Shifting Writing Styles and Sentiment Effects

Huiming (Minna) Zhang[1]

[a] Michael F. Price College of Business, University of Oklahoma, Norman OK 73019

Version: April 8th, 2026

Abstract

In this study, we investigate shifts in the writing style of financial news and the diminishing impact of information in financial news on security prices. First, we document that the writing style of financial news has gradually shifted toward a more analytical tone, with greater use of power-related words and stronger emphasis on authenticity. Despite these stylistic shifts, the impact of news sentiment on prices has diminished over time. Second, we propose one possible explanation: the adoption of AI writing tools may have contributed to these changes in writing style while reducing the credibility of financial news content as perceived by investors. Specifically, an average increase in monthly visits to the AI tool is associated with a 30% reduction in the impact of news tone, even though the circulation of financial news has continued to increase. We further show that this diminishing sentiment effect is less pronounced in the subsample of news columns subject to stricter writing formats. Overall, this study identifies a previously underexplored channel through which investor responses to financial news. Our findings suggest that caution should be exercised in the use of stylized writing in financial journalism.

Keywords: sentiment analysis; financial news; information; oil prices; oil shock

JEL Classification: G14, G17, G41, G53

[1] Corresponding author:
E-mail address: huimingzhang@ou.edu (H. Zhang).
Phone number: +1 (405)885-8548. Postal address: Michael F. Price College of Business, 307 West Brooks, Norman OK 73019

1. Introduction

The existing literature on sentiment analysis, including Tetlock (2007) and Loughran and McDonald (2011), documents a significant association between news reports and stock return changes by capturing sentiment as the tone of financial news articles. Another stream of literature examines factors that moderate the impact of sentiment, such as operational complexity in Loughran and McDonald (2024) and readability in financial disclosures in Loughran and McDonald (2014). While the existing literature on sentiment effects focuses primarily on whether news is more positive or more negative, much less is known about how the depiction of events and the narrative style of reporting, influence price responses. Collecting news articles from the most widely circulated financial and oil-market news sources, we study how shifts in writing styles influence the price response to news sentiment. Specifically, we examine whether changes in writing style alter the impact of oil news sentiment on oil prices and whether the adoption of AI writing tools has further weakened this effect. We propose one possible explanation for the phenomenon that the impact of news sentiment on prices has been diminishing: AI writing tools may have shifted financial journalism toward language that appears more authentic, analytical, and powerful, with greater signals of honesty, while at the same time reducing the credibility of news content as perceived by investors. Cao et al. (2023) show that the impact of sentiment in financial reports is mitigated by AI adoption. However, their study focuses on AI readership rather than AI-assisted writing by the news source.

We focus on oil-related news and oil markets. Oil prices have long been an important indicator of economic conditions (e.g., their predictive role during the 2008–2009 financial crisis), and the oil market is also a significant sector in global economies. We manually collect news articles from several of the most widely circulated financial news sources, including Dow Jones Institutional News, Financial Times, New York Times, the Wall Street Journal, and an oil-market-focused news feed, Energy Weekly News. From these sources, we construct a sample of news articles containing the keywords "oil" and "shock" from 1986 to 2023. We examine two main issues. First, we document shifts in the writing styles of oil-related news. We find that the price impact of news sentiment has declined over time. Second, we propose one possible explanation for this diminishing effect. While subscriptions to major financial news sources have continued to increase, the weakening sentiment effect may be related to the adoption of AI tools by major financial newspapers. In particular, AI-assisted writing at the same time reduces the credibility of the news content as perceived by investors.

To measure writing styles, we follow the conventional approach in social science research and use LIWC-22 as our text analysis tool. LIWC-22 provides article-level scores that capture a broad set of linguistic dimensions, including analytic and formal language, emotional content, and power-related words. To reduce the dimensionality of these style measures, we then apply principal component analysis (PCA) and extract three primary components that capture most of the common variation across the LIWC-based variables. We further show that these three components are strongly associated with the release of ChatGPT, as measured by monthly visits to the AI tool and Google search trends related to it. Prior studies, such as Buchanan and Hickman (2024) and Longoni et al. (2022), suggest that generative AI-assisted writing is perceived as less credible by human readers. Although traditional financial news sources such as the Wall Street Journal and the New York Times remain important channels for transmitting market-relevant information, as suggested by Tetlock (2011), it remains unclear whether these news sources continue to play the same role in the AI era. We find that, while sentiment measured by the Loughran-McDonald dictionary continues to significantly explain changes in daily oil returns, the release of the AI tool has weakened the impact of sentiment extracted from oil-shock news articles. Specifically, a 0.14 million increase in monthly visits to ChatGPT (the sample average monthly visits is associated with approximately a 30% reduction in the impact of sentiment on oil returns. For reference, in our sample, average monthly visits increased by about 0.14 million following the release of ChatGPT in late 2022.

We conduct a series of robustness tests to rule out alternative explanations. First, we focus on the *Market Talk* series in Dow Jones Institutional News. Compared with standard news articles, *Market Talk* follows a relatively fixed writing format and typically consists of brief comments or quotations from market professionals. We find that this series exhibits smaller shifts in writing style. Correspondingly, the decline in the sentiment effect on oil prices is also less affected as ChatGPT usage increases. Second, we show that article length, readability, and uncertainty do not account for the diminishing sentiment effect associated with greater ChatGPT usage. Third, our findings are robust to alternative measures of news tone, including sentiment indices weighted by article length. Finally, the results remain significant when we restrict the sample to the post-2020 period or exclude the year 2020, the onset of COVID-19, when oil futures prices turned negative in April 2020.

There is still an ongoing debate regarding whether AI tools improve information searches or whether they might have unintended negative consequences. Prior studies provide direct evidence that when readers believe or suspect that content is generated by AI, they tend to view such content as less credible and less trustworthy. For example, Altay and Gilardi (2024) show that headlines labeled as AI-generated are perceived as less accurate and are less likely to be shared, even when the headlines are true or originally written by humans. Similarly, Lermann Henestrosa and Kimmerle (2024) find that when the same GPT-generated science journalism article is labeled as AI-written rather than human-written, readers assign lower credibility source. Reis et al. (2024) show that identical medical advice is evaluated as less reliable and less empathetic, and is less likely to be followed, when people believe that AI is involved in generating the advice. Our results suggest caution should be taken in determining whether generative AI provides the best support for humans, particularly in facilitating the transfer of information to financial investors. These findings mirror the results of Buchanan and Hickman (2024), who highlight that AI-assisted content may be perceived as less trustworthy by people.

2. Existing Literature

2.1. Sentiment Analysis

Research on sentiment analysis is an important area in financial studies, primarily focusing on examining the impact of sentiment from financial news or financial statements on security prices. Merton (1987) argues that investors have limited attention, leading them to primarily focus on well-known stocks, suggesting that sentiment in news can attract investors' attention and consequently trigger price responses following the publication of financial news. Barber and Odean (2008) further demonstrate that this sentiment-driven attention effect is particularly significant among retail investors, whereas institutional investors are less likely to be influenced by attention-grabbing stocks.

Tetlock (2007, 2010) uses news articles from the Wall Street Journal or Dow Jones Institutional News to examine the impact of news releases on stock prices, demonstrating a reduction in information asymmetry through news dissemination. Similar results are presented by Tetlock (2011), Chan (2003), Klibanoff et al. (1998), Fang and Peress (2009), and Engelberg and Parsons (2011).

Several studies have used dictionaries to capture sentiment in financial news, such as Garcia (2013) and Jiang et al. (2019). The most commonly used dictionary is the Loughran and McDonald sentiment dictionary, originally introduced by Loughran and McDonald (2011) and subsequently used by Feldman et al. (2010), Solomon (2012), and Calomiris and Mamaysky (2019). These studies measure sentiment by calculating the relative percentage of negative to positive words in news articles and examining their impact on stock returns.

2.2. Writing Styles in Financial News

In this article, we use the Loughran-McDonald dictionary to capture sentiment in news articles related to “oil” and “shock.” However, our primary focus is on the impact of changes in the writing styles of these news articles and how such changes influence the sentiment effect. By focusing on the oil markets, we aim to exclude the influence of firm-specific noise on stock returns. The oil market is a key sector in the global economy. We examine whether the impact of oil-related news sentiment on oil returns is either diminished or amplified by changes in writing styles. Additionally, oil prices is an important indicator of overall economic conditions. Our dataset comprises 34,226 news articles, enabling us to closely examine shocks in the oil market and to analyze how various factors affect changes in oil returns.

There has been a growing trend in financial studies focusing on language styles, particularly the language used by CEOs and CFOs. For example, Shi et al. (2019) demonstrate that linguistic style matching between a company's CEO and CFO is associated with higher compensation for the CFO and an increased likelihood of being included on the board of directors by the CEO. Moreover, mergers and acquisitions (M&A) activity is more likely to occur when the CEO and CFO use similar language styles. Larcker and Zakolyukina (2012) utilize the LIWC language style analyzer to detect deceptive content in conference calls. Following a similar approach, we use LIWC to analyze changes in the writing styles of news articles related to oil shocks.

The rise of AI tools, particularly ChatGPT, has significantly influenced how people search for information, think, write, and edit language. Novy-Marx and Velikov (2025) highlight the power of AI in detecting trading signals and assisting in the writing of scholarly research reports. Malik et al. (2023) discuss the widespread use of generative AI tools in academic writing among higher education students, as well as students' perceptions of these tools. However, little attention has been paid to the use of generative AI in the writing of news articles - especially financial news articles published by popular news outlets. This study aims to address this gap and contribute to the existing literature.

Some of the existing literature suggests that AI-generated content is generally perceived as less trustworthy. For instance, Longoni et al. (2022) show that when a news article's title is generated by AI, readers are less likely to believe the article is authentic. Similarly, Buchanan and Hickman (2024) find that when readers suspect an article is written by AI, they are less likely to believe that the news accurately describes real-world facts. In this study, we examine whether this skepticism extends to financial news articles. Specifically, we investigate whether changes in writing style, potentially indicative of AI authorship, lead investors to trust the content less, thereby reducing their trading activity in the corresponding financial markets. This research addresses a gap in the literature on the intersection of AI-generated content and financial decision-making.

3. Theoretical Framework and Hypotheses

We follow the framework of De Long et al. (1990) to explain how noise traders' biased beliefs about expected payoffs affect asset prices. In their framework, sentiment-driven beliefs can move prices away from fundamentals when arbitrage is limited. We apply this idea to the oil market. In particular, we argue that oil prices respond to the sentiment in oil shock news because noise traders in the oil market form biased beliefs after reading such news.

Let ϕ_t denote the proportion of noise traders in the oil market at time t. Following the intuition of Peng and Xiong (2006), but adapting it to the oil market setting, we assume that sentiment in oil shock news shifts investors' expected oil returns. We write this belief revision as

$$\Delta \mathbb{E}_t[R_{t+1}] \; = \; \eta \, \phi_t \, q_t s_t$$

where η captures the oil market's exposure to the sentiment in oil shock news, s_t is the sentiment in oil shock news at time t, and q_t captures the extent to which the linguistic form of the news allows investors to process and trust that sentiment. In this sense, q_t reflects the transmission of sentiment from news text to investors' beliefs.

Under mean-variance preferences, investor demand is given by

$$X_t \; = \; \frac{1}{\gamma \, \sigma^2} \, \Delta \mathbb{E}_t[R_{t+1}] \; = \; \frac{1}{\gamma \, \sigma^2} \, \eta \, \phi_t q_t \, s_t$$

where γ is the coefficient of risk aversion and σ^2 is the variance of oil returns. Following Kyle (1985), price movements are driven by order flow under market depth $\Lambda > 0$. Therefore, the price change is

$$\Delta p_t \; = \; \Lambda \, \Delta X_t^n \; = \; \Lambda \cdot \frac{1}{\gamma \, \sigma^2} \, \eta \, \phi_t \, s_t,$$

Collecting terms, oil returns can be written as

$$R_t = \beta' f_t + \kappa \, \phi_t q_t \, s_t + \varepsilon_t$$

Where $\kappa = \frac{\Lambda \, \eta}{\gamma \, \sigma^2}$ and f_t denotes control factors related to oil market conditions and macroeconomic risks. This expression follows the key insight of De Long et al. (1990): when a nontrivial fraction of investors trade on sentiment-driven biased beliefs, prices may deviate from fundamentals because rational arbitrage is limited.

