Distortions in Financial Narratives: A ChatGPT Approach

Khaled Obaid and Kuntara Pukthuanthong

April 7, 2024

Abstract

We examine how financial news is distorted as it spreads across different news outlets, akin to the "telephone game". Using a sample of exclusive articles from the *Wall Street Journal*, we use ChatGPT-4 to quantify the distortion introduced in the information environment as competing news outlets retell these exclusive stories. We find strong evidence that retelling articles tend to be more opinionated and negative and less factual and appealing compared to the original story. Factors that influence the distortion include whether the stories are retold by news outlets that are less specialized, the time-lapse between the original story and its retelling, and the presence of competing narratives from other news outlets. Our findings indicate that media distortion affects asset prices, influences trading behaviors, and leads to disagreement among analysts.

JEL classification: G4, D8, M31

Keywords: Financial news, Telephone game, Distortion, Investor attention, ChatGPT, large language models.

Khaled Obaid is an Assistant Professor at the College of Business and Economics, California State University-East Bay, Hayward, CA 94542; email khaled.obaid@csueastbay.edu. Kuntara Pukthuanthong is the Robert J. Trulaske Jr. Professor of Finance at the Robert J. Trulaske Sr. College of Business, University of Missouri, Columbia, MO 65211; email pukthuanthongk@missouri.edu.

1. Introduction

In an era dominated by the rapid dissemination of information across many media platforms, financial news could be distorted as they permeate the media ecosystem. In this study, we utilize ChatGPT-4 to evaluate how financial content might be distorted as it spreads across news outlets and the financial implications of this distortion. We leverage a sample of exclusive news articles from the *Wall Street Journal* (WSJ) and track how these stories are retold in other news outlets. Oftentimes, investors learn about financial market developments from second- or third-hand sources. For instance, the WSJ reported on June 11, 2021, that Mudrick Capital Management LP experienced a 10% loss in its flagship fund due to an unexpected surge in AMC Entertainment Holdings' stock price.[1] This same story was retold by the *New York Post* (NYP) on June 14, 2021.[2] Some investors will learn about this story from the NYP instead of the WSJ.

These retelling articles could potentially distort the original story. News outlets have different biases and motives that can lead to subtle changes in tone in the retelling compared to the original account (e.g., Groseclose and Milyo, 2005; Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006). For example, it is possible that news outlets retelling the exclusive WSJ will add additional negative tone to the story in order to capture the attention of readers. In the example articles above, the percentage of negative words in the original WSJ article and the retelling according to the Loughran and McDonald (2011) lexicon is 2.80% and 5.03%, respectively. Another WSJ exclusive story reported that PepsiCo Inc. is implementing layoffs at the headquarters of its North American snacks and beverages divisions on December 5, 2022.[3] The *Washington Post* (WP) retold this same story on December 6, 2022.[4] Similarly, the original story has nearly half the proportion of negative words compared to the retelling (1.20% and 2.48%, respectively).

[1]Source: AMC Bet by Hedge Fund Unravels Thanks to Meme-Stock Traders
[2]Source: Mudrick Capital takes massive hit on AMC shares as short bet backfires
[3]Source: PepsiCo to Lay Off Hundreds of Workers in Headquarters Roles
[4]Source: PepsiCo layoffs could indicate broader economic woes

In light of these patterns, we draw parallels to the "telephone game" where the message erodes with each successive retelling. More specifically, we examine the following research questions: (1) How are these original news stories altered when retold (e.g., language style and level of detail)? (2) What news outlet and story characteristics influence this distortion? (3) Can advanced AI models help us quantify how news is distorted as it is retold across news outlets? and (4) Do retelling articles of news and distortion have implications on asset prices, trading behavior and the information environment (analyst forecasts dispersion)?

We focus our analysis on the full text of exclusive WSJ articles as our starting point. These stories are labelled "exclusive" by the WSJ, indicating that the WSJ has information, an interview, or a story that no other media outlet has at the time of publication or broadcast.[5] This sample serves as our starting point. Next, we create a sample of retelling articles, which is composed of articles from other major news outlets that retell the same story first released by the WSJ. The source of retelling articles comes from major US news outlets in Factiva, and we require articles to directly mention WSJ as the source of the story. We utilize Term Frequency-Inverse Document Frequency (TF-IDF) and cosine similarity to find potential matches and confirm the matches manually.[6]

Between 2013 and 2022, we identify 1,351 exclusive WSJ articles that we are able to match to at least one article from one of the major US news outlets in Factiva. On average, we find that each article is retold 1.40×. Our sample covers retelling articles from 18 news organizations including the *New York Times*, *Washington Post*, *New York Post*, *Investor's Business Daily*, *USA Today*, among others.[7]

After creating our sample of exclusive WSJ articles and respective retelling articles,

[5]Relying on a general news sample and using the publication date to determine the story's source is problematic for our study. While the publication date can indicate which news outlet was the quickest to release the story, it may not necessarily be the original source that other news outlets depend on.

[6]Term Frequency-Inverse Document Frequency is a statistic that gauges a word's significance within a document in a corpus, adjusting for its commonness across documents. Cosine similarity measures the similarity between two documents, regardless of size, by calculating the cosine of the angle between their vector representations in multidimensional space. In our study, we apply both methods to identify potential matches.

[7]Please see Appendix A for the complete list.

we utilize ChatGPT-4 to examine how the retelling differs from the original story in five dimensions of language style: Fact, Opinion, Negativity, Positivity, and Appeal. ChatGPT's advantage lies in its ability to understand and interpret complex text in a way that mimics human analysis, offering scalable, efficient, and nuanced insights. This makes it particularly suitable for analyzing the distortion of financial narratives as they spread across news outlets. Unlike traditional textual analysis or machine learning approaches like Word2Vec, ChatGPT can grasp the subtleties of language context and sentiment effectively (e.g., Lopez-Lira and Tang, 2023; Hansen and Kazinnik, 2023).

Our first empirical test is to examine how retelling articles differ from the original story in terms of language style. We document that retelling articles are deemed by ChatGPT-4 to be more opinionated and negative, and less factual, positive and appealing relative to the original story. We refer to the shift in writing as negative personalization. Moreover, we help validate our ChatGPT-4 variables by showing that automated text variables that capture tone (Loughran and McDonald, 2011), readability, and presence of numbers in the text relate to the ChatGPT-4 variables in the predicted direction. We control for these automated text variables in all our tests and show that the ChatGPT-4 language variables contain information beyond what these automated text variables contain.

Next, we explore the drivers behind the negative personalization in retelling articles. Specifically, we consider (1) the frequency news outlets retell stories, (2) the time-lapse between the retelling and the original story, and (3) the competition for retelling the story. We show that retelling articles that are retold later and retold by less specialized news outlets (i.e., news outlets that frequently retell stories) tend to be more opinionated and less factual. Moreover, we show that retelling articles that are retold by more specialized news outlets, published later, and retold in a more competitive environment tend to be more pessimistic. Stories that are retold by less specialized news outlets and in less competitive environments are deemed to be less appealing. Overall, we document clear patterns of negative personalization (i.e., the text becomes more opinionated, more negative, less factual, and less appealing) as news spreads and explore some of the factors that influence the distortion.

