.034***	-0.017**	-0.006	0.003	0.005
	(0.007)	(0.009)	(0.008)	(0.007)	(0.010)	(0.008)
Top 10% × D-Leaning		-0.049***	-0.040***		-0.017	-0.012
		(0.015)	(0.014)		(0.014)	(0.013)
Top 10% × R-Leaning		0.018	0.008		-0.020	-0.018
		(0.017)	(0.015)		(0.015)	(0.015)
Observations	284444	284444	284444	274000	274000	274000
Stations	697	697	697	699	699	699
DMAs (Clusters)	204	204	204	204	204	204
Mean Dep. Variable	1.966	1.966	1.966	1.896	1.896	1.896
Bottom 10% × D = Bottom 10% × R		0.035	0.175		0.248	0.950
10%-40% × D = 10%-40% × R		0.132	0.121		0.284	0.296
60%-90% × D = 60%-90% × R		0.275	0.618		0.005	0.015
Top 10% × D = Top 10% × R		0.001	0.006		0.864	0.677

Notes: This table shows the relationship between the salience of news coverage of local weather and weather events, on average and by media market political leaning. In columns (1) and (4), we regress the log minimum rank of segments about local weather (averaged across the different newscasts) on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, one lead and one lag of the same indicators, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.1). Columns (2) and (5) additionally include the weather events indicators interacted with interacted with dummies for the market being either Democratic- or Republican-leaning, with one lead and one lag of the same variables (see Equation 3.2). Columns (3) and (6), additionally control for number of local weather segments deciles. The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (columns (1) to (3)) and for winter (columns (4) to (6)), using OLS. Standard errors are clustered at the media market level.

Table 3.B.2: Publication Bias, CCES

	Local Weather Segments	
	Summer	Winter
	(1)	(2)
Bottom 10%	0.070***	0.109***
	(0.012)	(0.011)
Bottom 10% × Non-Sceptic	-0.031**	-0.040**
	(0.015)	(0.019)
Bottom 10% × Sceptic	-0.011	0.035
	(0.017)	(0.024)
10%-40%	0.005	0.022***
	(0.004)	(0.006)
10%-40% × Non-Sceptic	-0.006	0.000
	(0.007)	(0.011)
10%-40% × Sceptic	-0.009	0.016
	(0.009)	(0.011)
60%-90%	0.014***	-0.015***
	(0.004)	(0.005)
60%-90% × Non-Sceptic	0.013**	-0.003
	(0.006)	(0.009)
60%-90% × Sceptic	-0.002	0.030***
	(0.008)	(0.008)
Top 10%	0.028***	0.001
	(0.008)	(0.007)
Top 10% × Non-Sceptic	0.045***	-0.001
	(0.011)	(0.012)
Top 10% × Sceptic	0.007	0.033***
	(0.013)	(0.012)
Observations	297659	288726
Stations	698	699
DMAs (Clusters)	204	204
Mean Dep. Variable	27.905	27.915
Bottom 10% × D = Bottom 10% × R	0.205	0.005
10%-40% × D = 10%-40% × R	0.756	0.269
60%-90% × D = 60%-90% × R	0.109	0.001
Top 10% × D = Top 10% × R	0.005	0.019

Notes: This table shows the relationship between news coverage of local weather and weather events by media market climate change scepticism. We regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, the same indicators interacted with dummies for the market being either Democratic- or Republican-leaning, one lead and one lag of the same indicators, station fixed effects, day fixed effects, and number of segments decile fixed effects (similar to Equation 3.1). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (column (1)) and for winter (column (2)), using a Poisson model. Standard errors are clustered at the media market level.

Appendix 3.C Classification

We use a simple, but effective, dictionary method to identify whether a segment is about the weather. In particular, as we explain in the data section, we define a segment to be about weather if it contains at least *two* words from a dictionary of weather-related terms. We then define a segment to about local weather, if it additionally mentions at least one county or one municipality located in the media market the station operates in.

The dictionary we follow is from Baylis et al. (2019), who use to identify Tweets that are about weather, and contains the following terms:

aerovane air airstream altocumulus altostratus anemometer anemometers anticyclone anticyclones arctic arid aridity atmosphere atmospheric autumn autumnal balmy baroclinic barometer barometers barometric blizzard blizzards blustering blustery blustery breeze breezes breezy brisk calm celsius chill chilled chillier chilliest chilly chinook cirrocumulus cirrostratus cirrus climate climates cloud cloudburst cloudbursts cloudier cloudiest clouds cloudy cold colder coldest condensation contrail contrails cool cooled cooling cools cumulonimbus cumulus cyclone cyclones damp damp damper damper dampest dampest degree degrees deluge dew dews dewy doppler downburst downbursts downdraft downdrafts downpour downpours dried drier dries driest drizzle drizzled drizzles drizzly drought droughts dry dryline fall farenheit flood flooded flooding floods flurries flurry fog fogbow fogbows fogged fogging foggy fogs forecast forecasted forecasting forecasts freeze freezes freezing frigid frost frostier frostiest frosts frosty froze frozen gale gales galoshes gust gusting gusts gusty haboob haboobs hail hailed hailing hails haze hazes hazy heat heated heating heats hoarfrost hot hotter hottest humid humidity hurricane hurricanes ice iced ices icing icy inclement landspout landspouts lightning lightnings macroburst macrobursts maelstrom mercury meteorologic meteorologist meteorologists meteorology microburst microbursts microclimate microclimates millibar millibars mist misted mists misty moist moisture monsoon monsoons mugginess muggy nexrad nippy NOAA nor’easter nor’easters noreaster noreasters overcast ozone parched parching pollen precipitate precipitated precipitates precipitating precipitation psychrometer radar rain rainboots rainbow rainbows raincoat raincoats rained rainfall rainier rainiest raining rains rainy sandstorm sandstorms scorcher scorching searing shower showering showers skiff sleet slicker slickers slush slushy smog smoggier smoggiest smoggy snow snowed snowier snowiest snowing snowmageddon snowpocalypse snows snowy spring sprinkle sprinkles sprinkling squall squalls squally storm stormed stormier stormiest storming storms stormy stratocumulus stratus subtropical summer summery sun sunnier sunniest sunny temperate temperature tempest thaw thawed thawing thaws thermometer thunder thundered thundering thunders thunderstorm thunderstorms tornadic tornado tornadoes tropical troposphere tsunami turbulent twister twisters

Table 3.C.1: Weather classification performance measures

	One weather-related term	Two weather-related terms
	(1)	(2)
False positive rate	0.265	0.056
False negative rate	0.007	0.044
Precision	0.523	0.833
Recall	0.993	0.956
F-score	0.685	0.890
Accuracy	0.793	0.947

Notes: This table shows classification performance metrics for identifying weather-related segments using a keyword-based approach. Column (1) shows results when classifying segments containing at least one weather-related term, while column (2) restricts to segments containing at least two such terms. We report the false positive rate, false negative rate, precision, recall, F-score, and overall accuracy.

typhoon typhoons umbrella umbrellas vane warm warmed warming warms warmth waterspout waterspouts weather wet wetter wettest wind windchill windchills windier windiest windspeed windy winter wintery wintry

We explore whether our simple classification method is effective at identifying weather stories by implementing the following validation exercise. First, we randomly select 600 segments stratifying by political leaning of the DMA (in other words, we randomly 200 segments from Democratic-leaning DMAs, 200 segments from swing DMAs, and 200 segments from Republican-leaning DMAs). Second, we hand-classify these segments as being about weather or not through close reading. Third, we use our hand-classification as the "ground truth" and test the performance of our dictionary method against it.

Appendix Table 3.C.1 reports different measures of predictive performance using either one or two weather-related terms to identify weather segments. This exercise makes clear that using one weather-related term only captures virtually all weather-related segments, but yields a relatively high number of false positives. To the extent that several of the weather-related terms are polysemic, this is not surprising: think of "hot" or even "ice." Instead, requiring *two* weather words (rather than just one) to appear in a segment for it to be classified as weather-related greatly minimize this type of error, at the cost of a negligible increase in false negatives. Overall, the dictionary-based method we appears to have a good performance in this setting.

Our method to classify stories as local, instead, might yield both false positives (for example, if a county/municipality has a polysemic or common name) and false negatives (for example, if neighborhoods are mentioned instead of municipalities). However, this is only problematic to the extent that this type of measurement error varies depending on the size of weather events, which is not likely to be the case.

Finally, we might also be concerned that requiring both two weather terms and a county or municipality name to appear in the same 150-word segment might lead to further under-estimating coverage of local weather. The fact that our results are robust to using longer segments assuages this concern.

Appendix 3.D Additional Data Sources

Climate Regions. We follow the definition of climate regions from the National Centers for Environmental Information. We assign each DMA to the climate region that covers the majority of its area.

Additional DMA characteristics. Media market characteristics (population, population in urban versus rural areas, and population employed in different industries) are from NHGIS. In all cases, we start from county level data and aggregate them to the media market level. The area of each media market is from our own calculations, based on a shapefile of media markets.

Wildfires. Wildfire data is from the National Interagency Fire Center. The data reports the date of discovery and geographic coordinates for each wildfire event identified by the agency. We match each wildfire event to the relevant DMA using these geographic coordinates.

Natural Disasters. Information on natural disasters is from Federal Emergency Management Agency's (FEMA) Disaster Declarations Summaries dataset. We exclude incidents categorized as "fire" to avoid duplicates with wildfires. Natural disasters geographic information at the county level is then aggregated to the media market level.

Appendix 3.E Additional Analyses

In this Appendix we show that our results are robust to a number of concerns.

Heterogeneity in Shocks. A first concern is that the different responsiveness to weather events of stations that operate in Republican- and Democratic-leaning media markets might be explained by these markets experiencing substantively different events. This is possible in this case because we define weather events based on where a given day falls in the within-media market deviation distribution. If this distribution is systematically different in Republican- and Democratic-leaning media markets, then we might be treating as equivalent events that correspond to very different deviations.

However, this is unlikely to be the case in our setting. To start with, Figure 3.2 shows that the distribution of deviations within each weather event in the three types of media market is substantially overlapping, both in summer (panel (a)) and in winter (panel (b)). Nonetheless, it is also possible to see that the median deviation is at times slightly different across the different types of media markets.

In Appendix Table 3.E.1, we show that, even if we were to take these differences at face value, they would be too small to explain the effects we estimate through a simple back of the envelope calculation. We begin by computing the median deviation that correspond to each of the weather events we consider in the analysis. We report this in columns (1) for summer and column (4) for winter. From our usual specification, we can recover the percentage increase in coverage of local weather corresponding to each of these weather events. We use swing-DMAs as our benchmark. These are the estimates in columns (2) and (4). Finally, under a linearity assumption, we can estimate the predicted differential change in coverage that one would expect given the difference in median deviation in D-leaning and R-leaning markets for each event. The estimates reported in columns (3) and (6) are of one order of magnitude smaller than the ones we estimate in Table 1. This suggests that any differences in deviations across DMAs with different political leanings are unlikely to explain the patterns we estimate.

Finally, in line with these results, we can also show that the patterns we estimate are robust to using within-climate region variation only. Appendix Table 3.E.2 shows that our results are fully robust to estimating specifications that include climate region-by-year fixed effects (columns (1) and (2)).

DMA Characteristics. A second concern is that our publication bias estimates are measuring the effect of media market's characteristics that are different than ideology, and just happen to correlate with it. We explore this concern in Appendix Table 3.E.3, where we re-estimate our baseline specification but include different media market characteristics also interacted with the indicators for the weather events. The table clearly shows that our results are robust to controlling for media market population, area, level of urbanization, and industry composition. It is particularly interesting that most of our heterogeneity survives even a very stringent specification that includes all controls at the same time.

Wildfires and Natural Disasters. Third, one might be concerned that our extreme weather events (those in the top and bottom 10% of the deviation distribution, that imply the largest increases in coverage of local weather) are simply proxies for wildfires or other natural disasters. To test whether this is the case, we once again re-estimate our specification, but including interactions between whether the DMA is experiencing a wildfires or a natural disaster and DMA ideology. Appendix Table 3.E.4 shows that our results are virtually unchanged when including these controls, separately or together.

Robustness Checks. Finally, we also show that our results are robust to several perturbations to our measurement and modeling choices. In particular, in Appendix Table 3.E.5 we show that our findings for the summer are fully robust to using 300 word long segments (column (1)), estimating OLS specifications with the outcome expressed in logs (column (2)) or shares (column (3)), and splitting markets by ideology using 20th and 80th percentiles (column (4)) and the 30th and 70th percentiles (column (5)) as cutoffs. The heterogeneity in reporting in days in the bottom 10% of the deviation distribution in winter is instead not robust to defining the outcome as a share (while all other alternative specifications yield the same results).

Table 3.E.1: Back-of-the-Envelope Calculation

		Winter			Summer		
		Median Deviation	% Change in Swing DMA	Predicted % Change	Median Deviation	% Change in Swing DMA	Predicted % Change
		(1)	(2)	(3)	(4)	(5)	(6)
Bottom 10%	D-Leaning	-5.359		-0.004	-10.955		-0.004
	Swing	-5.707	0.061		-11.423	0.096	
	R-Leaning	-5.849		0.002	-11.893		0.004
10%-40%	D-Leaning	-1.947		0.000	-4.052		-0.002
	Swing	-2.059	-0.001		-4.443	0.018	
	R-Leaning	-2.021		0.000	-4.450		0.000
60%-90%	D-Leaning	2.218		0.002	3.942		0.002
	Swing	1.944	0.012		4.622	-0.015	
	R-Leaning	1.764		-0.001	4.927		-0.001
Top 10%	D-Leaning	5.336		0.004	10.023		
	Swing	4.789	0.032		10.271	0.003	
	R-Leaning	4.525		-0.002	10.440		0.000

Notes: This table performs a back-of-the-envelope calculation to understand whether the effects we estimate can be explained by differences in the median temperature deviation from the historical mean across media markets with different political leanings. Column (1) shows the median temperature deviation corresponding to different weather events by media market political leaning. Column (2) reports the increase in media coverage of local weather corresponding to a given weather event in swing media markets, that we estimate using Equation 3.2. Column (3) reports the predicted difference in reporting that we should expect in media markets with different political leanings based on (1) and (2).

Table 3.E.2: Publication Bias, Controls for Climate Region-Specific Shocks

	Local Weather Segments	
	Summer	Winter
	(1)	(2)
Bottom 10%	0.060***	0.096***
	(0.012)	(0.012)
Bottom 10% × D-Leaning	-0.015	-0.014
	(0.016)	(0.019)
Bottom 10% × R-Leaning	0.016	0.068***
	(0.017)	(0.020)
10%-40%	-0.001	0.016***
	(0.004)	(0.006)
10%-40% × D-Leaning	0.004	0.009
	(0.007)	(0.010)
10%-40% × R-Leaning	0.007	0.022**
	(0.008)	(0.010)
60%-90%	0.012***	-0.013**
	(0.004)	(0.005)
60%-90% × D-Leaning	0.017**	-0.002
	(0.007)	(0.009)
60%-90% × R-Leaning	-0.002	0.030***
	(0.006)	(0.008)
Top 10%	0.030***	0.005
	(0.008)	(0.007)
Top 10% × D-Leaning	0.034***	-0.003
	(0.012)	(0.011)
Top 10% × R-Leaning	-0.004	0.022
	(0.011)	(0.013)
Observations	297659	288726
Stations	698	699
DMAs (Clusters)	204	204
Mean Dep. Variable	27.905	27.915
Bottom 10% × D = Bottom 10% × R	0.048	0.000
10%-40% × D = 10%-40% × R	0.711	0.291
60%-90% × D = 60%-90% × R	0.011	0.001
Top 10% × D = Top 10% × R	0.003	0.111

Notes: This table shows the relationship between news coverage of local weather and weather events by media market political leaning, using only within climate region and year variation. We regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, the same indicators interacted with indicators for the market being either Democratic- or Republican-leaning, one lead and one lag of the same variables, climate region by year fixed effects, station fixed effects, day fixed effects, and number of segments decile fixed effects (similar to Equation 3.2). The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (column (1)) and for winter (column (2)), using a Poisson model. Standard errors are clustered at the media market level.

Table 3.E.3: Publication Bias, Controls for Media Market Characteristics

	Local Weather Segments									
	Summer					Winter				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Bottom 10%	0.058***	0.058***	0.057***	0.051***	0.049***	0.095***	0.097***	0.097***	0.089***	0.087***
	(0.011)	(0.011)	(0.010)	(0.008)	(0.008)	(0.012)	(0.012)	(0.012)	(0.010)	(0.010)
Bottom 10% × D-Leaning	-0.011	-0.037*	-0.033*	-0.001	-0.008	-0.008	-0.014	-0.013	0.005	0.012
	(0.015)	(0.020)	(0.017)	(0.013)	(0.013)	(0.019)	(0.020)	(0.019)	(0.018)	(0.017)
Bottom 10% × R-Leaning	0.013	0.034**	0.026	0.021	0.023	0.063***	0.071***	0.070***	0.075***	0.073***
	(0.019)	(0.016)	(0.018)	(0.015)	(0.015)	(0.021)	(0.021)	(0.020)	(0.022)	(0.022)
10%-40%	-0.003	-0.001	-0.001	-0.002	-0.005	0.019***	0.018***	0.018***	0.018***	0.019***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
10%-40% × D-Leaning	0.007	0.006	0.003	0.008	0.011	0.009	0.008	0.009	-0.005	-0.001
	(0.006)	(0.008)	(0.007)	(0.008)	(0.007)	(0.010)	(0.010)	(0.010)	(0.011)	(0.010)
10%-40% × R-Leaning	0.003	0.006	0.008	0.005	0.003	0.021*	0.022**	0.021*	0.031***	0.031***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.011)	(0.011)	(0.011)	(0.012)	(0.012)
60%-90%	0.011**	0.012***	0.012***	0.013***	0.011**	-0.017***	-0.014***	-0.014***	-0.016***	-0.018***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
60%-90% × D-Leaning	0.019***	0.023***	0.021***	0.012*	0.018**	0.000	0.008	0.007	0.004	0.011
	(0.007)	(0.008)	(0.007)	(0.007)	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
60%-90% × R-Leaning	-0.001	-0.003	-0.001	0.001	-0.001	0.028***	0.025***	0.027***	0.028***	0.026***
	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)
Top 10%	0.032***	0.031***	0.032***	0.034***	0.034***	0.001	0.003	0.003	0.000	-0.001
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Top 10% × D-Leaning	0.035***	0.027**	0.035***	0.035***	0.029**	0.002	-0.002	0.000	0.001	0.002
	(0.012)	(0.013)	(0.012)	(0.014)	(0.013)	(0.011)	(0.013)	(0.012)	(0.012)	(0.012)
Top 10% × R-Leaning	-0.005	-0.001	-0.005	-0.009	-0.004	0.024*	0.024*	0.022*	0.024	0.024*
	(0.012)	(0.011)	(0.011)	(0.012)	(0.013)	(0.013)	(0.014)	(0.013)	(0.015)	(0.015)
Log Area	✓				✓	✓				✓
Log Population		✓			✓		✓			✓
Urban Share			✓		✓			✓		✓
Industry Shares				✓	✓				✓	✓
Observations	297659	297659	297659	297659	297659	288726	288726	288726	288726	288726
Stations	698	698	698	698	698	699	699	699	699	699
DMAs (Clusters)	204	204	204	204	204	204	204	204	204	204
Mean Dep. Variable	27.905	27.905	27.905	27.905	27.905	27.915	27.915	27.915	27.915	27.915
Bottom 10% × D = Bottom 10% × R	0.170	0.002	0.002	0.223	0.101	0.003	0.001	0.000	0.009	0.022
10%-40% × D = 10%-40% × R	0.728	0.995	0.639	0.808	0.409	0.329	0.276	0.373	0.012	0.028
60%-90% × D = 60%-90% × R	0.016	0.007	0.008	0.218	0.081	0.005	0.122	0.033	0.035	0.208
Top 10% × D = Top 10% × R	0.003	0.074	0.002	0.007	0.049	0.167	0.126	0.147	0.208	0.219

Notes: This table shows the relationship between news coverage of local weather and weather events by media market political leaning, controlling for additional media market characteristics. We regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, the same indicators interacted with indicators for the market being either Democratic- or Republican-leaning, one lead and one lag of the same variables, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.2). The regression further includes the interaction between indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution and the control variable(s) noted in the respective column. The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (columns (1) to (5)) and for winter (columns (6) to (10)), using a Poisson model. Standard errors are clustered at the media market level.