A key feature of our framework is that writing style affects how sentiment is transmitted to market prices. We treat writing style as a moderator that changes how strongly sentiment in oil shock news influences oil returns. Specifically, we define q_t as a sentiment-transmission weight determined by writing style. Writing style may include several dimensions, such as readability, clarity, and machine parsability. We therefore write

$$q_t = q_0 + q_1 w_t$$

where $w_t$is the observed writing style of the news article. Substituting this into the return equation gives

$$R_t = \beta' f_t + \kappa\, \phi_t (q_0 + q_1 w_t)\, s_t + \varepsilon_t$$

This expression shows that the effect of sentiment on oil returns depends on writing style. When writing style changes, the sensitivity of oil returns to sentiment also changes. We further argue that the adoption of AI writing tools can affect oil returns indirectly through writing style. In particular, AI tools may change the observed style of oil shock news by changing the degree of formality and analytical language in the text. We model writing style as

$$w_t = w_0 + w_1 AI_t + \Omega' Z_t + u_t \tag{1}$$

where AI_t is a proxy for the adoption or usage of AI writing tools, and Z_t is a vector of other factors that may influence writing style. In this setting, AI changes the linguistic form of the news and moderates how sentiment affects oil returns. This framework leads to the following empirical specification:

$$R_t = \alpha + \beta' f_t + \omega_1 s_t + \omega_2 w_t + \omega_3 (s_t \times w_t) + \Gamma' Z_t + \varepsilon_t$$

In this equation, the interaction term $s_t \times w_t$ captures whether writing style moderates the impact of sentiment on oil returns. By substituting equation (1) into the above specification, we obtain the reduced-form model with the AI proxy:

$$R_t = \alpha + \beta' f_t + \delta_1 s_t + \delta_2 AI_t + \delta_3 (s_t \times AI_t) + \Gamma' Z_t + \varepsilon_t$$

The coefficient δ_3 captures whether the adoption of AI writing tools changes the relationship between oil returns and sentiment through its effect on writing style. Based on this framework, we propose the following hypotheses.

H1: Sentiment in oil shock news affects oil returns. More positive news is associated with higher oil returns, while more negative news is associated with lower oil returns.

H2: Writing style moderates the effect of sentiment in oil shock news on oil returns.

H3: The adoption of AI writing tools changes the writing style of oil shock news and, through this channel, moderates the effect of sentiment on oil returns.

4. Methodology and Data

4.1. Sample

We use the LIWC tool to analyze the writing styles of news articles related to oil shocks.[2] Previous studies have employed LIWC for various text analysis applications. For example, Shi et al. (2019) use LIWC to capture linguistic style matching between CEOs and CFOs, while Larcker and Zakolyukina (2012) analyze the narratives of CEOs and CFOs during conference calls. To measure disagreement or sentiment in oil-related news articles, we calculate sentiment as the percentage of negative words minus the percentage of positive words, using the Loughran and McDonald sentiment dictionary. Oil returns are calculated using WTI crude oil futures prices from 1986 to 2023. To account for the rise in generative AI tools, we include two proxy variables. First, *ChatGPT_Search* represents the Google Search trend for the keyword "ChatGPT," where the highest monthly search volume is benchmarked at 100, and all other months are relative values. Second, *Chat_Monthly_Visits* refers to the number of ChatGPT monthly visits, which is obtained from www.semrush.com.[3]

The manually collected news articles containing the keywords "oil" and "shock" are collected from ProQuest, from sources such as Dow Jones Institutional News, Financial Times, The New York Times, Energy Weekly News, and The Wall Street Journal (including Eastern, Asia, Europe, and Online editions) for the period 1986–2023. Table A presents a sample of the articles included in our dataset.

The three principal components are the orthogonalized first three components derived from a principal component analysis (PCA) conducted on 12 variables extracted using LIWC-22 from news articles containing the term "oil shock." LIWC-22 is a text analysis tool that analyzes linguistic features of language.[4] Among the variables, *wc* represents the word count, the remaining variables summarize linguistic styles, such as logical and formal thinking, leadership and status cues, perceived honesty and genuineness, tone (positive or negative), average words per sentence, percentage of words with six or more

[2] We use version liwc-22.
[3] We currently focus on ChatGPT because it held the dominant market share during our sample period and was also the earliest influential AI chatbot in widespread use. Other important AI tools, such as Claude, entered the market later, with Claude becoming available only in July 2023 while our sample ends in September 2023.
[4] See https://www.liwc.app/ for details on the text analysis tools. Refer to the LIWC-22 Psychometrics Manual for technical information about the variables.

letters, and the percentage of words captured by LIWC. The variables *function*, *pronoun*, and *ppron* represent total function words, total pronouns, and personal pronouns, respectively, within the text. The first three components together explain approximately 62.20% of the cumulative variance in the 12 LIWC-22 variables. Table B reports the loadings of the three principal components on these 12 LIWC-22 variables. We find that *F1* has the highest loadings on *tone, dictionary, linguistic*, and *function,* suggesting that *F1* best captures changes in writing style on emotional tone, vocabulary richness, and general linguistic usage. *F2* loads most heavily on *analytic, tone, pronoun,* and *ppron*, indicating that it reflects more logical, formal sentence and the use of personal references such as "I" and "you." *F3* shows the highest loadings on *clout* and *authentic*, suggesting authoritative language, perceived leadership, and genuineness.

Insert Table 1 about here

4.2. Statistical Summaries

Table 1 summarizes the statistics of the variables used in our analysis. The first 12 variables are LIWC writing style variables. On average, each news article contains approximately 1,360 words. The average number of words per sentence is around 33. The sentiment variable, *Disagreement*, is defined as the percentage of negative minus positive words, based on the Loughran-McDonald dictionary. Its average value is 0.025. A positive value indicates a generally negative tone in news articles related to oil shocks. The daily returns of WTI oil average -0.1%, with a volatility of 4.9%. The return distribution shows slight positive skewness (0.584) but exhibits a high kurtosis value of 24.617, indicating that oil returns are fat-tailed and contain extreme values during the sample period from 1986 to 2023. Figure 2 illustrates oil returns and sentiment over the sample period (1986-2023). Around 2020, there is a noticeable spike and a trough, reflecting the shocks to the oil market at the onset of the COVID-19 pandemic. The graphs reveal a slight negative correlation between oil returns and sentiment: when oil returns spike, sentiment tends to decrease, suggesting more positive tones during those periods.

Insert Figure 1 about here

Insert Figure 2 about here

In Panel B, most of our news article collection comes from Dow Jones Institutional News, which contributes 25,234 articles, accounting for 73.74% of our sample. Wall Street Journal Eastern Edition, the traditionally circulated printed version of the news, provides 7.49% of the articles. The Financial Times contributes 5.74%. The remaining news sources each account for less than 5% of the sample, is only a small portion of the overall dataset. This distribution is consistent with the article counts shown in Figure 1. Most news articles from Dow Jones Institutional News appear from around 2008 and continue through 2023, marking the period with the highest frequency of news coverage in our sample.

Insert Figure 3 about here

Insert Figure 4 about here

Panel C presents the correlation matrix, which shows that all principal components (*F1, F2*, and *F3*) are positively correlated with one another, indicating that they move together to some extent. Additionally, the three measures related to the release and popularity of ChatGPT exhibit high positive correlations, with all coefficients exceeding 94%. The correlation between ChatGPT search trends and monthly visits reaches 96.6%, suggesting a strong association between public interest (as measured by search volume) and actual usage of the tool. However, we observe that the principal components (*F1*, *F2*, and *F3*) are not strongly correlated with the ChatGPT-related measures. The highest correlation coefficient between the two groups is -0.113, indicating only a weak (and negative) relationship. Figure 3 displays the ChatGPT release-related measures starting from 2020. Prior to 2020, all three measures are either zero or missing. Following its release, monthly visits to ChatGPT increased abruptly, while Google search trends for ChatGPT surged significantly in 2021 and 2022 but began to decline thereafter. As shown in Figure 4, the second and third principal components do not exhibit any observable trends or patterns following the release of the AI tool. However, Panel A of Figure 4 reveals a clear decreasing trend in the first principal component (*F1*) after 2010, with another noticeable drop occurring after 2021.

Insert Table 2 about here

4.3. Empirical Model

Our main results estimate the following equation:

$$Oil\_Ret_t = \beta_0 + \sum_{\tau=1\ to\ 3} \beta_\tau\, Oil\_Ret_{t-\tau} + \beta_4\, Disagreement_{i,t} + \beta_5\, F1_{i,t} + \beta_6\, F2_{i,t} + \beta_7\, F3_{i,t} + \beta_8\, Disagreement_{i,t} \times F1_{i,t} + \beta_9\, Disagreement_{i,t} \times F2_{i,t} + \beta_{10}\, Disagreement_{i,t} \times F3_{i,t} + \beta_{11}\, GPT\_Release_t + \beta_{12}\, Disagreement_{i,t} \times GPT\_Release_t + \mu_j + \theta_y + \varepsilon_t \quad (2)$$

where Oil_Ret_t is the daily oil return on WTI on date t, $Disagreement_{i,t}$ is the Loughran-McDonald dictionary-based sentiment disagreement for news piece i on date t, $F1_{i,t}$, $F2_{i,t}$, and $F3_{i,t}$ are the first three principal components from the PCA analysis using the 12 LIWC variables on news piece i on date t, $GPT\_Release_t$ is a measure indicating ChatGPT-related events, such as a dummy for the release of ChatGPT, monthly visits to ChatGPT, or relative Google Trends search volume for "ChatGPT" on date t, μ_j is the publication fixed effect, θ_y is the year fixed effect, and $\varepsilon_{i,t}$ is the error term.

5. Results and Discussion

5.1. Writing Style Changes in Financial News

In Column (1), we examine the relationship between oil returns and news sentiment (measured by *Disagreement*). The results show a statistically significant coefficient of -0.0392 at the 10% level, indicating that a 1% increase in relative negative tone in the news is associated with a -0.0392% change in daily oil returns. Column (2) presents results controlling for three lags of oil returns. The relationship between sentiment and oil returns remains consistent, with a 1% increase in negative tone associated with a -0.0423% change in daily oil returns. In Column (3), the interaction between *F2* (the second principal component of writing style) and *Disagreement* shows a statistically significant and positive effect on oil returns. In Columns (4), we progressively control for up to three lags of oil returns. Even after accounting for these controls, we find consistent evidence that the interaction between *Disagreement* and *F2* remains significant, with a coefficient of 0.0218, statistically significant at the 10% level.

5.2. Release of ChatGPT and the Changes in the Sentiment Effect on Oil Returns

To investigate how changes in writing styles and the release of ChatGPT have influenced the effect of news sentiment on oil returns, we conduct regressions in Table 3. Specifically, we test the impact of sentiment and ChatGPT release measures on daily oil returns, as well as the interaction between sentiment and the release indicators. Column (1) reveals that *Disagreement* has a significant negative impact on oil returns: a 1% increase in relative negative tone leads to a -0.0427% change in daily oil returns. Additionally, the ChatGPT release, included as a dummy variable, reduces the strength of this impact. After the release of ChatGPT, the impact of negative sentiment decreases by an average of 0.124%, suggesting that the influence of sentiment on oil returns weakened post-release. This indicates that the release of ChatGPT mitigates the effect of news article sentiment on oil returns. In Column (2), the direct effect of *Disagreement* remains significant, with a coefficient of -0.0457, indicating that a 1% increase in negative sentiment is associated with a -0.0457% decline in daily oil returns. Furthermore, after the release of ChatGPT, this impact is reduced by 0.121%, confirming the dampening role of AI-generated language influence.

