

Successful Rhetorics: How Do Linguistic Dimensions Affect User Engagement with Different News Categories on Twitter?

Alireza Mohammadinodooshan, Niklas Carlsson

Linköping University, Sweden

{alireza.mohammadinodooshan, niklas.carlsson}@liu.se

Abstract

This paper analyzes how different rhetorical attributes in news tweets, specifically analytical, clout, perceptual, and risk language, influence user engagement across publishers with different bias and reliability ratings. Using the LIWC framework to quantify these linguistic dimensions in a 5.5 million tweets dataset covering 1,553 news publishers and capturing over 480 million tweet interactions, we perform and present a category-based analysis that captures the relative impact that such features have on the user engagement rates associated with different political bias and reliability categories. While highly biased and unreliable publishers saw increased engagement for clout and risk language, confirming audience biases, the least biased ones benefited more from analytical language. Perception language, on the other hand, uniformly reduced engagement. These insights not only further our understanding of persuasion tactics but also have implications for curbing misinformation by aligning recommendations with audience veracity and impartiality preferences.

1 Introduction

In the evolving landscape of news consumption, Twitter has emerged as a primary source, with 53% of its users relying on it for their news, surpassing other platforms like Facebook (43%), Reddit (38%), YouTube (32%), and Instagram (34%), according to a 2023 Pew Research Center study (Walker and Matsa 2023). However, the reliability and bias of news on Twitter vary, necessitating a closer examination of the factors that influence users' exposure and engagement with different content.

Previous studies have highlighted the pivotal role of user engagement in shaping content visibility (Ksiazek, Peer, and Lessard 2016), a sentiment echoed by Twitter's recent algorithm release (Twitter 2023), and examined factors affecting news engagement, including emotions (Molina et al. 2023), negativity (Kumar et al. 2018), cognitive inhibition (Bronstein et al. 2021), and analytical thinking (Pennycook and Rand 2020). However, the impact of the rhetorical attributes of the news posts themselves is less understood.

To address this void, in contrast to these works, we focus specifically on how four linguistic dimensions affect user engagement with different categories of news.

The four studied dimensions are analytic, clout, perceptual, and risk language. Encapsulated by the Linguistic Inquiry and Word Count (LIWC) framework, used by many prior works (Jaidka et al. 2020; Khalid and Srinivasan 2020; Beel et al. 2022), these attributes (see Sec. 3) represent credibility, authority projection, sensory grounding, and emotional provocation, respectively. For example, analytic language reflects logic, evidence, and critical analysis, all conveying credibility, appealing to rationality. Clout language denotes confidence and expertise for projecting credibility. Perceptual language grounds content in sensory observations rather than abstractions. Finally, risk language plays to emotions using defiant, shocking rhetoric, potentially increasing visceral appeal. Due to their important psychological effects, each of these dimensions has been the focus of other studies (Markowitz 2023; Mahmud, Chen, and Nichols 2014; Catherine A. Cherrstrom and Sherron 2023). However, no prior work has studied their effect on user engagement with different categories of news on social media platforms.

Motivated by the need for a nuanced understanding of how these linguistic attributes engage diverse ideological audiences, we address the following Research Question (RQ):

Within audiences with different political bias and reliability preferences, how does the usage of certain forms of rhetoric shift engagement rates compared to baseline levels without those attributes?

By engagement rate, here we refer to the number of interactions divided by impressions (views). Capitalizing on Twitter's recent release of impressions statistics for all tweets (not available until Dec. 2022), we first compiled a unique and comprehensive dataset covering 5.5 million tweets (responsible for 80.6 billion impressions and 483 million interactions) from 1,553 labelled U.S. news publishers over a six-month period. Using this dataset, which we annotate with LIWC rhetorical attributes, we then investigate how analytic, clout, perceptual, and risk language dimensions differentially engage audiences from distinct bias and reliability classes. Furthermore, by differentiating between *Misinformation* and *Non-Misinformation* audiences and adopting a granular five-tier bias taxonomy—*Left*, *Left-Center*, *Least Biased*, *Right-Center*, and *Right*—we capture several interesting dynamic effects in the modern media landscape, including the differentiated effects based on the reliability of the news sources, polarization effects based on both the level

of bias and the political alignment of the publisher.

Our analysis uncovers new insights, challenging common assumptions about the influence of rhetorical dimensions. For example, we find that analytical language typically reduces engagement, especially among more partisan groups. Conversely, inflammatory risk language disproportionately boosts engagement for unreliable and biased publishers. These findings have theoretical implications for understanding persuasive forces and practical applications for improving content legitimacy and balance. Beneficiaries include policymakers crafting anti-misinformation regulations, platform designers refining algorithms, journalistic institutions optimizing content strategies, and behavioural scientists unraveling drivers of political polarization and engagement in modern democracies.

This paper is organized as follows. The first two sections describe related work (Sec. 2) and the selected language dimensions (Sec. 3). The following sections describe our dataset (Sec. 4) and detail the metrics and statistical tests used to measure and compare distribution shifts in engagement rates (Sec. 5), before Sec. 6 presents our analysis and key insights. Finally, we present our conclusions (Sec. 7).

2 Related Works

While prior research has explored engagement differences across platforms and information types, no prior work has investigated how engagement with different classes of news is affected by the rhetoric of social media posts.

Engagement with Different News Classes: Several studies have compared interaction levels on mainstream, partisan, and unreliable content on different platforms, including on Facebook (Edelson et al. 2021; Hiaeshutter-Rice and Weeks 2021), Twitter (Spinde et al. 2023; Aldous, An, and Jansen 2022; Mohammadinooshan and Carlsson 2024; Vosoughi, Roy, and Aral 2018), and Reddit (Weld, Glenzki, and Althoff 2021). As an example, Weld et al. (2021) show that Reddit is more resilient to low factual content than Twitter; e.g., receiving 20% fewer upvotes for extremely biased and low factual content. Our work differs from this line of research as we focus on how rhetorical tactics encoded through analytic, clout, perceptual, and risk language dimensions affect user engagement across different bias and reliability audience categories.

Factors Impacting User Engagement: Other researchers have focused on factors affecting user engagement. For example, focusing on Twitter, Salehabadi et al. (2022) and Beknazari et al. (2022) show how toxicity in tweets can affect user engagement, and Antypas et al. (2023) show the positive effect of negativity on the virality and spread of tweets, with similar results reported on TikTok (Cheng and Li 2023) and Facebook (Rathje, Bavel, and van der Linden 2021), while some earlier research have reported a reverse trend (Trilling, Tolochko, and Burscher 2017) in that they observe that upbeat material tends to be more widely shared than negative material. Among the most closely related to our work are the works by Bellovary et al. (2021) and Aldous et al. (2019). Using tweets from 44 news organizations, Bellovary et al. show that negativity impact both the left- and right-party affiliated publishers, whereas

Aldous et al., using data from 53 news organizations, study the impact of topics on engagement with news posts. Neither of these works considers the rhetoric of the posts or the reliability of the publishers.

Impact of Rhetorics: Others have explored the rhetorical impact of LIWC features on engagement rates. For example, Mahmud et al. (2014) considered multiple dimensions, including perception, finding a significant (similar to our work) but positive effect on user engagement. Comparing with the consistently negative impact of perception in our context (c.f. Figure 7) suggests genre-specific influences (news-specific in our case). Lee et al. (2023) demonstrate the impact of anxiety, anger, and informal language from the LIWC dimensions on engagement with misinformation tweets. Others have highlighted the positive effect of the certainty language LIWC dimension on consumer engagement (Pezzuti, Leonhardt, and Warren 2021) or the usage effect of different pronouns (Hu, Farnham, and Talamadupula 2021).

Getting back to the news domain, Robertson et al. (2023) use the LIWC sentiment analysis to study the effect of negative language on news consumption. Somory et al. (2020) study the effect of different LIWC dimensions, including sentiment, use of pronouns, social words, cognitive mechanisms, perception, and time. Aldous et al. (2022) study the effect of nine emotions (anger, anticipation, anxiety, disgust, joy, fear, sadness, surprise, trust) extracted using LIWC on 133,487 posts by eight news organizations. Candia et al. (2022) study the link between the moral language dimension of LIWC and social news engagement. None of the above works have considered the impact of the four rhetorical features in our study on different (bias and reliability) classes of news posts.