Next, we examine whether distortion from retelling articles impacts asset prices. Past studies document how news coverage catches investors attention (e.g., Engelberg and Parsons, 2011; Solomon, Soltes, and Sosyura, 2014). and lead to temporary mispricing driven by biases such as overreaction and stale information (e.g., Dougal, Engelberg, Garcia, and Parsons, 2012; Engelberg, Sasseville, and Williams, 2012; Gurun and Butler, 2012). We find that retelling articles are associated with increasing contemporaneous abnormal returns and decreasing abnormal returns in the following trading month. The predictive effect of retelling articles is economically significant; a retelling article is linked to a 1.07% reduction in the stock return of the featured firm in the following month, which is substantial compared to the mean next-month abnormal return of -0.20%.

Our evidence is consistent with prior work documenting that investors react to stale information, leading to temporary mispricing (e.g., Huberman and Regev, 2001; Tetlock, 2011). However, harnessing the power of ChatGPT-4, we shed light on the drivers behind the reaction to stale news (i.e., stale news is not retelling stories in a neutral way). More specifically, retelling articles with more negative personalization, retold by less specialized news outlets, and retold more quickly are driving the positive relation with contemporaneous abnormal stock returns. However, the reversal is mostly driven by retelling articles with more negative personalization. We control for various article and stock characteristics, including sentiment of the text (e.g., Tetlock, 2007; Garcia, 2013). We argue that distortion in retelling articles attract more attention from readers and increase disagreement among investors, temporarily pushing up asset prices in presence of short-sale constraints.

We shift our focus to how retelling articles affects investor's trading. Huberman and Regev (2001) document a sharp increase in trading volume to no-new-news. We show that retelling articles are associated with increase in abnormal turnover on the publication date and the following trading month. The positive relation between retelling and abnormal trading volume is driven by more negative personalization and retelling articles by less specialized news outlets. Temporal gap and competition among news outlets are not driving the relation between retelling articles and abnormal turnover. Additionally, using retail trading data and 13F institutional holdings data, we show that retail traders are driven by retelling articles while institutional traders are not.

In addition, we have exclusive WSJ articles with a macro focus instead of company or sector specific articles. We narrow down the list of exclusive WSJ articles to those that do not contain any mentions of publically traded companies and average the article characteristics across these articles on a daily basis. We show a distinct return prediction pattern when using market-returns from the stock-level prediction we document earlier. We show that negative personalization positively predicts next trading month abnormal returns, but is not related to contemporaneous abnormal returns. Moreover, competition among media outlets to retell stories positively relates to contemporaneous and next trading month abnormal returns.

An important test in our paper is looking at how retelling articles relate to future analyst forecast dispersion. We find a strong positive relation between retelling articles and analyst forecast dispersion in the following trading quarter. This evidence suggests that distortion from retelling articless impacts the information environment of the firm and leads sophisticated market participants (i.e., financial analysts) to disagree. However, if retelling articles are simply a way to increase investor attention (e.g., Engelberg and Parsons, 2011; Hillert, Jacobs, and Müller, 2014) or sentiment (e.g., Tetlock, 2007; Hribar and McInnis, 2012; Garcia, 2013; Jiang, Lee, Martin, and Zhou, 2019), one would not expect to see an increase in analyst forecast disagreement.

This paper makes two important contributions to the literature. First, we carefully examine how financial news stories can be distorted as they spread in the market by different media outlets. Moreover, we examine how the distortion documented above impacts asset prices, trading behaviors, and the information environment. Prior studies argue that investors are acting on stale information due to behavioral biases like overconfidence (Tetlock, 2011). However, we show that retelling articles are not simply restating facts from the original source. There are clear patterns to how this information is sensationalized before it gets to the investor. This sensationalization has important financial market implications. In addition, we document some of the mechanisms behind the distortion, including the specialization of the news outlet, the temporal gap between the original story and the retelling article, and competition to retell the story among news outlets.

Second, we overcome a limitation in the literature of capturing distortion in news by utilizing advanced AI technology, ChatGPT-4. Related papers often use methods that do not take into account context to measure specific aspects of distortion. For example, prior work used the Loughran and McDonald (2011) dictionary to measure media slant based on tone (Gurun and Butler, 2012), the Jaccard (1901) index to capture text similarity (Tetlock, 2011), and counting the presence of political phrases to capture political slant (Gentzkow and Shapiro, 2010). In contrast, we use ChatGPT-4 to quantify how retelling articles differ from the original story in multiple dimensions. Unlike other machine learning approach, ChatGPT-4 has a deeper understanding of language context, enabling it to grasp the nuances of sentences and paragraphs. This makes ChatGPT-4 suitable for tasks that require an understanding of whole texts, such as summarizing news articles or answering questions based on them, which are the main tasks of this study. More recently, researcher apply ChatGPT for parsing news headlines (Lopez-Lira and Tang, 2023), Federal Reserve announcements (Hansen and Kazinnik, 2023), and forecasting innovation (Yang, 2023). Our findings demonstrate ChatGPT-4 can be a useful tool to help us understand how news is distorted as it spreads in financial markets.

Broadly, our paper is related to research focusing on news and financial markets (e.g., Tetlock, 2007; Fang and Peress, 2009; Engelberg and Parsons, 2011; Obaid and Pukthuanthong, 2022). More specifically, our paper is related to finance and economics studies examining bias or slant in news media (e.g., Gentzkow and Shapiro, 2006, 2008, 2010; Gurun and Butler, 2012). Gentzkow and Shapiro (2010) create an index for media political slant at the news outlet level by comparing the language used in news outlets and *Congressional Record*. They show that media outlets respond strongly to the political preferences of the readers. Gurun and Butler (2012) combine news stories for a given firm from a certain news outlet monthly and compute the proportion of negative words used in the text using a dictionary approach. They find that local media report news more positively about local companies. Our study provides an important contribution to this area in the following ways. First, we ensure that we examine distortion among articles strictly about the same story. This feature makes our setting resemble the *telephone game* where a message changes as it is passed from person to person. On the other hand, Gurun and Butler (2012) focus is more general on the coverage of firms in various news media. Second, we go beyond measuring differences in tone or political phrases to capture distortion. Our measure of distortion utilizes ChatGPT-4 takes into account differences in tone, factual details, opinions, and appeal of the writing style.