In Column (3), we replace the ChatGPT release dummy with the ChatGPT search trend as a continuous measure. The variable *Disagreement* remains negative and statistically significant, with a coefficient of -0.0447%, indicating that a 1% increase in negative sentiment leads to a corresponding decrease in daily oil returns. The ChatGPT search trend variable is also statistically significant at the 10% level, with a coefficient of 0.00260, suggesting that higher public interest in ChatGPT is associated with a reduction in the impact of negative sentiment on oil returns. To compare the relative impact of writing style and ChatGPT release measures, we conduct a horse race test in Columns (4) and (5). In Column (4), we use the ChatGPT release dummy, while in Column (5), we use the ChatGPT search trend. Both specifications yield results that are consistent with those in Columns (2) and (3), confirming the robustness of the findings. However, when writing style components are included in the model (Columns 4 and 5), the *Disagreement* sentiment measure is no longer statistically significant. Instead, the interaction between *Disagreement* and *F3* (the third principal component) becomes statistically significant at the 10% level, suggesting that changes in writing style may mediate the relationship between sentiment and oil returns. In summary, the results in Table 3 suggest that changes in writing style have a modest influence on the relationship between news sentiment and oil returns. More importantly, the release of the ChatGPT AI tool demonstrates a robust and statistically significant effect on how sentiment influences oil returns.

Insert Table 3 about here

5.3. The Sentiment Effect on Oil Returns and the Popularity of ChatGPT

From Table 3, the impact of ChatGPT release measures on the relationship between news sentiment and oil returns appears most significant when using the ChatGPT release dummy, which equals one for the period after the tool's launch on November 30, 2022. In Table 4, we replace the dummy variable with a more granular measure, ChatGPT monthly visit counts (in millions), to better capture the usage and popularity of the AI tool. First, Columns (1)–(3) report simple regressions of the first principal components on monthly ChatGPT visits. Since these specifications are based on time-series data rather than panel data, we do not include year or publication fixed effects, and standard errors are not clustered. The results show that F1 and F3 are significantly associated with monthly ChatGPT visits. In Column (4), an increase of 0.14 million in ChatGPT monthly visits (~0.14 million is average monthly increase in visits) is associated with a 0.0138 reduction in oil returns, suggesting a 30.6% reduction relative to the baseline coefficient on *Disagreement* of -0.0451. With sample mean of disagreement (0.025), this implies a smaller daily return response of $0.0983 \times 0.14 \times 0.025 = 0.000344$, or 0.0344% (3.44 basis points). By comparison, in the absence of ChatGPT visits, the baseline effect of mean disagreement on oil returns is $-0.0451 \times 0.025 = -0.001128$, or -0.1128% (-11.28 basis points). The economic magnitude of *Disagreement* is nontrivial: a 0.1 increase in disagreement (when relative negativity in news article tone increases by 10%) is associated with a 0.451% decline in daily oil returns. Relative to the standard deviation of daily oil returns (4.9%), this corresponds to about 9.2% of one standard deviation. In Column (5) of Table 4, we compare the relative influence of writing style changes and ChatGPT usage (measured by monthly visit counts) on the relationship between news sentiment and oil returns. We find that the interaction between *Disagreement* and the third principal component (*F3*) is statistically significant, with a coefficient of -0.0176%, significant at the 10% level. This indicates that changes in writing style captured by *F3* influence how sentiment affects oil returns. These findings suggest that the ChatGPT usage measure is more statistically significant and economically meaningful than the writing style component, implying that the growing use of ChatGPT has a stronger moderating effect on the influence of news sentiment in the oil market than shifts in writing style alone.

Insert Table 4 about here

6. Robustness

6.1. Covid-19

While the previous Tables 2-4 suggest that the release of ChatGPT has reduced the impact of news sentiment on daily oil returns, it is important to address potential sample period concerns. Specifically, the influence of ChatGPT primarily occurs in the post-2020 period, whereas our sample begins in the 1980s. Moreover, the year 2020 marks the onset of the COVID-19 pandemic, during which the oil market experienced heightened volatility, most notably the negative oil futures price in April 2020. To account for these factors, Table 5 presents robustness checks: In Columns (1), (3), and (5), the sample is restricted to post-2020 data. In Columns (2), (4), and (6), the year 2020 is excluded to eliminate the extreme shocks from the COVID-19 onset. In Column (1), the ChatGPT monthly visits variable still shows a significant moderating effect: a 1% increase in negative sentiment leads to a 0.0187% smaller impact on daily oil returns, with statistical significance at the 1% level. When we exclude 2020 in Column (2), the effect remains robust, with a 0.0105% reduction in the influence of sentiment, again significant at the 1% level. In Columns (3) and (4), we use an alternative sentiment measure, the daily aggregated sum of *Disagreement* across news articles. In

Columns (5) and (6), we use the daily average of sentiment scores. In Columns (7) and (8), we use article length-weighted average of sentiment scores. Across both sets of models, the interaction terms between the alternative sentiment measures and ChatGPT monthly visits remain statistically significant at the 1% level. These interactions consistently show a reduction in the impact of sentiment on oil returns, reinforcing our primary findings.

Insert Table 5 about here

In Columns (9) and (10), we report regression results that include both the ChatGPT monthly visit counts and the three principal components of writing style in the analysis. Additionally, we control for weekday fixed effects in both columns to account for variations across days of the week. In Column (10), we also cluster standard errors by weekday, as Wednesdays and Thursdays are often associated with major announcements in the oil market, and some macroeconomic news is typically released on specific days. The results remain consistent with our previous findings. The interaction between *Disagreement* and ChatGPT monthly visit counts continues to have a statistically significant effect on daily oil returns. Specifically, a 0.1 million increase in ChatGPT monthly visits (about the average monthly visits in our sample) is associated with a 0.0129% reduction in the impact of sentiment on oil returns. This effect is statistically significant at the 1% level in Column (9) and at the 5% level in Column (10). These findings reinforce the robustness of the relationship between ChatGPT usage and the diminished influence of sentiment in oil-related news on market behavior.

6.2. Alternative explanation: Supply side vs. demand side

One possible alternative explanation is that the diminishing effect of oil-shock news reflects a demand-side mechanism: investors may increasingly rely on ChatGPT to access financial information, rather than the writing style of financial news being altered by ChatGPT usage. In this study, we first document a decline in the sentiment effect and then propose one possible explanation related to AI tool adoption. Overall, however, our evidence is more supportive of the view that shifts in writing style contribute to the weakening of sentiment impact.

First, readership of major financial newspapers does not appear to be declining. As shown in Figure A in the Appendix, the number of paid subscriptions to *The Wall Street Journal* and *The New York Times* has continued to rise over time. We manually collect subscription data for these two major financial newspapers, which publicly disclose such information in their financial reports (10-K and 10-Q filings). For all years in which the data are available, subscriptions to both newspapers display an upward rather than downward trend. This pattern is notable because internet-based news has made financial information more accessible, yet both newspapers have successfully expanded their digital versions, and a large share of *The Wall Street Journal* subscriptions now comes from digital subscriptions. Taken together, this evidence suggests that the circulation of major financial newspapers has not been declining.

Second, internet access was not available in ChatGPT until March 2023, with only very limited access at first. Web browsing was later expanded to paid users in May 2023 and was temporarily shut down from July to September 2023, while our sample ends in September 2023.[5] Before the release of the web-browsing plugin, ChatGPT was not able to access same-day information, since the model was typically trained on earlier data. Because our analysis focuses on contemporaneous oil return responses, the effect is unlikely to be driven by investors obtaining same-day oil-shock news or related information through ChatGPT.[6] For most readers, access to internet searching through ChatGPT would only have become available after May 2023.

Third, an alternative explanation is that the diminishing sentiment effect is driven by factors other than writing style, such as macroeconomic events, geopolitical events, or fundamental shocks that dominate sentiment signals. To address this concern, we perform a DiD robustness check by focusing on a unique series of news articles, the *Market Talk* series in *Dow Jones Institutional News*, which follows a more fixed and less flexible reporting template. The *Market Talk* series follows a relatively standardized brief-news format. A typical item contains a timestamp, a short summary of a market-relevant development, and an attributed statement from a cited professional source.[7] Because these pieces rely heavily on brief, formulaic structure and attributed quotations, they leave comparatively less room for idiosyncratic reporter writing style than longer narrative articles. We therefore classify *Market Talk* articles as standard news and the remaining articles in *Dow Jones Institutional News* as adaptive news, since their writing styles are less restricted by a fixed reporting template.

We then separately report the first three principal components from the PCA of the LIWC-22 writing scores and plot them in Figure 5, with F1, F2, and F3 shown in Panels A, B, and C, respectively. The graphs show the daily averages of the adaptive principal components and the standard principal components for the writing-style measures. The standard news is relatively more stable than the adaptive news across F1,

[5] See ChatGPT release notes: https://openai.com/index/chatgpt-plugins/.
[6] As a robustness check, we focus only on *Dow Jones Institutional News* articles published before March 23, 2023. The results continue to show a diminished sentiment effect and are reported in Column (1) of Table C in the Appendix.
[7] The following stylized example illustrates the template and is not reproduced from any specific article.
10:15 AM: "Oil prices rose after supply concerns intensified", according to a market analyst.

F2, and F3. Note that LIWC-22 is measured by the percentage of words in a given category, so article length would only potentially reinforce the volatility of the principal components. However, we find that the standard news is more stable, and we do not observe any significant structural or gradual changes in the writing-style components after 2022.

We then formally test whether there is heterogeneity in the sentiment effect moderated by ChatGPT usage. We run the following regression:

$$
\begin{aligned}
Oil_R\_Ret_t = \beta_0 &+ \sum\nolimits_{\tau=1\ to\ 3} \beta_\tau\, Oil_R\_Ret_{t-\tau} + \beta_4\, Disagreement_{i,t} + \beta_5\, Standard_{i,t} \\
&+ \beta_6\, Disagreement_{i,t} \times\ Standard_{i,t} + \beta_7 ChatGPT\_Monthly\_Visits_t \\
&+ \beta_8\, Standard_{i,t} \times\ ChatGPT\_Monthly\_Visits_t \\
&+ \beta_9\, Disagreement_{i,t} \times ChatGPT\_Monthly\_Visits_t \\
&+ \beta_{10}\, Standard_{i,t} \times Disagreement_{i,t} \times ChatGPT\_Monthly\_Visit_t +\ \mu_j + \theta_y + \varepsilon_t
\end{aligned}
$$

The three-way interaction term, $\beta_{10}\, Standard_{i,t} \times Disagreement_{i,t} \times ChatGPT\_Monthly\_Visits_t$, tests whether the effect of ChatGPT usage on the sentiment-return relation differs between standard news and other news, where $Standard_{i,t} = 1$ for standard news, defined as the *Market Talk* series. We report the results in Table 6. Since *Market Talk* first became available on December 13, 2016, we repeat the regression in column (2) using the subsample after December 13, 2016. In Table 6, we include lags of oil returns because oil returns are highly serially correlated. We continue to find that disagreement is significantly related to oil returns: more negative tone is associated with lower oil returns in column (1) for the full sample and remains negative in the post-2016 sample. We also continue to find that the sentiment effect diminishes as ChatGPT usage increases, as the coefficient on $Disagreement_{i,t} \times ChatGPT\_Monthly\_Visits_t$ is positive and significant at the 5% and 10% levels. Finally, the three-way interaction term, $Standard_{i,t} \times Disagreement_{i,t} \times ChatGPT\_Monthly\_Visits_t$, tests whether the standard news series experiences a smaller decline in sentiment impact as ChatGPT usage grows. The results support this view. The coefficient is negative and significant at the 1% level in both the full sample and the post-2016 subsample. The magnitude of the three-way interaction almost fully offsets the effect of $Disagreement_{i,t} \times ChatGPT\_Monthly\_Visits_t$ in the full sample and offsets about 70% of it in the post-2016 subsample. These estimates indicate that, relative to other news articles, standard news experiences little or no decline in sentiment impact as ChatGPT usage grows. If the diminished sentiment effect were driven primarily by a demand-side channel, that is, if readers increasingly used ChatGPT to access financial information rather than reading the original news, we would not expect to observe a significantly weaker impact in the diminishing effect for news written in a relatively fixed reporting format.