Table 3.E.4: Publication Bias, Controls for Wildfires and Natural Disasters

	Local Weather Segments					
	Summer			Winter		
	(1)	(2)	(3)	(4)	(5)	(6)
Bottom 10%	0.061***	0.058***	0.058***	0.096***	0.095***	0.095***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.011)	(0.011)
Bottom 10% × D-Leaning	-0.015	-0.012	-0.012	-0.012	-0.010	-0.010
	(0.016)	(0.015)	(0.015)	(0.019)	(0.019)	(0.019)
Bottom 10% × R-Leaning	0.018	0.017	0.017	0.069***	0.068***	0.068***
	(0.018)	(0.017)	(0.017)	(0.020)	(0.020)	(0.020)
10%-40%	-0.001	-0.001	-0.001	0.018***	0.018***	0.018***
	(0.004)	(0.004)	(0.004)	(0.006)	(0.006)	(0.006)
10%-40% × D-Leaning	0.003	0.006	0.005	0.009	0.008	0.008
	(0.007)	(0.007)	(0.007)	(0.010)	(0.010)	(0.010)
10%-40% × R-Leaning	0.008	0.010	0.010	0.021**	0.018*	0.018*
	(0.008)	(0.008)	(0.008)	(0.011)	(0.010)	(0.011)
60%-90%	0.012***	0.011**	0.011**	-0.015***	-0.014***	-0.014***
	(0.004)	(0.004)	(0.004)	(0.006)	(0.005)	(0.005)
60%-90% × D-Leaning	0.017**	0.019***	0.018***	-0.001	-0.001	-0.001
	(0.007)	(0.007)	(0.007)	(0.009)	(0.009)	(0.009)
60%-90% × R-Leaning	0.000	0.001	0.001	0.031***	0.029***	0.029***
	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)
Top 10%	0.032***	0.032***	0.032***	0.003	0.004	0.004
	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)
Top 10% × D-Leaning	0.035***	0.037***	0.038***	0.002	-0.000	0.000
	(0.012)	(0.012)	(0.012)	(0.011)	(0.011)	(0.011)
Top 10% × R-Leaning	-0.005	-0.003	-0.003	0.023*	0.021	0.020
	(0.011)	(0.011)	(0.011)	(0.013)	(0.013)	(0.013)
Wildfires	✓		✓	✓		✓
Disasters		✓	✓		✓	✓
Observations	297659	297659	297659	288726	288726	288726
Stations	698	698	698	699	699	699
DMAs (Clusters)	204	204	204	204	204	204
Mean Dep. Variable	27.905	27.905	27.905	27.915	27.915	27.915
Bottom 10% × D = Bottom 10% × R	0.047	0.070	0.066	0.000	0.001	0.001
10%-40% × D = 10%-40% × R	0.590	0.657	0.614	0.335	0.391	0.392
60%-90% × D = 60%-90% × R	0.027	0.018	0.021	0.001	0.001	0.001
Top 10% × D = Top 10% × R	0.001	0.001	0.001	0.157	0.158	0.172

Notes: This table shows the relationship between news coverage of local weather and weather events by media market political leaning, controlling for wildfires and natural disasters. We regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, the same indicators interacted with indicators for the market being either Democratic- or Republican-leaning, one lead and one lag of the same variables, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.2). Columns (1) and (4) additionally control for presence of an active wildfire in the media market, columns (2) and (5) for presence of a declared natural disaster in the media market, and columns (3) and (6) for both. The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (columns (1) to (3)) and for winter (columns (4) to (6)), using a Poisson model. Standard errors are clustered at the media market level.

Table 3.E.5: Publication Bias, Robustness

	Local Weather Segments									
	Outcome Def.			DMAs Split		Outcome Def.			DMAs Split	
	300	Log	Share	20-80	30-70	300	Log	Share	20-80	30-70
	Summer					Winter				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Bottom 10%	0.038***	0.049***	0.005***	0.059***	0.071***	0.067***	0.087***	0.011***	0.105***	0.101***
	(0.008)	(0.009)	(0.001)	(0.011)	(0.015)	(0.008)	(0.011)	(0.001)	(0.011)	(0.012)
Bottom 10% × D-Leaning	-0.007	-0.010	-0.001	-0.009	-0.029*	-0.008	-0.012	0.000	-0.033*	-0.023
	(0.011)	(0.013)	(0.002)	(0.016)	(0.017)	(0.013)	(0.018)	(0.003)	(0.020)	(0.018)
Bottom 10% × R-Leaning	0.014	0.025*	0.001	0.016	0.002	0.041***	0.050***	0.003	0.057**	0.056***
	(0.011)	(0.013)	(0.001)	(0.018)	(0.018)	(0.013)	(0.018)	(0.003)	(0.023)	(0.021)
10%-40%	-0.002	-0.002	-0.000	-0.000	0.000	0.013***	0.016**	0.002***	0.018***	0.019***
	(0.003)	(0.004)	(0.000)	(0.004)	(0.004)	(0.004)	(0.006)	(0.000)	(0.006)	(0.006)
10%-40% × D-Leaning	0.003	0.002	0.000	0.005	0.000	0.005	-0.005	0.001	0.011	0.004
	(0.005)	(0.008)	(0.001)	(0.007)	(0.007)	(0.007)	(0.009)	(0.001)	(0.011)	(0.010)
10%-40% × R-Leaning	0.004	-0.001	0.001	0.002	0.006	0.009	0.006	0.001	0.022*	0.020**
	(0.006)	(0.009)	(0.001)	(0.009)	(0.008)	(0.007)	(0.011)	(0.001)	(0.011)	(0.010)
60%-90%	0.007**	0.011**	0.001**	0.014***	0.011**	-0.007*	-0.012**	-0.001**	-0.013***	-0.017***
	(0.003)	(0.005)	(0.000)	(0.004)	(0.005)	(0.004)	(0.006)	(0.001)	(0.005)	(0.006)
60%-90% × D-Leaning	0.011**	0.014*	0.002**	0.015**	0.017***	-0.002	0.001	0.000	-0.004	-0.001
	(0.005)	(0.008)	(0.001)	(0.007)	(0.007)	(0.006)	(0.008)	(0.001)	(0.009)	(0.008)
60%-90% × R-Leaning	-0.001	-0.005	0.000	-0.003	-0.000	0.017***	0.018*	0.002***	0.029***	0.034***
	(0.005)	(0.008)	(0.001)	(0.007)	(0.008)	(0.006)	(0.009)	(0.001)	(0.008)	(0.008)
Top 10%	0.022***	0.039***	0.003***	0.036***	0.031***	0.005	-0.004	0.000	0.004	0.004
	(0.005)	(0.007)	(0.001)	(0.007)	(0.009)	(0.005)	(0.009)	(0.001)	(0.006)	(0.008)
Top 10% × D-Leaning	0.022***	0.023**	0.003***	0.033***	0.033***	-0.001	0.014	0.001	-0.000	-0.002
	(0.008)	(0.011)	(0.001)	(0.012)	(0.012)	(0.008)	(0.012)	(0.001)	(0.012)	(0.011)
Top 10% × R-Leaning	-0.011	-0.021	-0.001	-0.011	-0.005	0.010	0.010	0.001	0.021	0.019
	(0.007)	(0.013)	(0.001)	(0.012)	(0.012)	(0.009)	(0.014)	(0.001)	(0.014)	(0.012)
Observations	297659	284444	297714	297659	297659	288726	274000	288731	288726	288726
Stations	698	697	698	698	698	699	699	699	699	699
DMAs (Clusters)	204	204	204	204	204	204	204	204	204	204
Mean Dep. Variable	30.406	2.957	0.092	27.905	27.905	29.375	2.909	0.096	27.915	27.915
Bottom 10% × D = Bottom 10% × R	0.058	0.014	0.134	0.185	0.027	0.002	0.002	0.499	0.001	0.001
10%-40% × D = 10%-40% × R	0.951	0.821	0.670	0.774	0.471	0.609	0.309	0.920	0.432	0.173
60%-90% × D = 60%-90% × R	0.021	0.043	0.031	0.026	0.024	0.005	0.095	0.055	0.002	0.000
Top 10% × D = Top 10% × R	0.000	0.002	0.001	0.001	0.002	0.275	0.796	0.886	0.198	0.111

Notes: This table shows robustness of the relationship between news coverage of local weather and weather events by media market political leaning. Unless otherwise specified, we regress the number of segments about local weather on indicator variables for the deviation from the historical mean falling in a given bin of the within-media market deviation distribution, the same indicators interacted with indicators for the market being either Democratic- or Republican-leaning, one lead and one lag of the same variables, station fixed effects, day fixed effects, and number of segments decile fixed effects (see Equation 3.2). Columns (1) and (6) use variables defined based on 300-word long segments. Columns (2) and (7) report estimates from an OLS regression using the log number of local weather segments as the outcome, and columns (3) and (8) using the share of segments that are about local weather. Columns (4) and (9) define as Democratic-leaning (Republican-leaning) media markets in the bottom (top) 20% of the Republican vote share in the 2008 Presidential election and columns (5) and (10) media markets in the bottom (top) 30% of the same variable. The omitted category is the 40%-60% bin, which approximately corresponds to days in which temperatures are in line with the historical mean. We estimate the regression separately for summer (columns (1) to (5)) and for winter (columns (6) to (10)), again using a Poisson model unless otherwise specified. Standard errors are clustered at the media market level.

Appendix 3.F Weather Events and Climate Change Beliefs

A large body of works investigate how temperatures affect climate change beliefs. This is usually done using individual-level survey data and correlating respondents' perspective on climate change with the temperatures they experience (objectively and subjectively). In an excellent review, Howe et al. (2019) summarize a few patterns emerging from the literature. First, individuals seem to be much more sensitive to short-term variations than long-term trends: temperature anomalies on the day or in the temporal vicinity of an individual's interview tends to affect the respondent's belief on climate change. Second, individuals tend to believe more in climate change or become more supportive of actions to mitigate climate change following hotter than usual days, whereas the effect of days colder than usual is more ambiguous (e.g., Joireman, Truelove, and Duell (2010); Hamilton and Stampone (2013); Brooks et al. (2014); Di Leo and Midões, 2023).

While several studies have addressed the question of how temperatures influence individuals' states positions over climate change, there is no agreement on how to measure temperature anomalies or over which geography (Howe et al. (2019)). This makes it important for us to show that indeed weather events, as we define them, also induce the same shifts in beliefs. To do so, we use questions about individuals' stated position on whether action against climate change is required, which were included in the CCES 2009 to 2013 waves.

To get identification, we exploit the fact that CCES is a large scale survey that is fielded over several days, even within the same media market. This gives us quasi-random within-market variation in exposure to weather events on the data individuals responded to the interview. We focus in particular on exposure to weather events in the top or bottom 10% of the within-market deviation distribution. To allow short-term effects of exposure, we consider a respondent exposed if they experienced such weather events in the seven days prior to the interview. Because we are interested in the political dimension of the relationship between weather events and beliefs, we always estimating heterogeneous effects based on individuals' stated ideology.[19] One important caveat is that the survey is fielded in autumn, when the weather tends to be slightly less predictable. However, by making it harder for individuals to benchmark the temperatures they experience against the seasonal norms, this higher variability should, if anything, bias our estimates downwards.

[19] A potential concern with this approach is that ideology could also be post-treatment. However, we believe this concern to be unwarranted. We are, after all, looking at relatively small events: the weather.

More precisely, we estimate the following specification:

$$\begin{aligned} Y_j = & \sum_{i\in\{L,M,C\}} \beta_i^{top10\%}\{\text{\# days in top } 10\%_{m(j)d(j)}\} \times \mathbb{I}\{ideology_j = i\} \\ & + \sum_{i\in\{L,M,C\}} \beta_i^{bottom10\%}\{\text{\# days in bottom } 10\%_{m(j)d(j)}\} \times \mathbb{I}\{ideology_j = i\} \\ & + \sum_{i\in\{L,C\}} \lambda_i \mathbb{I}\{ideology_j = i\} + X_j'\gamma + \delta_{m(j)t(j)} + \epsilon_i, \end{aligned} \tag{3.F.1}$$

where Y_j is an indicator variable equal to 1 if individual j reports that climate change requires action; # days in top $10\%_{m(j)}$ measures the number of days in the top 10% of the within-market deviation distribution that individual j in media market $m(j)$ was exposed to in the week prior to the interview;[20] $\mathbb{I}\{ideology_j = i\}$ are indicator variables for the self-reported ideology of individual j (which can be liberal, moderate, or conservative); X_j is a matrix of individual-level controls (namely, age, gender, race, employment status, education, marital status, homeowner status); $\delta_{m(j)t(j)}$ are media market by year fixed effects. Standard error are clustered at the media market level.

Table 3.F.1 reports the estimates from regressions that progressively build to our preferred specification. In particular, column (1) reports estimates from a regression that only includes media market and year fixed effects (separately); column (2) includes media market by year fixed effects; column (3) further includes individual-level controls (equation 3.F.1). In line with the existing evidence, we document a small effect of short-term temperature anomalies on climate change beliefs. Moderates, in particular, tend to be more (less) supportive of actions against climate change after experiencing temperatures above (below) the historical mean. The same is true to a lesser extent for liberals, who already believe more in climate change (see Appendix Figure 3.A.5). Conservatives, in contrast, always react negatively: they become less supportive of climate change action following any type of weather events.

This pattern should be interpreted with caution. It could indicate that conservatives engage in motivated reasoning discarding evidence going against their beliefs. Alternatively, it could be that weather events increase the salience of climate change issue and motivate conservatives to respond more negatively to a (somewhat) politically charged question. As we do not have data to adjudicate between these possible rationales (they survey, after all, is not incentivized), we leave a further exploration of these findings to future research. Nonetheless, we

[20] Since the survey waves we use correspond to the period 2009-2013, we define the historical mean on the period 2000-2008.

believe that there are several lessons we can learn from this exercise. In particular, the evidence points to: i) individuals having a (small) propensity to react to weather events; and ii) those reactions being different according to the baseline ideology of the individual. As such, individuals' interpretation of weather events may provide one rationale for why TV stations cover the same event differently depending on the media market they operate in.[21]

Table 3.F.1: Weather Events and Support for Climate Change Action

	Climate Change Requires Action		
	(1)	(2)	(3)
Liberal	0.204***	0.204***	0.193***
	(0.005)	(0.005)	(0.005)
Conservative	-0.470***	-0.469***	-0.447***
	(0.006)	(0.006)	(0.007)
# Top 10% Days × Liberal	0.003	0.003**	0.003*
	(0.002)	(0.002)	(0.002)
# Top 10% Days × Moderate	0.004*	0.005**	0.004*
	(0.002)	(0.002)	(0.002)
# Top 10% Days × Conservative	-0.007***	-0.007***	-0.008***
	(0.002)	(0.002)	(0.002)
# Bottom 10% Days × Liberal	-0.002	-0.003	-0.002
	(0.003)	(0.003)	(0.003)
# Bottom 10% Days × Moderate	-0.004**	-0.005**	-0.004*
	(0.002)	(0.002)	(0.002)
# Bottom 10% Days × Conservative	-0.006**	-0.007**	-0.007**
	(0.003)	(0.003)	(0.003)
DMA FEs	✓		
Year FEs	✓		
DMA-By-Year FEs		✓	✓
Additional Controls			✓
Observations	146375	146374	144463
DMAs (Clusters)	204	204	204
Mean Dep. Variable if Moderate	0.694	0.694	0.694

Notes: This table shows the relationship between climate change beliefs and weather events, by individual ideology. In columns (1), we regress an indicator variable equal to one if the respondent reports that climate change requires action on indicator variables for whether the individual is liberal or conservative, the same indicator variables interacted with the number of days in the top 10% of the within-media market deviation distribution that the respondent was exposed to in the week prior to the interview, media market fixed effects and year fixed effects. Column (2) includes media market by year fixed effects. Column (3) reports estimates from our preferred specification (see 3.F.1), that further includes individual level controls (namely: age, gender, race, employment status, education, marital status, and homeownership status). Standard errors are clustered at the media market level.

[21] One limitation of our approach here is that the lack of available data (questions about local TV viewership are asked inconsistently) do not allow us to test for the impact of local TV newscasts on individuals' interpretation of weather events. Yet, others have documented an effect of cable news (Ash et al. 2023) and local media outlets (Andrews et al., 2023) on how citizens interpret weather events or natural disasters, which would be in line with the mechanisms we will develop later in the paper.

Appendix 3.G Topic Modelling

Our topic modelling approach leverages BERTopic, a framework designed to extract meaningful themes from large collections of text by combining transformer-based embeddings, dimensionality reduction, clustering, and topic representation techniques. The main challenge in topic modelling lies in effectively capturing the semantic meaning of text and grouping similar documents into coherent topics without predefined labels. Traditional methods often rely on simple word frequency counts, which fail to account for the complex relationships between words and their contextual meanings. To overcome this limitation, we use a BERT-based sentence embedding model, which transforms text into high-dimensional numerical representations that encode rich contextual and semantic information. Unlike earlier approaches based on word co-occurrence (e.g., LDA), transformer models like BERT allow us to understand how words are used in different contexts, leading to more accurate topic clustering.

Since these embeddings exist in a high-dimensional space, we apply UMAP (Uniform Manifold Approximation and Projection for Dimension Reduction) for dimensionality reduction. UMAP preserves the structure of the data while projecting it into a lower-dimensional space, where similar documents remain close together. We then apply HDBSCAN, a hierarchical density-based clustering algorithm, to identify groups of documents that share similar topics. Unlike traditional clustering methods that require specifying the number of clusters (topics) in advance, HDBSCAN dynamically determines the number of topics based on the density of the data, making it more flexible for tasks where the number of topics is not known beforehand.