Our research differs from prior studies (aggregate rhetorical effects) by considering how different groups of audiences (by controlling the impressions variable) from different bias and reliability preferences (differential responses) engage with four important (as we discuss in the next section), yet, underrepresented LIWC dimensions.

3 LIWC and Selected Language Dimensions

LIWC: Lexicon-based analysis is a well-established method for quantifying psychological and linguistic attributes in text. Among existing lexicons, the Linguistic Inquiry and Word Count (LIWC) framework stands out for its alignment with psychological and sociological theories (Tausczik and Pennebaker 2010), as well as its extensive validation and widespread use in prior ICWSM research (Jaidka et al. 2020; Kim et al. 2021; Beel et al. 2022; Haworth et al. 2021). In this study, we use the recently released LIWC2022 version, which includes a rebuilt processing engine and an expanded dictionary better suited to informal netspeak and social media language (Boyd et al. 2022b).

Dimension Selection: LIWC supports two groups of linguistic dimensions: (1) the standard dimensions, which provide the keywords on one specific psychological dimension of the text (e.g., sad word groups), and (2) summary dimensions, which is a broader measure of a particular linguistic or psychological construct. Due to space constraints,

we chose to focus on the two dimensions in each group that we deemed most relevant to the language used in news posts associated with different news categories: *analytic* and *clout* from the summary dimensions, and *perception* and *risk* from the standard dimensions. While this set differs from those studied (in other contexts) by prior work, the critical importance of studying these dimensions is supported by classic theoretical frameworks such as Media Richness Theory (Daft and Lengel 1986) and Social Presence Theory (Short, Williams, and Christie 1976).

1) Analytic: This summary dimension reflects logical and formal thinking, aligns well with the cognitive dimension of communication, and is of special interest in environments like news engagement, where the accuracy and depth of content are valued. For example, the Elaboration Likelihood Model (ELM) (Petty and Briñol 2011) suggests that analytical language facilitates the central route of processing, where audiences are more likely to engage deeply with the content, leading to more stable attitude changes. Here, we study to what extent this applies to social media news engagement and (more broadly) shed light on how the presentation of logical arguments and information affects user engagement with different news on Twitter.

2) Clout: This summary dimension encapsulates language that conveys expertise, confidence, and leadership. Being indicative of expertise and authority, it resonates with the concept of authority bias (Cialdini 2006), which refers to the tendency of individuals to attribute greater accuracy and credibility to the opinions of authority figures. In the context of social media news, clout language can trigger this bias, leading users to engage more with content that demonstrates authority and expertise. When news sources use language that conveys confidence, leadership, and knowledge, users may perceive the information as more trustworthy and valuable, even without critically evaluating the content itself.

3) Perception: This standard dimension combines features, aggregating sensory descriptors pertaining to visual, spatial, kinesthetic, and auditory observations. With media richness theory arguing that richer media, capable of conveying more sensory information, are more effective for complex message delivery, we ask whether perception language, which often includes sensory details (making news stories more vivid and engaging) potentially lead to higher reader involvement (e.g., due to emotional response).

4) Risk: This standard dimension measures defiance of norms and conventions through daring, shocking, rule-breaking rhetoric designed to grab attention. Previous research supports the fact that risk language, which often incorporates elements of danger or uncertainty, can significantly heighten engagement with news content as audiences are drawn to assess threats and risks (Rozin and Royzman 2001). Accordingly, we wanted to measure this effect among different groups of news followers.

4 Dataset Compilation

We start by outlining our multi-step methodology used for data collection and labelling.

Publishers Selection (Step 1): We selected to use the list of U.S publishers provided by Media Bias Fact

Check (MBFC 2023), as gathered in Feb. 2023. To focus on the original publishers, we removed any news aggregation websites (such as alternativenews.com) from the list by examining the description page provided by MBFC for each publisher. After this exclusion, the dataset consisted of 4,109 original news publishers.

While there are other independent organizations assessing and labeling news publishers, including NewsGuard (NewsGuard 2023), Adfontes Media (Ad Fontes Media 2023), and Allsides (AllSides 2023), we have chosen to primarily use MBFC for the following reasons. (1) MBFC is used by numerous academic studies (Sharma, Ferrara, and Liu 2022; Papadogiannakis et al. 2023). (2) The more than 4K news publishers in the dataset allow a reliable statistical analysis. (3) Both the bias and the reliability dimensions are evaluated. (4) Ratings are up-to-date and freely accessible, allowing others to more easily and accurately replicate our results. (5) Employ a transparent methodology for their labeling process, detailed in (MBFC 2023). (6) Lin et al. (2023) have recently demonstrated a strong correlation between their rating and other ratings, including NewsGuard (which is not free). As a delimitation, our study does not include “social media only” news channels that might elude traditional media watchdog groups, as we could not find a comprehensive list of these outlets. This limitation may result in an underrepresentation of the prevalence of synthetic or algorithmically generated news articles.

Augmenting Bias Labels (Step 2): Like many prior works (Edelson et al. 2021; Huszár et al. 2022) we utilize the *Bias* labels typically provided by MBFC (from left-to-right): *Left*, *Left-Center*, *Least Biased*, *Right-Center*, and *Right*.

For publishers already having one of these labels as their main bias label, we utilized the corresponding label. However, some publishers have one of the following main labels: “Pro-Science”, “Conspiracy-Pseudoscience”, and “Questionable Source”. In these cases, similar to other works (CITAP 2022), we use one of the five main bias labels if such label is included in the descriptive text, tags, or bias-meter icons for the publishers (manually identified by visiting each publisher’s MBFC page). For the 304 publishers without explicit bias classifications, we utilized Robertson et al.’s (2018) bias scores alongside MBFC labels to construct non-parametric Kernel Density Estimator (KDE) models for each of the five bias groups. Gaussian kernel functions were fitted, and the bandwidth hyperparameter was optimized using Silverman’s rule and 5-fold cross-validation across 50 values. The tuned KDEs were employed to infer bias labels for unclassified publishers. For the remaining cases, the same approach was applied using (Ad Fontes Media 2023) bias scores. The efficacy of KDEs in capturing bias stratification (with an average 0.78 macro-averaged F1-score over the cross validations), is demonstrated in Figure 1(a), with 112 publishers excluded due to uninference bias labels.

Augmenting (Mis)information Labels (Step 3): The iffy index (Josef Verbanac 2023), widely used in scholarly studies (Pieri et al. 2023; Hanley, Kumar, and Durumeric 2023; Broniatowski et al. 2023), is a prominent method for assigning reliability labels to publishers. Leveraging this index, which relies on MBFC labels, we adopted a consistent




Figure 1: Tuned KDEs. Bias scores by Adfontes and reliability scores by Lin et al. (2023).

approach in our study. Given the gap between their latest update and our Feb. 2023 MBFC data collection, we replicated their methodology. Publishers marked as “Conspiracy/Pseudoscience” and “Questionable Source/Fake News” with a “Low Credibility” rating by MBFC were classified as misinformation (“iffy”). For publishers lacking some of these labels, we employed a training KDE approach on reliability (PC1) values from (Lin et al. 2023) to infer their reliability class. Figure 1(b) visualizes these KDEs.

Twitter Accounts Mapping (Step 4): We next identified publishers Twitter accounts (mainly by visiting their website) and discarded 407 publishers without a Twitter account.

Tweets, Engagement and Followers Data Collection (Step 5): We recorded the number of followers for each account in our publisher set as of Mar. 2023. Next, we excluded accounts with less than 10K followers. This decision was made because these smaller accounts have disproportionately skewed statistics and generally generate minimal interest compared to the “average” account (with 562,663 followers). For the remaining 1,553 accounts, we collected all tweets from Dec. 15, 2022 to June 2023 and their impressions and interactions statistics. Dec. 15, 2022, was selected as starting point since it is the earliest date from which tweets are accompanied by view count statistics (crucial metric for our analysis), and we were able to collect data for all publishers to June 2023, providing us a comprehensive dataset spanning six months for all publishers. Throughout this period, the impressions and interactions statistics for each tweet were collected simultaneously, ensuring a consistent and synchronized dataset for our analysis. By ensuring that there was at least one month delay between the posting dates and the retrieval of data for each tweet, we ensure that most totals have converged. Prior research has also shown that 95% of tweets cease to receive additional impressions after 24 hours of posting (Pfeffer, Matter, and Sargsyan 2023).