Insert Table 6 about here

6.3. Alternative explanation: Changes in article characteristics

An alternative explanation for the diminished sentiment effect is that the characteristics of oil-market news articles may have changed over time. For example, major market events may have been more concentrated in earlier periods, whereas more recent oil-market news may be shorter and contain less information. In addition, recent market conditions may be more uncertain and ambiguous, which could mechanically weaken the price impact of news after 2022. In other words, the significantly weakened effect of disagreement as ChatGPT usage increases may reflect changes in underlying market conditions or article characteristics, rather than the influence of AI tool adoption itself.

To address this concern, we conduct a robustness analysis by splitting the sample along three article characteristics. We extend the sample through the end of 2024 to capture a fuller picture of these shifts over time. In addition, we focus on *Dow Jones Institutional News* to mitigate variation in article characteristics across publications. The three article characteristics are: (1) longer versus shorter news articles; (2) harder-to-read versus easier-to-read articles, proxied by a readability index; and (3) high-uncertainty versus low-uncertainty articles, measured using the fraction of words classified as "uncertainty" terms in the Loughran and McDonald dictionary. Article length is measured by word count, and *Long_Article* is an indicator equal to one if an article's word count is above the median word count in the pre-ChatGPT period. Readability difficulty is constructed as the sum of the standardized values of *Words per Sentence* and *BigWords*, where *Words per Sentence* is the average number of words per sentence in article *i*. The standardization uses the mean and standard deviation from the pre-ChatGPT period. *BigWords* is the percentage of words containing seven or more letters.

$$Readability = \frac{Word\ per\ Sentence_i - \overline{Word\ per\ Sentence}}{\sigma(Word\ per\ Sentence)} + \frac{Big\ Word_i - \overline{Big\ Word}}{\sigma(Bid\ Word)}$$

Hard_Read is an indicator equal to one if the readability-difficulty index is above the pre-ChatGPT median. Finally, article-level uncertainty is measured as the percentage of words matched to the Loughran-McDonald uncertainty dictionary, and *High_Uncertainty* equals one if this measure is above the pre-ChatGPT median. Thus, all of our article-characteristic dummies are defined relative to the median values in the pre-ChatGPT period. This construction allows us to identify subsamples of news articles that are directly comparable with the pre-ChatGPT benchmark and that are relatively more "information-rich" than the pre-ChatGPT sample median.

If changes in market conditions are driving variation in article length, readability, or uncertainty in the news, then ChatGPT usage may simply proxy for these article characteristics. Formally, suppose that

$$\text{Characteristic Measure}_t = \alpha + \beta \text{ Market Conditions}_t + u_t$$

and

$$\text{ChatGPT Visits}_t = \partial + \gamma \text{ Market Conditions}_t + e_t$$

Then it follows that

$$\text{Characteristic Measure}_t = \eta + \delta \text{ ChatGPT Visits}_t + \theta_t$$

Under this explanation, once we add these article-characteristic measures to the regression, the significance of ChatGPT visits should disappear, as these measures would absorb the explanatory power associated with market conditions. However, we do not find such evidence.

We report the results in Table 7. The estimates show that Disagreement continues to have a negative effect on oil returns, while the interaction term between Disagreement and Chat_Monthly_Visits remains positive and significant in Columns (1) and (2), suggesting that greater ChatGPT usage is associated with a weaker sentiment effect. However, we do not find evidence that this moderating effect differs systematically across article characteristics. Specifically, the three-way interaction terms, *Characteristic Dummy × Disagreement × Chat_Monthly_Visits*, are never statistically significant. Overall, these results suggest that the diminished sentiment effect associated with ChatGPT usage is unlikely to be driven by changes in article characteristics or by fundamental shifts in market conditions that in turn affect those characteristics.

To summarize, our evidence is more consistent with the view that the diminished sentiment effect is driven by changes in the depiction and narration of news, rather than by declining newspaper readership, the availability of web browsing after May 2023, or other macroeconomic events during the sample period.

Insert Table 7 about here

6.4. Expanded Dow Jones Institutional News Sample

Our baseline results use the full newspaper sample. Because Dow Jones Institutional News is only available after 2008 but accounts for the majority of observations, we report DJI-only results as an additional robustness check. Column (2) of Table C in the Appendix re-estimates the main specification using only DJI articles and an expanded sample through the end of 2024, which increases the number of observations after the release of ChatGPT. We continue to find that the sentiment effect weakens as ChatGPT usage grows.

7. Conclusion

In this study, we examine the impact of ChatGPT's release on writing style changes in news articles covering oil shocks. Our dataset consists of articles collected from multiple popular sources, including Dow Jones Institutional News and the Wall Street Journal. We find that the release of ChatGPT has had a significant influence on news writing styles. These changes include the use of more analytical language and logical reasoning, increased use of power-related words, and more frequent use of words indicating honesty and authenticity. Sentiment analysis has long been an established field in financial literature, with studies demonstrating the impact of tone in news articles on security prices. We contribute to this literature by showing the influence of AI tool, ChatGPT, on news content creation. Our results show that AI writing tool release is associated with reduced sensitivity of financial markets to news sentiment. In other words, the release of AI tools like ChatGPT appear to weaken the link between news tone and asset prices.

We further find that the greater the public engagement with ChatGPT, as measured by its monthly visit counts, the smaller the effect of negative sentiment in news articles on daily oil returns. This suggests that the increased reliance on AI-generated or AI-edited news content may lead markets to respond less sharply to sentiment cues, possibly due to perceived changes in the authenticity or trustworthiness of such content. Our findings are aligned with those of Buchanan & Hickman (2024), who showed that experimental participants were less likely to trust articles they believed were written by AI. This raises important considerations, not only about the benefits of AI in enhancing productivity and writing efficiency, but also about the potential risks to traditional publishing models and the flow of information in financial markets.

References

Altay, S., & Gilardi, F. (2024). People are skeptical of headlines labeled as AI-generated, even if true or human-made, because they assume full AI automation. *PNAS nexus*, *3*(10), pgae403.

Lermann Henestrosa, A., & Kimmerle, J. (2024). The effects of assumed AI vs. human authorship on the perception of a GPT-generated text. *Journalism and Media*, *5*(3), 1085-1097.

Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The Review of Financial studies*, *21*(2), 785-818.

Buchanan, J., & Hickman, W. (2024). Do people trust humans more than ChatGPT?. *Journal of Behavioral and Experimental Economics*, 112, 102239.

Cao, S., Jiang, W., Yang, B., & Zhang, A. L. (2023). How to talk when a machine is listening: Corporate disclosure in the age of AI. *The Review of Financial Studies*, *36*(9), 3603-3642.

Calomiris, C. W., & Mamaysky, H. (2019). How news and its context drive risk and returns around the world. *Journal of Financial Economics*, *133*(2), 299-336.

Chan, W. S. (2003). Stock price reaction to news and no-news: drift and reversal after headlines. *Journal of Financial Economics*, *70*(2), 223-260.

Engelberg, J. E., & Parsons, C. A. (2011). The causal impact of media in financial markets. *the Journal of Finance*, *66*(1), 67-97.

Fang, L., & Peress, J. (2009). Media coverage and the cross-section of stock returns. *The Journal of Finance*, *64*(5), 2023-2052.

Feldman, R., Govindaraj, S., Livnat, J., & Segal, B. (2010). Management's tone change, post earnings announcement drift and accruals. *Review of Accounting Studies*, *15*, 915-953.

Garcia, D. (2013). Sentiment during recessions. *The Journal of Finance*, *68*(3), 1267-1300.

Jiang, F., Lee, J., Martin, X., & Zhou, G. (2019). Manager sentiment and stock returns. *Journal of Financial Economics*, *132*(1), 126-149.

Klibanoff, P., Lamont, O., & Wizman, T. A. (1998). Investor reaction to salient news in closed-end country funds. *The Journal of Finance*, 53(2), 673-699.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315-1335.

Larcker, D. F., & Zakolyukina, A. A. (2012). Detecting deceptive discussions in conference calls. *Journal of Accounting Research,* 50(2), 495-540.

Longoni, C., Fradkin, A., Cian, L., & Pennycook, G. (2022, June). News from generative artificial intelligence is believed less. *In Proceedings of the 2022 ACM Conference on Fairness, Accountability*, and Transparency (pp. 97-106).

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, *66*(1), 35-65

Loughran, T., & McDonald, B. (2014). Measuring readability in financial disclosures. *the Journal of Finance,* 69(4), 1643-1671.

Loughran, T., & McDonald, B. (2024). Measuring firm complexity. *Journal of Financial and Quantitative Analysis,* 59(6), 2487-2514.

Malik, A. R., Pratiwi, Y., Andajani, K., Numertayasa, I. W., Suharti, S., & Darwis, A. (2023). Exploring artificial intelligence in academic essay: higher education student's perspective. *International Journal of Educational Research Open,* 5, 100296.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, 42, 483-510.

Novy-Marx, R., & Velikov, M. Z. (2025). AI-powered (finance) scholarship (No. w33363). *National Bureau of Economic Research.*

Peng, L., & Xiong, W. (2006). Investor attention, overconfidence and category learning. *Journal of financial Economics, 80(3), 563-602.*

Reis, M., Reis, F., & Kunde, W. (2024). Influence of believed AI involvement on the perception of digital medical advice. *Nature Medicine*, *30*(11), 3098-3100.

Shi, W., Zhang, Y., & Hoskisson, R. E. (2019). Examination of CEO–CFO social interaction through language style matching: Outcomes for the CFO and the organization. *Academy of Management Journal*, 62(2), 383-414.

Solomon, D. H. (2012). Selective publicity and stock prices. *The Journal of Finance*, *67*(2), 599-638.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, *62*(3), 1139

Tetlock, P. C. (2010). Does public financial news resolve asymmetric information?. *The Review of Financial Studies*, *23*(9), 3520-3557.

Tetlock, P. C. (2011). All the news that's fit to reprint: Do investors react to stale information?. *The Review of Financial Studies*, *24*(5), 1481-1512.