Our paper is related to Tetlock (2011). The author calculates the text similarity using the Jaccard (1901) index between news text and the previous ten news stories related to the same firm and argue that this measures stale news and show that individual investors overreact to stale news. Even though the approach is straightforward and does not require training data, one major limitation of this approach is that the Jaccard (1901) index does not take context and semantic meaning into consideration when computing similarity. You can have two articles with high Jaccard scores that contain new information. Compared to Tetlock (2011), we carefully verify retelling articles to make sure that they are retelling the same story as the original article. Our starting point is identifying articles that retell existing stories; however, unlike Tetlock (2011), our focus is examining how the retelling stories distort the message in the original story.

Our study is related to studies focusing on the dissemination of information. For example, Cookson, Engelberg, and Mullins (2023) document patterns consistent with echo chambers among users of a popular finance social media platform. The concept of an echo chamber refers to a situation where individuals are exposed only to information that reinforce their existing beliefs, leading to a narrow perspective. While echo chambers and telephone game both involve the dissemination of information, our study examines the transformation of content as it passes through various channels, whereas the echo chamber concept centers on the selective exposure and reinforcement of ideas within closed systems.

2. Data

In this section, we discuss the various data we use in this study. First, we discuss the sample of exclusive WSJ articles. Next, we explain the sample of retelling articles from Factiva and the method of identifying which retelling articles are related to the exclusive WSJ articles. We also discuss key variables, including those created based on ChatGPT-4.

2.1. Exclusive WSJ sample

Our study utilizes the full text of all the exclusive WSJ articles. These stories have been tagged “exclusive” by the journal to indicate that the WSJ has information, an interview, or a story that no other media outlet has at the time of publication or broadcast. This sample ensures that, up to the point of publication, the public was not aware of the story in the exclusive WSJ article. We use the “PMDM” tag to isolate the WSJ articles that are labelled exclusive by Dow Jones. We collect all 28,875 exclusive WSJ articles that can potentially be retold by other news outlets.

2.2. Retelling articles sample

We compile a sample of articles that can be considered candidates for retelling exclusive WSJ articles from Factiva. We narrow our corpus to include only articles published by news outlets listed under “Newspapers: Top US newspapers” to avoid articles published by minor blogs or lesser known news outlets. Moreover, we make sure to exclude any article published in "The Wall Street Journal" itself, or any of the publications by Dow Jones like "Barron's" and "MarketWatch". To ensure consistency and relevance, we stipulate that all articles must be in English and contain explicit references to the WSJ within the first paragraph of the full text.[8] After applying the filters above, we collect a sample of 3,113 articles between 2013 and 2022 across all the major newspapers in Factiva that can possibly match with exclusive WSJ articles from our sample.

2.3. *Matching method*

We perform document similarity matching between the exclusive WSJ articles and the Factiva sample of retelling articles using TF-IDF and cosine similarity. TF-IDF allows for the weighting of terms based on their importance, reducing the impact of commonly used words that might not be meaningful for similarity analysis. We combine the text of all articles from both groups for TF-IDF vectorization. This process transforms the articles into vectors, allowing the calculation of cosine similarities between each article in the exclusive WSJ articles group and the Factiva retellings group. Term Frequency (TF) measures how frequently a term appears in a document. It is calculated as:

$$TF(t,d) = \frac{\text{“Number of times term } t \text{ appears in document } d\text{”}}{\text{“Total number of terms in document } d\text{”}} \tag{1}$$

Inverse Document Frequency (IDF) assesses the importance of a term t across all documents, formulated as:

$$IDF(t,D) = \log\left(\frac{N}{|\{d \in D : t \in d\}|}\right) \tag{2}$$

where N is the total number of documents in the corpus D, and $|\{d \in D : t \in d\}|$ is the

[8]For a comprehensive enumeration of the search phrases considered, refer to Appendix B

count of documents containing term t. The TF-IDF score combines these two measurements:

$$TFIDF(t, d, D) = TF(t, d) \times IDF(t, D) \tag{3}$$

After transforming text into TF-IDF vectors, then we calculate the cosine similarity between vectors from the exclusive WSJ articles group and all vectors from Factiva retellings group. Cosine similarity is defined as:

$$\text{cosine\_similarity}(\mathbf{E}, \mathbf{R}) = \frac{\mathbf{E} \cdot \mathbf{R}}{\|\mathbf{E}\| \cdot \|\mathbf{R}\|} \tag{4}$$

where E and R are the TF-IDF vectors of the two documents, $(E \cdot R)$ is the dot product of the vectors, and $||E||$ and $||R||$ are the norms of the vectors E and R, respectively.

For each article in the exclusive WSJ group, we identify the top five most similar articles from the Factiva retellings group (with replacement) within 14 days from the publication date of the exclusive WSJ article. We sort these articles by cosine similarity score. The potential matches are verified manually. We review the content of each WSJ article to grasp its primary focus, and then evaluate the relatedness of each Factiva retelling article sequentially. A binary system is employed for assessment; 1 if a Factiva retellings group article directly referenced the content of the WSJ article in a meaningful way, such as mentioning a specific event or detail initially reported by the WSJ, and 0 for unrelated articles.

2.4. OpenAI ChatGPT-4 variables

ChatGPT, developed by Open AI, is a large language model that can take sophisticated tasks and provide detailed and clear answers at a level similar to human experts. The model is based on the Generative Pre-trained Transformer (GPT) series of large language models. The GPT framework utilizes transformer architectures. Transformers are adept at processing sequential data, notably text. More specifically, transformers are characterized by their ability discern intricate word relationships within sentences. Pioneering this approach, Google's BERT (Bidirectional Encoder Representations from Transformers) emerged in 2018 as a foundational transformer-based model, garnering significant recognition (Devlin, Chang, Lee, and Toutanova, 2018). Subsequently, OpenAI's release of GPT-3 in June 2020, with its unprecedented 175 billion parameters and trained on 45TB of data. In 2022, Open AI released ChatGPT-3 with its proficient generation of coherent and comprehensive responses across a multitude of knowledge areas. Continuing the evolution of transformer-based architectures, OpenAI introduced GPT-4, an even more advanced iteration of the Generative Pre-trained Transformer series. With a capacity exceeding its predecessor, GPT-4 is distinguished by its larger parameter count, enhanced training dataset, and improved fine-tuning techniques, resulting in superior understanding and generation of text (Wiggers, 2023).[9]

We analyze the language style of exclusive WSJ news and matched retelling articles using ChatGPT-4. We focus on the following five writing style dimensions: Fact, Opinion, Negativity, Positivity, and Appeal. Motivated by Melumad, Meyer, and Kim (2021), we create a prompt for ChatGPT-4 comprising 14 questions. In each of the 14 questions, we use a scale of 1 = "not at all", and 7 = "very much so". The prompt is provided in Appendix C. we follow Open AI's prompt best practices guidelines when designing the prompt.[10] We utilize OpenAI's API to request ChatGPT-4 (version GPT-4-0613) to evaluate the articles in our sample. We send the requests between December 10, 2023, to December 12, 2023. Next, we use the responses for each pair of articles to the 14 questions from ChatGPT-4 to score the text on the five dimensions.