Table 1 provides a statistical summary of our final dataset, categorized based on bias and reliability categories. For each category, the table shows: (1) the number of outlets (N), (2) the aggregate number of tweets in the dataset (Twts.), (3) the combined sum of all interaction types in the dataset (Ints.), including likes, retweets, quotes, and replies, and (4) the number of impressions recorded in each category (Imprs.). Overall, our dataset consists of 483 million interactions from 5.5 million tweets across 1,553 news outlets, which are followed in total by 874 million users. This significant level of engagement highlights the extensive reach and influence of

	Class	N	Twts.	Ints.	Impr.
Bias	Left	200	464.4K	126.2M	15.3B
	Left-Center	408	1.3M	84.4M	28.3B
	Least Biased	582	2.4M	27.4M	8.2B
	Right-Center	170	781.5K	38.4M	9.8B
	Right	193	497.8K	206.5M	19.1B
Rel.	Non-misinfo.	1,419	5.1M	266.9M	61.4B
	Misinfo.	134	343.7K	216.1M	19.2B
Total		1,553	5.5M	483.0M	80.6B

Table 1: Dataset summary split per bias and reliability type.

news outlets on Twitter and underscores their crucial role in the news consumption landscape and dissemination on social media. Finally, considering all the publishers (last row), we note an overall engagement rate of 0.6% ($\frac{483M}{80.6B}$), near to what Twitter’s CEO (Musk 2022) has reported.

Compiling the LIWC Values (Step 6): Using LIWC2022, for all tweets in our dataset, we then computed the LIWC features for the four dimensions of interest.

The code and supplementary materials of this study are public at: <https://github.com/alireza-mon/liwc-news>.

5 Metrics and Statistical Tests

5.1 Measuring Publisher-Level Effects

To assess the influence of each of the LIWC-determined linguistic features on user engagement of different publishers, it is important to quantify variations in the engagement rates related to each specific LIWC feature. For example, when considering the effects of using analytical language for a publisher such as @nytimes, one needs to quantitatively compare the engagement distribution of the subset of @nytimes tweets that exhibit higher levels of analytical language with the overall engagement distribution of @nytimes tweets.

Challenges with Related Metrics: While various methods, such as Cumulative Distribution Function (CDF)-based or information theory-based approaches, exist for comparing distributions, they each present limitations in our context. Traditional methods like the Kolmogorov-Smirnov test lack directionality, hindering the interpretation of linguistic feature impact on engagement. Information theory-based approaches, such as Jensen-Shannon divergence, lack intuitive interpretability. Additionally, tests like the Mann-Whitney U test are sensitive to the range of values, potentially biasing results across publishers with varying engagement rates. Emphasizing maximum differences, as seen in, for example, the Kolmogorov-Smirnov test, may overlook other significant aspects of distribution differences. Moreover, methods like the Paired t-test may not be suitable for the long-tail nature of social media engagement data.

Median Percentile Shift (MPS): To overcome the above challenges, we define and utilize “Median Percentile Shift” (MPS) as our chosen metric. MPS is direction-sensitive, intuitively interpretable, less influenced by the range of original values, and emphasizes central tendencies over ex-




Figure 2: Distribution of the engagement rates for all tweets published by @nytimes and @foxnews and their analytical tweets. We limit the x-range to provide a better resolution.

treme differences. Its applicability to long-tailed distributions makes it suitable for equitable comparisons across publishers, irrespective of their baseline engagement rates.

Consider the illustrative example in Figure 2. Here, solid lines show the overall user engagement rate distributions across all tweets associated with @nytimes and @foxnews, and the dashed lines show the corresponding distributions over only the analytical tweets. By analytical tweets, we refer to the ones that have an analytic score (according to LIWC) higher than the median in our dataset. Interestingly, while @nytimes sees right-shifted distributions signifying engagement boosts, @foxnews exhibits left-shifts denoting declines. This polarization underscores analytic rhetoric’s fickle relationship with audience engagement.

Now, let us try to quantify this shift, marked as a black triangle in the figure, for @nytimes. Here, we first determine the median engagement rate of @nytimes tweets classified within the *high-analytic* feature subset. This median is then located within the total engagement distribution for @nytimes, culminating in a percentile rank. Here, the median of the high-analytic tweets falls at the 55.8th percentile of the total distribution. Finally, to calculate the *MPS* for @nytimes concerning the analytic feature ($MPS_{@nytimes}^{analytic}$) we report the difference from the median percentile (50%), resulting in an *MPS* of 5.8. In simpler terms, $MPS_{@nytimes}^{analytic} = 5.8$.

More formally, for each publisher (twitter account) a and LIWC feature F , we calculate the *MPS* as follows. (1) Let S_a^F represent the engagement rates of publisher a ’s tweets scoring high on feature F . The rest are considered in the low category. (2) Let the median of S_a^F be denoted as $M_{S_a^F}$. (3) Let S_a denote the full set of engagement rates for publisher a (encompassing both the high and low groups). Finally, the “Median Percentile Shift” for feature F for publisher a is calculated as: $MPS_a^F = (P_{S_a}(M_{S_a^F}) - 50)$, where $P_{S_a}(M_{S_a^F})$ is the percentile rank of $M_{S_a^F}$ within S_a .

By design, *MPS* values always fall within a normalized range between $\pm 50\%$, regardless of the original magnitude of the engagement rate, making it agnostic to the original engagement levels. This facilitates equitable comparison across publishers with vastly different engagement rates. Positive *MPS* values indicate increased engagement relative

to the publisher’s baseline for the feature (e.g., analytic language), while negative values indicate reduced engagement.

Threshold Selection and Validation: To define high feature scores, we adopt a robust measure of central tendency. While LIWC norms exist for general tweet domains (Boyd et al. 2022a), they are not tailored to the news-specific context of our dataset. Given the skewed distribution of engagement rates, we use the median (rather than the mean) as a more representative threshold. Therefore, tweets with scores above the median for feature F are classified as high in that feature. This choice helps mitigate the influence of outliers, common in skewed distributions. We also validated this approach by testing alternative thresholds (e.g., 50–70th percentiles in 5% steps) and the results remained consistent. We also validated the *MPS* approach by varying percentile thresholds (Appendix A), observing consistent effect directions. Robustness checks across topics and time (Appendix B) further confirmed the stability of our findings.

5.2 Statistical Tests

In social media, including for our *MPS* distributions, pronounced skewness and deviation from normality, including long tails, are common. Therefore, to compare distribution locations across multiple groups, we leveraged the non-parametric Kruskal-Wallis test, which quantifies discrepancy in median (and also distribution) trends. Throughout the paper, when comparing a set of distributions, we use a significance threshold of 0.01 and conduct post-hoc Dunn tests on all pairwise couplings in the cases where the omnibus Kruskal-Wallis metric flagged significant differences. Following the same reasoning (i.e., accounting for non-normal, skewed distributions), for the case when we compare distributions against a constant (baseline) value, we use the Wilcoxon signed rank test.

6 Results and Discussions

Using the *MPS* metric, we next analyze the effect of each LIWC feature on the user engagement rates.

6.1 Analytic Language

The analytic dimension within the LIWC framework provides a lens through which cognitive processes in communication can be examined. This dimension focuses on the presence of language indicative of higher-order thinking, including aspects like analysis, reasoning, and use of evidence. High analytical scores are typically reflective of texts that not only delve into a subject with depth but also demonstrate intellectual rigour in their structure and content.

The Effect of Analytic Language on Engagement Rates: To see how users engage with analytical language in tweets, Figure 3 shows the distribution of the Median Percentile Shift (*MPS*) values $D_{category}^{analytic}$ for different categories of publishers (e.g., *Left*), where $D_{category}^{analytic}$ consists of the set of $\{MPS_{a_i}^{analytic}\}$ for all publishers (accounts) a_i belonging to the category of interest. Here, we include results for all five *bias* categories, the two *reliability* categories, and for the full set of publishers calculated across *All* publishers (irrespective of their bias and reliability category).




Figure 3: Analytical language effect on engagement rates: negative values indicate that analytical posts inversely affect engagement rates.