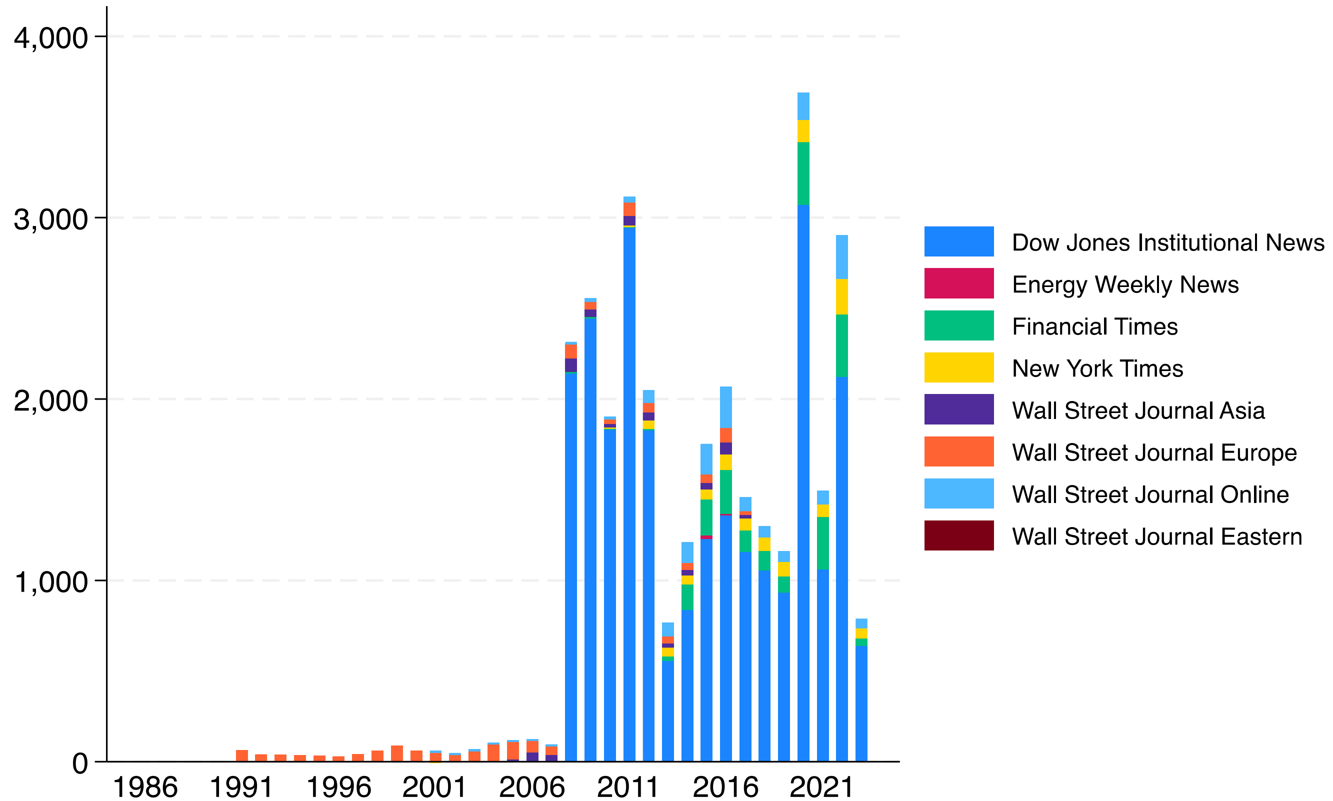

Figure 1. Count of News Articles on Oil Shocks by Publication Source and Year

Figure 1 shows the number of news articles containing the keywords "oil" and "shock," collected from various publication sources and plotted over time by year.

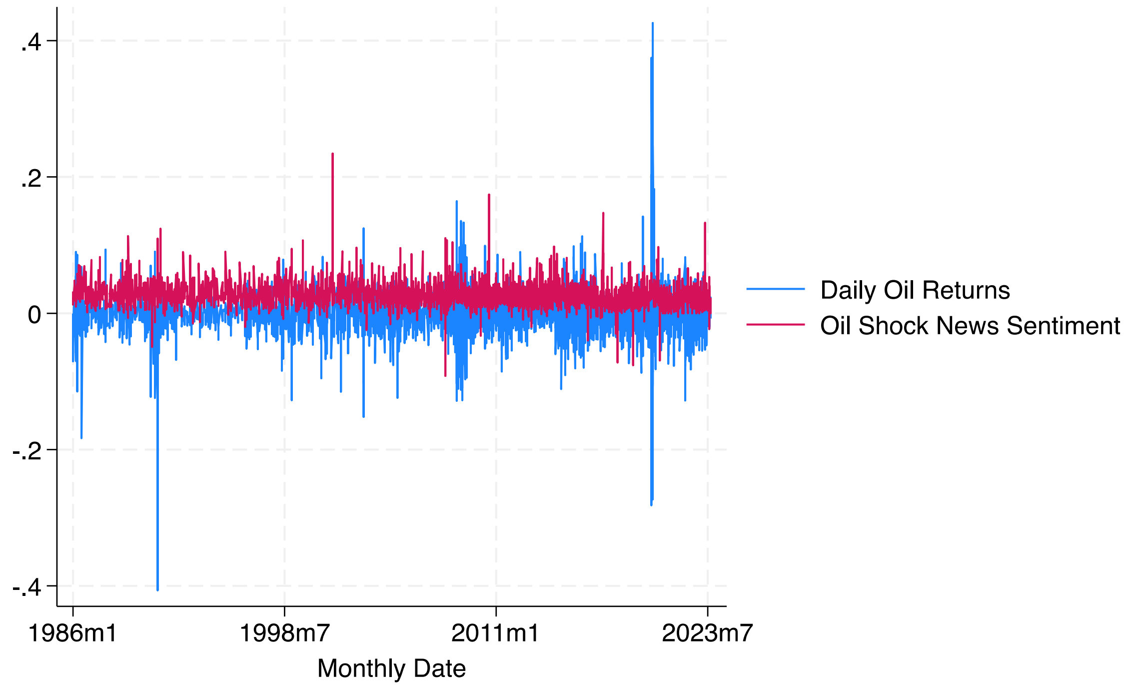

Figure 2. WTI Daily Oil Returns and Loughran-McDonald Dictionary-Based Sentiment on Oil Shock News

Figure 2 plots WTI daily oil returns, calculated from WTI futures prices, and sentiment scores on oil shock news. Sentiment is measured as the percentage of negative to positive words using the Loughran-McDonald dictionary. News articles were collected from Dow Jones Institutional News, Financial Times, The New York Times, Energy Weekly News, and The Wall Street Journal. Both variables are displayed as monthly averages.

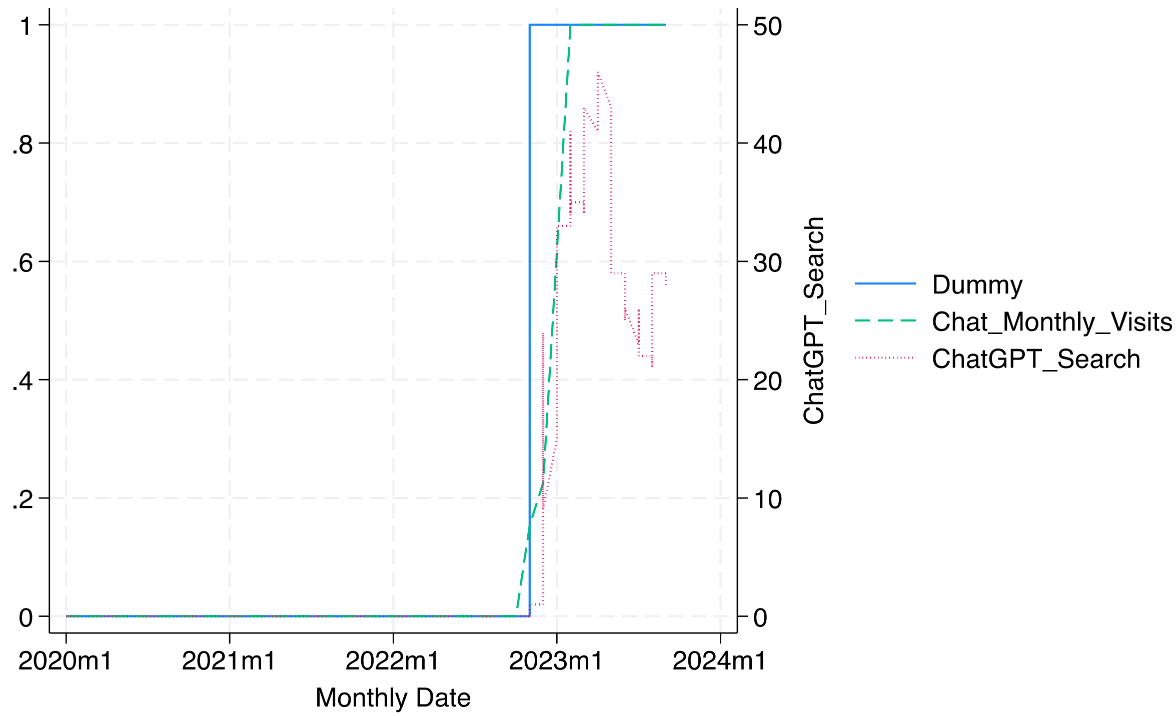

Figure 3. ChatGPT Release Date and Trends

This figure plots three variables: a dummy variable equal to one on the release date of ChatGPT (*Dummy*), *Chat_Monthly_Visits* is the number of visits to ChatGPT (in millions), and *ChatGPT_Search*, which is the Google relative search trend where 100 indicates the peak search interest in history. All variables are plotted as monthly averages.

Panel A: First Principal Component of LIWC Writing Style Measures

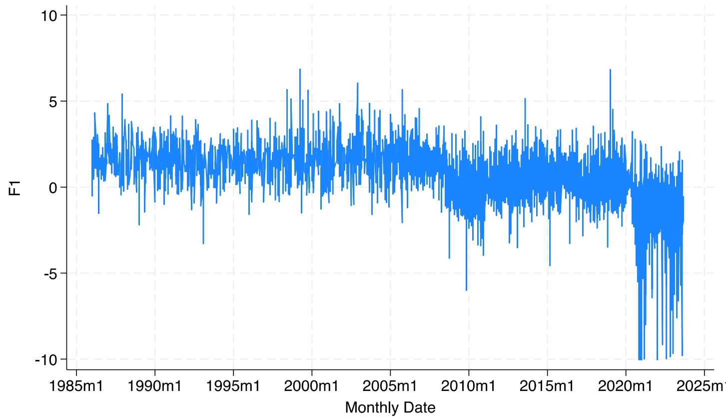

Panel B: Second Principal Component of LIWC Writing Style Measures

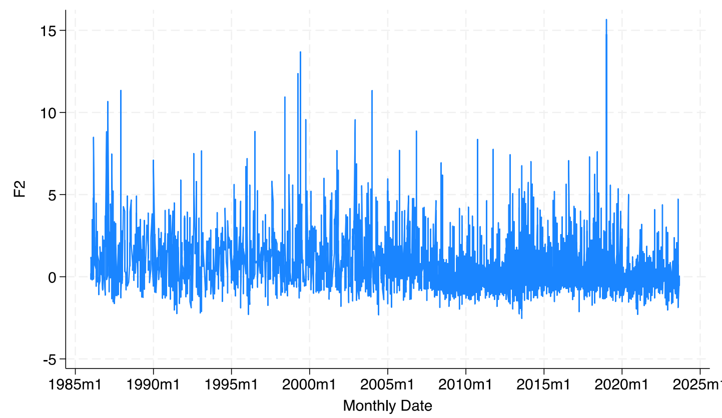

Panel C: Third Principal Component of LIWC Writing Style Measures

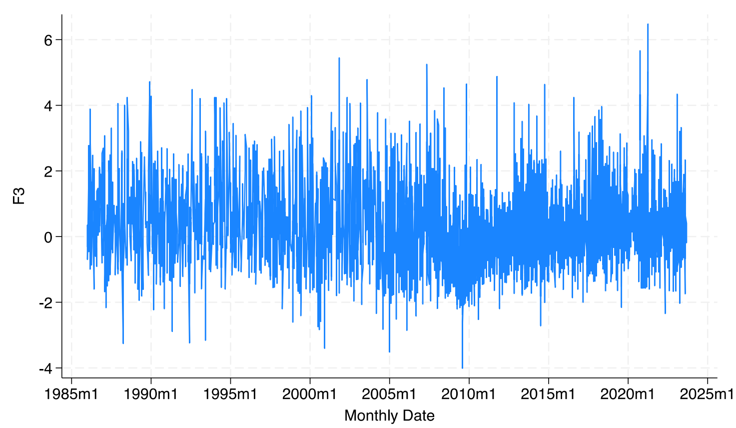

Figure 4. Three Principal Components of LIWC Variables in Oil Shock News Articles

Figures 4.1, 4.2, and 4.3 plot the first, second, and third orthogonal components extracted from principal component analysis (PCA) of 12 LIWC variables that capture writing styles in news articles containing the keywords "oil" and "shock." The articles were collected from Dow Jones Institutional News, Financial Times, The New York Times, Energy Weekly News, and The Wall Street Journal, covering the period from 1986 to 2023.