Clustering alone does not produce interpretable topics. To address this, we apply different representation techniques to refine topic descriptions. Specifically, we use KeyBERT inspired keyword extraction, a method that builds on BERT embeddings to generate keywords and key phrases with minimal computational overhead. Although many existing methods for keyword extraction are available (e.g., RAKE, YAKE!, TF-IDF), KeyBERT offers a simple yet effective approach using pre-trained BERT embeddings and cosine similarity. The process begins by extracting document embeddings with BERT to obtain a document-level representation. Then, word embeddings are generated for N-gram words and phrases. Finally, cosine similarity is used to determine which words and phrases are most similar to the document as a whole. The highest-ranked words are identified as the most representative of the document's content. To further enhance topic interpretability, we integrate Maximal Marginal Relevance (MMR), which balances relevance and diversity when selecting topic keywords. This ensures that extracted terms are not only representative but also varied, preventing redundancy in topic descriptions.

One limitation of clustering-based topic modelling is that some documents may not fit neatly into a single topic. To improve accuracy, we apply an outlier reduction strategy, reassigning ambiguous documents based on their similarity to existing topic embeddings. This ensures that documents are categorised into the most appropriate topics, improving the overall coherence of the model.

In the final step, we update the topic model with these refined assignments and generate a structured dataset containing document-topic distributions and keyword representations. By leveraging BERT-based embeddings, dimensionality reduction, density-based clustering, and topic representation techniques, our approach provides a robust and scalable method for discovering latent topics in large text corpora.

We apply the topic model to segments about local weather. Because the number of segments about local weather is large, running the topic model on the entire set would be computationally expensive. To address this, we focus on segments about local weather in five randomly selected days per year per season (summer and winter). Even after this initial sampling, we are left with over one million stories, so we further reduce the data by selecting a random subsample of 250,000 observations per season.

Table 3.G.1: Topics of Segments about Local Weather, Winter

	Top Words	Segments Count
Topic 1	morning, snow, today, rain, weather, temperatures, day, right, cold, look	248053
Topic 2	pope, vatican, morning, cardinals, new, today, church, day, conclave, new pope	1947

Notes: This table shows the top words generated by a BERTopic model trained on text segments about local weather in winter. We construct the dataset by randomly selecting five winter days per year and sampling 250,000 observations. Text embeddings are generated using the pre-trained sentence transformer model all-MiniLM-L6-v2. To group similar texts, we reduce the dimensionality of the embeddings using UMAP and then apply HDBSCAN clustering with a minimum cluster size of 1,000. Texts are represented using a bag-of-words approach via CountVectorizer, which removes English stopwords, includes bigrams, and filters out rare terms. Topic representations are based on a combination of KeyBERT-inspired keywords and Maximal Marginal Relevance (MMR). Outliers are reassigned using an embedding-based similarity strategy.

Table 3.G.2: Topics of Segments about Local Weather, Summer

	Top Words	Segment Count
Topic 1	rain, showers, temperatures, afternoon, degrees, tomorrow, day, today, going, morning	46407
Topic 2	police, morning, man, say, old, year old, year, county, news, today	32770
Topic 3	new, good, today, morning, day, like, summer, just, right, time	23288
Topic 4	water, flooding, people, flood, help, harvey, hurricane, storm, florence, county	17260
Topic 5	traffic, road, morning, right, delays, look, lane, good, southbound, accident	15451
Topic 6	game, team, sports, season, football, today, day, tonight, high, just	14582
Topic 7	trump, president, clinton, morning, today, donald, new, donald trump, news, hillary	12370
Topic 8	firefighters, acres, fires, crews, burning, flames, homes, burned, contained, morning	11552
Topic 9	rain, showers, storms, beach, coast, florida, west, county, right, miami	10717
Topic 10	san, degrees, temperatures, bay, inland, low, today, high, going, valley	10495
Topic 11	school, students, kids, schools, summer, year, district, new, day, today	10157
Topic 12	power, tornado, damage, trees, tree, storm, county, people, just, weather	7906
Topic 13	rain, texas, temperatures, heat, showers, today, degrees, going, afternoon, dallas	5785
Topic 14	dog, dogs, animal, animals, pet, pets, zoo, shelter, just, today	5359
Topic 15	iowa, storms, dakota, weather, morning, low, highs, showers, south, today	5162
Topic 16	oklahoma, kansas, storms, rain, tulsa, heat, morning, chance, showers, hot	5094
Topic 17	seattle, clouds, degrees, temperatures, coast, tomorrow, low, mid, cascades, morning	4300
Topic 18	denver, colorado, storms, tomorrow, today, rain, afternoon, plains, pueblo, 80s	3963
Topic 19	virus, mosquito, zika, mosquitoes, nile, west nile, health, nile virus, mosquitos, west	2635
Topic 20	fireworks, july, fourth, fourth july, 4th, weekend, day, 4th july, city, weather	2507
Topic 21	utah, salt lake, salt, wasatch, st george, northern utah, george, lake, northern, southern utah	2240

Notes: This table shows the top words generated by a BERTopic model trained on text segments about local weather in summer. We construct the dataset by randomly selecting five summer days per year and sampling 250,000 observations. Text embeddings are generated using the pre-trained sentence transformer model all-MiniLM-L6-v2. To group similar texts, we reduce the dimensionality of the embeddings using UMAP and then apply HDBSCAN clustering with a minimum cluster size of 1,000. Texts are represented using a bag-of-words approach via CountVectorizer, which removes English stopwords, includes bigrams, and filters out rare terms. Topic representations are based on a combination of KeyBERT-inspired keywords and Maximal Marginal Relevance (MMR). Outliers are reassigned using an embedding-based similarity strategy.

Appendix 3.H Proofs of the Formal Model

Proof of Proposition 4

As we noted in the text, the outlet maximizes pointwise $\frac{1}{\delta}\int_0^1 w(\tilde{c})\big(v(\tilde{c})+\alpha z_D(\rho(w(\tilde{c})))+(1-\alpha)z_R(\rho(w(\tilde{c})))\big)+g(1-w(\tilde{c}))\underline{u}\, dF^e(\tilde{c})$.

Notice that if for a given c, the outlet chooses $w(c)=0$, then the utility of the "average" citizen of watching the newscast conditional on the realization of this event is $g(1)\underline{u}$. Hence, we can define an alternative objective function for the outlet: $w(c)\big(v(c)+\alpha z_D(\mu(c))+(1-\alpha)z_R(\mu(c))\big)+g(1-w(c))\underline{u}$ for any realization c. This function removes the discontinuity in posterior at $w(c)=0$ compared to the original objective function. Yet, point by point (relative to $w(c)$), it equals $w(c)\big(v(c)+\alpha z_D(\rho(w(c)))+(1-\alpha)z_R(\rho(w(c)))\big)+g(1-w(c))\underline{u}$. Hence, any maximum of our alternative function is a maximum of the function we study. Under the assumption of the model (especially $g'(1)=0$ and $g(\cdot)$ strictly concave), the maximum is unique and interior and satisfies:

$$g'(1-w^{dem}(c;\alpha)) = \frac{v(c)+\alpha z_D(\mu(c))+(1-\alpha)z_R(\mu(c))}{\underline{u}} \tag{3.H.1}$$

Proof of Corollary 1

As all functions are differentiable, we obtain (with subscript denoting the partial derivative with respect to the variable):

$$w^{dem}_c(c;\alpha)(-g''(1-w^{dem}(c;\alpha))) = \frac{v'(c)+\mu'(c)\big(\alpha z'_D(\mu(c))+(1-\alpha)z'_R(\mu(c))\big)}{\underline{u}} \tag{3.H.2}$$

$$w^{dem}_\alpha(c;\alpha)(-g''(1-w^{dem}(c;\alpha))) = \frac{z_D(\mu(c))-z_R(\mu(c)}{\underline{u}} \tag{3.H.3}$$

Given that $g''<0$, $z_D(\mu(c))\geq 0\geq Z_R(\mu(c))$ and our assumptions $v'(c)>-\mu'(c)z_R(\mu'(c))$ for all $c>0$ and $v'(c)<-\mu'(c)z_D(\mu'(c))$ for all $c<0$, the result follows.

Proof of Proposition 5

Define $\overline{g} = \min_{c\in[-1,1],\alpha\in[0,1]} \frac{1}{\underline{u}} \frac{\mu'(c)\big(z'_D(\mu(c))-z'_R(\mu(c))\big)}{w^{dem}_c(c;\alpha)w^{dem}_\alpha(c;\alpha)}$. Notice that since all functions are continuous, the minimum is well defined. Further, observe that $\overline{g}>0$. We will first show that if $g'''(c)\leq\overline{g}$ for all c, then $w^{dem}_{\alpha c}(c;\alpha)>0$ for almost all c and all α.

Using Equation 3.H.3 (noting that all functions are again differentiable by assumption), we obtain:

$$w_{c\alpha}^{dem}(c;\alpha)(-g''(1-w^{dem}(c;\alpha)))+w_{c}^{dem}(c;\alpha)w_{\alpha}^{dem}(c;\alpha)g'''(1-w^{dem}(c;\alpha)) = \frac{\mu'(c)\big(z_D'(\mu(c)) - z_R'(\mu(c))\big)}{\underline{u}} \tag{3.H.4}$$

Quite directly, if $g'''(c) \leq \overline{g}$, then $w_{\alpha c}^{dem}(c;\alpha) > 0$ for almost all c.
With this, we can prove the proposition. As noted in the text, the first point of the proposition follows from Corollary 1. We then prove the third point. Notice that $\frac{\partial \Delta^{dem}(c;\alpha^D) - \Delta^{dem}(c;\alpha^R)}{\partial c} = w_c^{dem}(c;\alpha^D) - w^{dem}(c;\alpha^R) = \int_{\alpha^R}^{\alpha^D} w_{c\alpha}^{dem}(c;\widetilde{\alpha})d\widetilde{\alpha} > 0$ since $w_{c\alpha}^{dem}(c;\alpha) > 0$ for almost all c and all α. The second point of the proposition follows from the previous result by noting that $\Delta^{dem}(0;\alpha) = 0$ for all α by construction.

Proof of Proposition 6

Consider the following three sets:

$$\begin{aligned} H_\tau :=& \big\{c \in [-1,1] : w_\tau^{sup}(c;\alpha) = 0\big\} \\ F_\tau :=& \big\{c \in [-1,1] : w_\tau^{sup}(c;\alpha) = 1\big\} \\ I_\tau :=& \big\{c \in [-1,1] : w_\tau^{sup}(c;\alpha) \in (0,1)\big\} \end{aligned}$$

The first is the set of events that a biased outlet $\tau \in \{b, s\}$ hides. F is the set of events that take up the whole programme. I is the set of events that have interior coverage. We suppose that none of these sets is empty (and not singleton). We will show then that the outlet is better off by deviating for each $c \in [-1, 1]$. Recall that A is the audience of the outlet, when all sets are non-empty, the expected utility of the outlet τ is:

$$F^e(H_\tau)V_\tau(A\mu(\emptyset) + (1-A)\pi) + \int_{\widetilde{c}\in F_\tau \cup I_\tau} V_\tau(A\mu(\widetilde{c}) + (1-A)\pi)dF^e(\widetilde{c})$$

For any $c \in I_\tau$, the first derivative satisfies:

$$\frac{\partial A}{\partial w(c)}\Big[F^e(H_\tau)V_\tau'(A\mu(\emptyset) + (1-A)\pi)(\mu(\emptyset) - \pi) + \int_{\widetilde{c}\in F_\tau \cup I_\tau} V_\tau'(A\mu(\widetilde{c}) + (1-A)\pi)(\mu(\widetilde{c}) - \pi)dF^e(\widetilde{c})\Big]$$

Define $c_\tau^* = arg\max_{c\in F_\tau \cup I_\tau} \mu(c)\ s.t.\ \mu(c) \leq \pi$ and

$$\underline{V}_\tau' = \begin{cases} V_\tau'(A\mu(\emptyset) + (1-A)\pi) & \text{if } \pi \geq \mu(\emptyset) \geq \mu(c_\tau^*) \\ V_\tau'(A\mu(c_\tau^*) + (1-A)\pi) & \text{otherwise} \end{cases}$$

We now show that $V_\tau'(A\mu(\emptyset) + (1-A)\pi)(\mu(\emptyset) - \pi) \leq \underline{V}_\tau'(\mu(\emptyset) - \pi)$. In the first case ($\pi \geq \mu(\emptyset) \geq \mu(c_\tau^*)$), it is immediate as the two are equal. Suppose $\mu(\emptyset) \geq \pi$

(recall that $\mu(\pi) \geq \mu(c^*_\tau)$ by definition). Then, given the concavity of $V_\tau(\cdot)$, $V'_\tau(A\mu(\emptyset) + (1-A)\pi) \leq V'_\tau(A\mu(c^*_\tau) + (1-A)\pi)$. Since $\mu(\emptyset) - \pi > 0$, we get $V'_\tau(A\mu(\emptyset) + (1-A)\pi)(\mu(\emptyset) - \pi) < \underline{V}'_\tau(\mu(\emptyset) - \pi)$. Suppose $\mu(\emptyset) < \mu(c^*_\tau)$ (which implies $\mu(\emptyset) < \pi$ since $\mu(\pi) \geq \mu(c^*_\tau)$ by definition). Then, given the concavity of $V_\tau(\cdot)$, $V'_\tau(A\mu(\emptyset) + (1-A)\pi) \geq V'_\tau(A\mu(c^*_\tau) + (1-A)\pi)$. Since $\mu(\emptyset) - \pi < 0$, we get $V'_\tau(A\mu(\emptyset) + (1-A)\pi)(\mu(\emptyset) - \pi) < \underline{V}'_\tau(\mu(\emptyset) - \pi)$.
We now show that $V'_\tau(A\mu(c) + (1-A)\pi)(\mu(c) - \pi) \leq \underline{V}'_\tau(\mu(c) - \pi)$ for all $c \in F_\tau \cup I_\tau$ with strict inequality for a non-empty, non-singleton subset. To do so, we divide the set $F_\tau \cup I_\tau$ into two subsets: $\left(F_\tau \cup I_\tau\right)^+ = \{c \in F_\tau \cup I_\tau : \mu(c) > \pi\}$ and $\left(F_\tau \cup I_\tau\right)^- = \{c \in F_\tau \cup I_\tau : \mu(c) \leq \pi\}$. Note that by definition $c^*_\tau \in \left(F_\tau \cup I_\tau\right)^-$.
Take any $c \in \left(F_\tau \cup I_\tau\right)^+$, by concavity of $V_\tau(\cdot)$, $V'_\tau(A\mu(c) + (1-A)\pi) < V'_\tau(A\mu(\emptyset) + (1-A)\pi)$ if $\pi \geq \mu(\emptyset) \geq \mu(c^*_\tau)$ (as $\mu(c) > \pi$) and $V'_\tau(A\mu(c) + (1-A)\pi) < V'_\tau(A\mu(c^*_1) + (1-A)\pi)$ (as $\mu(c) > \pi \geq \mu(c^*_\tau)$ by definition of c^*_τ). Hence, $V'_\tau(A\mu(c) + (1-A)\pi) < \underline{V}'_\tau$ and since $\mu(c) > \pi$, $V'_\tau(A\mu(c) + (1-A)\pi)(\mu(c) - \pi) < \underline{V}'_\tau(\mu(c) - \pi)$.
Take any $c \in \left(F_\tau \cup I_\tau\right)^-$, by concavity of $V_\tau(\cdot)$, $V'_\tau(A\mu(c) + (1-A)\pi) \geq V'_\tau(A\mu(\emptyset) + (1-A)\pi)$ if $\pi \geq \mu(\emptyset) \geq \mu(c^*_\tau)$ (as $\mu(c) > \pi$) and $V'_\tau(A\mu(c) + (1-A)\pi) \geq V'_\tau(A\mu(c^*_1) + (1-A)\pi)$ (by definition of c^*_τ). Hence, $V'_\tau(A\mu(c) + (1-A)\pi) \geq \underline{V}'_\tau$ and since $\mu(c) \leq \pi$, $V'_\tau(A\mu(c) + (1-A)\pi)(\mu(c) - \pi) \leq \underline{V}'_\tau(\mu(c) - \pi)$. Using all these results, we obtain that

$$\begin{aligned}
&F^e(H_\tau)V'_\tau(A\mu(\emptyset) + (1-A)\pi)(\mu(\emptyset) - \pi) + \int_{\widetilde{c} \in F_\tau \cup I_\tau} V'_\tau(A\mu(\widetilde{c}) + (1-A)\pi)(\mu(\widetilde{c}) - \pi) dF^e(\widetilde{c}) \\
&\quad < \underline{V}'_\tau\left(F^e(H_\tau)(\mu(\emptyset) - \pi) + \int_{\widetilde{c} \in F_\tau \cup I_\tau} (\mu(\widetilde{c}) - \pi) dF^e(\widetilde{c})\right)
\end{aligned}$$

Since citizens are Bayesian:

$$\begin{aligned}
\mu(\emptyset) =& \frac{\pi F_1(H_\tau)}{F^e(H_\tau)} \\
\mu(c) =& \frac{\pi f_1(c)}{f^e(c)}
\end{aligned}$$

Then,

$$\begin{aligned}
&F^e(H_\tau)(\mu(\emptyset) - \pi) + \int_{\widetilde{c} \in F_\tau \cup I_\tau} (\mu(\widetilde{c}) - \pi) dF^e(\widetilde{c}) \\
=& \pi F_1(H_\tau) + \int_{\widetilde{c} \in F_\tau \cup I_\tau} \pi f_1(\widetilde{c}) d\widetilde{c} - \pi(F^e(H_\tau) + F^e(H_\tau \cup I_\tau) = 0
\end{aligned}$$

This implies that

$$F^e(H_\tau)V'_\tau(A\mu(\emptyset) + (1-A)\pi)(\mu(\emptyset) - \pi) + \int_{\widetilde{c} \in F_\tau \cup I_\tau} V'_\tau(A\mu(\widetilde{c}) + (1-A)\pi)(\mu(\widetilde{c}) - \pi) dF^e(\widetilde{c}) < 0$$

By Proposition 4, we know that $\frac{\partial A}{\partial w(c)} > 0$ for all $w(c) < w^{dem}(c)$ and $\frac{\partial A}{\partial w(c)} < 0$ for all $w(c) > w^{dem}(c)$. Hence, the first derivative

$$\frac{\partial A}{\partial w(c)}\Big[F^e(H_\tau)V'_\tau(A\mu(\emptyset) + (1-A)\pi)(\mu(\emptyset) - \pi) + \int_{\widetilde{c}\in F_\tau \cup I_\tau} V'_\tau(A\mu(\widetilde{c}) + (1-A)\pi)(\mu(\widetilde{c}) - \pi)dF^e(\widetilde{c})\Big]$$

is strictly negative for all $w(c) < w^{dem}(c)$ and strictly positive for $w(c) > w^{dem}(c)$. This means that for any $w(c) \in (0,1)$, the outlet has a profitable deviation by either increasing the coverage of weather event or decreasing it. Hence, it cannot be that there is an equilibrium in which all three sets (H_τ, F_τ, I_τ) are non-empty and non-singleton. We can then perform the same analysis with two of the sets including I_τ being non-empty, non-singleton or I_τ to encompass all events (indeed, the proof above does not require all sets to be non-empty, non-singleton, just I_τ) and obtain the same profitable deviation. Hence, there is no equilibrium in which I_τ is non-empty, non-singleton. We can then exclude that it is a singleton using the same as above. As a result, I_τ must be empty, which completes the proof.