Throughout the paper, each class is shown as a boxplot, with horizontal markers indicating the median MPS values, boxes representing the interquartile ranges, and whiskers extending to the 20th and 80th percentiles. Outliers beyond these percentiles are omitted to enhance the visual resolution and a reference line at the zero mark is included to aid in a comparative analysis of the distributions. In general, negative values in the distributions of each class (lower than this baseline) show the negative effect of the analytical language on the tweets of the publishers belonging to that class.

Several observations can be drawn from Figure 3. First, focusing on the rightmost boxplot, showing the distribution of MPS over all 1,553 publishers when considering the usage of the analytical language feature ($\mathcal{D}_{\text{All}}^{\text{analytic}}$), we observe a clear negative effect on user engagement rates when analytic language is used. Although this distribution is biased by the number of samples in each bias and reliability category (e.g., the *Non-misinformation* has more samples than the *Misinformation* class) it offers insights into the news domain landscape without selection bias in publisher choice.

Furthermore, the central tendency of $\mathcal{D}_{\text{All}}^{\text{analytic}}$ is notably below zero, with a median of -0.9 . The negative median is statistically significant, confirmed by a p-value of 2.9×10^{-42} using the Wilcoxon signed-rank test. In fact, this negative effect of using analytical language on the users engagement is statistically supported for all categories considered, irrespective of bias or reliability category. The largest (least significant) p-value here belongs to the *Left-Center* group which has a p-value of 2.4×10^{-6} , which also has the smallest absolute deviations from 0 (median of -0.4). The largest deviations belong to the *Right* and *Misinformation* classes, with medians of -2.2 and -1.8 , respectively. Finally, we should mention that the difference between the left party (combining the two left classes) and the right party (combining the two right classes) is also significant at the p-value of 1.4×10^{-3} according to the Kruskal-Wallis test, indicating that right-aligned publishers (as an aggregate) see relatively lower user engagement when using analytic language.




Figure 4: Percentage of tweets with analytical language content across various categories.

Key Observation: Analytical language negatively impacts user engagement rate across all classes, with the *Right* and *Misinformation* groups exhibiting the strongest declines, and the *Left-Center* demonstrating the least pronounced effect.

Need for Studying the Prevalence: To better understand the pronounced negative impact of analytical language on engagement, especially pronounced within the *Right* class, for example, we must explore how frequently it is used. The scarcity principle, derived from the “novelty effect” theory in communications, suggests that users may exhibit heightened responsiveness to content that is less prevalent within their typical information ecosystem. To further our understanding of the interplay between this linguistic style, its prevalence, and audience interaction, and to distinguish the potential impact of the “novelty effect” from other engagement drivers, we next analyze the prevalence of analytical language. If a type of content is already prevalent among a certain class of publishers, its (positive) impact on engagement might be less attributed to its novelty and more to other factors. The inverse case also holds.

Prevalence of Analytical Language: Figure 4 shows the distribution of (highly) analytic tweet percentages across different publisher classes. These distributions are obtained by calculating, for each publisher, the percentage of tweets classified as using analytic language, and then aggregating these percentages for all publishers within a class.

Before comparing the categories against each other, we should point out a notable distinction in distribution characteristics when compared to the engagement rate distributions. While engagement rates typically exhibit long-tailed distributions, the prevalence distributions for most groups conform to normality tests (the means are marked for all distributions). This contrast in distribution shapes is significant, suggesting that the use of analytical language (and the other features) across different publishers is more evenly spread and less skewed than the engagement patterns observed.

The figure reveals several interesting key trends and patterns. First, we note the marked and statistically signifi-

cant disparity in the prevalence of analytical language between the *Misinformation* and *Non-misinformation* classes. Notably, the median prevalence of analytic tweets for *Misinformation* publishers is 30.7%, while *Non-misinformation* outlets exhibit a higher median of 47.0%. This discrepancy suggests that misinformation sources are less likely to use analytical language, potentially favoring emotive or assertive rhetoric over analytical depth. This trend aligns with the commonly recognized tactics of misinformation sources, which often prioritize persuasion over analytical rigour.

Second, turning our attention to the bias classes, a Kruskal-Wallis test (p-value of 2.6×10^{-84}) and subsequent post-hoc Dunn tests, reveals that almost all pairwise comparisons between the bias classes to be statistically significant, with the only two exceptions being between (1) *Left* and *Right* and (2) *Least Biased* and *Right-Center*. This observation suggests a correlation between the level of bias in publishers and their use of analytical language. Specifically, publishers exhibiting the highest levels of bias, regardless of their political orientation, tend to use analytical language less frequently. Conversely, the *Least Biased* publishers are characterized by a higher rate of analytical language usage. This inverse relationship may reflect a strategic choice by highly biased publishers to engage their audience with more ideologically driven content over analytical discourse.

Discussion and Combined Insights: Considering the trends observed in Figures 3 and 4, interesting insights emerge regarding the interplay between the frequency of analytical language and user engagement across different classes. For example, while *Right* publishers feature less frequent use of analytic tweets, their followers concurrently exhibit the least engagement with such content. This pattern effectively negates the “novelty effect”. Instead, it suggests a congruence between the publisher’s content strategy and their audience’s preferences. The followers of *Right* publishers, it appears, are less inclined towards engaging with analytic content, possibly favoring narratives that align more closely with their ideological stances or emotional appeal.

Similarly, in the realm of reliability, *Misinformation* publishers and comparing to the *Non-misinformation* publishers, despite their infrequent use of analytic tweets, do not see an increase in user engagement when such language is employed. This observation again counters the “Novelty Effect” theory. Instead, it indicates that the audience of misinformation sources might be less influenced by the scarcity of analytic content, perhaps due to a predisposition towards narratives that confirm pre-existing beliefs or biases, rather than those offering analytical depth.

Key Observation: *Misinformation* publishers use analytical language notably less compared to *Non-misinformation* outlets. In terms of political bias, highly biased publishers demonstrate lower analytical language usage than their least biased counterparts. This trend, negating the “novelty effect”, shows the interest of audiences of *Right* and *Misinformation* publishers favoring less analytic content.

6.2 Clout, Perception, and Risk Language

We next look closer at three additional LIWC dimensions, each offering complementing insights to our observations regarding the prevalence and effects of analytic language.

The Effect of Clout Language: The clout dimension quantifies linguistic markers of authority, assertiveness, and influence. Specifically, the prevalence of clout language in social media spans expressions of institutional authority to populist rhetoric resonating with the reader’s experiences.

To assess the impact of clout language on user engagement, Figure 5 shows the *MPS* distributions across different publisher categories. Several key observations are notable, here split into three types. First, when contrasted with analytic language, clout language exerts a less distinct impact on the composite engagement rate (*All* category) and most of the individual categories. For example, whereas the median of $D_{All}^{analytic}$ for analytic language was -0.9 , the median of D_{All}^{clout} for clout language is -0.2 . This shift, while less pronounced, remains statistically significant with a p-value of 2.8×10^{-8} , as affirmed by the Wilcoxon signed-rank test. This finding indicates that the clout language retains a salient influence on user engagement, albeit to a varied extent across different classes.

Second, when comparing the bias categories, a more granular assessment reveals a heterogeneous response to clout language among the publisher classes. For example, the *Right* class again stands out and demonstrates a positive median shift in engagement rates with D_{Right}^{clout} having a median of 1.3 percentiles, implying that such language positively resonates with the user base of these publishers. Conversely, the left party publisher classes (*Left* and *Left-Center*) exhibit negative effect with $D_{Left-Center}^{clout}$ and D_{Left}^{clout} having median values of -0.6 and -0.5 percentiles, respectively. Although small, these deviations from the baseline are statistically significant, with the least significant p-value (8×10^{-4}) found in the comparison between the *Left* class and the baseline.

In summary, the positive engagement shift in the *Right* class may suggest a preference for assertive and dominant content, aligning with conservative rhetorical styles. The negative shift among the *Left* and *Left-Center* classes might indicate a skepticism towards authoritative discourse.

Third, turning our focus to the reliability factor, the differential impact of clout language on user engagement is striking. The *Misinformation* class demonstrates a significant increase in engagement rates when (news-related) tweets use clout-themed language, a trend not seen in the *Non-misinformation* class. For example, the medians of the clout distributions for the two classes are 1.2 and -0.3 percentiles, respectively. These observations suggest that audiences engaging with *Misinformation* sources may be particularly receptive to, or even seeking, content that projects confidence and authority, characteristics often associated with clout language. This inclination could reflect a cognitive bias wherein assertive content is perceived as more credible. Conversely, the smaller but still significant negative shift within the *Non-misinformation* class may indicate a preference for more nuanced or less assertive communication, aligned with a more critical engagement with content.