Panel A. First Principal Component (F1)

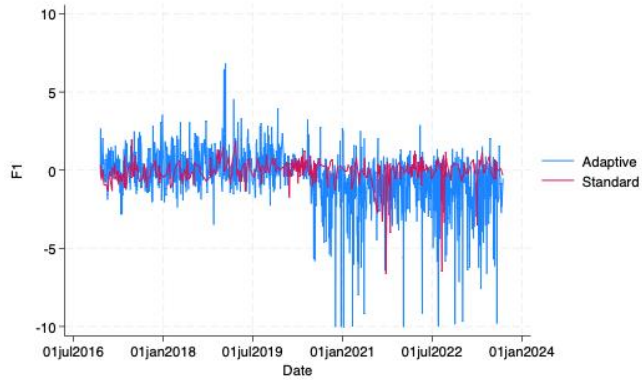

Panel B. Second Principal Component (F2)

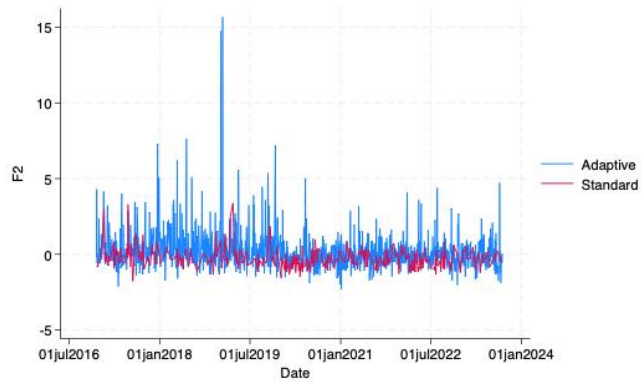

Panel C. Third Principal Component (F3)

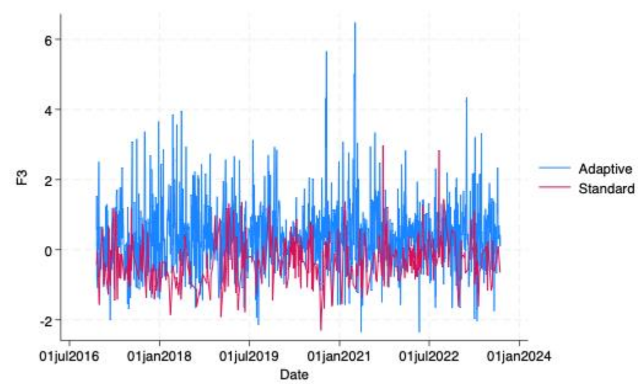

Figure 5. Writing Style Dynamics in Standard and Adaptive Oil News

This figure compares the time-series variation in the first three principal components of writing style, F1, F2, and F3, between standard and adaptive oil-related news articles. Standard news refers to articles in the recurring *Market Talk* format, which follows a relatively fixed and standardized template. Adaptive news refers to other articles that do not follow this preset format and therefore allow more flexible language. The three principal components are constructed from the LIWC-22.

Table 1. Summary Statistics

Table 1 summarizes the variables used in our tests. The first 12 variables are derived from LIWC. *Disagreement* is based on the Loughran-McDonald dictionary sentiment scores, calculated as the percentage of negative words minus positive words. *Agg_Disagreement* is the sum of all news articles' disagreement scores on the trading date. *Mean_Disagreement* is the daily average of news article disagreements. *F1*, *F2*, and *F3* are the first three principal components extracted from the principal component analysis of the 12 LIWC variables. *Oil_Ret* refers to the daily returns of WTI crude oil. *Oil_Ret_L1*, *Oil_Ret_L2*, and *Oil_Ret_L3* are the first-, second-, and third-order lags of the oil returns, respectively. *Dummy* equals 1 for dates after November 30, 2022. *ChatGPT_Search* represents the Google relative search trend for the keyword "ChatGPT." *Chat_Monthly_Visits* is the number of visits to ChatGPT (in millions), collected from www.semrush.com. The number of news articles containing the keywords "oil" and "shock" from various sources is also reported.

Panel A: Variable	Obs	Mean	Std. dev.	Min	Max	Skewness	Kurtosis	5% Percentile	95% Percentile
wc	34,226	1360.475	1694.193	3.000	54220.000	13.250	302.633	300.000	3079.000
analytic	34,226	94.413	5.270	10.190	99.000	-4.657	40.572	85.700	98.550
clout	34,210	41.487	11.512	1.000	98.220	0.940	4.981	25.400	62.980
authentic	34,226	45.226	18.184	1.000	99.000	0.179	2.526	16.600	77.030
tone	34,196	26.101	14.045	1.000	99.000	1.695	9.470	7.110	47.440
wps	34,226	33.026	29.080	3.000	804.000	9.232	133.343	19.040	69.190
bigwords	34,226	26.698	4.681	8.220	66.670	0.825	4.159	20.340	35.340
dic	34,226	77.835	4.407	37.140	100.000	-0.838	6.436	71.540	84.690
linguistic	34,226	52.596	6.589	12.660	75.780	-1.246	8.495	44.190	62.330
function	34,226	38.156	6.260	4.670	60.940	-1.227	8.448	29.960	47.320
pronoun	34,226	4.002	1.844	0.000	21.370	1.630	8.483	1.700	7.250
ppron	34,226	1.435	1.199	0.000	17.640	2.795	15.510	0.250	3.610
Disagreement	34,226	0.025	0.021	-0.500	0.628	3.363	69.698	0.002	0.058
Agg_Disagreement	34,226	0.460	0.586	-0.295	5.557	4.007	28.198	0.034	1.481
Mean_Disagreement	34,226	0.025	0.011	-0.092	0.234	1.509	18.490	0.011	0.043
F1	34,181	-0.001	1.847	-10.156	7.319	-1.647	11.428	-2.227	2.605
F2	34,181	-0.003	1.515	-3.532	18.124	3.041	20.206	-1.584	2.556
F3	34,181	-0.004	1.314	-4.092	13.621	0.989	6.652	-1.877	2.279
Oil_Ret	34,226	-0.001	0.049	-0.406	0.426	0.584	24.617	-0.061	0.054
Oil_Ret_L1	34,226	-0.002	0.047	-0.406	0.426	0.343	24.186	-0.063	0.048
Oil_Ret_L2	34,226	-0.001	0.045	-0.281	0.426	0.905	28.040	-0.057	0.047
Oil_Ret_L3	34,226	-0.001	0.043	-0.281	0.426	0.488	26.006	-0.056	0.045
Dummy	34,226	0.027	0.163	0.000	1.000	5.799	34.633	0.000	0.000
ChatGPT_Search	34,226	0.832	5.253	0.000	46.000	6.585	46.533	0.000	0.000
Chat_Monthly_Visits	34,226	0.024	0.147	0.000	1.000	6.256	40.834	0.000	0.000

Panel B: Publication	Freq.	Percent	Cum.	Panel C:	F1	F2	F3	Dummy	ChatGPT _Search
Dow Jones Institutional News	25,234	73.74	73.74	F1	1				
Energy Weekly News	24	0.07	73.81	F2	0.425	1			
Financial Times	1,963	5.74	79.55	F3	0.078	0.295	1		
New York Times	962	2.81	82.36	Dummy	-0.113	-0.014	0.044	1	
WSJ Asia	512	1.50	83.86	ChatGPT_Search	-0.112	-0.009	0.046	0.945	1
WSJ Eastern	2,562	7.49	91.35	Chat_Monthly_Visits	-0.110	-0.015	0.045	0.951	0.966
WSJ Europe	1,445	4.22	95.57						
WSJ Online	1,516	4.43	100.00						

Table 2. The Impact of Writing Styles on the Relationship Between Oil News Sentiment and Oil Returns

This table presents regression results examining how changes in writing styles within oil shock news articles affect the relationship between news sentiment and daily WTI oil returns. The sentiment measure (Disagreement) is based on the Loughran-McDonald dictionary and is calculated as the percentage difference between negative and positive words in the articles. News articles on "oil" and "shock" were collected from Dow Jones Institutional News, The Wall Street Journal, The New York Times, The Financial Times, and Energy Weekly News, covering the period from 1986 to 2023. Writing styles are captured by F1, F2, and F3, which are the first three orthogonal components extracted via principal component analysis of the 12 LIWC variables. We include publication-year fixed effects. Standard errors (S.E.) are clustered at the level of the corresponding fixed effects and are reported in parentheses. *, **, and *** indicate that the coefficients are significant at the 10%, 5%, and 1% significance levels, respectively.

VARIABLES	(1) Oil_Ret	(2) Oil_Ret	(3) Oil_Ret	(4) Oil_Ret
Oil_Ret_L1		-0.0457***		-0.0453***
		(0.009)		(0.009)
Oil_Ret_L2		0.0935**		0.0938**
		(0.037)		(0.037)
Oil_Ret_L3		-0.119***		-0.120***
		(0.031)		(0.032)
Disagreement	-0.0392*	-0.0423*	-0.00927	-0.00894
	(0.017)	(0.018)	(0.010)	(0.009)
F1			-0.000377	-0.00041
			(0.000)	(0.001)
F2			0.000201	0.000136
			(0.000)	(0.000)
F3			0.000207	0.000217
			(0.000)	(0.000)
Disagreement× F1			-0.0113	-0.0139
			(0.015)	(0.017)
Disagreement× F2			0.0179*	0.0218*
			(0.008)	(0.011)
Disagreement× F3			-0.015	-0.0163
			(0.009)	(0.009)
Constant	-0.00615***	-0.00635**	-0.00686***	-0.00712**
	(0.002)	(0.002)	(0.002)	(0.002)
Publication FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Clustered S.E.	Publication × Year	Publication × Year	Publication × Year	Publication × Year
Observations	34,218	34,218	34,173	34,173
R-squared	0.008	0.030	0.009	0.031

Table 3. The Relationship Between Writing Styles, the Release of ChatGPT, and the Impact of Oil News Sentiment on Oil Returns

This table reports the regression results examining the relationship between ChatGPT-related variables, Google search trends (*ChatGPT_Search*) and a dummy variable (*Dummy*) equal to one for the release of ChatGPT, and the three principal components (*F1*, *F2*, and *F3*) extracted from the principal component analysis of the 12 LIWC variables. We present the regression results of oil returns (*Oil_Ret*) on *Disagreement*, which is the sentiment score based on the Loughran-McDonald dictionary on oil shock news. We include publication-year fixed effects. Standard errors (S.E.) are clustered at the level of the corresponding fixed effects and are reported in parentheses. *, **, and *** indicate that the coefficients are significant at the 10%, 5%, and 1% significance levels, respectively.