Proof of Corollary 2

Immediate from Proposition 6.

Chapter 4

Uncovering Disagreement in Central Bank Communication: Social Media and Monetary Policy Surprises

with Donato Masciandaro (Bocconi), Conor Parle (Fidelity International), and Davide Romelli (TCD).

4.1 Introduction

Central Bank communication can move markets through multiple channels, and exploring its mechanisms has sparked a growing academic literature since the early 2000s. A central question in this literature is whether monetary policy statements and other forms of central bank communications move markets through the message itself. Such messaging can work in two ways. First, it can add new information beyond the market's current information set, prompting investors to update their beliefs. Second, it can deliver a monetary policy surprise by signaling a different reaction function or an unexpected path for monetary policy.

One channel that has received relatively little attention is the anticipation effects: part of the market reaction may reflect the build-up to a decision itself, rather than its substance. There is a reasonable case that the more a decision is expected and discussed in advance, the more prices can move simply because it arrives, regardless of the announcement's content. In addition, when pre-decision opinions are more dispersed, the subsequent adjustment is likely to be larger, again regardless of the actual content of the communication.

This links to the "information effect" literature, which suggests that monetary policy announcements move prices by shifting not only the policy rate, but also beliefs. Nakamura and Steinsson (2018) show that FOMC announcements influence not only beliefs about monetary policy but also about other economic fundamentals, while Bauer and Swanson (2023) warn that failing to purge the "Fed response to news" channel can misstate policy effects. D'Amico and King (2023) add that anticipated monetary policy is a noisy signal and that credibility and horizon shape the response. Read together, these results suggest a simple test for anticipation effects: when pre-meeting attention is high or disagreement is wide, market moves should be larger on impact, even holding the content of the monetary policy announcement roughly constant.

Most empirical work on anticipated monetary policy and the information effect has focused on professional forecaster surveys, especially the Blue Chip survey, leaving pre-meeting public attention on social media largely unexplored (Bauer and Swanson, 2023; D'Amico and King, 2023). We take a different approach and construct a high-frequency Twitter/X dataset around the policy meetings for the Bank of England (BoE), the European Central Bank (ECB), and the US Federal Reserve (Fed). Using the (now-retired) Twitter Academic API, we retrieve the full archive of around 125 million tweets that match a comprehensive keyword set spanning official names, common abbreviations, and governors' surnames, and then focus on the subset of 65 million tweets posted within a ±7-day window around each decision. We translate non-English tweets to English, filter out irrelevant posts with a supervised classifier into unrelated, broadly related, and strictly monetary-policy tweets, and then quantify pre-decision "attention" (volume and engagement) and "disagreement" using our embedding-based dispersion metric and stance splits (hawkish, dovish, neutral). Our disagreement measure represents a methodological innovation in capturing the heterogeneity of monetary policy discourse. Unlike traditional sentiment-based or volume-based metrics, we construct a geometric measure of narrative dispersion by treating each tweet as a point in high-dimensional semantic space and computing the average distance from the discussion's centroid. This approach captures something fundamentally different from tone or activity levels: it quantifies how tightly public narratives cohere around a shared interpretation versus how widely they fragment across competing framings. By combining this dispersion metric with directional stance detection, we can simultaneously measure both the spread of opinions (how much views diverge) and their distribution along the hawkish-dovish spectrum, providing a richer characterization of pre-announcement expectations.

Our empirical analysis focuses on the ECB to test anticipation effects via social media. We focus on a single central bank both to keep a consistent institutional setting and to avoid cross-regime heterogeneity in communication styles and policy frameworks. Our results suggest that monetary policy announcement days with higher pre-meeting user disagreement are followed by stronger asset price moves at the announcement. This is consistent with an anticipation effect in which the decision resolves competing narratives visible in advance. This suggests that part of the asset price volatility which follows monetary policy announcements builds before the decision itself.

These results provide a policy-relevant takeaway for policymakers. Clear and consistent communication ahead of monetary policy decisions can help markets adjust in an orderly way. To reduce noise from pre-meeting disagreement and the associated uncertainty, between-meeting updates can clarify the outlook and support the transmission of monetary policy.

4.1.1 Related Literature

The impact of monetary policy communication on financial markets has been extensively studied, with text-based methods now routinely used in the literature.

The foundational research in this field focused primarily on measuring either monetary policy communication directly or quantifying the impact of monetary policy announcements. Early approaches employed econometric techniques to isolate the "predictable component" of announcements or used high-frequency data to construct measures of "monetary policy surprises".

A seminal contribution to this literature is Romer and Romer (2004), who construct a novel measure of US monetary policy responses that eliminates endogenous and anticipatory movements. They purge their measure of systematic responses to information about future economic developments to obtain a clean identification of monetary policy shocks. Their results show that output responds significantly more strongly to the unanticipated component of monetary policy.

An alternative approach leverages high-frequency information from financial market movements to obtain more direct measures of monetary policy communication effects. Kuttner (2001) pioneered this methodology by exploiting high-frequency movements in federal funds futures rates to disentangle expected from unexpected policy actions. Building on this framework, Bernanke and Kuttner (2005) examine stock market reactions to unanticipated Fed monetary policy actions, finding that an unexpected 25 basis point cut in the federal funds target rate is associated with approximately a 1% increase in broad stock market indices.

Numerous subsequent studies have adopted event study and high-frequency data approaches to measuring monetary policy effects. Gürkaynak, Sack, and Swanson (2005) extends this methodology by decomposing US monetary policy effects on asset prices through the extraction of two distinct factors: a "current federal funds rate target" factor and a "future path of policy" one. This factor-based approach has been extensively developed and applied to other central banks. For example, Altavilla, Brugnolini, et al. (2019) employ similar factor analysis to measure euro area monetary policy effects.

A critical distinction that has emerged in the literature concerns the difference between "policy" and "information" shocks, representing a key innovation beyond earlier work that focused solely on measuring anticipated versus unanticipated effects. This framework recognizes two distinct channels through which monetary policy announcements affect markets: a market reaction to pure policy surprises, and a reaction to information surprises where central banks communicate information about the broader economic outlook through their policy decisions.

This distinction connects directly to recent work by Nakamura and Steinsson (2018), Bauer and Swanson (2023), and D'Amico and King (2023) on information effects, and provides theoretical grounding for our central thesis regarding the importance of pre-announcement anticipation effects. When disagreement among market participants is elevated, they are likely to turn to central banks as authoritative sources to resolve informational ambiguities and clarify economic conditions. Our paper contributes to this literature by introducing a novel methodology that leverages social media data to quantify existing uncertainties and disagreement prior to communication events, and by empirically showing how these pre-existing information gaps are associated with asset price volatility following monetary policy announcements.

Jarociński and Karadi (2020) implement a structural VAR methodology to decompose high-frequency monetary policy shocks into two distinct components: a pure policy component and a component reflecting new information about the economic outlook conveyed through central bank announcements. They show that these components generate different and distinct macroeconomic and financial market effects. Miranda-Agrippino and Ricco (2021) build on this framework by providing a comprehensive decomposition of policy shocks and central bank information effects. Byrne, Goodhead, Mcmahon, et al. (2023) introduce a further innovation by demonstrating that markets can be moved not only by new information about the economic outlook, but also by central banks' alternative interpretations of commonly available information relative to market participants' views.

Beyond measuring monetary policy shocks, a natural progression in the literature has been to examine what specific elements within announcements drive market reactions. This development marks the convergence of natural language processing techniques with traditional monetary policy research, as scholars began quantifying communication content to identify which types of messages most effectively move markets.

Hansen and McMahon (2016) employ NLP (Natural Language Processing) methods to analyze Fed communication, decomposing it into signals regarding forward guidance and economic conditions. Their analysis suggests that forward guidance generates more pronounced market movements, as it conveys information about the likely future path of interest rates, while communication about economic conditions has more limited effects on asset prices. Similarly, Hansen, McMahon, and Tong (2019) examine the Bank of England context, measuring the long-run information effects of central bank communication by extracting topic-based signals from the BoE's Inflation Report.

A central focus in this literature has been the extraction of signals to quantitatively measure tone and uncertainty in central bank communication, an approach we extend through our hawkish/dovish classification methodology. Tone analysis has been extensively explored across multiple studies, including Parle (2022), Tobback, Nardelli, and Martens (2017), Dossani (2021), Picault and Renault (2017), Apel, Blix Grimaldi, and Hull (2019), and Apel and Blix Grimaldi (2014). Complementary research has focused on quantifying other dimensions of information content, particularly uncertainty (Fadda et al., 2025) and communication complexity (Byrne, Goodhead, McMahon, et al., 2025).

Research on monetary policy communication has moved beyond official statements and press conferences. Swanson (2023) shows that central bankers' speeches, especially those by the Fed Chair, can move asset prices even when they aren't tied to scheduled announcements. Altavilla, Gürkaynak, et al. (2025) provide systematic evidence on the market impact of ECB speeches. Building on this, Campiglio et al. (2025) focus on climate-related content in central bank speeches, using data from a broad set of countries.

Recent advances have moved beyond traditional text-based methods to incorporate AI techniques analyzing multimodal communication channels. Gorodnichenko, Pham, and Talavera (2023) employ deep learning methods to detect monetary policy tone from audio recordings, finding that shifting the tone of a Fed press conference from negative to positive can increase S&P 500 returns by approximately 200 basis points. Barry et al. (2025) further extend this approach by using AI to measure facial and vocal expressions of ECB presidents at minute-level frequency, demonstrating that facial and vocal expressions can amplify hawkish policy messages, while positive emotions can enhance the effectiveness of forward guidance statements.

Our paper employs text-based techniques applied to social media data, contributing to the rapidly expanding literature in this domain. Gorodnichenko, Pham, and Talavera (2025) similarly utilize Twitter/X data about the US Federal Reserve to examine what central banks communicate, identify who engages with such communications, and analyze their reactions. Their analysis yields real-time measures of inflation expectations and tracks how these metrics respond to central bank communications and policy actions. Masciandaro, Peia, and Romelli (2024) explore the characteristics of central bank Twitter communication, while Ehrmann and Wabitsch (2022) use Twitter data about the European Central Bank to distinguish between communications directed at experts versus non-experts.

Our paper's distinctive contribution lies in leveraging social media data to quantify disagreement among economic actors regarding central banks and their expected actions. While previous research has examined disagreement, we focus on the micro-level behavior of external participants rather than policymaker-level dynamics. Existing studies have predominantly analyzed disagreement at the policymaker level, such as Tillmann and Walter (2019), who examine divergences between ECB and Bundesbank policy communications. Masciandaro, Romelli, and Rubera (2023) employ Twitter data to explore how similarity differences between tweets and monetary policy statements explain market reactions, an approach closely related to our framework. However, our core innovation focuses on disagreement among Twitter users themselves rather than comparing user content to official central bank communications.

Our paper contributes to the literature across four key dimensions. First, we advance the application of language models to analyze social media reactions and anticipation surrounding central bank decisions. Second, we develop a novel measure of disagreement and show that inter-agent disagreement is associated with larger market reactions to monetary policy announcements. Third, we introduce a comprehensive tone classification system that provides new measures of perceived hawkishness and dovishness in sentiment surrounding central bank decisions. Fourth, we document patterns consistent with the thesis that central bank communication serves to fill information gaps, by identifying and measuring instances where the largest "information gaps" exist prior to monetary policy decisions.

The remainder of the paper is structured as follows. Section 4.2 discusses our data construction and methodology, including the comprehensive tweet extraction strategy across three major central banks, our novel translation pipeline for multilingual content, and the two-stage AI-based classification system that filters tweets by relevance and stance. We also introduce our embedding-based disagreement measure and complementary directional stance indices. Section 4.3 presents descriptive evidence on social media activity patterns around monetary policy announcements, documenting both temporal dynamics and the evolution of discourse across different policy episodes. Section 4.4 provides stylized facts linking pre-announcement social media disagreement to asset price volatility following ECB monetary policy decisions. Section 4.5 concludes.

4.2 Data & Methodology

This section discusses the data and the methodology used in the paper. In particular, we detail the construction of our dataset, the preprocessing and classification pipelines, and the empirical measures used in the analysis.

4.2.1 Twitter/X data

We use the Twitter Academic API, which provided researchers free large-scale access to the full archive of public tweets until its closure in April 2023, to extract every tweet from any user that contains keywords associated with the three major central banks: the European Central Bank (ECB), the Bank of England (BoE), and the Federal Reserve (Fed). To maximize our tweet coverage for these institutions, we implement a keyword-based matching strategy that goes well beyond official names and includes common abbreviations, language variants, and the names of past and present central bank governors. For instance, in the case of the ECB, we include terms such as "European Central Bank", "ECB", and translated or abbreviated forms like "bce" (Italian/French/Spanish), "ezb" (German), or "EKT" (Greek), along with references to presidents such as "Trichet", "Draghi", and "Lagarde". Similar approaches are followed for the Bank of England (e.g., "boe", "bank of england", "Carney", "Bailey") and the Federal Reserve (e.g., "federal reserve", "federalreserve", "Bernanke", "Yellen", "Powell"). Table 4.1 reports a summary of the keyword taxonomy.

This approach ensures that our dataset captures tweets referring to these central banks regardless of regional language differences, spelling variations, or

Table 4.1: Overview of the keywords used for Twitter message extraction

Central Bank	User Handles	Abbreviations	Full Names / Variants	President's Surname
ECB	@ecb, @lagarde*	bce, ecb, ekp, esb, ekt, ezb	banca centrale europea, banco central europeo, banco central europeu, banque centrale europeenne, banque centrale européenne, eiropas centrala banka, eiropas centrālā banka, euroopa keskpank, euroopan keskuspankki, europaische zentralbank, europäische zentralbank, european central bank, europese centrale bank, europos centrinis bankas, europska centralna banka, európska centrálna banka, europska sredisnja banka, europska središnja banka, evropska centralna banka, ευρωπαικη κεντρικη τραπεζα	trichet, draghi, lagarde
BoE	@bankofengland	boe	bank of england, bankofengland, bank_of_england	mervyn king**, carney, bailey
Fed	@federalreserve	fed	federal reserve, federalreserve, federal_reserve	bernanke, yellen, powell

Notes: The table reports the taxonomy of keywords used to retrieve central bank–related tweets. For each institution, we include official handles, common abbreviations, full name variants across major European languages, and the surnames of past and present governors. This broad coverage ensures the capture of both formal and colloquial references.
*Christine Lagarde is the only central bank leader among these institutions with an official personal X (formerly Twitter) account (@lagarde), which is included in our search strategy.
**To reduce noise from unrelated tweets, we extract only those referring explicitly to "mervyn king", given the ambiguity of the term "king".

the use of social media shorthand[1], and also includes discourse involving the key individuals associated with these institutions. We are, therefore, able to gather the complete universe of tweets that mention these keywords, including original tweets, replies, quoted retweets, and retweets, covering the period from 2007 (X's launch year) to May 2023, when the API was discontinued. Along with the tweet text, we collect all associated metadata, including engagement metrics and comprehensive user-level information, *e.g.* the user's name, profile picture, bio, and follower and following counts). The resulting dataset comprises more than 125 million tweets across the three central banks.

Given the fast nature of online discourse and the vast volume of available data, we chose to concentrate exclusively on the online activity during the 7 days surrounding central bank monetary policy announcements (± 7 days from each announcement). This approach reduces the number of tweets while preserving the most meaningful content, allowing us to filter out non-essential data while

[1]This consideration is especially important for the ECB, given the larger number of European language variants in which it is referenced.

capturing the most significant online discussions that emerge before, during, and after central bank communications and announcements. Table 4.2 summarizes these numbers by central bank. The BoE records the highest volume of X activity in our dataset, with over 47 million tweets in the full sample and approximately 23 million tweets during the 7-day meeting windows. The ECB follows with around 40 million tweets in the full sample, while the Fed shows the lowest activity with approximately 38 million tweets. It is important to note that these figures represent noisy data, as a part of it consists of tweets that match our pre-specified keywords but are not actually related to central bank activities or monetary policy discussions. This noise reflects the inherent challenge of keyword-based data collection in capturing only relevant content from the broader X and social media ecosystem. We address this measurement problem in Section 4.2.1.1.

Table 4.2: Tweet Volume by Central Bank

	Full Sample			**± 7 Days Around Meetings**		
Central Bank	English	Non-English	Total	English	Non-English	Total
BoE	29,583,576	17,832,960	47,416,536	13,720,039	9,233,823	22,953,862
ECB	19,095,159	21,511,941	40,607,100	8,356,080	9,232,853	17,588,933
Fed	33,970,513	3,979,846	37,950,359	22,225,513	2,615,237	24,840,750
Total	**82,649,248**	**43,324,747**	**125,973,995**	**44,301,632**	**21,081,913**	**65,383,545**

Notes: The table shows the number of tweets collected for each central bank (BoE, ECB, and Fed) divided by language (english and non-english). The full sample includes all tweets collected in the dataset (i.e. tweets, retweets, replies and quote retweets), while the 7 Days Around Meetings restricts the sample to the X activity within ±7 days of central bank meeting announcements.

4.2.1.1 Measurement Problems

One of the main challenges in using a keyword-matching approach for tweet extraction is filtering out irrelevant tweets that, despite matching the keywords based strategy, are not actually related to the central bank of interest. While this issue might be less relevant when using specific, unambiguous combination of words (such as "European Central Bank," "Federal Reserve," or "Bank of England"), it becomes considerably more problematic when using ambiguous acronyms like "ECB," "Fed," or "BoE."[2] Furthermore, even targeting the names of central bank governors can capture tweets where these names are mentioned in contexts unrelated to central banking, *e.g.*, "Draghi" (former ECB president), which in Italian also means "dragons." This propensity for false positives is a well-documented challenge in the literature on central banking communication. This propensity for false positives is a well-documented challenge in the literature on central banking communication (Ehrmann and Wabitsch, 2022; Gorodnichenko, Pham, and Talavera, 2023; Masciandaro, Romelli, and Rubera, 2023; Grebe, Kandemir, and Tillmann, 2024a,b; Campiglio et al., 2025).

To overcome this, Ehrmann and Wabitsch (2022) follow a rule-based filtering method. They begin by manually inspecting random samples of tweets to identify

[2] Often confused spuriously with the England Cricket Board.

sources of noise, such as sports-related tweets referencing the England and Wales Cricket Board. They then iteratively apply exclusion criteria based on specific keywords (e.g., "cricket," "batsman," "coach," etc.) and unrelated usernames. To refine their filtering, they use visual tools like word clouds to compare the vocabulary of included versus excluded tweets.

As an alternative to the rule-based filtering procedure employed by Ehrmann and Wabitsch (2022), Masciandaro, Peia, and Romelli (2024) adopt a supervised machine learning approach. Starting from a manually labelled sample of 3,000 tweets, independently coded by multiple annotators for relevance to monetary policy announcements, they train an algorithm to classify tweets in a larger corpus of over 467,000 messages. The resulting classifier achieves an accuracy rate of 79% and successfully identifies clear spikes in relevant tweet activity around scheduled central bank announcements.