Figure 5: Clout language effect on engagement rates: mixed effects on engagement rates when using clout language.




Figure 6: Percentage of tweets belonging to clout language along different classes.

Key Observation: The clout dimension exerts a more differentiated impact on user engagement than the analytic one: positively influencing the *Right* and *Misinformation* classes, suggesting their audiences' preference for assertive content, while eliciting a negative response from the left party classes.

The Prevalence of Clout Language: Figure 6 shows the clout prevalence distributions associated with the different classes. When considering the bias classes, we observe significant differences in the use of clout language between the different groups (validated using the Kruskal-Wallis test). The *Right-Center* class exhibits a significantly lower median usage of clout language compared to the other four classes, suggesting that *Right-Center* publishers might adopt a more moderated or less assertive tone in their content. Another insight emerges when we combine the extreme classes (*Right* and *Left*) and compare them as a single entity against the other three less biased classes. In all comparisons, the extreme classes demonstrate higher use of clout language, with the largest p-value being 0.06. This observation indicates that publishers with more pronounced political biases tend to use clout language more frequently.

The Effect of Perception Language: Next, we focus on the dimension of perception. The perception category within the LIWC framework encompasses language pertaining to sensory and physical experiences, spatial orientation, and motion. It is a linguistic marker of how content is grounded in the tangible and observable world. Our goal in analyzing the prevalence of perceptual words in tweets is to provide insights into how sensory-rich language influences user engagement across various types of publishers followers.

The mediated effect of perceptual language is visualized in Figure 7, depicting *MPS* distributions for different groups. Notably, the median of $D_{All}^{\text{Perception}}$ at -1.0 indicates a significant downward shift in engagement when tweets are more perceptually descriptive (p-value: 3.1×10^{-62}). Perception language shows a more pronounced correlation with reduced engagement compared to analytical (2.9×10^{-42}) or clout (2.8×10^{-8}) language. This suggests potential user pref-

erences for abstract or narrative-driven content on Twitter, questioning the integration of sensory language (which, while enriching, could lead to higher cognitive load) within the fast-paced environment of social media.

Turning attention to the impact of perception language on engagement across various classes, significant shifts from baseline rates are observed for all classes (p-values ranging from 7.8×10^{-32} for *Least Biased* to 2.8×10^{-4} for *Left*). This statistical underpinning confirms the pervasive influence of perceptual language across publisher biases. Detailed examination of bias classes, in comparison to the analytic and clout dimensions (Figures 3 and 5), reveals a uniform effect on engagement rates. Unlike clout and analytical dimensions, the Kruskal-Wallis test does not reach significance, suggesting homogeneity in how perceptual content affects user engagement (compared to analytic and clout language), irrespective of political orientation.

Additionally, the test does not substantiate a significant difference between reliability classes' distributions (*Misinformation* vs. *Non-misinformation*), aligning with visual observations. This lack of distinction implies that perception language universally affects user engagement without introducing biases based on the publisher's perceived reliability.

These insights suggest that, while perception language uniformly decreases engagement across publisher classes, the reasons behind this phenomenon may be more related to content style preferences and processing behaviours rather than to political or reliability classifications. This encourages a deeper exploration of how sensory language is used, which we shall expand upon next.

Key Observation: Perception language consistently decreases user engagement across all classes of publishers, with statistically significant while uniform shifts regardless of political bias or reliability.

The Prevalence of Perception Language: Figure 8 offers a comprehensive view of the use of perception language across different groups, leading to several insights. First, the Kruskal-Wallis test statistically confirms the notably lower




Figure 7: Perception language effect on engagement rates: negative values show that perception-related posts have an inverse effect on engagement rates.

use of perception language of the *Misinformation* class (median of 39.7%) compared to the *Non-misinformation* class (median of 49.5%). This pronounced difference indicates that the perception language is more frequently employed by *Non-Misinformation* publishers. This trend might suggest a strategic emphasis on concrete, experiential content in credible news sources as opposed to the potentially abstract or emotionally driven narratives in misinformation outlets.

Second, examining bias classes, Kruskal-Wallis test confirms significant variations, with the Dunn tests highlighting significant differences between the less biased classes and the more biased classes. For example, the pairwise comparisons with the *Least Biased* class (showing the highest prevalence of perception language) are all significant. Additionally, a statistically significant higher prevalence is observed in the *Right-Center* class compared to *Right*, and a similar pattern is observed between *Left-Center* and *Left*. The shift towards increasing usage for the less biased classes is also observed on the median percentages (from Left to Right): 42.4%, 48.5%, 52.1%, 49.9%, and 39.5%. This gradient suggests that more centrist classes are more likely to use perception language, indicating a preference for tangible, relatable content in less ideologically extreme publications.

In summary, the varied use of perception language across publisher classes reveals nuanced content strategies. Notably, less biased publishers, especially the *Least Biased* class, employ more perceptual language while others choose alternative narrative styles.

Combined Prevalence Observations: While Figure 8 shows that the usage of perception language is less frequent in both *Misinformation* and highly biased classes, Figure 7 reveals that these ones yield similar levels of engagement when employing perception language as other classes. This observation suggests that, despite the lower prevalence of sensory and tangible language in *Misinformation* and highly biased content, its impact on user engagement does not differ significantly from other classes. This, combined with the “novelty effect”, suggests that followers of these classes may engage less with perception-related content.

The Effect of Risk Language: The linguistic attributes




Figure 8: Percentage of tweets belonging to perception language along different classes.

explored thus far – analytic reasoning, clout authority, and perception – represent pillars of credibility, social influence, and connection with observable reality. However, the pressure to attract attention and shape beliefs may also incentivize publishers to employ sensational rhetoric that defies these very attributes, instead playing on emotions and identities. The risk category within LIWC encapsulates such language, words that express defiance, rule-breaking, daring, and extremism. By quantifying the engagement with risk language, we aim to understand the extent to which different groups respond to such provocation and whether distinct dynamics underlie these responses.

Figure 9 shows the *MPS* distributions of risk language across various publisher classes. First, considering the aggregate effects across all 1,553 publishers (right-most boxplot), risk language only moderately boosts engagement rates. The distribution $\mathcal{D}_{All}^{\text{risk}}$ has a median of 0.7, indicating a 0.7 percentile increase in the median engagement rate. This effect is the smallest among the studied dimensions, supported by a non-significant p-value of 0.03 compared to the baseline. However, similar to clout language, the observed effects vary substantially among the classes, emphasizing the importance of class-specific analysis.

Second, starting with the reliability classes, we observe big differences, with $\mathcal{D}_{Misinfo}^{\text{risk}}$ and $\mathcal{D}_{Non-misinfo}^{\text{risk}}$ having medians of 3.0 and 0.6, respectively. These differences are supported by the Kruskal-Wallis test (p-value of 0.01), suggesting that *Misinformation* sources benefit more from the use of risk language and that the use of such language may help them obtain increased user engagement. This effect may potentially be due to its alignment with sensationalized content that reinforces pre-existing beliefs or biases.

Third, considering bias, a correlation emerges between the degree of bias—regardless of political orientation—and *MPS* values. Highly biased *Left* and *Right* classes show the highest *MPS* values, indicating receptivity to risk language aligned with the extremity of their positions. Conversely, the *Right-Center* and *Left-Center* classes exhibit a non-significant deviation from the baseline. Specifically, highly biased *Right* and *Left* classes statistically benefit from risk language, with medians of 3.0 and 3.5 and p-values of




Figure 9: Risk language effect on engagement rates: positive values show that risk-related posts positively increase the engagement rates.

2.0×10^{-6} and 7.5×10^{-7} respectively. In contrast, the *Least Biased* class shows a statistically significant negative effect, with a median shift of -0.3 and a p-value of 6.8×10^{-5} . This suggests that extreme publishers’ audiences are drawn to risk language, while audiences of *Least Biased* publishers prefer more measured and less sensational content.