VARIABLES	(1) Oil_Ret	(2) Oil_Ret	(3) Oil_Ret	(4) Oil_Ret	(5) Oil_Ret
Oil_Ret_L1		-0.0457***	-0.0457***	-0.0453***	-0.0453***
		(0.010)	(0.009)	(0.010)	(0.009)
Oil_Ret_L2		0.0935**	0.0936**	0.0938**	0.0939**
		(0.037)	(0.037)	(0.038)	(0.037)
Oil_Ret_L3		-0.119***	-0.119***	-0.120***	-0.120***
		(0.032)	(0.031)	(0.032)	(0.032)
Disagreement	-0.0427**	-0.0457**	-0.0447**	-0.0123	-0.0115
	(0.018)	(0.019)	(0.019)	(0.010)	(0.010)
F1				-0.000454	-0.000448
				(0.001)	(0.001)
F2				0.000161	0.000159
				(0.000)	(0.000)
F3				0.000251	0.000245
				(0.000)	(0.000)
Disagreement× F1				-0.013	-0.0131
				(0.018)	(0.017)
Disagreement× F2				0.0211	0.0211
				(0.011)	(0.011)
Disagreement× F3				-0.0181*	-0.0176*
				(0.009)	(0.009)
Dummy	-0.00741***	-0.00766**		-0.00830***	
	(0.002)	(0.002)		(0.002)	
Disagreement× Dummy	0.124***	0.121***		0.142***	
	(0.005)	(0.015)		(0.017)	
ChatGPT_Search			-0.000271*		-0.000292*
			(0.000)		(0.000)
Disagreement× ChatGPT_Search			0.00260*		0.00318*
			(0.001)		(0.001)
Constant	-0.00601***	-0.00621**	-0.00624**	-0.00697**	-0.00700**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Publication FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES
Clustered S.E.	Publication × Year	Publication × Year	Publication × Year	Publication × Year	Publication × Year
Observations	34,218	34,218	34,218	34,173	34,173
R-squared	0.009	0.03	0.03	0.031	0.031

Table 4. The Impact of ChatGPT's Release on the Relationship Between Oil News Sentiment and Oil Returns

Table 4 reports the regression results of daily WTI oil returns on monthly ChatGPT visit counts (in millions) and their interaction with sentiment in oil shock news, as measured using the Loughran-McDonald dictionary. Columns (1)–(3) report simple regressions of the first principal components on monthly ChatGPT visits. Column (5) compares the effect of ChatGPT visit counts with writing styles, the first three principal components extracted from the 12 LIWC variables (*F1, F2, F3*). Standard errors (S.E.) are clustered at the level of the corresponding fixed effects and are reported in parentheses. *, **, and *** indicate that the coefficients are significant at the 10%, 5%, and 1% significance levels, respectively.

VARIABLES	(1) F1	(2) F2	(3) F3	(4) Oil_Ret	(5) Oil_Ret
Oil_Ret_L1				-0.0457***	-0.0453***
				(0.0115)	(0.011)
Oil_Ret_L2				0.0935**	0.0938**
				(0.0371)	(0.037)
Oil_Ret_L3				-0.119***	-0.120***
				(0.0322)	(0.032)
Disagreement				-0.0451**	-0.0119
				(0.0187)	(0.010)
F1					-0.000451
					(0.001)
F2					0.000155
					(0.000)
F3					0.000246
					(0.000)
Disagreement× F1					-0.013
					(0.018)
Disagreement× F2					0.021
					(0.013)
Disagreement× F3					-0.0176*
					(0.009)
Chat_Monthly_Visits	0.0607**	0.0202	0.0516***	-0.0271**	-0.0275**
	(0.0274)	(0.0203)	(0.0152)	(0.00919)	(0.009)
Disagreement× Chat_Monthly_Visits				0.0983*** (0.0249)	0.120*** (0.026)
Constant	-0.242***	0.385***	-0.205***	-0.00623**	-0.00698**
	(0.0254)	(0.0188)	(0.0141)	(0.00217)	(0.002)
Publication FE	NO	NO	NO	YES	YES
Year FE	NO	NO	NO	YES	YES
Clustered S.E.	None	None	None	Publication × Year	Publication × Year
Observations	8,025	8,025	8,025	34,218	34,173
R-squared	0.001	0.000	0.001	0.030	0.031

Table 5. Alternative Specifications for the Impact of ChatGPT's Release on the Relationship Between Oil News Sentiment and Oil Returns

Table 5 reports the regression results on the impact of oil shock news sentiment, measured using the Loughran–McDonald dictionary, and its interaction with monthly ChatGPT visit counts (in millions) on daily WTI oil returns. Columns (1), (3), and (5) examine subsamples restricted to the post-2020 period. Columns (2), (4), and (6) exclude the year 2020, the year with significant oil futures price shocks due to the onset of COVID-19. Columns (3) and (4) use an alternative measure of sentiment: the daily aggregated sentiment across all news articles. Columns (5) and (6) use the daily average sentiment instead. Columns (7) and (8) use the article length-weighted average daily sentiment instead. Column (9) includes weekday fixed effects. Column (10) includes weekday fixed effects and clusters standard errors by weekday. Standard errors (S.E.) are clustered at the level of the corresponding fixed effects and are reported in parentheses. *, **, and *** indicate that the coefficients are significant at the 10%, 5%, and 1% significance levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
VARIABLES	Oil_Ret	Oil_Ret	Oil_Ret	Oil_Ret	Oil_Ret	Oil_Ret	Oil_Ret	Oil_Ret	Oil_Ret	Oil_Ret
		Excluding		Excluding		Excluding		Excluding		
	After 2020	2020	After 2020	2020	After 2020	2020	After 2020	2020		
Oil_Ret_L1	-0.0357**	0.00804***	-0.105***	0.00838***	-0.0388**	0.00705***	-0.0423**	0.00781***	-0.0451***	-0.0451
	(0.010)	(0.002)	(0.014)	(0.002)	(0.010)	(0.002)	(0.0108)	(0.00167)	(0.011)	(0.045)
Oil_Ret_L2	0.154***	-0.0306*	0.111***	-0.0297*	0.151***	-0.0311**	0.150***	-0.0314**	0.0958**	0.0958
	(0.014)	(0.013)	(0.018)	(0.014)	(0.014)	(0.013)	(0.0143)	(0.0122)	(0.039)	(0.055)
Oil_Ret_L3	-0.178***	0.0338	-0.222***	0.0342	-0.183***	0.0334	-0.187***	0.0325	-0.120***	-0.120*
	(0.011)	(0.020)	(0.011)	(0.020)	(0.011)	(0.020)	(0.0115)	(0.0193)	(0.031)	(0.056)
Disagreement	-0.122***	-0.0353*							-0.0086	-0.0086
	(0.023)	(0.016)							(0.009)	(0.020)
F1									-0.000403	-0.000403
									(0.001)	(0.001)
F2									0.000218	0.000218
									(0.000)	(0.000)
F3									0.000255	0.000255
									(0.000)	(0.000)
Disagreement×F1									-0.0124	-0.0124
									(0.018)	(0.020)
Disagreement×F2									0.0186	0.0186
									(0.011)	(0.018)
Disagreement×F3									-0.0177*	-0.0177
									(0.009)	(0.010)
Chat_Monthly_Visits	-0.00313***	-0.000567	-0.0216***	-0.00237**	-0.0177***	-0.00483***	-0.0366***	-0.00467**	-0.0262**	-0.0262*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.00493)	(0.00173)	(0.010)	(0.010)
Disagreement×Chat_Monthly_Visits	0.187**	0.105***							0.129***	0.129**
	(0.058)	(0.023)							(0.030)	(0.041)
Agg_Disagreement			-0.0252***	-0.00176						
			(0.002)	(0.002)						
Agg_Disagreement×Chat_Monthly_Visits			0.0355***	0.0196***						
			(0.003)	(0.003)						
Mean_Disagreement					-0.570***	-0.118**				
					(0.076)	(0.045)				
Mean_Disagreement×Chat_Monthly_Visits					0.747***	0.286***				
					(0.078)	(0.050)				
Weighted_Disagreement							-1.433***	-0.171*		
							(0.211)	(0.0808)		
Weighted_Disagreement×Chat_Monthly_Visits							1.610***	0.300**		
							(0.198)	(0.0932)		
Constant	0.00344***	-0.000114	0.0208***	-0.000205	0.0155***	0.00189	0.0345***	0.00270	-0.0182**	-0.0182
	(0.001)	(0.000)	(0.001)	(0.001)	(0.002)	(0.001)	(0.00502)	(0.00173)	(0.006)	(0.009)
Publication FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Weekday FE	NO	NO	NO	NO	NO	NO	NO	NO	YES	YES
Clustered S.E.	Publication × Year	Publication × Year	Publication × Year	Publication × Year	Publication × Year	Publication × Year	Publication × Year	Publication × Year	Publication × Year	Publication × Year × Weekday
Observations	9,139	30,498	9,139	30,498	9,139	30,498	9,139	30,498	34,173	34,173
R-squared	0.048	0.004	0.116	0.004	0.052	0.005	0.061	0.006	0.040	0.040

Table 6 Differential Effects of Oil News Sentiment on Oil Returns in Standard versus Adaptive News as ChatGPT Usage Changes

This table reports regression results comparing how oil news sentiment affects daily WTI oil returns across *standard* versus *adaptive* news articles as ChatGPT usage changes. The dependent variable is *Oil_Ret*. *Disagreement* measures news sentiment using the Loughran-McDonald dictionary. *Chat_Monthly_Visits* denotes monthly ChatGPT visits. *Standard* is an indicator equal to one for articles in the recurring "Market Talk" format, which follows a more fixed, standardized template. By contrast, adaptive news refers to other articles that do not follow this format. in column (2) using the subsample after December 13, 2016 since first "Market Talk" news. The regressions include up to three lags of oil returns, publication fixed effects, and year fixed effects. Standard errors are two-way clustered by publication and year, and are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

VARIABLES	(1) Oil_Ret Full Sample	(2) Oil_Ret After 2016
Oil_Ret_L1	-0.0457***	-0.0472**
	(0.0113)	(0.0137)
Oil_Ret_L2	0.0936**	0.144***
	(0.0372)	(0.0251)
Oil_Ret_L3	-0.119***	-0.176***
	(0.0311)	(0.0177)
Standard	-0.00437*	-0.00563
	(0.00224)	(0.00342)
Disagreement	-0.0458**	-0.0740
	(0.0193)	(0.0415)
Standard × Disagreement	0.164*	0.178
	(0.0776)	(0.106)
Chat_Monthly_Visits	-0.000283*	-0.000320*
	(0.000131)	(0.000150)
Standard × Chat_Monthly_Visits	0.000693***	0.000703***
	(5.73e-05)	(8.68e-05)
Disagreement × Chat_Monthly_Visits	0.00299**	0.00381*
	(0.00115)	(0.00192)
Standard × Disagreement × Chat_Monthly_Visits	-0.0262***	-0.0265***
	(0.00291)	(0.00383)
Constant	0.00055	0.000524
	(0.001)	(0.000)
Publication FE	YES	YES
Year FE	YES	YES
Clustered S.E.	Publication × Year	Publication × Year
Observations	34,218	13,226
R-squared	0.030	0.048

Table 7. Heterogeneity by Article Length, Readability, and Uncertainty

This table reports heterogeneity tests on whether article length, readability difficulty, and uncertainty affect the moderating role of ChatGPT usage in the relation between oil news sentiment and daily WTI oil returns. The dependent variable is *Oil_Ret*. *Disagreement* is the Loughran-McDonald dictionary-based sentiment measure. *Chat_Monthly_Visits* denotes monthly ChatGPT visits. In Column (1), *Long_Article* is an indicator equal to one if the article's word count is above the median word count in the pre-ChatGPT period. In Column (2), *Hard_Read* is an indicator equal to one if the article's readability difficulty is above the pre-ChatGPT median, where readability difficulty is constructed as the sum of the standardized values of words per sentence (*wps*) and the share of big words (*bigwords*), using pre-ChatGPT means and standard deviations. In Column (3), *High_Uncertainty* is an indicator equal to one if article-level uncertainty, measured as uncertainty-word counts divided by total word count, is above the pre-ChatGPT median. The pre-ChatGPT period is defined as dates before November 30, 2022. All regressions include lagged oil returns and year fixed effects. Standard errors are clustered by year and reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