While previous studies focus exclusively on English tweets or, at most, a limited set of major European languages, our paper expands the analysis to include tweets written not only in English but also in the official languages of all 20 euro area countries. We extend our keyword search to encompass these official languages but do not restrict our analysis exclusively to them. Instead, we process all tweets that mention these keywords regardless of the language they are written in - whether Japanese, Cyrillic, Chinese, or any other language. For instance, while searching for "banca centrale europea" will primarily yield Italian tweets, we also capture and analyze any tweets containing this combination of keywords written in Chinese or any other language. This broader linguistic coverage provides a more comprehensive view of public discourse surrounding the European Central Bank's communication.

However, this multilingual scope introduces additional challenges. Because our approach centers on fine-tuning an English-language model, which has demonstrated superior performance in prior tests compared to multilingual alternatives (Devlin et al., 2019; Conneau & Lample, 2019), non-English tweets may introduce classification errors (Conneau et al., 2020).

To overcome this challenge, we apply an open source Large Langage Model to convert non-English tweets into English prior to classification. This allows us to leverage the superior performance of English-language models while preserving the multilingual scope of our dataset. In the next subsection, we first describe the translation pipeline used to process non-English content and ensure a more robust and consistent filtering process. We then introduce our classification strategy, one that, to our knowledge, has not yet been applied in this context, which consists of two sequential steps. The first step systematically filters out false positives by distinguishing between tweets that are unrelated to central banking, broadly related to monetary policy, and those strictly focused on it. The second step focuses on the strictly related tweets, identifying their stance by classifying them as hawkish, dovish, or neutral.

4.2.1.2 Tweets Translation & Classification

Tweets Translation. To minimize the risk of misclassifying or discarding potentially relevant tweets while maintaining a consistent classification framework, we translate all non-English tweets into English before classification. For this, we use *Mistral24B*, an open-source large language model that ensures full replicability of our results. Large language models have emerged as highly effective translation tools due to their ability to understand context, handle idiomatic expressions, and preserve meaning across languages, capabilities that are particularly valuable when dealing with the informal, abbreviated, and often colloquial nature of social media content. Unlike traditional neural machine translation models that are trained specifically for translation tasks, LLMs benefit from their broader training on diverse multilingual text, enabling them to better capture nuances, slang, and domain-specific terminology that frequently appear in tweets about financial and monetary policy topics.

Our translation process uses the following system prompt:

```
Translate the tweet below into fluent English.
Return a JSON object with:
- "tweet_id": The original ID,
- "translation": The translated text
Example: {"tweet_id": 123, "translation": "Hello world"}
```

This approach enables us to process the entire dataset with a single English-focused classification model, ensuring that relevant content is preserved and classification is applied uniformly across tweets. For instance, an Italian tweet reading "#ANataleNonMancaMai il tentato suicidio di Boe" is accurately translated to "#ChristmasIsNeverMissing the attempted suicide of Boe," demonstrating the model's ability to handle both hashtags and colloquial expressions where "Boe" appears in a non-central banking context. Similarly, a Japanese tweet UTF8min " - " is correctly translated to "BoE Deputy Governor: No support for raising rates until sustained wage growth is confirmed – Reuters", showing the model's capacity to handle technical monetary policy terminology across different writing systems. The model also captures informal contexts effectively, as seen in a Spanish tweet "@KarenHerlein no, me sacaste a mis amigos y despues venis con el caballo cansado, mami retirate. Boe" translated to "@KarenHerlein no, you took away my friends and then you come with the tired horse, mami back off. Boe," where "Boe" appears as a colloquial expression rather than referring to the Bank of England.

While some non-English tweets remain untranslated due to model limitations, we successfully translate close to 95% of them. To avoid translating identical content multiple times and to ensure that our classification focuses on unique tweet id - text pair, we exclude all retweets from the translation process while keeping original tweets, replies and quote retweets. As a result, the number of tweets at this stage is substantially smaller than in Table 4.2. Table 4.3 summarizes translation outcomes by central bank within the ± 7 days around policy meetings, reporting the number of tweets originally in English, those requiring translation, and those successfully translated by our *Mistral24B* model. The "Success Rate" corresponds to the share of non-English tweets translated successfully, and the "Total" column aggregates English and translated tweets. Results show consistently high success rates, 98.29% for the ECB, 95.13% for the Fed, and 92.21% for the BoE, for an overall rate of 94.92%. The lower counts in Table 4.3 compared to Table 4.2 are explained by the exclusion of approximately 32 millions retweets, which make up the majority of the raw dataset relative to original tweets, replies, and quote retweets.

Table 4.3: Tweet Translation Statistics ± 7 Days Around Meetings

Central Bank	English	To Translate	Translated	Success Rate	Total English + Translated
BoE	5,868,565	4,436,892	4,118,600	92.83%	9,987,165
ECB	4,360,237	4,978,560	4,892,666	98.29%	9,252,903
Fed	4,961,816	606,846	573,640	95.13%	5,535,456
Total	**15,190,618**	**10,022,298**	**9,584,906**	**(Avg.) 95.41%**	**24,775,524**

Notes: These statistics cover only tweets within ± 7 days around central bank monetary policy announcements and include original tweets, replies to tweets, and quote retweets. Regular retweets are excluded to avoid translating identical content multiple times.

Tweets Classification. Following the translation of all non-English tweets into English, we proceed with the classification of tweet content. The goal of this step is to systematically filter out irrelevant messages and capture only those tweets that are meaningfully connected to central bank policy discussions. Our classification framework builds on, but extends beyond, existing approaches in the literature.

Masciandaro, Romelli, and Rubera (2023) employ a binary classification distinguishing monetary from non-monetary tweets, our approach adopts a more nuanced taxonomy that reflects the evolving role of central banks since the global financial crisis. In line with recent work highlighting the emergence of secondary and ancillary mandates in central bank communication (Moschella and Pinto, 2019; Campiglio et al., 2025), we classify tweets into three distinct categories: (i) unrelated to monetary policy tweets, (ii) tweets broadly related to monetary policy, e.g., addressing financial supervision, climate change, digital currencies, etc, and (iii) tweets strictly related to monetary policy in its traditional sense, e.g., discussions of inflation or interest rates.

To implement this classification, we randomly selected 35,600 tweets from the combined corpora of tweets. These were manually labelled into the three categories described above: 18,510 tweets were classified as unrelated to non-monetary policy, 7,090 as broadly related, and 10,000 as strictly related to monetary policy. This labelled dataset serves as the foundation for fine-tuning a state-of-the-art zero-shot classifier.[3] By leveraging the general language understanding capabilities of the pre-trained model and tailoring it to our domain-specific

[3]Best models per architectures: `https://ibm.github.io/model-recycling/`. See Grebe, Kandemir, and Tillmann (2024b) for an economics application.

classification task, we are able to achieve strong performance while limiting the amount of manually annotated data required.

We split the labeled tweets into a training set (55%) and test set (45%), preserving class balance across the subsets. After fine-tuning, our best model achieves a balanced accuracy score and F1-micro score of 88%.

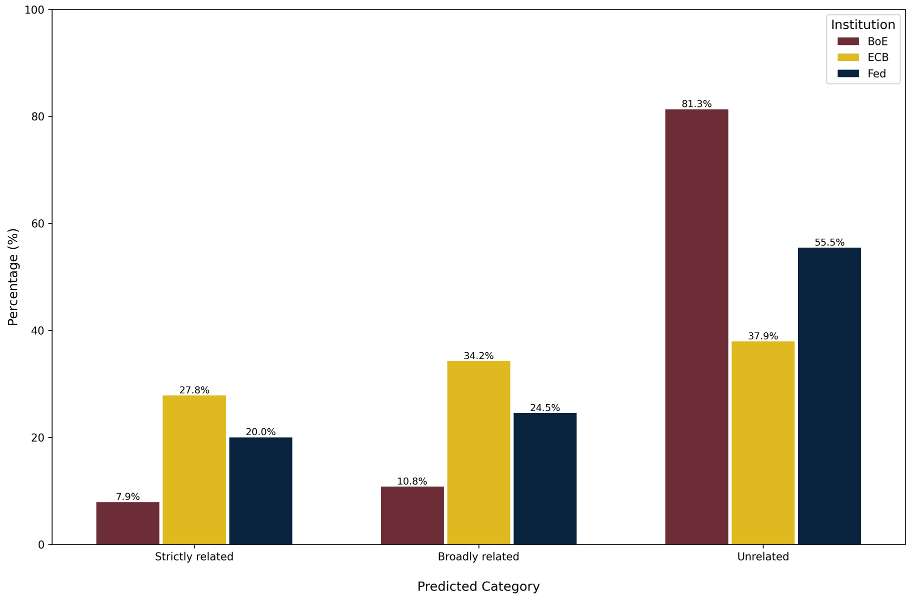

Figure 4.1: Tweet Classification Results Across Central Banks

Notes: The figure reports the percentage distribution of tweets across the three categories, strictly related to monetary policy, broadly related, and unrelated, separately for each institution. The classification is obtained using a fine-tuned BERT model.

Figure 4.1 presents the results of our three-way classification of tweets across the ECB, BoE, and Fed. The classification reveals the substantial amount of noise present in social media data collected through keyword-based approaches, demonstrating the critical importance of systematic content filtering.

The results demonstrate that a large proportion of tweets collected through our keyword-based approach constitute noise that must be filtered out. The "Unrelated" category represents irrelevant content that, while containing our target keywords, bears no substantive relation to central bank policy discussions.

The BoE exhibits the highest noise levels, with 81.3% of collected tweets classified as unrelated to monetary policy. This result is not surprising, as the keyword "boe", the most frequently occurring keyword in our dataset[4], is widely used in contexts unrelated to the Bank of England. In Spanish, it most often refers to the Boletín Oficial del Estado (the official state gazette), but it can also denote other institutions (e.g., Banco de Occidente), appears as an acronym in education or engineering contexts (e.g., Bases de Orientación Educativa, Balance of Energy), or simply occur as a personal name or in general social media usage. In our

[4]Nearly 47% of all tweets extracted under the BoE keyword set.

dataset, 'boe' emerges as the most frequently occurring keyword. This highlights a fundamental challenge in keyword-based data collection and the importance of employing a systematic classification approach to ensure data quality and analytical validity.

The European Central Bank and Federal Reserve show comparatively lower but still substantial noise levels, with 37.9% and 55.5% of tweets respectively classified as unrelated. The relatively better performance of ECB and Fed keywords likely reflects lower susceptibility of their keywords to alternative interpretations.

After filtering, we retain varying proportions of the original samples across institutions. For the BoE, only 18.7% of collected tweets (combining both strict and broad monetary policy categories) are retained for our analysis. The ECB performs better with 62.1% retention, while the Fed retains 44.5% of the originally collected content. In absolute terms, this filtering process eliminates over 8 million potentially misleading tweets from the BoE sample alone, while retaining approximately 1.9 million relevant tweets. For the ECB and Fed, we filter out approximately 3.5 million and 3.1 million irrelevant tweets respectively, while preserving approximately 5.7 million and 2.5 million policy-relevant tweets. Across all three institutions, our classification identifies a total of 10.08 million tweets as relevant to monetary policy analysis (5.6 million broad monetary policy and 4.47 million strict monetary policy), while filtering out the remaining irrelevant content.

These results underscore the necessity of moving beyond simple keyword matching when analyzing institutional communication on social media platforms, as the raw data contains substantial amounts of noise that would severely compromise any subsequent analysis of central bank communication patterns.

Hawkish and Dovish Classification. Building on the first classification step, we then focus on the tone of monetary policy communication. Specifically, among the tweets previously labelled as strictly related to monetary policy, we implement a second classification layer that identifies whether a message conveys a hawkish (indicative of a tendency towards higher rates), dovish (indicative of a tendency towards lower rates and more accommodative monetary policies in general), or neutral stance.

To construct the training set for this tone classification, we randomly selected a balanced sample of tweets, equally drawn from the ECB, BoE, and Fed corpora, and from both originally English and translated messages. Each tweet was independently labelled by the co-authors based on a set of qualitative criteria. Tweets referring to higher inflation, stronger growth, policy tightening, or the downplaying of rate cuts were categorized as hawkish. Conversely, references to easing measures such as rate cuts, increased QE, or weak macroeconomic indicators were labeled as dovish. Tweets that did not signal a clear directional tilt were marked as neutral.

Once the initial set was labeled, disagreements between annotators were resolved through discussion, producing the hand-labelled dataset used for model training. To strengthen classifier robustness, we included both strictly related tweets (9,161) and a subset of broadly related tweets (4,562) in the training process. The inclusion of broadly related content, while not part of our main econometric analysis, helps improve generalization and sensitivity to tone variation across topics. We then fine-tuned a pre-trained zero-shot classifier on this combined set and applied it only to the strictly related monetary policy tweets. Out of 13,576 classified tweets, we identified 4,510 as dovish, 4,773 as hawkish, and 4,293 as neutral. Given the relatively limited sample size available for training, we opt for a 65%-35% training-test split to provide the model with additional learning data while maintaining sufficient evaluation samples. We preserve class balance across both subsets to prevent overfitting to any particular category. After fine-tuning, our best model achieves a balanced accuracy score and F1-micro score of 86%.

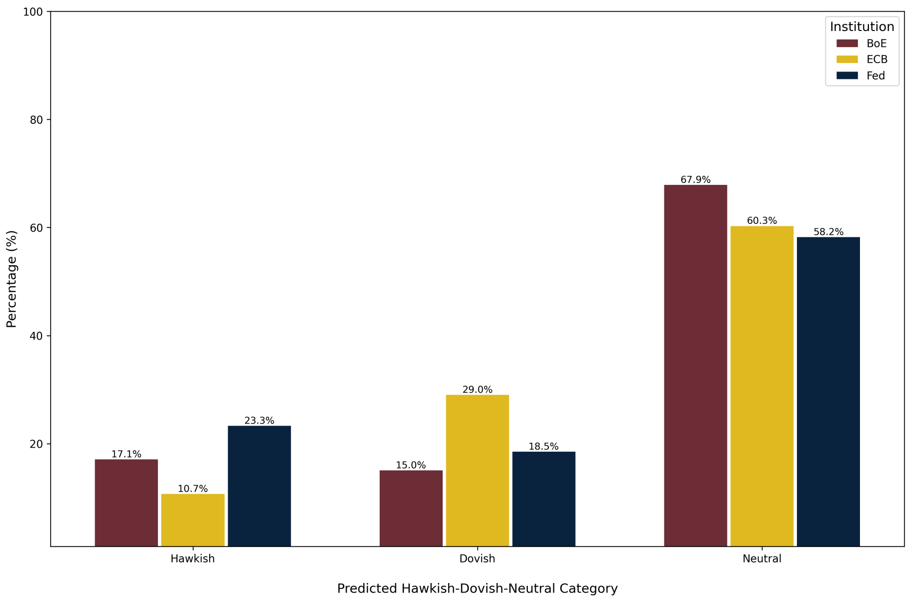

Figure 4.2: Hawkish, Dovish, Neutral Classification Results by Institution

Notes: This figure shows the percentage distribution of hawkish, dovish, and neutral sentiment classifications for monetary policy-relevant tweets across the BoE, ECB, and Fed.

Figure 4.2 presents the hawkish-dovish-neutral (HDN) classification results for tweets strictly related to monetary policy across the three central banks. The analysis reveals a predominance of neutral classification across all institutions, accounting for the majority of classified content: 67.9% for the BoE, 60.3% for the ECB, and 58.2% for the Fed.

The ECB exhibits the highest proportion of dovish content at 29.0%, compared to 18.5% for the Fed and 15.0% for the BoE. Conversely, hawkish classification is most pronounced in Fed related tweets at 23.3%, while the BoE shows 17.1% and the ECB demonstrates the lowest hawkish classification at 10.7%. The patterns in tweets tone reflect distinct monetary policy cycles during the sample period. The ECB's prolonged accommodative stance, including negative interest rates and extensive quantitative easing programs, most likely drove the dovish sentiment in public discourse, exemplified by Draghi's "whatever it takes" claim. The Fed's earlier rate hiking cycles and aggressive tightening to combat inflation most likely fueled hawkish tweets, amplified by Powell's focus on price stability.

4.2.1.3 A measure of Disagreement and Stance/Tone Detection

While the hawkish, dovish, neutral classification captures the directional of tweets discussing monetary policy, it does not fully describe the distribution of views within the debate. A sequence of tweets can all be dovish (hawkish) yet tightly clustered around the same narrative, or they can diverge in how they frame easing (tightening) policies. To move beyond the tone dimension and assess the extent to which public discussion reflects consensus or fragmentation, we introduce a measure of disagreement. At its core, disagreement is a property of a discussion in which participants' beliefs, expectations, or interpretations do not align (De Kock and Vlachos, 2021). It is not simply noise or emotion; it is about how far views diverge. Two conversations can be equally loud yet differ in disagreement: in one, participants reiterate a common narrative (low disagreement); in the other, they advance distinct, incompatible narratives (high disagreement). Conceptually, then, disagreement is about dispersion in the space of ideas: how spread out the positions are and not how positive, negative (sentiment) or intense (volume) the conversation becomes.

Building on this logic, we apply the same lens to discussions around central bank announcements (Gorodnichenko, Pham, and Talavera, 2025; Camous and Matveev, 2021). Before a decision, users can diverge on the expected action and the macro backdrop; after the decision, they can split over how to read the statement and press conference, or whether the news constitutes a surprise relative to priors. In both stages, our focus is not sentiment or sheer volume but the dispersion of narratives—how tightly conversations cohere around a shared framing versus how widely they fragment across competing interpretations.

However, disagreement is not directly observable, and textual data are short, high-dimensional, and noisy. An operational proxy must (i) respect the geometry of short texts; (ii) scale to event-time windows before and after meetings; and (iii) be interpretable as a dispersion statistic (rather than sentiment or activity per se). Although no single metric is perfect, our approach is deliberately transparent: we treat each tweet as a point in a semantic space and summarize dispersion from the center of the discussion within a time window.

The construction is intuitive. For a given window around a meeting, we map each tweet to a vector that captures its semantic content, compute the window's "center" as the average of those vectors (the centroid), and then take the average Euclidean distance of tweets from that center. If users cluster around a shared framing (e.g., "hold with hawkish tilt"), distances are small and disagreement is low. If users pull in different directions, distances grow and disagreement is high. This is the same logic as looking at how tightly points bunch around their own center focusing on the conversation's internal coherence rather than where it sits in the broader space.

Each tweet is represented by a sentence level embedding that maps short texts into a high dimensional vector space designed so that semantically similar messages are close together even when they use different words. In practice, this representation aggregates topical framing (e.g., "hawkish pause" vs. "policy error"), relational cues (e.g., references to inflation, growth, or guidance), and coarse stance like signals that appear in the wording. This is precisely what allows distances to meaningfully reflect whether the conversation converges on a single interpretation or splits across competing ones.