The Prevalence of Risk Language: Our attention now turns to the prevalence of risk language across different publishers, as depicted in Figure 10. This analysis reveals several key observations. First, compared to the other linguistic dimensions there is no statistically significant difference in the prevalence of risk language between the different groups (e.g., p-value of 0.87 when comparing the *Left-Center* and the *Right-Center* with medians of 8.6% and 10.1%).

Second, we observe greater variations (e.g., inter-quartile differences) in the distribution of risk language usage for the more biased classes (*Left* and *Right*), suggesting that publishers in these categories possibly may apply more diverse editorial strategies when it comes to using risk language. In contrast, the *Least Biased* and *Right-Center* groups see substantially smaller variations, suggesting that publishers within these groups may use more homogeneous strategies regarding their use of such language.

Combined Risk Language Observations: Now, by factoring out the influence of the “Novelty Effect”, and comparing these insights with the engagement patterns seen in Figure 9, an intriguing narrative emerges. While the more biased publishers exhibit higher engagement with risk language, and *Misinformation* sources similarly see greater engagement shifts, the relatively uniform prevalence of risk language suggests that engagement spikes are not driven by its rarity. Instead, this indicates that the audiences of more biased or misinformation-leaning publishers might have a particular affinity for content that incorporates elements of risk, regardless of its frequency of use.

6.3 Robustness of Findings and Interplay between Dimensions

To ensure robustness of our findings, we performed tests with alternative percentile thresholds and conducted addi-




Figure 10: Percentage of tweets belonging to risk language along different classes.

tional analyses examining the potential confounding effects of topic coverage and temporal variations on the observed patterns. These analyses, detailed in Appendix A and Appendix B, respectively, showed that the main conclusions of our study were supported using other percentile thresholds and across a diverse range of topics, and remained consistent over time, strengthening the validity of our findings.

Finally, we conducted a preliminary investigation of the interplay between analytic and risk language, as detailed in Appendix C. While this paper’s focus is on the individual factors, these results suggest that the interplay between analytic and risk language yields diverse effects on user engagement across different publisher categories, opening the door for interesting follow-up work studying the combined effects of multiple dimensions.

7 Conclusion

This work took a LIWC-based approach to uncover how rhetorical attributes in tweets affect audience engagement, specifically investigating analytical, clout, perceptual, and risk language usage by news publishers across the political spectrum. We derived several key findings.

First, analytical language typically decreased engagement, especially among right-leaning and less reliable publishers. Combined with lower analytic tweets prevalence for these groups, this nullifies the novelty effect and confirms audience biases favoring ideologically aligned narratives over analytic discourse. Second, assertive clout language resonated more positively with right-leaning and unreliable publishers, suggesting a preference for authoritative rhetoric for these groups. Third, concrete perceptual language consistently reduced engagement across the categories, suggesting its effects stem more from perceptual style preferences independent of partisan bias or reliability.

Finally, sensational risk language increased engagement for politically extreme and unreliable publishers, indicative of inflammatory content confirming biases. Yet, uniform prevalence hinted that rarity was not the underlying driver.

While these insights can assist publishers, including those in misinformation, in strategically targeting audiences for higher engagement, the benefits of scientifically quantify-

ing these dynamics outweigh the potential downsides. Our granular analysis provides valuable intelligence for multiple stakeholders to shape ethical and responsible content strategies. Policymakers can derive directives to prioritize mitigating provocative appeals for vulnerable groups. Platforms can align algorithms to nourish credibility. Publishers can compete fairly by tailoring public discourse to constructive preferences. Most importantly, these findings highlight pathways for redirecting individual users through cognitive targeting, fostering truth-aligned recommendations resistant to manipulation. Overall, despite risks of misuse, decoding the intricate science of persuasion has profound potential to advance personal agency and digital public welfare.

Acknowledgments

The authors thank the reviewers and the associate editor for their insightful comments that helped improve the paper. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

Ad Fontes Media. 2023. Ad Fontes Media Bias Chart. <https://adfontesmedia.com/interactive-media-bias-chart/>. Accessed: 2023-09-10.

Aldous, K. K.; An, J.; and Jansen, B. J. 2019. View, Like, Comment, Post: Analyzing User Engagement by Topic at 4 Levels across 5 Social Media Platforms for 53 News Organizations. *Proc. ICWSM*, 13(01): 47–57.

Aldous, K. K.; An, J.; and Jansen, B. J. 2022. Measuring 9 Emotions of News Posts from 8 News Organizations across 4 Social Media Platforms for 8 months. *Trans. Soc. Comput.*, 4(4).

AllSides. 2023. AllSides. <https://www.allsides.com>. Accessed: 2023-09-10.

Antypas, D.; Preece, A.; and Camacho-Collados, J. 2023. Negativity spreads faster: A large-scale multilingual twitter analysis on the role of sentiment in political communication. *J. Online Soc. Netw. Media.*, 33: 100242.

Beel, J.; Xiang, T.; Soni, S.; and Yang, D. 2022. Linguistic Characterization of Divisive Topics Online: Case Studies on Contentiousness in Abortion, Climate Change, and Gun Control. *Proc. ICWSM*, 16(1): 32–42.

Beknazár-Yuzbashev, G.; Jiménez Durán, R.; McCrosky, J.; and Stalinski, M. 2022. Toxic content and user engagement on social media: Evidence from a field experiment. *SSRN*.

Bellovary, A. K.; Young, N. A.; and Goldenberg, A. 2021. Left-and right-leaning news organizations use negative emotional content and elicit user engagement similarly. *Affective Science*, 2: 391–396.

Boyd, R.; Ashokkumar, A.; Seraj, S.; and Pennebaker, J. 2022a. LIWC-22: Descriptive Statistics and Norms.

Boyd, R. L.; Ashokkumar, A.; Seraj, S.; and Pennebaker, J. W. 2022b. The development and psychometric properties of LIWC-22. The University of Texas at Austin.

Broniatowski, D. A.; Simons, J. R.; Gu, J.; Jamison, A. M.; and Abroms, L. C. 2023. The efficacy of Facebook’s vaccine misinformation policies and architecture during the COVID-19 pandemic. *Science Advances*, 9.

Bronstein, M. V.; Pennycook, G.; Buonomano, L.; and Cannon, T. D. 2021. Belief in fake news, responsiveness to cognitive conflict, and analytic reasoning engagement. *Thinking & Reasoning*, 27(4): 510–535.

Candia, C.; Atari, M.; Kteily, N.; and Uzzi, B. 2022. Overuse of moral language dampens content engagement on social media.

Catherine A. Cherrstrom, C. J. B.; and Sherron, T. 2023. Nontraditional Students and Credit for Prior Learning—Analytical Thinking, Clout, Drives, and Motives. *The J. Contin. High. Educ.*, 1–16.

Cheng, Z.; and Li, Y. 2023. Like, Comment, and Share on TikTok: Exploring the Effect of Sentiment and Second-Person View on the User Engagement with TikTok News Videos. *Social Science Computer Review*.

Cialdini, R. 2006. Influence: The psychology of persuasion, revised edition of Harper Business. ISBN-13, 978-0061241895.

CITAP. 2022. Methods and Data. https://pcad.ils.unc.edu/methods_and_data. Center for Information and Public Life at the University of North Carolina.

Daft, R. L.; and Lengel, R. H. 1986. Organizational information requirements, media richness and structural design. *Management science*, 32(5): 554–571.

Edelson, L.; Nguyen, M. K.; Goldstein, I.; Goga, O.; McCoy, D.; and Lauinger, T. 2021. Understanding engagement with U.S. (mis)information news sources on Facebook. *Proc. ACM IMC*, 444–463.

Egger, R.; and Yu, J. 2022. A topic modeling comparison between lda, nmf, top2vec, and bertopic to demystify twitter posts. *Frontiers in sociology*, 7: 886498.

Hanley, H. W. A.; Kumar, D.; and Durumeric, Z. 2023. A Golden Age: Conspiracy Theories’ Relationship with Misinformation Outlets, News Media, and the Wider Internet. *Proc. ACM Hum.-Comput. Interact.*, 7(CSCW2).

Haworth, E.; Grover, T.; Langston, J.; Patel, A.; West, J.; and Williams, A. C. 2021. Classifying Reasonability in Retellings of Personal Events Shared on Social Media: A Preliminary Case Study with /r/AmITheAsshole. *Proc. ICWSM*, 15(1): 1075–1079.