The first index captures the degree of specific factual details in the text. The first two questions in the prompt capture the degree of specific factual details (7 = "very much so", and 1 = "not at all") and the degree of vagueness (reverse-coded, 7 = "not at all", and 1 = "very much so") in the text, respectively. The sum of the ratings on the first two questions

[9] Despite its astonishing performance, ChatGPT has some limitations when used in academic research. First, its responses can change over time as it is updated with new data and algorithms. However, this limitation can be mitigated as we are transparent about the version and date of the model we use.Second, the answer also depends on the prompt. We follow best practices when constructing the prompt and include it in Appendix C.

[10] Source: Open AI Prompt Engineering

is our *Fact* variable. We construct $Fact_A$ and $Fact_{Match}$ that correspond to the exclusive WSJ articles and the matched retelling articles (range between 2 and 14), respectively.

The second index captures the presence of opinions in the text. For this index, we sum ratings on questions 8, 11 (reverse-coded), and 14 (reverse-coded) as our *Opinion* variable. We construct $Opinion_A$ and $Opinion_{Match}$ that correspond to the exclusive WSJ articles and the matched retelling articles (range between 3 and 21), respectively.

The third and fourth indices capture negativity and positivity in the text. Negativity (positivity) is the degree of opposition (support) or disagreement (agreement) conveyed in the text. *Neg* is the sum of questions 6, 10, and 13. *Pos* is the sum of elements 5, 9, and 12. We construct Neg_A (Pos_A) and Neg_{Match} (Pos_{Match}) that correspond to the exclusive WSJ articles and the matched retelling articles (range between 3 and 21), respectively.

The final index is capturing the appeal of the text. In other words, how interesting the text is to read and the quality of the writing. We create *Appeal* as the sum of questions 4, 5, and 7. We construct $Appeal_A$ and $Appeal_{Match}$ that correspond to the exclusive WSJ articles and the matched retelling articles (range between 3 and 21), respectively.

2.5. Automated text variables

We compute various automated text variables to help us validate our responses from ChatGPT-4 and control for article characteristics in our analysis. To assess the level of complexity and specificity in the text, we compute the following variables. First, we compute Flesch Reading Ease (*readability*), which is a readability test designed to assess the clarity of English writing (Kincaid, Fishburne Jr, Rogers, and Chissom, 1975). A higher score indicates easier readability, with scores typically ranging between 0 to 100. Second, we count the numbers in the text (*numbers*) to assess the level of details and specificity in the text. Third, we calculate the percentage of words in the text that are complex (*complex*). We define complex words as those with two more syllables. To capture negativity/positivity, we use the 2018 version of the Loughran and McDonald (2011) lexicon. We count the number of positive words and negative words for each article. We then scale these numbers by the total number of words in the document to get the percentage of positive words and the percentage of negative words (positive, negative). We focus on the net of the negative and positive variables ($neg - pos_{lm}$).

2.6. CRSP variables

We collect returns and volume data from CRSP for companies mentioned by the exclusive WSJ articles and respective retelling articles. We are able to identify 1,330 unique stocks in our sample across 17,984 article-stock observations with return and price data available in CRSP. We calculate the market capitalization (mve_m) of the stock by multiplying the price of the share by the shares outstanding. We calculate the market beta ($beta$) for the stock by estimating the Fama and MacBeth (1973) regressions using weekly returns and equal weighted market returns for three years. Idiosyncratic volatility ($idiovol$) is calculated following Ali, Hwang, and Trombley (2003) by taking the standard deviation of residuals of weekly returns on weekly equal weighted market returns for three year window. We calculate the Amihud (2002) illiquidity measure (ill) by taking the average of daily absolute returns divided by dollar volume. We measure the leverage ratio (lev) following Bhandari (1988) by taking total liabilities and divide by the market capitalization. We measure the profitability of companies using the definition in Brown and Rowe (2007) by taking the annual earnings before interest and taxes net of non-operating income and scale by non-cash enterprise value ($roic$). Finally, we capture market expectations by calculating book-to-market ratio following Rosenberg, Reid, and Lanstein (1985) as the book value of equity divided by market capitalization (bm).

We calculate the compounded abnormal returns ($abret$) for the referenced stocks over some window (e.g., $t - 1$ and t) using the Fama-French Plus Momentum model (Fama and French, 1993; Carhart, 1997) estimated using a 100-day trading window with a 50-day gap (70 day minimum window). Motivated by Cookson et al. (2023), $AbLogTurnover$ is the difference between average daily log turnover between some window (e.g., $t-1$ and t) and the average daily log turnover between t–140 and t–20 (6-month period, skipping the most recent month).

2.7. *Institutional and retail trading data*

We use retail trading data based on the method in Boehmer, Jones, Zhang, and Zhang (2021) using TAQ data. The measure takes advantage of a unique aspect of U.S. financial markets regulation: only retail orders, but not institutional orders, can receive price improvement relative to the National Best Bid or Offer (NBBO). We identify orders executed through internalization in TAQ "consolidated tape" as retail buy (sell) if the price is slightly below (above) the round penny. One limitation of this approach is that we are only able to identify marketable retail orders, but not limit orders, since Reg NMS requires limit orders to be traded at round pennies (Boehmer et al., 2021). Once we calculate the shares of retail orders on a given day, we scale by the shares outstanding on a particular stock to get retail turnover.

We use 13F Holdings Data to calculate the trading by institutions in stocks in our sample. We aggregate all the shares held by institutions by firm and date across all the institutions. Next, we calculate the quarterly change in shares. Next, we scale by the shares outstanding on a particular stock to get institutional turnover.

2.8. *I/B/E/S*

We use Institutional Brokers Estimate System (I/B/E/S) database to obtain the analysts' earnings estimates. Dispersion of analyst forecasts is computed as the standard deviation of EPS forecasts divided by the absolute value of the mean EPS forecast. We use the last estimate for a firm-date observation within broker-analyst group.

3. Results

In this section, we provide our main empirical results. First, we report descriptive statistics for the key variables. Second, we report evidence of disagreeable personalization as exclusive WSJ articles are spread and retold in various news outlets. Third, We examine the mechanism behind the disagreeable personalization. Fourth, we examine how retelling articles can have asset pricing implications. Finally, we examine how retelling articles affects trading behavior and the information environment.

3.1. Descriptive statistics

Table 1 displays descriptive data on exclusive news articles from the WSJ and the subsequent retelling articles by different news organizations. The dataset contains 1,940 observations of exclusive WSJ articles matched to a retelling article.[11] There are a total of 1,351 unique exclusive WSJ articles that are retold. The dataset also contains 27,524 exclusive WSJ articles that are not retold.[12] The sample spans between 2013 and 2022, with a peak in 2014 and the fewest observations in 2017. Out of the 1,351 unique exclusive WSJ articles that are retold, 332 of those articles mention specific tickers that can be matched to CRSP and Compustat variables. In contrast, out of the 27,524 exclusive WSJ articles that are not retold, 9,131 of those articles mention tickers that match to CRSP and Compustat variables.[13]

The retelling articles are sourced from 18 different news organizations. The organizations with the most retelling articles are 'INVD (*Investor's Business Daily*)' with 589 retelling articles, followed by 'NYPO (New York Post)' with 356, 'WPCO (*Washington Post*)' with

[11] An original story could be retold multiple times.