Let $\mathscr{W}$ denote a time window around a meeting, with $n_{\mathscr{W}} \geq 3$ tweets. Let $\mathbf{x}_i \in \mathbb{R}^d$ be the embedding of tweet i in that window, obtained from a sentence-level model tailored to short texts (Fang, Nguyen, and Oberski, 2022; Chen et al., 2024).[5] Define the centroid of the window as

$$\bar{\mathbf{x}} = \frac{1}{n_{\mathscr{W}}} \sum_{i \in \mathscr{W}} \mathbf{x}_i .$$

Our baseline disagreement index is the mean Euclidean distance from each tweet to the centroid:

$$D_{\mathscr{W}} = \frac{1}{n_{\mathscr{W}}} \sum_{i \in \mathscr{W}} \left\| \mathbf{x}_i - \bar{\mathbf{x}} \right\|_2 .$$

To reflect the idea that dispersion measured from broader participation is both more informative and more likely to matter for market attention, we also report a tweet count weighted version,

$$D^{\star}_{\mathscr{W}} = D_{\mathscr{W}} \times \sqrt{n_{\mathscr{W}}} ,$$

which preserves the interpretation as dispersion while increasing the influence of windows with more X activity (tweets). We evaluate a window only if it contains at least three tweets, given that with a single tweet, the centroid coincides with the tweet and the dispersion is mechanically zero (uninformative). With two tweets, the centroid sits at the midpoint and the dispersion reduces to half of their pairwise distance, which is extremely sensitive to noise and idiosyncratic wording. Panel A of Figure 4.3 shows a case where the discussion scores lower disagreement, with tweets clustering more closely together, while Panel B shows a case of higher disagreement, with tweets more widely dispersed.

To understand the evolution of disagreement around monetary policy announcements, we construct non overlapping windows of 15, 30, 60, and 120 minutes before and after each meeting's decision time, defined as the scheduled public release of the monetary policy decision (statement/press release). For every meeting of the ECB, BoE, and the Fed, we set $t = 0$ at this timestamp and compute the indices separately by meeting and by central bank[6]. This yields a consistent set of pre–post pairs for every meeting and window length. For each window, we compute $D_{\mathscr{W}}$ and $D^{\star}_{\mathscr{W}}$ exactly as described above. The pre-decision

[5] We employ the *bge-large-en-v1.5* model for sentence embeddings.

[6] The ECB announces its monetary policy decision via a press release at 13:45 CET, followed by a press conference beginning at 14:30 CET. We set $t = 0$ at the press release time (13:45 CET) for the press-release analysis and at the press conference start time (14:30 CET) for the press-conference analysis. For the BoE and the Fed, $t = 0$ is set at the scheduled time of the policy announcement and, where applicable, the start of the press conference.

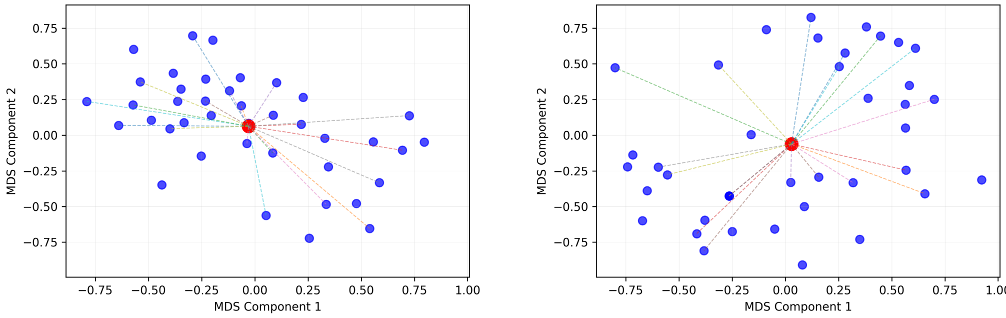

Panel A. Low disagreement **Panel B.** High disagreement

Figure 4.3: Low vs. high disagreement as semantic dispersion.

Notes: Each dot is a tweet embedding projected into two dimensions using a shared metric MDS layout computed on equal-sized samples from the two windows; the red dot is the within-window centroid. Panel A shows a relatively coherent window (low dispersion); Panel B shows a more fragmented window (high dispersion). Distances in the figure are for visualization; our empirical disagreement index is the mean Euclidean distance to the centroid computed in the original embedding space.

indices capture an anticipation phase in which users debate the likely action and its justification; the post-decision indices capture an interpretation phase in which users parse the statement and press conference relative to their priors. Constructed symmetrically, these measures allow us to compare disagreement before and after the same meeting and relate each component to high-frequency market reactions (Adams et al., 2023).

Because $D_{\mathcal{W}}$ summarizes the *spread* of messages, a uniformly negative or uniformly positive conversation can yield a low score if users express similar narratives, whereas a small window can yield a high score if the few messages present are far apart in meaning. The weighted index $D^{\star}_{\mathcal{W}}$ makes this dependence on participation explicit through the $\sqrt{n_{\mathcal{W}}}$ factor.

After finalizing the disagreement measure based on overall dispersion in the embedding space, we also want to capture another, complementary dimension: the extent to which the conversation *tilts* toward hawkish or dovish positions. While the embedding-based measure reflects how dispersed opinions are in general, regardless of their substantive orientation, this second step focuses directly on directional tone/stance within a window.

In particular, we summarize how the discussion divides between *hawkish* and *dovish* stances. Let h and d be the counts of hawkish and dovish tweets, and let u be the number of *neutral* tweets (no clear directional stance). We define the *directional stance index conditionally on active stances* as

$$S_{\text{dir}} = \frac{h-d}{h+d}, \qquad \text{for } h+d>0.$$

This index ranges from −1 (all active tweets dovish) to +1 (all active tweets hawkish), with 0 indicating a balanced split. Compared with the dispersion measure, which quantifies how far tweets are spread in the semantic embedding space without using stance labels, the signed stance index S_{dir} adds direction along the hawkish–dovish axis. Dispersion is stance-agnostic and multi-dimensional; S_{dir} is stance-specific and one-dimensional, indicating whether the balance of expressed stances leans hawkish or dovish.

Precision-weighted direction (active stances). While S_{dir} captures direction, it treats a split based on a handful of *active* tweets the same as one based on many. To reflect that the directional signal becomes more stable as *active stances* accumulate, we report an active-volume–adjusted score:

$$S^{\star}_{\text{dir}} = \sqrt{h+d} \cdot S_{\text{dir}}.$$

Relative to S_{dir}, $S^{\star}_{\text{dir}}$ gives more influence to windows with *more hawkish and dovish tweets*. It is a precision-weighted *score* (not bounded in $[-1, 1]$) and does not incorporate neutral tweets.

Participation-aware direction (all tweets). Because $S^{\star}_{\text{dir}}$ ignores neutrality, we also report a measure that explicitly accounts for *overall participation* in the window:

$$\widetilde{S}_{\text{dir}} = \frac{h+d}{n} \cdot S_{\text{dir}}, \qquad n = h + d + u.$$

Relative to $S^{\star}_{\text{dir}}$, $\widetilde{S}_{\text{dir}}$ down-weights windows in which few tweets take a side, incorporates neutral tweets via $(h+d)/n$, and preserves the $[-1, 1]$ scale.

Taken together, these measures provide two clean, separable lenses on disagreement around monetary policy decisions. The embedding–based dispersion index captures *narrative spread*, that is, how far users' meanings fan out within a window, whereas the stance–based indices capture *policy-tilt polarization*: the balance between hawkish and dovish interpretations.

4.2.2 Other data sources

To complement our social media measures, we incorporate high-frequency financial market data from the Euro Area Monetary Policy Event-Study Database (EA-MPD), published by the European Central Bank.[7] The EA-MPD provides standardized intraday return data around ECB monetary policy events, calculated from one-minute price observations.

We focus on six financial instruments that are highly liquid and directly sensitive to monetary policy expectations: the EURO STOXX 50 index and the EURO STOXX Banks Index (SX7E), which capture broad equity and banking sector performance; ten-year sovereign bond yields for France, Germany, and Spain; and the ten-year overnight index swap (OIS) rate, a widely used proxy for long-term expectations of short-term interest rates.

The EA-MPD computes returns as the percentage change between median prices in narrow pre- and post-event windows, using medians over short intervals

[7] Available at: https://www.ecb.europa.eu/pub/pdf/annex/Dataset_EA-MPD.xlsx.

rather than single timestamps to minimize the influence of transient noise. The timing conventions are as follows:
Press release (decision announced at 13:45 CET): pre-event median over 13:25–13:35; post-event median over 14:00–14:15.
Press conference (begins at 14:30 CET): pre-event median over 14:15–14:25; post-event median over 15:40–15:50.
Overall monetary policy event: pre-press-release median (13:25–13:35) compared to post-press-conference median (15:40–15:50).

Following Masciandaro, Romelli, and Rubera (2023), we take the absolute value of these returns. This choice is motivated by the nature of our social media measures: disagreement captures the dispersion of views, not their direction. High disagreement may precede large market moves in either direction, so the relevant question is how much prices move, not which way. Using absolute returns isolates the intensity of the market reaction.

4.3 Descriptive Evidence

This section documents patterns in social media discussions surrounding monetary policy announcements across our three central banks. We begin by examining the temporal evolution of Twitter/X activity over our entire sample period (2007-2023), focusing on the seven-day windows around each monetary policy meeting to identify long-term trends and major policy episodes. We then provide a granular view of social media dynamics by analyzing tweet frequency in 15-minute intervals spanning the sixteen hours around each announcement, revealing how social media discussions build up before, peak during, and dissipate after policy decisions.

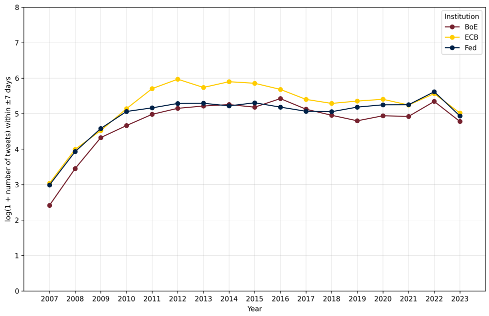

Panel A. Tweet volume around monetary policy meetings by institution, yearly.

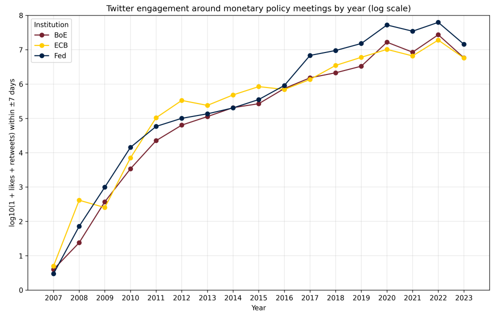

Panel B. Engagement around monetary policy meetings by institution, yearly.

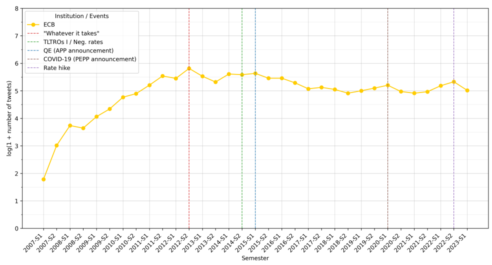

Panel C. ECB tweet volume by semester with major policy events.

Figure 4.4: Twitter activity and engagement around monetary policy meetings.

Notes: Panel A plots log(1 + tweets) within ±7-day window around each monetary policy announcement meetings for the BoE, ECB, and Fed (2007–2023). Panel B reports log(1 + likes + retweets) over the same window and period. Panel C shows ECB tweet volume by semester with major policy events marked by dashed vertical lines (e.g., "Whatever it takes," TLTROs/negative rates, APP, PEPP, first rate hike). Axes are on log scales to compress right tails and improve readability.

The temporal evolution of social media activity reveals several striking patterns. Panel A of Figure 4.4 tracks yearly Twitter/X activity across the three central banks, with the most dramatic feature being a pronounced spike in ECB-related discussions during 2012, i.e. the year of Draghi's pivotal "whatever it takes" intervention. Following this crisis period, all three institutions exhibit remarkably similar and stable posting patterns, suggesting that while the initial adoption of social media varied, the core community of users engaged with monetary policy discussions has remained relatively consistent. The ECB maintains the highest overall volume throughout the sample, likely reflecting both the complexity of multi-country policy coordination and the institution's higher public profile during the European sovereign debt crisis.

More intriguing is the divergence between posting behavior and audience engagement shown in Panel B. While active posting stabilizes after the initial platform adoption phase, passive engagement—measured through likes and retweets—continues its steady upward trajectory throughout the entire sample period. This pattern suggests that while the number of users actively contributing to monetary policy discussions remains relatively constant, the audience consuming and amplifying this content grows substantially over time, indicating the increasing influence of social media discussion on broader public awareness of central banking.

Panel C provides a granular view of social media activity about the ECB that maps perfectly onto major policy innovations: the period around the 2012 "whatever it takes" speech stands out as the single largest spike, followed by notable increases during the introduction of targeted longer-term refinancing operations (TLTROs) and negative interest rates, the Asset Purchase Programme launch, the Pandemic Emergency Purchase Programme (PEPP), and the recent policy normalization. This tight correspondence between policy innovation and social media activity suggests that Twitter/X served as a real-time barometer of perceived policy significance.

Examining the high-frequency dynamics around individual policy announcements (Figure 4.5) reveals a consistent temporal pattern that transcends the institutional differences across all three central banks. This common pattern emerges despite the substantial differences in communication protocols, most prominently for the ECB's unique dual-event framework that consists of an initial press, followed by a press conference which starts 45 minutes after the first communication channel. The evolution of social media discussion follows a predictable three-phase evolution: anticipatory discussions gradually intensify in the pre-announcement period, reach explosive peaks during the policy release itself, and subsequently decay in a measured return to baseline activity levels.

The ECB's communication structure offers particularly rich insights into how policy information cascades through social media channels. Its characteristic dual-spike pattern, featuring an initial surge accompanying the 13:45 press release followed by an even more pronounced peak during the 14:30 press conference, illuminates the real-time information processing dynamics of both financial markets and social media participants. This sequential reaction pattern suggests that while the press release provides the fundamental policy decision, the press conference delivers crucial additional context that substantially amplifies public engagement, indicating that the interactive Q&A format that follows the introductory statement of the press conference adds significant explanatory value beyond the initial policy announcement.

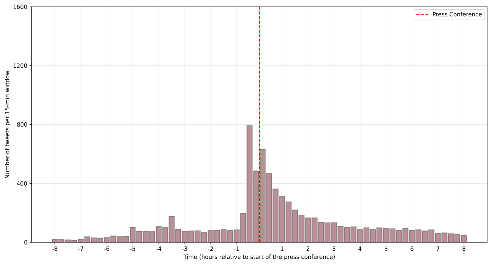

Panel A. BoE

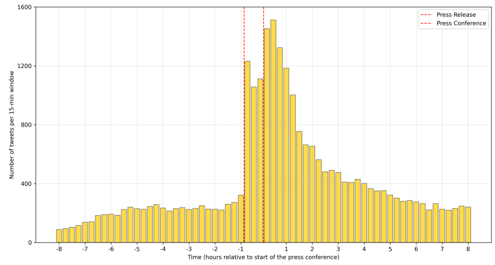

Panel B. ECB

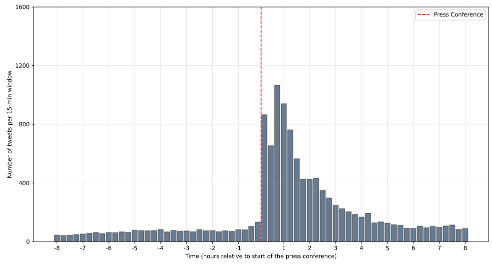

Panel C. Fed

Figure 4.5: Average posting intensity around policy announcements (15-minute bins)

Notes: Bars show the average number of tweets per 15-minute bin within ±8 hours of the press conference, pooled across meetings for the whole sample period (2007-2023). For the ECB (Panel B) we also indicate the press release time; dashed vertical lines mark the release (ECB only) and the start of the press conference.

4.4 Stylized Facts: Disagreement, Sentiment, and Asset Price Movements

This section explores how social media activity before ECB announcements is associated with the magnitude of high-frequency market reactions. We focus on the ECB, which offers a particularly informative setting for two reasons. First, the ECB's dual-event communication structure, a press release at 13:45 CET announcing the policy decision, followed by a press conference beginning at 14:30 CET, allows us to examine whether different types of social media signals matter at different stages of the communication process. Second, restricting attention to a single institution avoids cross-regime heterogeneity in communication style and policy framework. Our analysis examines three dimensions of pre-announcement online discourse: the dispersion of narratives (disagreement), the directional tilt of the conversation (hawkish vs. dovish stance), and simple activity levels (tweet volume and engagement).

We stress at the outset that these results document systematic associations between social media measures and market reactions; they are not intended to establish a causal link. The timing structure of the exercise, however, is informative: social media variables are measured *before* each event, so any associations reflect pre-announcement expectations and positioning rather than direct reactions to the policy decision.

4.4.1 Specification and Timing

Our sample covers all ECB monetary policy meeting dates over the period 2007–2023.[8] For each meeting date t and each asset τ, we estimate the following bivariate regression:

$$|r^{e}_{\tau,t}| = \alpha + \beta \, X^{\text{pre-}e}_{t} + \varepsilon_{\tau,t}, \tag{4.4.1}$$

where the superscript $e \in \{\text{PR}, \text{PC}\}$ indexes the two ECB communication events, the press release (PR) and the press conference (PC), and $X^{\text{pre-}e}_{t}$ is a social media measure computed in the 15 minutes immediately before event e.

Dependent variable. The dependent variable $|r^{e}_{\tau,t}|$ is the absolute percentage change in asset τ around event e on date t. We take absolute values because our social media measures capture dispersion and intensity of discussion, not its direction; high disagreement may precede large moves in either direction. Returns are computed from the Euro Area Monetary Policy Event-Study Database (EA-MPD), which calculates percentage changes between median prices in narrow intraday windows to minimize the influence of transient noise:[9]

- **Press release window:** the pre-event price is the median over 13:25–13:35 CET; the post-event price is the median over 14:00–14:15 CET.
- **Press conference window:** the pre-event price is the median over 14:15–14:25 CET; the post-event price is the median over 15:40–15:50 CET.

Assets. We consider six euro-area financial instruments that are liquid and directly sensitive to monetary policy expectations: the EURO STOXX 50 index, the EURO STOXX Banks Index (SX7E), ten-year sovereign bond yields for France, Germany, and Spain, and the ten-year overnight index swap (OIS) rate.

Social media regressors. The right-hand-side variable $X^{\text{pre-}e}_{t}$ is one of the social media measures described in Section 4.2.1.3, computed over the 15-minute window ending at the start of event e, that is, 13:30–13:45 CET for the press release and 14:15–14:30 CET for the press conference. We consider three measures in turn:

1. $D_{\mathscr{W}}$: the baseline (unweighted) embedding-based disagreement index;
2. $D^{\star}_{\mathscr{W}}$: the tweet-count weighted disagreement index, which scales $D_{\mathscr{W}}$ by $\sqrt{n_{\mathscr{W}}}$ to give more weight to windows with greater participation;

[8] The exact number of meetings included varies slightly across assets depending on data availability in the EA-MPD database.

[9] Available at: https://www.ecb.europa.eu/pub/pdf/annex/Dataset_EA-MPD.xlsx.