Hiaeshutter-Rice, D.; and Weeks, B. 2021. Understanding Audience Engagement with Mainstream and Alternative News Posts on Facebook. *Digital Journalism*, 9(5).

Hu, Y.; Farnham, S.; and Talamadupula, K. 2021. Predicting User Engagement on Twitter with Real-World Events. *Proc. ICWSM*, 9(1): 168–177.

Huszár, F.; Ktena, S. I.; O’Brien, C.; Belli, L.; Schlaikjer, A.; and Hardt, M. 2022. Algorithmic amplification of politics on Twitter. *Proc. of the National Academy of Sciences*, 119(1).

Jaidka, K.; Chhaya, N.; Mumick, S.; Killingsworth, M.; Halevy, A.; and Ungar, L. 2020. Beyond Positive Emotion:

Deconstructing Happy Moments Based on Writing Prompts. *Proc. ICWSM*, 14(1): 294–302.

Jiyoung Lee, Y. K.; and Zhu, X. 2023. Liked and shared tweets during the pandemic: the relationship between intrinsic message features and (mis)information engagement. *Behaviour & Information Technology*, 0(0): 1–18.

Josef Verbanac. 2023. Iffy news. <https://www.iffy.news/>. Accessed: 2023-09-10.

Khalid, O.; and Srinivasan, P. 2020. Style Matters! Investigating Linguistic Style in Online Communities. *Proc. ICWSM*, 14(1): 360–369.

Kim, S.; Razi, A.; Stringhini, G.; Wisniewski, P. J.; and De Choudhury, M. 2021. You Don't Know How I Feel: Insider-Outsider Perspective Gaps in Cyberbullying Risk Detection. *Proc. ICWSM*, 15(1): 290–302.

Ksiazek, T. B.; Peer, L.; and Lessard, K. 2016. User engagement with online news: Conceptualizing interactivity and exploring the relationship between online news videos and user comments. *New media & society*, 18(3): 502–520.

Kumar, N.; Nagalla, R.; Marwah, T.; and Singh, M. 2018. Sentiment dynamics in social media news channels. *J. Online Soc. Netw. Media.*, 8: 42–54.

Lin, H.; Lasser, J.; Lewandowsky, S.; Cole, R.; Gully, A.; Rand, D. G.; and Pennycook, G. 2023. High level of correspondence across different news domain quality rating sets. *PNAS Nexus*, 2(9).

Mahmud, J.; Chen, J.; and Nichols, J. 2014. Why are you more engaged? predicting social engagement from word use. *arXiv:1402.6690*.

Markowitz, D. M. 2023. Analytic thinking as revealed by function words: What does language really measure? *Applied Cognitive Psychology*, 37(3): 643–650.

MBFC. 2023. MBFC. <https://mediabiasfactcheck.com>. Accessed: 2023-09-10.

Mohammadinodooshan, A.; and Carlsson, N. 2024. Understanding Engagement Dynamics with (Un) Reliable News Publishers on Twitter. In *Proc. ASONAM*, 36–47. Springer-Verlag.

Molina, M. D.; Wang, J.; Sundar, S. S.; Le, T.; and DiRusso, C. 2023. Reading, commenting and sharing of fake news: How online bandwagons and bots dictate user engagement. *Communication Research*, 50(6): 667–694.

Musk, E. 2022. Tweets are read 100 times more than they are liked. Twitter. Tweet ID: 1606107465174503424.

NewsGuard. 2023. NewsGuard. <https://newsguardtech.com>. Accessed: 2023-09-10.

Papadogiannakis, E.; Papadopoulos, P.; Markatos, E. P.; and Kourtellis, N. 2023. Who Funds Misinformation? A Systematic Analysis of the Ad-Related Profit Routines of Fake News Sites. *Proc. ACM Web conf 2023*.

Pennycook, G.; and Rand, D. G. 2020. Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *J. personality*, 88(2): 185–200.

Petty, R. E.; and Briñol, P. 2011. The elaboration likelihood model. *Handbook of theories of Soc. Psychol.*, 1: 224–245.

Pezzuti, T.; Leonhardt, J. M.; and Warren, C. 2021. Certainty in Language Increases Consumer Engagement on Social Media. *Journal of Interactive Marketing*, 53(1): 32–46.

Pfeffer, J.; Matter, D.; and Sargsyan, A. 2023. The Half-Life of a Tweet. *Proc. ICWSM*, 17(1): 1163–1167.

Pierri, F.; Luceri, L.; Jindal, N.; and Ferrara, E. 2023. Propaganda and Misinformation on Facebook and Twitter during the Russian Invasion of Ukraine. *Proc. Websci'23*.

Rathje, S.; Bavel, J. J. V.; and van der Linden, S. 2021. Out-group animosity drives engagement on social media. *Proc. National Academy of Sciences*, 118(26).

Robertson, C. E.; Pröllochs, N.; Schwarzenegger, K.; Pärnamets, P.; Bavel, J. J. V.; and Feuerriegel, S. 2023. Negativity drives online news consumption. *Nature Human Behaviour*, 7: 812–822.

Robertson, R. E.; Jiang, S.; Joseph, K.; Friedland, L.; Lazer, D.; and Wilson, C. 2018. Auditing Partisan Audience Bias within Google Search. *Proc. ACM on Human-Computer Interaction*, 2(CSCW): 148.

Rozin, P.; and Royzman, E. B. 2001. Negativity Bias, Negativity Dominance, and Contagion. *Personality and Social Psychology Review*, 5(4): 296–320.

Salehabadi, N.; Groggel, A.; Singhal, M.; Roy, S. S.; and Nilizadeh, S. 2022. User Engagement and the Toxicity of Tweets. *arXiv:2211.03856*.

Samory, M.; Abnousi, V. K.; and Mitra, T. 2020. Characterizing the social media news sphere through user co-sharing practices. In *Proc. ICWSM*, volume 14, 602–613.

Sharma, K.; Ferrara, E.; and Liu, Y. 2022. Characterizing Online Engagement with Disinformation and Conspiracies in the 2020 U.S. Presidential Election. *Proc. ICWSM*, 16(1).

Short, J.; Williams, E.; and Christie, B. 1976. *The Social Psychology of Telecommunications*. Wiley. ISBN 0471015814.

Spinde, T.; Richter, E.; Wessel, M.; Kulshrestha, J.; and Donnay, K. 2023. What do Twitter comments tell about news article bias? Assessing the impact of news article bias on its perception on Twitter. *J. Online Soc. Netw. Media*.

Tausczik, Y. R.; and Pennebaker, J. W. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *J. Lang. Soc. Psychol.*, 29(1): 24–54.

Trilling, D.; Tolochko, P.; and Burscher, B. 2017. From Newsworthiness to Shareworthiness. *J. Mass Commun. Q.*, 94(1): 38–60.

Twitter. 2023. Twitter's Recommendation Algorithm. <https://github.com/twitter/the-algorithm/>. Accessed: 2024-12-10.

Vosoughi, S.; Roy, D.; and Aral, S. 2018. The spread of true and false news online. *Science*, 359(6380): 1146–1151.

Walker, M.; and Matsa, K. E. 2023. News Consumption Across Social Media in 2023. Pew Research Center.

Weld, G.; Glenski, M.; and Althoff, T. 2021. Political bias and factualness in news sharing across more than 100,000 online communities. *Proc. ICWSM*, 15: 796–807.




Figure 11: Analytical language effect on engagement rates computed using the Average Percentiles Shifts (APS).