[12] Our aim in constructing the dataset is to accurately track the retelling of exclusive WSJ articles. By focusing on retelling articles in only the major US newspapers, ensuring retelling articles explicitly reference the WSJ, and carefully confirming the matches, we ensure that our dataset is accurately capturing retelling articles of the same exclusive WSJ story. If we relax the criteria, we increase our sample of exclusive WSJ articles with a potential matched retelling article; however, the match might discuss the same topic as the exclusive WSJ article, but not exactly the same story and thus not a retelling.

[13] We use articles with no ticker in one of our empirical tests where we create a retelling variable based on macro news, or news with no ticker mentioned.

226, 'USAT (*USA Today*)' with 212, and 'NYDN (*New York Daily News*)' with 98.[14]

The number of companies featured in each exclusive WSJ article that are retold is 2.654 on average with a median of 2, and a standard deviation of 3.572. For exclusive WSJ articles that are not retold, the average number of tickers mentioned is 1.596 articles, with a median of 1 and standard deviation of 1.147. The average number of retelling articles per exclusive WSJ article is 1.404× with a standard deviation of 0.818. The average number of days between the publication of an exclusive WSJ article and a retelling article is 2.635, with a median of 1, and a standard deviation of 3.38.

Table 2, Panel A, displays descriptive statistics for the ChatGPT-4 indices of language style. Panel B reports the mean difference between the exclusive WSJ articles sample and the retelling articles sample. We note that exclusive WSJ articles are deemed to contain more factual details ($Fact$) and are less vague than the matched retelling articles. Retelling articles are judged to be more opinionated ($Opinion$) than the exclusive WSJ articles. Moreover, retelling articles tend to have more negativity (Neg) and less positivity (Pos) than the exclusive WSJ articles. Finally, the exclusive WSJ articles have stronger appeal ($Appeal$) than the retelling articles. All the differences are statistically significant at the 1% level except for Pos which is significant at the 10% level. The findings in Table 2 support the hypothesis that retelling articles go beyond relaying the key facts in the original story. Retelling articles attempt to provide guidance persuasively by becoming more opinionated and negative and less factual and appealing (as the text becomes more personalized and disagreeable).

In Table 3, we look at the descriptive statistics for the automated text variables that help capture some of the writing style of the articles. We compute the Flesch Reading Ease of the text to capture linguistic complexity (*readability*) and find that the average readability for the exclusive WSJ articles sample and retelling sample are 57.225 and 51.108, respectively. Moreover, we look at the presence of numbers (*numbers*) and complex (*complex*) words in the text. Exclusive WSJ articles use an average of 14.814 numbers in the text, whereas the

[14]Please see Appendix A for a complete list.

retelling sample uses an average of 9.880 numbers in the text. Exclusive WSJ articles contain 0.17% of complex words in the text, whereas retelling articles contain 0.162% of complex words. Text in both samples tend to have more negative than positive words ($neg - pos$), but the retelling articles sample has a slightly more negative tilt than the exclusive sample. The average *cosinesimilarity* between the articles in both samples is 0.704.

Next, we report the descriptive statistics for the characteristics of the firms mentioned in the exclusive WSJ articles in Table 4, Panel A. We have 16,863 stock-exclusive WSJ article or -retelling article observations in our sample. 6.1% of those article-stock observations are retold by other newspapers. 7.1% represent the retelling articles by newspapers other than the WSJ. We note that the average market beta (*beta*) of stocks in our sample is 1.118 and idiosyncratic volatility (*idiovol*) is 3.5%. The average market capitalization (mve_m) of companies in our sample is $139 billion. Firms in our sample have an average return on invested capital (*roic*) of 9.5% and book-to-market ratio (*bm*) of 0.481. The average leverage ratio (*lev*) is 2.807%. The average illiquidity using the Amihud (2002) measure is 0.002.

In Table 4, Panel B, we report the descriptive statistics for the stock returns and turnover of firms mentioned in our sample of exclusive WSJ articles. The returns and abnormal returns corresponding to the date (t–1, t) of the exclusive WSJ article or retelling article is 0.59 bps and 0.47 bps, respectively. In contrast, the monthly returns and abnormal returns after (t+1, $t + 21$) the exclusive WSJ article or retelling article is 1.12% and –0.23 bps, respectively. We also report the average abnormal log turnover of stocks in our sample ($AbLogTurnover$) during the time of the article and the month after.

In Table 4, Panels C and D, we compare the firm characteristics and market variables between the stocks mentioned in exclusive WSJ articles that were retold and those that were not. We note that firms in exclusive WSJ articles that are not retold tend to be riskier, less liquid, more levered, have lower profitability, have lower market expectations, and are smaller than firms that are mentioned in exclusive WSJ article that are retold. Firms mentioned in exclusive WSJ articles that are not retold have lower returns between $t - 1$ and t, but higher returns between $t+1$ and $t+21$ compared to firms mentioned in exclusive WSJ articles that are retold.

3.2. Disagreeable personalization

We examine whether information in exclusive WSJ articles is systematically altered in a strategic manner as it is retold across news outlets. On the one hand, news outlets who retell exclusive WSJ articles might simply retell what they deem is an important story that is beyond their purview to their audience. In this case, one would not expect the writing in the retelling article to be noticeably different from the original story. On the other hand, news outlets might attempt to guide and persuade readers in a certain direction. In this case, the retelling article will contain more opinions and guidance. Moreover, the news outlet might emphasize the negative aspects of the content when retelling news, driven by the news outlet's desire to stand out and persuade the audience to value their guidance. In the latter case, the final narrative received by the end audience is noticeably different from the original news, characterized by less factual details, more opinions, and a more pessimistic tone.

Several factors can influence how information passing across various news outlets is altered. First, memory can influence how stories are recounted, with research suggesting that unusual or unexpected details are more easily remembered and retold (Barrett and Nyhof, 2001). Second, bias can be generated from social and motivational factors (Hyman, 1994; Barasch and Berger, 2014; Melumad et al., 2021. More specifically related to news outlets, advertising pressure (Reuter and Zitzewitz, 2006), media ownership (Besley and Prat, 2006), and competition (Baron, 2005; Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006) can all influence bias of news outlets. Finally, the specialization of news outlets can shape their retelling of financial stories. For example, more specialized news outlets tend to provide in-depth analysis, while general news outlets simplify financial information for broader appeal.