3. $\widetilde{S}_{\text{dir}}$: the participation-aware directional stance index, which measures the hawkish–dovish tilt of the conversation while down-weighting windows dominated by neutral tweets.

Each regression in equation (4.4.1) includes a single social media regressor at a time. We estimate the equation separately for each asset and each event (press release or press conference), yielding a total of $6 \times 2 = 12$ regressions per social media measure. This deliberately simple specification serves our purpose of documenting stylized facts. Confidence intervals are shown at the 90% level.

Timing summary. The crucial feature of this design is that all social media variables are measured *before* the event window opens. Any relationship we detect therefore reflects pre-announcement expectations, not a mechanical reaction to the policy decision itself.

4.4.2 Results: Disagreement and Stance

Figure 4.6 presents coefficient estimates from equation (4.4.1) across our six assets and three social media measures, separately for press releases (Panel A) and press conferences (Panel B).

Weighted disagreement around press conferences. The most robust finding involves the weighted disagreement measure $D^{\star}_{\mathscr{W}}$ around press conferences. Across several assets, the coefficient is positive and precisely estimated, indicating that meetings preceded by more dispersed social media narratives are followed by larger absolute price moves during the press conference window. The unweighted measure $D_{\mathscr{W}}$, by contrast, yields coefficients that are generally positive but imprecisely estimated. The difference is informative: the $\sqrt{n_{\mathscr{W}}}$ scaling in $D^{\star}_{\mathscr{W}}$ gives more weight to meetings that attract substantial pre-announcement discussion, precisely the meetings where the disagreement signal is estimated with less noise and where broader public attention is more likely to matter for market dynamics.

The stronger performance of disagreement measures around press conferences rather than press releases also has an institutional explanation. The press conference, with its extended duration, introductory statement, and interactive Q&A session, generates substantially more social media engagement than the brief initial press release. This higher participation allows us to construct more informative disagreement measures, and it also represents the stage of communication where broader interpretive uncertainty is resolved. The press conference is where the ECB explains the rationale, risk assessment, and forward guidance behind the decision, precisely the elements over which pre-announcement narratives are most likely to diverge.

Directional stance around press releases. The directional stance measure $\widetilde{S}_{\text{dir}}$ reveals a complementary pattern: it shows more precisely estimated coefficients around press releases than around press conferences. This finding is intuitive. The press release is the moment when the core policy decision, the key rate change or the announcement of unconventional measures, is revealed. This is precisely when the hawkish–dovish dimension of the pre-announcement debate is most relevant. The directional stance index captures the focused discussion about whether the central bank will tighten or ease, and its informational content peaks at the moment the decision is announced.

Why different measures matter at different stages. The complementarity between these two findings is not an artefact; it reflects the distinct informational content of the two ECB communication events.

The press release conveys the *decision*: the policy rate level and any changes to asset purchase programs. Before this announcement, the most salient dimension of social media discussion is directional, "will they hike or hold?" This maps directly onto $\widetilde{S}_{\text{dir}}$. Once the decision is known, the directional debate is largely resolved, and $\widetilde{S}_{\text{dir}}$ loses predictive content.

The press conference conveys the *rationale*: the economic assessment, the balance of risks, and signals about the future policy path. Before this event, the relevant dimension of social media discussion is broader, users debate not just the direction of the next move but how to interpret the decision, what it implies for the medium-term outlook, and whether the communication shifts the ECB's reaction function. This interpretive uncertainty is multi-dimensional and maps naturally onto the embedding-based dispersion measure $D^{\star}_{\mathscr{W}}$, which captures how widely narratives fragment across semantic space regardless of their hawkish or dovish orientation.

This pattern also has implications for understanding central bank communication more broadly. It suggests that central banks operate in an information environment with two layers of anticipation. The first layer concerns the immediate policy action and is resolved by the press release; the associated social media signal is directional (stance). The second layer concerns the interpretation and forward-looking implications of the decision and is resolved by the press conference; the associated signal is dispersion (disagreement). Markets appear sensitive to both, but at different moments in the communication process.

4.4.3 Cross-Asset Heterogeneity

A notable feature of Figure 4.6 is the heterogeneity in coefficient magnitude and precision across asset classes.

Equities. The EURO STOXX 50 and, to a lesser extent, the EURO STOXX Banks Index (SX7E) consistently yield some of the largest and most precisely estimated coefficients across specifications and events. This is consistent with the high liquidity and rapid price adjustment of equity markets, which may make them particularly responsive to shifts in the information environment ahead of monetary policy announcements. The banking index's sensitivity is further explained by the direct exposure of euro-area banks to interest-rate changes via net interest margins.

Sovereign bond yields. Among sovereign bond yields, the results are more heterogeneous. French and German ten-year yields tend to show precisely estimated positive associations with $D^{\star}_{\mathscr{W}}$, particularly around press conferences. Spanish ten-year yields (ES10Y) display large point estimates in several specifications, but the confidence intervals are wider and often include zero, reflecting greater imprecision. This imprecision likely stems from the higher volatility of Spanish yields over the sample period, which spans the European sovereign debt crisis. During this period, Spanish sovereign spreads were at the epicentre of the crisis and remained sensitive to ECB communication. The large point estimates are consistent with the idea that peripheral bond yields, which carried higher credit risk and redenomination risk relative to core issuers, are more responsive to shifts in perceived ECB stance. However, the wider confidence intervals mean we cannot draw strong conclusions about whether the association is systematically larger for Spain than for other sovereign yields. The OIS rate shows mixed results, with coefficients close to zero in most specifications.

Overall, the cross-asset patterns suggest that the association between pre-announcement disagreement and market reactions is not confined to a single asset class. Equity indices, which adjust quickly and are highly liquid, show some of the strongest results, while sovereign bond yields display more heterogeneity that partly reflects differences in volatility and risk characteristics across issuers.

4.4.4 The Role of the Sample Period

Our sample covers a period of considerable heterogeneity in the monetary policy environment: the global financial crisis (2007–2009), the European sovereign debt crisis (2010–2012), the extended low-rate and QE era (2013–2019), the pandemic response (2020–2021), and the rapid tightening cycle (2022–2023). A natural question is whether the associations we document vary across these regimes.

We note that our current analysis does not split the sample by sub-period. The number of ECB monetary policy meetings per year is limited, implying that sub-period regressions would be estimated on very small samples, sharply reducing statistical power. Moreover, crisis periods coincide with changes in the ECB's communication format and the introduction of new policy instruments, so any sub-period differences would confound regime effects with institutional changes.

Nonetheless, several features of our results are informative about the time dimension. The descriptive evidence in Section 4.3 shows that social media activity and engagement spike during major policy episodes, the "whatever it takes" speech, the introduction of TLTROs and negative rates, the PEPP, and the 2022 tightening, suggesting that our disagreement measures are naturally higher during periods of elevated policy uncertainty. The large point estimates for peripheral bond yields, discussed above, are in line with the idea that these are the countries that are more exposed to shifts in ECB stance, even though the wider confidence intervals prevent us from drawing firm conclusions.

A systematic exploration of time-varying effects, potentially using interaction terms with crisis indicators or rolling-window estimates, is a natural extension that we leave for future work as the dataset expands to include additional central banks and a longer post-2023 sample.

4.4.5 Ruling Out Simple Activity Effects

To confirm that our results are not driven mechanically by higher volumes of social media activity, Figure 4.7 reports coefficient estimates from equation (4.4.1) replacing the disagreement and stance measures with basic engagement metrics: total engagement (sum of tweets, retweets, and likes) and the count of hawkish plus dovish tweets in the 15-minute pre-event window. Most coefficients are statistically indistinguishable from zero for both press releases and press conferences. This null result indicates that sheer volume of discussion does not predict the magnitude of market reactions, and that the associations documented above reflect the specific informational content captured by our disagreement and stance measures, narrative dispersion and directional tilt, rather than general social media activity levels.

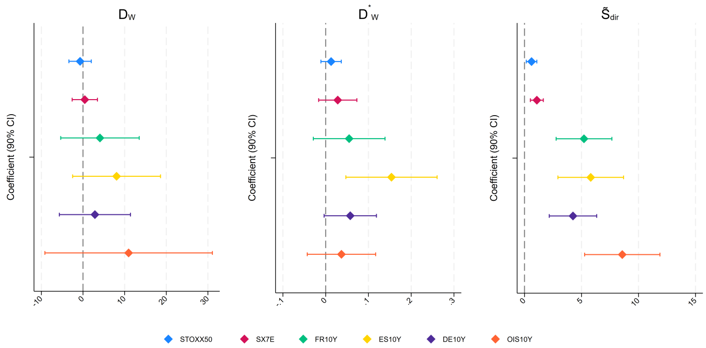

((a)) Press release

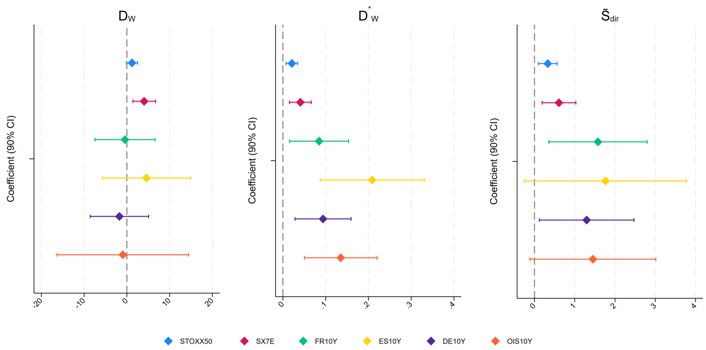

((b)) Press conference

Figure 4.6: Social media disagreement and euro area asset price returns

Notes: Coefficient estimates from equation (4.4.1): bivariate regressions of absolute asset price returns on social media measures computed in the 15 minutes before each ECB communication event. Assets include the EURO STOXX 50, EURO STOXX Banks Index, ten-year sovereign bond yields for France, Germany, and Spain, and the ten-year OIS rate. Panel (a) reports results for press releases (return window: 13:25–14:15 CET; social media window: 13:30–13:45 CET). Panel (b) reports results for press conferences (return window: 14:15–15:50 CET; social media window: 14:15–14:30 CET). $D_{\mathscr{W}}$: baseline disagreement; $D^{\star}_{\mathscr{W}}$: tweet-count weighted disagreement; $\widetilde{S}_{\text{dir}}$: participation-aware directional stance. Bars show 90% confidence intervals.

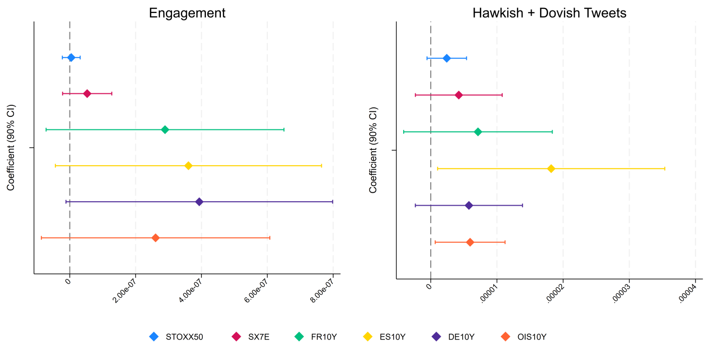

((a)) Press release

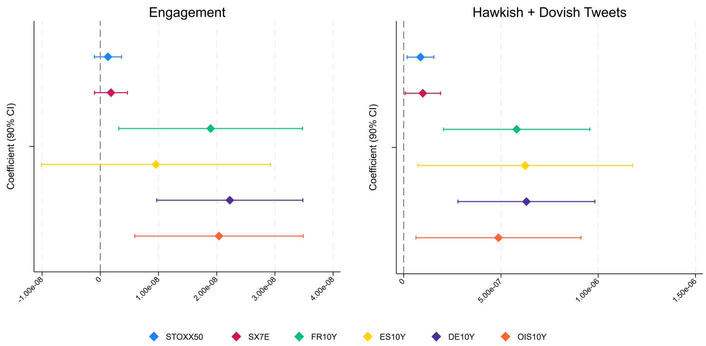

((b)) Press conference

Figure 4.7: Social media engagement and euro area asset price returns

Notes: Coefficient estimates from equation (4.4.1) with basic engagement metrics replacing the disagreement and stance measures. Total engagement is the sum of tweets, retweets, and likes; H+D count is the number of hawkish plus dovish tweets. Both measures are computed in the 15 minutes before each event. Panel (a): press releases; Panel (b): press conferences. Bars show 90% confidence intervals.

4.5 Conclusion

This paper examines how social media discussions around central bank monetary policy announcements relate to high-frequency market reactions. Using a comprehensive dataset of millions of tweets surrounding policy meetings of the BoE, ECB, and Fed, we develop novel measures of pre-decision disagreement and directional stance to test for anticipation effects in financial markets.

Our primary methodological contribution is an embedding-based disagreement measure that captures narrative dispersion in monetary policy discourse. By treating each tweet as a point in high-dimensional semantic space and computing average distances from the discussion's centroid, we quantify how tightly public narratives cohere around shared interpretations versus how widely they fragment across competing framings. This approach captures something fundamentally different from traditional sentiment or volume metrics: it measures the spread of ideas themselves, not just their tone or intensity. Combined with our hawkish–

dovish stance detection, we can simultaneously track both the magnitude of disagreement and its directional distribution along the policy spectrum.

Our empirical analysis of the ECB reveals that different social media measures capture distinct aspects of monetary policy communication at different stages. The weighted disagreement measure shows stronger associations with market volatility around press conferences, when broader participation generates more precise estimates of narrative fragmentation and when markets process the rationale, risk assessment, and forward guidance behind the decision. The directional stance measure, by contrast, performs better around press releases, when the focused debate about the policy rate decision is most relevant. These complementary patterns suggest that social media discourse contains multiple layers of information that markets process at different moments of the communication cycle.

Across asset classes, the associations are not confined to a single market segment. Equity indices show some of the strongest and most precisely estimated results, while sovereign bond yields display more heterogeneity. Peripheral yields such as Spain exhibit large point estimates, in line with the idea that these are the countries more exposed to shifts in ECB stance, though the wider confidence intervals call for caution in interpretation.

These findings have implications for central bank communication strategy. High pre-announcement disagreement signals areas where additional clarity might improve policy transmission. The distinct roles of the press release and press conference in resolving different types of uncertainty highlight the value of the ECB's dual-event communication structure. More broadly, our results suggest that monitoring social media discourse can provide central banks with a real-time gauge of communication effectiveness that complements traditional survey-based measures. As the information environment grows increasingly fragmented, understanding how public discourse forms and resolves around monetary policy decisions becomes critical for successful policy implementation.

A natural extension is to examine whether the associations we document vary across monetary policy regimes and to expand the analysis to the Bank of England and the Federal Reserve. We leave these exercises for future work as the dataset grows and additional high-frequency financial data become available.

Chapter 5

General Conclusion

This thesis presented three chapters examining how gatekeepers translate observable signals of demand into offline choices across politics, news production, and financial markets. In this concluding chapter, I summarise the main findings and outline the implications for platform design, media governance, and policy communication.

The first chapter studies how social-media virality operates as a booking signal for television producers. Using tweets by Italian MPs and a content-agnostic virality threshold, it shows that when politicians get extremely viral, they receive more next-day TV invitations and speaking time. A simple model clarifies why: producers treat virality as a noisy proxy for audience interest, and politicians anticipate this, shifting rhetoric toward attention-grabbing extremes.

The second chapter shifts from "who to invite" to "what to cover," analysing local TV news decisions around weather, a situation where underlying facts are observable. Coverage rises with event rarity ("man bites dog") but varies with audience partisanship: markets emphasise temperature deviations that fit local ideological priors. Ownership changes to a politically conservative group do not meaningfully alter these patterns, pointing to demand-driven, not supply-driven, selection.

The third chapter turns to financial markets and the timing of reactions, showing that pre-announcement social-media disagreement around central-bank meetings foreshadows larger announcement-day asset-price moves. Disagreement, capturing dispersion of narratives rather than sentiment alone, functions as an anticipation signal that markets resolve when policy information arrives.

Taken together, the three essays map complementary facets of gatekeeping: *who* is amplified (politicians selected via virality), *what* is covered (events filtered through audience-shaped newsworthiness), and *when* reactions occur (markets responding to the resolution of public disagreement). A unifying lesson is that measurable signals, engagement metrics, audience composition, and narrative dispersion, do not simply mirror demand; they feed back into editorial and policy choices.

Bibliography

Adams, Travis, Andrea Ajello, Diego Silva, and Francisco Vazquez-Grande (2023). *More than Words: Twitter Chatter and Financial Market Sentiment*. Finance and Economics Discussion Series 2023-034. Board of Governors of the Federal Reserve System (U.S.)

Allcott, Hunt and Matthew Gentzkow (2017). "Social media and fake news in the 2016 election". *Journal of economic perspectives* 31.2, pp. 211–236.

Altavilla, Carlo, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Giuseppe Ragusa (2019). "Measuring euro area monetary policy". *Journal of Monetary Economics* 108, pp. 162–179.

Altavilla, Carlo, Refet Gürkaynak, Tobias Kind, and Luc Laeven (2025). *Monetary Transmission with Frequent Policy Events*. Tech. rep. CEPR Discussion Paper DP20196, Centre for Economic Policy Research (CEPR . . .

Anand, Bharat, Rafael Di Tella, and Alexander Galetovic (2007). "Information or opinion? Media bias as product differentiation". *Journal of Economics & Management Strategy* 16.3, pp. 635–682.

Apel, Mikael and Marianna Blix Grimaldi (2014). "How Informative Are Central Bank Minutes?" *Review of Economics* 65.1, pp. 53–76.

Apel, Mikael, Marianna Blix Grimaldi, and Isaiah Hull (2019). *How Much Information Do Monetary Policy Committees Disclose? Evidence from the FOMC's Minutes and Transcripts*. Working Paper Series 381. Sveriges Riksbank.

Armona, Luis, Matthew Gentzkow, Emir Kamenica, and Jesse M Shapiro (2024). "What is Newsworthy? Theory and Evidence". NBER WP 32512.

Ash, Elliott, Anton Boltachka, Sergio Galletta, and Matteo Pinna (2023). "Media Bias and Climate Change Skepticism". Available at SSRN: http://dx.doi.org/10.2139/ssrn.4632854.

Bail, Christopher A, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky (2018). "Exposure to opposing views on social media can increase political polarization". *Proceedings of the National Academy of Sciences* 115.37, pp. 9216–9221.

Bär, Dominik, Nicolas Pröllochs, and Stefan Feuerriegel (2025). "The role of social media ads for election outcomes: Evidence from the 2021 German election". *PNAS nexus* 4.3, pgaf073.

Baron, David P (2006). "Persistent media bias". *Journal of Public Economics* 90.1-2, pp. 1–36.

Barry, Mamadou-Lamine, Brenton Joey Bruns, Sinem Kandemir, Jens Klose, Victor Smirnov, and Peter Tillmann (2025). *The Emotions of Monetary Policy*. SSRN Working Paper Series.

Bauer, Michael D and Eric T Swanson (2023). "An alternative explanation for the "fed information effect"". *American Economic Review* 113.3, pp. 664–700.