A Robustness Against Alternative Percentile Shifts

To address the potential limitations of the Median Percentile Shift (*MPS*) metric and demonstrate the robustness of our findings, we performed additional analyses using alternative percentile shifts. While *MPS* is a useful metric for comparing the central tendencies of engagement rate distributions, it has some inherent limitations. First, *MPS* focuses solely on the median, which may not capture other essential aspects of the distribution, such as variance or skewness. Second, *MPS* assumes that the compared distributions have similar shapes, which may not always be the case. Finally, *MPS* is a descriptive metric and does not provide information about the statistical significance of the observed differences. To overcome these limitations and validate the reliability of the insights presented based on *MPS*, we performed supplementary analysis using Average Percentile Shifts (*APS*) across alternative percentiles. Specifically and using the definitions in Sec. 5.1, for each candidate threshold $p \in \{10, 20, \dots, 90\}$, we computed: (1) The p^{th} percentile of the high feature group's engagement rates (S_a^F). (2) Its percentile rank within the overall rates (S_a). (3) The deviation between the percentile rank from Step 2 and the original threshold p . Finally, we averaged these deviations to obtain the *APS* as follows:

$$\frac{1}{9} \sum_{p \in 10, \dots, 90} [P_{S_a}(\text{percentile}(S_a^F, p)) - p].$$

APS confirms the validity of all observed *MPS* trends presented in this paper. As an example, we plot these distributions for the analytic language effect in Figure 11. Compared to Figure 3, we see the same trends repeating here with some variation in magnitudes. For example the median of $D_{\text{All}}^{\text{analytic}}$ has decreased from -0.9 to -0.78 here. This consistency despite using alternative percentile calculations, affirms the reliability of our methodology and findings.

B Robustness Against Topic and Time

Two potential concerns or confounding factors that could influence the observed patterns for the four linguistic dimensions are that maybe a few topics within those dimensions affect the patterns, not the actual linguistic dimensions, and

second, the temporal variations in the dataset. To address these concerns, we conducted a series of robustness checks using topic modeling techniques and temporal segmentation.

Topic-based Analysis: For the topic analysis, we employed BERTopic. It was selected for its superior performance in extracting coherent and meaningful topics, as well as for its ability to handle the unstructured and brief nature of tweets effectively (Egger and Yu 2022). For each linguistic dimension (e.g., perception), we followed a four-step process:

1. We selected a subset of tweets belonging to each linguistic dimension (e.g., perception tweets).
2. We applied BERTopic to these subsets, using default embedding computations and allowing it to automatically determine the optimal number of topics, to identify the underlying topics within that dimension.
3. For each identified topic, we created a new dataset consisting of tweets related to that specific topic and tweets not belonging to the linguistic dimension under consideration (e.g., for a topic in the perception dimension, the dataset would include tweets related to that topic and all non-perception tweets).
4. We computed the Median Percentile Shift (*MPS*) (the same metric as our main analysis) on this dataset to assess the robustness of our original findings.

The results of our topic modeling analysis revealed that the observed patterns for each linguistic dimension were not limited to a one or small number of specific topics. Notably, for all four dimensions, we found that at least 7 out of the top 10 topics (sorted by the number of tweets they cover) supported the original findings¹, indicating a high degree of consistency across a diverse range of subjects. This suggests that the engagement patterns are driven by the broader linguistic characteristics captured by each dimension, rather than a narrow set of topics.

However, it is important to acknowledge that the topics identified within each linguistic dimension may collectively play a role in shaping user engagement. While individual topics may not solely account for the observed patterns, the combination of topics falling under the umbrella of a given linguistic dimension may contribute to the overall effect.

To illustrate this, let us consider some concrete examples. In the analytic language dimension, we identified a clear climate change topic (with top keywords: “climate”, “energy”, “solar”, “carbon”, “fossil”) among the top 10 topics. However, we also found more general topics, such as the largest topic in this group, which was food-related (with keywords: “restaurants”, “menu”, “cream”, “food”, “flavor”)². Notably, both of these topics, along with 7 out of the top 10 topics,

¹This holds true when we relaxed the significance level of the p-values (from 0.01) to 0.1.

²While initially unexpected, our manual analysis revealed that news tweets involving these keywords often tapped into current societal interests and lifestyle choices. Specifically, discussions about “restaurants” and “food” can reflect community engagement and public interest, highlighting the cultural relevance of these topics in the context of current events.

supported most of the patterns observed in the overall analytic language (Figure 3).

Similarly, in the perception dimension, we found event-related topics like (“music”, “band”, “concert”, “tour”, “festival”) as the 3rd largest topic. Surprisingly, an initially unexpected topic about pets (“dog”, “pet”, “animal”, “puppy”, “adoption”) emerged as the largest topic in this group. Both of these topic categories supported most of the patterns in the overall perception language (Figure 9).

Finally, it is worth noting that many of the topics in each group aligned with the expected news-related topics for that dimension. For example, we found topics related to abortion in the analytic dimension, COVID-19 in the risk dimension, the TikTok (ban) in clout, and music in the perception dimension. This consistency across multiple topics within each dimension suggests that the observed patterns are not limited to a small subset of specific topics.

Time-based Analysis: Time is another potential confounding factor that could influence the observed patterns. To study the generalizability of our findings with regard to the time window studied, we divided our data collection period (from Dec. 15, 2022, to June 2023) into three buckets and evaluated the tweets in each bucket individually. Using three buckets allowed the results to remain significant, whereas using more buckets led to a number of samples in some buckets becoming too small, compromising the significance of those results.

To assess the stability of the linguistic dimensions over time, we first computed the percentage of tweets belonging to each dimension relative to the total number of tweets in each bucket. Interestingly, we observed no significant differences between the percentages in the buckets and the overall timeline. The percentage of tweets containing analytic language was consistently between 46-47% across all buckets, while perceptual language was present in 49% of tweets in each bucket. Similarly, the percentage of tweets containing risk language remained stable at 9% across all buckets, and the percentage of tweets with clout language ranged from 32-33%. These findings suggest a consistent prevalence of the four linguistic dimensions over time.

Second, we investigated whether the patterns observed for each of the four dimensions in the original analysis held true when limiting the timeline of the study to each bucket. To account for the reduced sample size in each bucket, we relaxed the significance level of the p-value (from 0.01) to 0.1. Notably, the main findings of the paper (reported in Figures 3-9) remained statistically significant for all four dimensions in all three buckets. This indicates that the observed relationships between the linguistic dimensions and user engagement are robust and persist over time, even when considering shorter time periods.

The results of our time-based analysis provide strong evidence for the stability and consistency of the observed patterns across the four linguistic dimensions. The fact that the prevalence of each dimension remains relatively constant over time and that the main findings hold true in each time bucket suggests that the relationships between language use and user engagement are not significantly influenced by temporal factors. This robustness to time-based variations




Figure 12: Combined Analytic and Risk language effect on engagement rates.

strengthens the validity of our conclusions and highlights the enduring nature of the language-engagement dynamics we have uncovered here.

C Exploring the Interplay Between Analytic and Risk Language

While our main analysis focused on the individual effects of the four linguistic dimensions (analytic, clout, perceptual, and risk language) on user engagement, it is also interesting to consider the potential interplay between these dimensions. As a starting point, we chose to study the combined effect of analytic and risk language. For example, analytic language is associated with logical and rational thinking, while risk language is often characterized by emotional and provocative content. Investigating the interplay between these two seemingly opposing dimensions can provide valuable insights into how their combination influences user engagement across different publisher categories.

The results of our analysis when combining these dimensions are shown in Figure 12. When considering the individual effects of these two dimensions, as seen in Figure 3 (analytic language effect) and Figure 9 (risk language effect), we observe that they have an inverse effect on most groups. For example, while risk language increases engagement among followers of biased and misinformation publishers, analytic language decreases it.

Interestingly, when we examine the combined effect of analytic and risk language, we observe a diverse set of effects on different groups. Notably, this combination appears to cancel out the individual effects of both dimensions for the *Right* group. In other words, risk language that also contains analytic themes does not significantly affect the engagement of *Right* followers. The same pattern is observed for the *Non-misinformation* group and when considering all publishers (the *All* group).

The most pronounced effect of this combined language is seen in the *Misinformation* group, with a median effect of 2.62 percentiles, followed by the *Left* group, with a median effect of 1.63 percentiles. These results suggest that the combination of analytic and risk language resonates differently with various audience segments. For the *Misinformation* and

Left groups, the presence of both dimensions seems to amplify engagement, possibly by making topics appealing to a mix of rational and emotional triggers.

The findings from this preliminary analysis highlight the complex interplay between linguistic dimensions and their impact on user engagement. The fact that the combination of analytic and risk language produces different effects across publisher categories underscores the importance of future work also considering the interaction between dimensions. Future research can extend this analysis to other combinations of linguistic dimensions, such as the interplay between clout and perceptual language or the interaction of all four dimensions. Additionally, investigating the psychological mechanisms underlying these combined effects could provide a deeper understanding of how language influences user engagement and how these dynamics vary across different audience segments.