To examine how retelling articles might change details and language style, we regress the

ChatGPT-4 indices of language style on an indicator variable for whether the story is an exclusive WSJ story or a retelling article (*retelling*). Each observation in our regression is an article (either an exclusive WSJ article or a retelling article). We create two additional variables based on the ChatGPT-4 language indices. First, $Opinion - Fact$, which is the difference between *Opinion* and *Fact* indices described in Section 2.4 and captures the level of opinions less level of factual details in the text. Second, $Neg - Pos$, which is the difference between *Neg* and *Pos* and captures the net pessimism in the text. Finally, we use the logarithm of *Appeal* to capture how appealing the text is perceived.

We also include controls for article characteristics in our regression. First, we control for the number of tickers mentioned in the article ($log(TickerCount)$). It is possible that articles mentioning many tickers are discussing an industry broadly and thus are less focused and detailed, whereas articles mentioning few tickers are focusing on a specific company and thus are more specialized and detailed. We also control for readability using the Flesch Reading Ease (*readability*), the presence of numbers in the text (*numbers*) to assess the level of details and specificity in the text, the percentage of words in the text that are complex (*complex*), and the net % of negative words minus positive words according to the 2018 version of the Loughran and McDonald (2011) lexicon ($neg - pos_{lm}$). We run the following regression:

$$LanguageIndex_{i,t} = \alpha + \beta_1 \times retelling_{i,t} + \beta_2 \times STORY_{i,t} \tag{5}$$

where *LanguageIndex* is one of the three ChatGPT-4 language style variables for article i at time t: $Opinion - Fact$, $Neg - Pos$, and $log(Appeal)$. *STORY* is a vector of story characteristics including number of tickers mentioned, readability score, % of negative net positive words in the text, count of numbers in the text, and % of complex words. We include time fixed effects and cluster standard errors by article id. We report the results in Table 5.

In the first specification, we find a strong and significant (at the 1% level) positive relation between *retelling* and $Opinion - Fact$ suggesting that the retelling article tends to, on average, contain more opinions and less factual details compared to the original story.

This relation is consistent with alterations or exaggerations being introduced as news outlets add their interpretations or perspectives to the original story. In the second specification, we include story-level controls and continue to find a positive and significant relation between *retelling* and $Opinion - Fact$. Moreover, we find a negative and significant coefficient (at the 1% level) on $numbers$ suggesting that more opinionated and less factual text contains fewer numbers. This relation is intuitive, as one would expect more opinionated writing to contain fewer numbers. Hence, the negative coefficient on $numbers$ helps validate our $Opinion - Fact$ from ChatGPT-4.

In the third specification, we find a positive and significant (at the 1% level) relation between *retelling* and $Neg{-}Pos$ indicating that retelling articles tend to be more pessimistic. In the fourth specification, we include story-level characteristics as controls and continue to find a positive and significant (at the 1% level) relation between *retelling* and $Neg{-}Pos$. The coefficient on $neg - pos_{lm}$ is positive and highly significant, suggesting that the automated pessimism measure from the 2018 version of the Loughran and McDonald (2011) lexicon and the ChatGPT-4 index are related in the predicted way. This relation helps validate our $Neg - Pos$ ChatGPT-4 variable.

In the fifth specification, we focus on the natural logarithm of *Appeal* (*log*(*Appeal*)), which captures how interesting the text is to read and the quality of the writing. We find a negative and significant (at the 1% level) coefficient on *retelling* suggesting that retelling articles are deemed to be less interesting to read with lower writing quality compared to the original exclusive WSJ articles. In the sixth specification, we add story-level characteristics. We continue to find a negative and significant coefficient on *retelling*. Moreover, we show a negative and significant (at the 5% level) relation between *readability* and *log*(*Appeal*) which suggests that text that is more challenging to read is less interesting to read with lower writing quality, on average. One would expect that the readability score (lower score suggest the text is more difficult to read) will be lower for text that is less appealing to read. Hence, the negative relation between *readability* and *log*(*Appeal*) helps validate our *Appeal*

ChatGPT-4 measure.

Overall, in Table 5, we provide evidence that retelling articles tend to be more disagreeable, opinionated, and less interesting to readers compared to the original. The results are statistically strong and hold after controlling for various article-level controls. Our evidence suggests that memory is not the only factor present in how news is retold. Our evidence suggests that news is retold systematically in a strategic way. Moreover, we provide evidence helping validate our ChatGPT-4 language variables.

3.3. Mechanisms: Specialization, memory, and motivation

In this section, our goal is to explore the underlying mechanisms behind the systematic shift in content of news as it is retold. We explore how specialization, memory, and social and motivational factors can play a role in how the content might shift as it is retold.

First, we study how news outlets that rarely versus those that frequently retell stories by the WSJ shift the content of the original story. We postulate that news outlets that rarely retell news from other organizations are likely more specialized and thus rely on their own reporting for the most part. In contrast, news outlets that frequently retell news from the WSJ are possibly not specialized and regularly depend on outside sources to cover financial topics. When news outlets perceive themselves as more knowledgeable and specialized about a specific topic, they might be more inclined to include their own opinions and interpretations.

Second, we explore how the time between the original story and the time of the retelling impacts how the story shifts compared to the original. A longer temporal gap creates more room for distortion as the story starts to fade from the memory of the public (Barrett and Nyhof, 2001). This temporal gap allows for the accumulation of errors, misinterpretations, and the insertion of bias as the story is passed from one journalist to another. Furthermore, as time progresses, the urgency to publish may lead secondary sources to rely on incomplete information or unchecked sources, compounding inaccuracies.

Third, we explore how distortion varies between exclusive WSJ articles that are retold by many news outlets versus those that are only retold by a few. Competition between news outlets can bias their reporting (Baron, 2005; Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006). When an exclusive article from WSJ is retold by several other news outlets at the same time, the potential for distortion increases due to the tendency of reporters to sensationalize the story to gain more readers in a competitive environment.

In Table 6, we regress the ChatGPT-4 language style variables on three key motivation variables. First, to differentiate between news outlets that rarely or frequently retell news by WSJ, we use $log(NewsOrgFreq)$ which is the natural logarithm of the number of retelling articles by the retelling news outlet in the past. Second, to capture how time since the original article can impact distortion, we use $log(DaysBetween)$ which is the natural logarithm of the days between the date of the original exclusive WSJ article and the retelling article. Third, to capture how the number of competing retelling articles of the same story on the same day can impact distortion, we use $log(RetellingsSameDay)$ which is the natural logarithm of the number of competing retelling articles on the same exclusive WSJ article on the same date.

In the first specification of Table 6, the dependent variable is $Opinion - Fact$. We find a positive and significant (at the 1% level) coefficient on $log(NewsOrgFreq)$ suggesting that articles retelling exclusive WSJ stories by news outlets who frequently retell WSJ stories tend to be more opinionated and less factual. This finding is consistent with less specialized news outlets (those that retell news often) injecting more opinions and removing factual details in their retelling. Next, we find a positive and significant (at the 1% level) coefficient on $log(DaysBetween)$. This finding suggests that the larger the time-lapse between the original and the retelling article, the retelling articles tend to be more opinionated and less factual. This finding is consistent with the temporal gap allowing for accumulation of bias. We do not find a significant relation between competition among retelling stories and the level of opinions and factual details in the writing of retelling articles.