VARIABLES	(1) Oil_Ret Article Length: Characteristic_Dummy = Long_Article	(2) Oil_Ret Difficulty to Read: Characteristic_Dummy = Hard_Read	(3) Oil_Ret Uncertainty: Characteristic_Dummy = High_Uncertainty
Oil_Ret_L1	-0.0481**	-0.0501**	-0.0498**
	(0.0160)	(0.0168)	(0.0174)
Oil_Ret_L2	0.140***	0.140***	0.142***
	(0.0178)	(0.0185)	(0.0197)
Oil_Ret_L3	-0.182***	-0.182***	-0.181***
	(0.0220)	(0.0206)	(0.0201)
Disagreement	-0.0546**	-0.194***	-0.294
	(0.0133)	(0.0331)	(0.215)
Chat_Monthly_Visits	-0.000249	0.00208	-0.00150
	(0.00164)	(0.00230)	(0.00236)
Disagreement×Chat_Monthly_Visits	0.0699***	0.0996**	0.178
	(0.0112)	(0.0353)	(0.109)
Characteristic_Dummy	0.0117	0.00884	-0.00407
	(0.0114)	(0.00954)	(0.00317)
Characteristic_Dummy × Disagreement	-0.713	0.150**	0.287
	(0.558)	(0.0528)	(0.238)
Characteristic_Dummy × Chat_Monthly_Visits	-0.00423	-0.00516	0.00130
	(0.00498)	(0.00438)	(0.00162)
Characteristic_Dummy × Disagreement × Chat_Monthly_Visits	0.253	-0.0209	-0.137
	(0.253)	(0.0333)	(0.117)
Constant	0.000749	-0.00324	0.00302
	(0.001)	(0.000)	(0.001)
Year FE	YES	YES	YES
Clustered S.E.	Year	Year	Year
Observations	8,037	8,037	8,037
R-squared	0.051	0.054	0.050

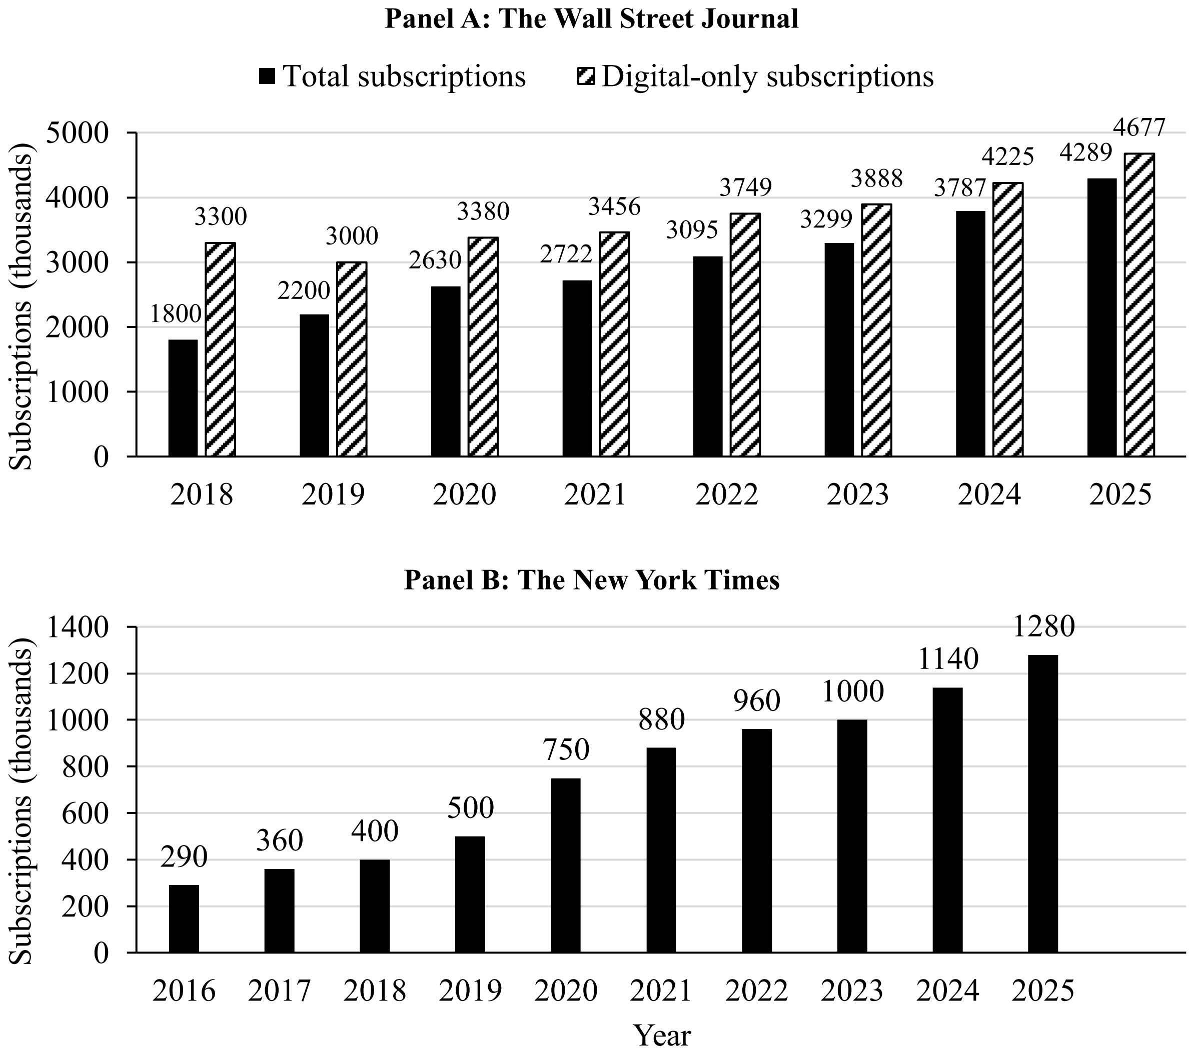

Figure A. Paid Subscriptions of The Wall Street Journal and The New York Times

Figure A plots annual paid subscription counts for The Wall Street Journal and The New York Times. The New York Times series is collected from annual reports, while the Wall Street Journal series is collected from News Corp annual and quarterly reports. Panel A reports The Wall Street Journal's total subscriptions and digital-only subscriptions. Panel B reports total paid subscriptions for The New York Times. Subscription counts are reported in thousands. [8]

[8] Some early observations are approximate because the underlying reports describe the values as "about" or "more than" specific thresholds, and total subscription values are unavailable before 2017 for Wall Street Journal and before 2015 for The New York Times.

Table A. A Sample of "Oil Shock" News

Table A provides an example of a news article from our Dow Jones Institutional Newswire collection, showing coverage of "oil shock" events. This sample was sourced from ProQuest for the period 2010–2020. The figure includes a screenshot capturing the article's title, publication info, a partial excerpt of its content, and associated metadata such as mentioned companies.

European Midday Briefing: Strong Earnings Boost European Stocks

Publication info: Dow Jones Institutional News ; New York [New York]08 May 2019.

Full text: European Markets:
Strong first quarter earnings reports helped markets in Europe shake off the pull of Tuesday's sharp losses.
The Stoxx 600 dipped 0.2%. The U.K.'s FTSE 100 was a rare index in the red on Wednesday, down 0.2%.
In Germany, the DAX increased to a gain of 0.6% following Tuesday's 1.6% loss.
Italy's FTSE MIB was flat and France's CAC 40 was climbing 0.3% but not quite reversing Tuesday's fall.
In European indexes earnings appear to be dominating, rather than news about trade or oil.
Iran announced that it would stop implementing some aspects of its 2015 nuclear deal, citing frustrations with signatory countries' failure to help them sell oil as well as the U.S.'s withdrawal from the agreement in 2018. But the oil price didn't react dramatically, nor did oil producers' stocks.
U.S. negotiators had been touting the progress of trade talks with China for weeks, when U.S. President Donald Trump's announcement on Saturday of fresh punitive measures sharply reset investors' expectations.
Confirmation from U.S. officials that it wasn't mere bluster firmly sank in Tuesday, when a trickle of losses turned into a selloff in U.S. markets. The S&P 500 fell 1.7%, the largest loss since March 22.
Jasper Lawler, head of research at London Capital Group, said: "We are starting to see the market price in the U.S.-China trade dispute as both an immediate and long-term risk factor, rather than an issue that was on the brink of being resolved. With this in mind, the S&P looks mispriced close to all-time highs."
In Germany, industrial production figures for March rose unexpectedly, coming in at 0.5% month over month versus -0.5% expected. Chinese data was less sunny, with exports for April sinking 2.7% year over year versus

May 08, 2019 06:19 ET (10:19 GMT)

Subject: Corporate profits; Insurance industry; Dow Jones averages; EU membership; Brand loyalty; Industrial production; Interest rates; Central banks; Stock exchanges; Spinoffs

Location: Italy Beijing China Denmark United Kingdom--UK Iran France Europe United States--US Germany China

People: Trump, Donald J Powell, Jerome

Company / organization: Name: Walmart Inc; NAICS: 452112, 452910, 454111; Name: Wall Street Journal; NAICS: 511110; Name: Dow Jones & Co Inc; NAICS: 511110, 511120, 511130, 519130, 523210; Name: Congress; NAICS: 921120; Name: London Capital Group Ltd; NAICS: 523110; Name: European Central Bank; NAICS: 521110

Title: European Midday Briefing: Strong Earnings Boost European Stocks

Publication title: Dow Jones Institutional News; New York

Publication year: 2019

Publication date: May 8, 2019

Publisher: Dow Jones & Company Inc

Place of publication: New York

Country of publication: United States, New York

Publication subject: Business And Economics

Source type: Wire Feeds

Language of publication: English

Document type: News

ProQuest document ID: 2221655354

Table B. Principal Component Analysis of LIWC-22 Writing Style Variables

This table reports the coefficients of the LIWC-22 variables on the first three principal components derived from the principal component analysis of the writing style variables.

Variable	Comp1	Comp2	Comp3	Unexplained
wc	-0.056	0.164	-0.031	0.953
analytic	0.011	-0.569	0.077	0.305
clout	0.069	0.114	0.529	0.372
authentic	0.025	0.046	-0.604	0.403
tone	-0.431	0.445	-0.260	0.341
wps	-0.149	-0.071	0.307	0.769
bigwords	-0.212	-0.069	0.314	0.681
dic	0.477	-0.005	-0.083	0.229
linguistic	0.493	0.053	-0.076	0.110
function	0.473	0.073	0.019	0.134
pronoun	0.205	0.426	0.154	0.096
ppron	0.057	0.488	0.229	0.145

Table C. Web-Browsing Plugin and Expanded Dow Jones Institutional News Data

This table extends our main specification in Table 4 using an expanded Dow Jones Institutional News sample. First, we focus on the period before March 23, 2023, when internet browsing first became available in ChatGPT. Second, we expand the Dow Jones Institutional News sample through the end of 2024 and estimate the regressions using only Dow Jones Institutional News articles. The table reports regression results for daily WTI oil returns on monthly ChatGPT visit counts (in millions) and their interaction with sentiment in oil-shock news, measured using the Loughran-McDonald dictionary. Standard errors are clustered at the level of the corresponding fixed effects and are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)
VARIABLES	Oil_Ret Before Web Browsing	Oil_Ret DJI Expanded Sample
Oil_Ret_L1	-0.0557***	-0.0550***
	(0.0163)	(0.0164)
Oil_Ret_L2	0.0950*	0.0941*
	(0.0492)	(0.0492)
Oil_Ret_L3	-0.119*	-0.118*
	(0.0610)	(0.0610)
Disagreement	-0.00750	-0.0103
	(0.0164)	(0.0157)
F1	-0.000565	-0.000524
	(0.000637)	(0.000590)
F2	-9.15e-06	-5.42e-05
	(0.000328)	(0.000304)
F3	0.00166	0.00164
	(0.00100)	(0.000955)
Disagreement× F1	0.0209	0.0183
	(0.0162)	(0.0151)
Disagreement× F2	-0.00307	-0.00193
	(0.0148)	(0.0137)
Disagreement× F3	0.0255*	0.0204*
	(0.0127)	(0.0116)
Chat_Monthly_Visits	-0.000644***	-7.91e-05
	(0.000126)	(0.000118)
Disagreement× Chat_Monthly_Visits	0.00443**	0.00379***
	(0.00174)	(0.00116)
Constant	-0.0158***	-0.0158***
	(0.000537)	(0.000526)
Year FE	YES	YES
Clustered S.E.	Year	Year
Observations	24,813	26,343
R-squared	0.034	0.034