Baylis, Patrick, Nick Obradovich, Yury Kryvasheyeu, Haohui Chen, Lorenzo Coviello, Esteban Moro, Manuel Cebrian, and James H Fowler (2018). "Weather impacts expressed sentiment". *PloS one* 13.4, e0195750.

Bell, Samuel J, Mariano Coria Meglioli, Megan Richards, Eduardo Sánchez, Christophe Ropers, Skyler Wang, Adina Williams, Levent Sagun, and Marta R Costa-jussà (2024). "On the role of speech data in reducing toxicity detection bias". *arXiv preprint arXiv:2411.08135*.

Bernanke, Ben S and Kenneth N Kuttner (2005). "What explains the stock market's reaction to federal reserve policy?" *Journal of Finance* 60.3, pp. 1221–1257.

Boken, Johannes, Mirko Draca, Nicola Mastrorocco, and Arianna Ornaghi (2023). "The returns to viral media: the case of US campaign contributions".

Brady, William J, Julian A Wills, John T Jost, Joshua A Tucker, and Jay J Van Bavel (2017). "Emotion shapes the diffusion of moralized content in social networks". *Proceedings of the National Academy of Sciences* 114.28, pp. 7313–7318.

Brooks, Jeremy, Douglas Oxley, Arnold Vedlitz, Sammy Zahran, and Charles Lindsey (2014). "Abnormal Daily Temperature and Concern about Climate Change Across the United States". *Review of Policy Research* 31.3, pp. 199–217.

Buchanan, Tom, Rotem Perach, Deborah Husbands, Amber F Tout, Ekaterina Kostyuk, James Kempley, and Laura Joyner (2024). "Individual differences in sharing false political information on social media: Deliberate and accidental sharing, motivations and positive schizotypy". *Plos One* 19.6, e0304855.

Byrne, David, Robert Goodhead, Michael McMahon, Matthew Naylor, and Conor Parle (2025). "From Forward Guidance to Data Dependence: Temporality and Complexity in ECB Communication After the Pandemic". *Central Bank of Ireland, Economic Letter* Vol. 2025, No. 2.

Byrne, David, Robert Goodhead, Michael Mcmahon, and Conor Parle (2023). *The Central Bank Crystal Ball: Temporal information in monetary policy communication*. CEPR Discussion Papers 17930. C.E.P.R. Discussion Papers.

Cagé, Julia, Nicolas Hervé, and Béatrice Mazoyer (2022). "Social media and newsroom production decisions".

Camous, Antoine and Dmitry Matveev (2021). "Furor over the fed: A president's tweets and central bank independence". *CESifo Economic Studies* 67.1, pp. 106–127.

Campiglio, Emanuele, Jerome Deyris, Davide Romelli, and Ginevra Scalisi (2025). "Warning words in a warming world: Central bank communication and climate change". *European Economic Review* 178, p. 105101.

Caprini, Giulia (2023). "Does candidates' media exposure affect vote shares? Evidence from Pope breaking news". *Journal of Public Economics* 220, p. 104847.

Carey, John M, Gretchen Helmke, Brendan Nyhan, Mitchell Sanders, and Susan Stokes (2019). "Searching for bright lines in the Trump presidency". *Perspectives on Politics* 17.3, pp. 699–718.

Chen, Jianlv, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu (2024). "Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation". *arXiv preprint arXiv:2402.03216*.

Chinoy, Sahil, Nathan Nunn, Sandra Sequeira, and Stefanie Stantcheva (2025). *Zero-sum Thinking and the Roots of US Political Differences*. National Bureau of Economic Research.

Chopra, Felix, Ingar Haaland, and Christopher Roth (2024). "The demand for news: Accuracy concerns versus belief confirmation motives". *The Economic Journal* 134.661, pp. 1806–1834.

D'Amico, Stefania and Thomas B King (2023). "What does anticipated monetary policy do?" *Journal of Monetary Economics* 138, pp. 123–139.

De Kock, Christine and Andreas Vlachos (2021). "I beg to differ: A study of constructive disagreement in online conversations". *arXiv preprint arXiv:2101.10917*.

DellaVigna, Stefano and Eliana La Ferrara (2015). "Economic and social impacts of the media". *Handbook of Media Economics*. Vol. 1. Elsevier, pp. 723–768.

Demir, Vedat (2025). "Erdougan's populist rhetoric and hate speech: Anti-opposition discourse and the polarization of Turkish politics". *American Journal of Qualitative Research* 9.2, pp. 62–78.

DeVerna, Matthew R, Andrew M Guess, Adam J Berinsky, Joshua A Tucker, and John T Jost (2024). "Rumors in retweet: Ideological asymmetry in the failure to correct misinformation". *Personality and Social Psychology Bulletin* 50.1, pp. 3–17.

Di Leo, Riccardo and Catarina Midões (2023). "Good COP/Bad COP: Estimating the Impact of Global Political Initiatives on Climate Change Attitudes". *Available at SSRN 4624696*.

Djourelova, Milena, Ruben Durante, Elliot Motte, and Eleonora Patacchini (2024). "Experience, Narratives, and Climate Change Beliefs". CEPR DP 18738.

Dossani, Asad (2021). "Central Bank Tone and Currency Risk Premia". *Journal of International Money and Finance* 117, p. 102424.

Durante, Ruben and Brian Knight (2012). "Partisan control, media bias, and viewer responses: Evidence from Berlusconi's Italy". *Journal of the European Economic Association* 10.3, pp. 451–481.

Durante, Ruben, Paolo Pinotti, and Andrea Tesei (2019). "The political legacy of entertainment TV". *American Economic Review* 109.7, pp. 2497–2530.

Ehrmann, Michael and Alena Wabitsch (2022). "Central bank communication with non-experts–A road to nowhere?" *Journal of Monetary Economics* 127, pp. 69–85.

Eliaz, Kfir and Ran Spiegler (2024). "News Media as Suppliers of Narratives (and Information)". *arXiv preprint arXiv:2403.09155*.

Ely, Jeffrey, Alexander Frankel, and Emir Kamenica (2015). "Suspense and surprise". *Journal of Political Economy* 123.1, pp. 215–260.

Fadda, Pietro, Rayane Hanifi, Klodiana Istrefi, and Adrian Penalver (2025). "Central bank communication of uncertainty". *Journal of International Money and Finance* 157, p. 103385.

Fang, Qixiang, Dong Nguyen, and Daniel L Oberski (2022). "Evaluating the construct validity of text embeddings with application to survey questions". *EPJ Data Science* 11.1, p. 39.

Fetzer, Thiemo and Prashant Garg (2025). "Network determinants of cross-border media coverage of natural disasters". Working Paper.

Gambaro, Marco, Valentino Larcinese, Riccardo Puglisi, and James M Snyder Jr (2025). "The Revealed Demand for Hard Versus Soft News: Evidence from Italian TV Viewership". *The Journal of Politics* 87.4, pp. 000–000.

Gentzkow, Matthew and Jesse M Shapiro (2006). "Media bias and reputation". *Journal of Political Economy* 114.2, pp. 280–316.

— (2010). "What drives media slant? Evidence from US daily newspapers". *Econometrica* 78.1, pp. 35–71.

Gentzkow, Matthew, Jesse M Shapiro, and Daniel F Stone (2015). "Media bias in the marketplace: Theory". *Handbook of Media Economics*. Vol. 1. Elsevier, pp. 623–645.

Gervais, Bryan T (2015). "Incivility online: Affective and behavioral reactions to uncivil political posts in a web-based experiment". *Journal of Information Technology & Politics* 12.2, pp. 167–185.

Goodman-Bacon, Andrew (2021). "Difference-in-differences with variation in treatment timing". *Journal of econometrics* 225.2, pp. 254–277.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera (2023). "The Voice of Monetary Policy". *American Economic Review* 113.2, pp. 548–584.

— (2025). "Central bank communication on social media: What, to whom, and how?" *Journal of Econometrics* 249, p. 105869.

Gottfried, Jeffrey and Elisa Shearer (2017). "Americans' Online News Use is Closing in on TV News Use". Pew Research Center.

Grainger, Kimberly Cogdell (2018). "Political Rhetoric And Minority Health: Introducing The Rhetoric-Policy-Health Paradigm". *. Louis UJ Health L. & Pol'y* 12, p. 121.

Grebe, Moritz, Sinem Kandemir, and Peter Tillmann (2024a). "Geopolitics in the Boardroom: How German Managers Respond to the War in Ukraine". *Available at SSRN 5098484*.

— (2024b). "Uncertainty about the war in Ukraine: Measurement and effects on the German economy". *Journal of Economic Behavior & Organization* 217, pp. 493–506.

Gürkaynak, Refet S, Brian Sack, and Eric T Swanson (2005). "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements". *International Journal of Central Banking* 1.2, pp. 55–93.

Hagen, Sal and Tommaso Venturini (2024). "Memecry: tracing the repetition-with-variation of formulas on 4chan/pol". *Information, Communication & Society* 27.3, pp. 466–497.

Hamilton, Lawrence C and Mary D Stampone (2013). "Blowin' in the wind: Short-term weather and belief in anthropogenic climate change". *Weather, Climate, and Society* 5.2, pp. 112–119.

Hansen, Stephen, Michael McMahon, and Matthew Tong (2019). "The long-run information effect of central bank communication". *Journal of Monetary Economics* 108, pp. 185–202.

Hansen, Stephen and Michael F. McMahon (2016). "Shocking language: Understanding the macroeconomic effects of central bank communication". *Journal of International Economics* 99, S114–S133.

Hatte, Sophie, Etienne Madinier, and Ekaterina Zhuravskaya (2021). "Reading twitter in the newsroom: How social media affects traditional-media reporting of conflicts".

Hilden, David, Michael R Kistner, et al. (2025). "Promoting Conspiracy Theories Strategically". *Quarterly Journal of Political Science* 20.3, pp. 307–337.

Holmström, Bengt (1999). "Managerial incentive problems: A dynamic perspective". *The review of Economic studies* 66.1, pp. 169–182.

Howe, Peter D, Jennifer R Marlon, Matto Mildenberger, and Brittany S Shield (2019). "How will climate change shape climate opinion?" *Environmental Research Letters* 14.11, p. 113001.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood (2019). "The origins and consequences of affective polarization in the United States". *Annual review of political science* 22.1, pp. 129–146.

Iyengar, Shanto and Sean J Westwood (2015). "Fear and loathing across party lines: New evidence on group polarization". *American journal of political science* 59.3, pp. 690–707.

Jarociński, Marek and Peter Karadi (2020). "Deconstructing monetary policy surprises—The role of information shocks". *American Economic Journal: Macroeconomics* 12.2, pp. 1–43.

Jensema, Carl, Ralph McCann, and Scott Ramsey (1996). "Closed-captioned television presentation speed and vocabulary". *American Annals of the Deaf*, pp. 284–292.

Joireman, Jeff, Heather Barnes Truelove, and Blythe Duell (2010). "Effect of outdoor temperature, heat primes and anchoring on belief in global warming". *Journal of Environmental Psychology* 30.4, pp. 358–367.

Kemp, Jeremy, Richard Milne, and Dave S Reay (2010). "Sceptics and deniers of climate change not to be confused". *Nature* 464.7289, pp. 673–673.

Kowal, Michael (2023). "The value of a like: Facebook, viral posts, and campaign finance in US Congressional elections". *Media and Communication* 11.3, pp. 153–163.

Kuttner, Kenneth (2001). "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market". *Journal of Monetary Economics* 47.3, pp. 523–544.

Larcinese, Valentino, Riccardo Puglisi, and James M Snyder Jr (2011). "Partisan bias in economic news: Evidence on the agenda-setting behavior of US newspapers". *Journal of Public Economics* 95.9-10, pp. 1178–1189.

Levitsky, Steven and Daniel Ziblatt (2024). "How democracies die". *Ideals and Ideologies*. Routledge, pp. 73–80.

Mamakos, Michalis and Eli J Finkel (2023). "The social media discourse of engaged partisans is toxic even when politics are irrelevant". *PNAS nexus* 2.10, pgad325.

Mangani, Andrea and Elisa Tarrini (2018). "Social pluralism in public and private television broadcasting". *Javnost-The Public* 25.3, pp. 282–297.

Martin, Gregory J, Nicola Mastrorocco, Joshua McCrain, and Arianna Ornaghi (2024). "Media Consolidation".

Martin, Gregory J and Joshua McCrain (2019). "Local news and national politics". *American Political Science Review* 113.2, pp. 372–384.

Masciandaro, Donato, Oana Peia, and Davide Romelli (2024). "Central bank communication and social media: From silence to Twitter". *Journal of Economic Surveys* 38.2, pp. 365–388.

Masciandaro, Donato, Davide Romelli, and Gaia Rubera (2023). "Monetary policy and financial markets: evidence from Twitter traffic". *Trinity Economic Papers (TEP) Working Paper* 1023.

Mason, Lilliana (2018). *Uncivil agreement: How politics became our identity*. University of Chicago Press.

Mastrorocco, Nicola and Arianna Ornaghi (Forthcoming). "Who watches the watchmen? Local news and police behavior in the United States".

Matsa, Katerina Eva (2018). "Fewer Americans Rely on TV News; What Type They Watch Varies by Who They Are". Pew Research Center.

Mendelberg, Tali (2001). *The race card: Campaign strategy, implicit messages, and the norm of equality*. Princeton University Press.

Miho, Antonela (2024). "Small screen, big echo? Political persuasion of local TV news: evidence from Sinclair". Available at SSRN: http://dx.doi.org/10.2139/ssrn.3266755.

Miranda-Agrippino, Silvia and Giovanni Ricco (2021). "The transmission of monetary policy shocks". *American Economic Journal: Macroeconomics* 13.3, pp. 74–107.

Moore, Frances C, Nick Obradovich, Flavio Lehner, and Patrick Baylis (2019). "Rapidly declining remarkability of temperature anomalies may obscure public perception of climate change". *Proceedings of the National Academy of Sciences* 116.11, pp. 4905–4910.

Moschella, Manuela and Luca Pinto (2019). "Central banks' communication as reputation management: How the Fed talks under uncertainty". *Public administration* 97.3, pp. 513–529.

Moskowitz, Daniel (2021). "Local News, Information, and the Nationalization of U.S. Elections". *American Political Science Review* 115.1, pp. 114–129.

Mudde, Cas and Cristóbal Rovira Kaltwasser (2017). *Populism: A very short introduction*. Oxford University Press.

Mullainathan, Sendhil and Andrei Shleifer (2005). "The market for news". *American Economic Review* 95.4, pp. 1031–1053.

Nakamura, Emi and Jón Steinsson (2018). "High-frequency identification of monetary non-neutrality: the information effect". *The Quarterly Journal of Economics* 133.3, pp. 1283–1330.

Newman, Nic, A Ross Arguedas, Craig T Robertson, Rasmus Kleis Nielsen, and Richard Fletcher (2025). *Digital news report 2025*. Reuters Institute for the study of Journalism.

Nielsen (2019). "Local Reference Supplement 2019-2020".

Norris, Pippa (2019). "Do perceptions of electoral malpractice undermine democratic satisfaction? The US in comparative perspective". *International Political Science Review* 40.1, pp. 5–22.

Papper, Bob (2017). "Local News by the Numbers". Radio Television Digital News Assocation Research Reports.

Parle, Conor (2022). “The financial market impact of ECB monetary policy press conferences — A text-based approach”. *European Journal of Political Economy* 74.

Petrova, Maria, Ananya Sen, and Pinar Yildirim (2021). “Social media and political contributions: The impact of new technology on political competition”. *Management Science* 67.5, pp. 2997–3021.

Pianta, Silvia and Matthew R Sisco (2020). “A hot topic in hot times: how media coverage of climate change is affected by temperature abnormalities”. *Environmental Research Letters* 15.11, p. 114038.

Picault, Matthieu and Thomas Renault (2017). “Words are not all created equal: A new measure of ECB communication”. *Journal of International Money and Finance* 79, pp. 136–156.

Puglisi, Riccardo and James M Snyder Jr (2011). “Newspaper coverage of political scandals”. *The Journal of Politics* 73.3, pp. 931–950.

Rathje, Steve, Jay J Van Bavel, and Sander Van Der Linden (2021). “Out-group animosity drives engagement on social media”. *Proceedings of the national academy of sciences* 118.26, e2024292118.

Roets, Arne et al. (2019). “Is there an ideological asymmetry in the moral approval of spreading misinformation by politicians?” *Personality and Individual Differences* 143, pp. 165–169.

Romer, Christina D and David H Romer (2004). “A new measure of monetary shocks: Derivation and implications”. *American Economic Review* 94.4, pp. 1055–1084.

Sen, Ananya and Pinar Yildirim (2015). “Clicks bias in editorial decisions: How does popularity shape online news coverage?” *Available at SSRN 2619440.*

Strömberg, David (2004). “Mass media competition, political competition, and public policy”. *The Review of Economic Studies* 71.1, pp. 265–284.

Sun, Liyang and Sarah Abraham (2021). “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects”. *Journal of econometrics* 225.2, pp. 175–199.

Swanson, Eric T (2023). “The importance of fed chair speeches as a monetary policy tool”. *AEA papers and proceedings*. Vol. 113. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 394–400.

Tajfel, Henri, John Turner, William G Austin, and Stephen Worchel (2001). “An integrative theory of intergroup conflict”. *Intergroup relations: Essential readings*, pp. 94–109.

Theocharis, Yannis, Pablo Barberá, Zoltán Fazekas, Sebastian Adrian Popa, and Olivier Parnet (2016). “A bad workman blames his tweets: The consequences of citizens’ uncivil Twitter use when interacting with party candidates”. *Journal of communication* 66.6, pp. 1007–1031.

Tillmann, Peter and Andreas Walter (2019). “The effect of diverging communication: The case of the ECB and the Bundesbank”. *Economics Letters* 176, pp. 68–74.

Tobback, Ellen, Stefano Nardelli, and David Martens (2017). “Between hawks and doves: measuring central bank communication”. *European Central Bank, Working Paper No. 2085.*

Van Erkel, Patrick FA and Peter Van Aelst (2021). "Why don't we learn from social media? Studying effects of and mechanisms behind social media news use on general surveillance political knowledge". *Political Communication* 38.4, pp. 407–425.

Vosoughi, Soroush, Deb Roy, and Sinan Aral (2018). "The spread of true and false news online". *science* 359.6380, pp. 1146–1151.

Washington, Haydn (2013). *Climate change denial: Heads in the sand*. Routledge.

Wasow, Omar (2020). "Agenda seeding: How 1960s black protests moved elites, public opinion and voting". *American Political Science Review* 114.3, pp. 638–659.

Wolton, Stephane (2019). "Are biased media bad for democracy?" *American Journal of Political Science* 63.3, pp. 548–562.

Wu, Tim (2017). *The attention merchants: The epic scramble to get inside our heads*. Vintage.

27670.4cm

This thesis was typeset using the LaTeX typesetting system created by Leslie Lamport and Donald Knuth, and the memoir class. The body text is set 11pt with Palatino designed by Hermann Zapf, which includes italics and small caps. Other fonts include `Monospace` from Young Ryu's TX Fonts family, Sans from TeX User Group Poland's Latin Modern family, and Helvetica by Max Miedinger and Eduard Hoffmann.