In the second specification of Table 6, we focus on the level of disagreeableness in the text. We find a negative and significant (at the 10% level) on $log(NewsOrgFreq)$ which suggests that news outlets that frequently retell exclusive WSJ articles (less specialized) tend to write articles that are less disagreeable. Moreover, we document a significant (at the 1% level) and positive relation between the time-lapse between the retelling article and the original story and the level of disagreeableness in the retelling article. Next, we look at $log(RetellingsSameDay)$ and find a positive and significant (at the 1% level) relation with $Neg - Pos$. This finding is consistent with news outlets relying on negativity to capture the attention of readers in a more competitive environment.

In the third specification of Table 6, the main dependent variable is $log(Appeal)$. The coefficient on $log(NewsOrgFreq)$ is negative and significant at the 1% level. This result is consistent with news outlets that retell often writing articles that are less appealing to read. The coefficient on $log(DaysBetween)$ is not statistically significant. We find a positive and significant coefficient on $log(RetellingsSameDay)$ which is consistent with competition encouraging news outlets that want to retell the exclusive WSJ story to write a more interesting and appealing article.

In Table 6, we document that news outlets that frequently retell exclusive WSJ stories tend to write more opinionated, less disagreeable, and less appealing articles. We document that the larger the time-lapse between the retelling article and the original results in retelling articles that are more opinionated and disagreeable. The temporal gap can create more opportunities for bias and misinterpretation to enter the retelling. Finally, when there is competition to retell exclusive WSJ stories, retelling articles tend to be, on average, more disagreeable in tone and more appealing in hope to catch attention.

3.4. Asset pricing implications

The retelling of financial news, especially when news is distorted or interpreted differently by various newspapers, can impact asset returns. Fang and Peress (2009) argue that media coverage reaches a broad audience and impacts expected returns. Moreover, Tetlock (2010) show that media coverage reduces information asymmetries. In this section, we explore how retelling articles can impact asset returns. We regress abnormal returns on the date (t–1, 0) of the exclusive WSJ article or retelling article on indicators for whether the article has been retold ($retold$) or is a retelling of an exclusive WSJ article ($retelling$). Each observation in our regression is an article-firm. Moreover, we add 1) $STORY$ which is a vector of story characteristics including number of tickers mentioned, readability score, % of negative words in the text, count of numbers in the text, and % of complex words; and 2) $FIRM$ which is a vector of observable firm characteristics including industry fixed effects and past returns. We cluster standard errors by firm.

$$log(abret_{i,j,t-1,t}) = \alpha + \beta_1 \times retold_{i,j,t} + \beta_2 \times retelling_{i,t} + \beta_3 \times STORY_{i,t} + \beta_4 \times FIRM_{i,t-1} \quad (6)$$

where $abret_{i,j,t-1,t}$ is the compounded abnormal returns between t–1 and t adjusted using the Fama-French three factor model plus momentum for article i and stock j. We use 100 days estimation window 50 days prior to time t and require a minimum of 70 valid returns to estimate the expected returns used to calculate $abret_{i,j,t-1,t}$.

We report the regression results in Table 7, Panel A. In specification (4), instead of $retelling$ variable, we add more nuanced variables to capture specific characteristics of the retelling article including: 1) $NegPer$; 2) $log(NewsOrgFreq)$, 3) $log(DaysBetween)$; and 4) $log(RetellingsSameDay)$. These variables are 0 for all articles except those that are retelling articles of an exclusive WSJ article. $NegPer$ is the difference between the negative personalization scores between the original story and the retelling. The negative personalization score is calculated using the ChatGPT-4 language style indices as: $Neg - Pos + Opinion - Fact - Appeal$. A higher value of $NegPer$ indicates that a retelling article has more negative personalization compared to the original exclusive WSJ article (more opinionated, more negative tone, less factual, less positive tone, and less appealing). $log(NewsOrgFreq)$ captures the frequency of retelling articles by various news outlets. We postulate that a higher value means that the news outlet is less specialized and thus rely less on their own reporting. To differentiate between retelling articles that occur soon after the initial report and those that come later, we include $log(DaysBetween)$, which is the log of the number of days between the retelling article and the exclusive WSJ article. Finally, we capture whether there are competing retelling articles about the same exclusive WSJ article on the same day using the $log(RetellingsSameDay)$.

We find a positive and significant (at the 1% level) coefficient on $retelling$. This finding suggests that retelling articles of exclusive WSJ articles are positively associated with contemporaneous abnormal returns after controlling for past returns, story characteristics, and firm characteristics. In specification (4), we find a positive and significant (at the 1% level) relation between contemporaneous abnormal returns and negative personalization of articles ($NegPer$). In other words, $retelling$ articles with higher negative personalization are associated with higher abnormal returns. Retelling articles with higher negative personalization scores can attract greater attention from investors and the media, increasing the visibility of the stock and disagreement among investors. Disagreement among investors can push prices up especially in light of short-sale constraints (e.g., Diether, Malloy, and Scherbina, 2002; Jones and Lamont, 2002.

We find a positive and significant (at the 1% level) coefficient on $log(NewsOrgFreq)$ which suggests that when news are retold by news outlets that are less specialized (frequently retell stories instead of generating their own reporting), garner more attention driving up prices. It is possible that less specialized news outlets reach a larger crowd and increase the attention on the company.

We find a negative and significant (at the 1% level) coefficient on $log(DaysBetween)$ suggesting that the longer the time between the retelling article and the exclusive WSJ article, the lower the contemporaneous abnormal returns. When more time has passed, investors and journalists might lose memory of the events and more distortion enters the story. Moreover, the retelling of news after a significant delay could suggest that the company is struggling to move past its challenges. If the market perceives the retelling as indicative of ongoing or unresolved issues, it may lead to doubts about the company's future prospects, resulting in negative abnormal returns.

We find a positive and significant (at the 10% level) coefficient on $log(RetellingsSameDay)$. Multiple retelling articles by different outlets may serve as a form of validation or confirmation of the original news. This can reduce uncertainty or skepticism among investors about the news's validity, leading to increased confidence. As confidence grows, so might the stock's price, reflecting the perceived reduction in informational risk.

Overall, we find that retelling articles are associated with higher contemporaneous abnormal returns. More specifically, we find that the more negative personalization in the retelling article compared to the original story and the more competing retelling articles of the same story, the higher the contemporaneous abnormal returns. We argue that negative personalization and competition among news outlets increase attention on the firm and confidence among investors, leading to higher prices. In contrast, the time between the original story and the retelling could indicate that the company is struggling to move past its challenges. We show that the time between the retelling and the original story is negatively related to contemporaneous abnormal returns.

Next, we shift focus on return predictability. In Panel B of Table 7, we examine how retelling articles can impact future abnormal returns. We run the following regression:

$$log(abret_{i,j,t,t+21}) = \alpha + \beta_1 \times retold_{i,j,t} + \beta_2 \times retelling_{i,t} + \beta_3 \times STORY_{i,t} + \beta_4 \times FIRM_{i,t-1} \quad (7)$$

where $abret_{i,j,t,t+21}$ is the compounded abnormal returns between t and $t+21$ adjusted using the Fama-French three-factor model plus momentum for stock j mentioned in article i. If retelling articles do not contain new material information not mentioned previously in the original exclusive WSJ article, we expect to see return reversal.

We find that the coefficients on *retelling* are all negative and significant at the 1% level in the first three specifications. This finding suggests that the contemporaneous relation documented in Table 7, Panel A, reverses in the following month. This finding suggests a behavioral explanation to our results on abnormal return. In specification (4), the reversal is mostly driven by negative personalization ($NegPer$). In contrast, we note no reversal for $log(NewsOrgFreq)$, $log(DaysBetween)$, and $log(RetellingsSameDay)$, suggesting that some forms of distortion are lasting.

3.5. Turnover

Next, we examine how retelling articles can generate trading. As news stories are circulated and altered, they capture the attention of a broader audience and increase disagreement. Disagreement among investors can lead to an increase in trading (Karpoff, 1986). In Table 8, Panel A, we focus on the relation between retelling and contemporaneous abnormal shares turnover. To examine this relation, we run the following regression:

$$AbLogTurnover_{i,j,t-1,t} = \alpha + \beta_1 \times retold_{i,j,t} + \beta_2 \times retelling_{i,t} + \beta_3 \times STORY_{i,t} + \beta_4 \times FIRM_{i,t-1} \tag{8}$$

where $AbLogTurnover_{i,j,t-1,t}$ is the difference between average daily log turnover between $t-1$ and t and the average daily log turnover between t–140 and t–20. We also control for firm and story characteristics. If retelling articles are responsible for generating trading, we expect to find $\beta_2 > 0$. In specification (4), instead of *retelling* variable, we add: 1) $NegPer$; 2) $log(NewsOrgFreq)$, 3) $log(DaysBetween)$; and 4) $log(RetellingsSameDay)$.

In Panel A of Table 8, we find a positive and statistically significant (at the 1% level) relation between *retelling* and abnormal shares turnover around the time of the exclusive WSJ article. This finding is consistent with retelling articles generating more disagreement among traders and leading to an increase in trading. In specification (4), we examine which characteristics of the retelling article are driving the positive relation between *retelling* and abnormal trading volume. We note a positive and significant (at the 1% level) coefficient on $NegPer$ highlighting that the negative personalization is driving much of the abnormal trading volume. Moreover, retelling articles by news outlets that retell more often (less specialized) are associated with higher abnormal trading volume.

In Panel B of Table 8, we examine how retelling articles can generate future trading. We run the following regression:

$$AbLogTurnover_{i,j,t+1,t+21} = \alpha + \beta_1 \times retold_{i,j,t} + \beta_2 \times retelling_{i,t} + \beta_3 \times STORY_{i,t} + \beta_4 \times FIRM_{i,t-1} \tag{9}$$

We note a positive and significant coefficient (at the 1% level) on *retelling*. The positive coefficient on *retelling* indicates that articles retelling exclusive WSJ articles generate additional abnormal turnover in the following month. The finding suggest that the impact of retelling on abnormal trading volume persists in the following month. In specification (4), we show that the relation between *retelling* and future abnormal turnover is mostly driven by negative personalization ($NegPer$) in the retelling article and media that frequently retell exclusive WSJ stories ($log(NewsOrgFreq$).

3.6. Retail and institutional trading

In this section, we explore how retail and institutional traders behave differently in response to retelling news. Consistent with papers portraying retail traders as unsophisticated "noise" trades in the behavioral finance literature (Barber and Odean, 2008; Barber, Odean, and Zhu, 2008; Barber and Odean, 2000; Barber, Lee, Liu, and Odean, 2009), one would expect that retail traders will respond more strongly to distortion in retelling articles. Whereas institutions are more sophisticated with better tools and thus are less likely to act on distortions in retelling articles. However, if retelling articles are contributing to a more nuanced and complete information environment, one can expect that institutions will respond to retelling articles.

In Table 9, we look at the relation between *retelling* and contemporaneous and future abnormal shares turnover for retail and institutional traders. To examine this relation, we run the following regressions:

$$\begin{aligned} AbLogTurnoverRetail_{i,j,t-62,t} \text{ or } AbLogTurnoverInst_{i,j,t-62,t} &= \alpha + \beta_1 \times retold_{i,j,t} \\ &+ \beta_2 \times retelling_{i,t} + \beta_3 \times STORY_{i,t} + \beta_4 \times FIRM_{i,t-1} \end{aligned} \tag{10}$$

$$\begin{aligned} AbLogTurnoverRetail_{i,j,t+1,t+63} \text{ or } AbLogTurnoverInst_{i,j,t+1,t+63} &= \alpha + \beta_1 \times retold_{i,j,t} \\ &+ \beta_2 \times retelling_{i,t} + \beta_3 \times STORY_{i,t} + \beta_4 \times FIRM_{i,t-1} \end{aligned} \tag{11}$$

where $AbLogTurnoverRetail_{i,j,t-62,t}$ ($AbLogTurnoverInst_{i,j,t-62,t}$) is the difference between average daily log turnover for retail (institutional) traders between $t-62$ and t and the average daily log turnover for retail (institutional) traders between t–252 and t–63 (9-month period, skipping most recent quarter).[15] $AbLogTurnoverRetail_{i,j,t+1,t+63}$ ($AbLogTurnoverInst_{i,j,t+1,t+63}$) is the difference between average daily log turnover for retail (institutional) traders between $t+1$ and $t+63$ and the average daily log turnover for retail (institutional) traders between t–252 and t–63. We also control for firm and story characteristics.

In specification (1) of Table 9, we focus on contemporaneous abnormal turnover for retail investors and find a positive and significant coefficient on *retelling* (at the 5% level). In specification (2), we focus on institutional trading and find no significant relation between turnover and *retelling*. In specification (3), we shift to future abnormal turnover. Again, we find that retail traders trade more frequently in the quarter following publication of retelling articles that are retold. The relation is significant at the 5% level. In specification (4), we find no relation between institutional trading and *retelling*. Overall, we find that retail traders are influenced by retelling articles, but not institutional trading.

[15] The reason we use quarterly windows instead of monthly or daily is that we are limited by the quarterly frequency in the 13F holdings data.

3.7. Macro news and market returns

In this section, we explore how the retelling of exclusive WSJ macro articles relate to market returns. Up to this point, we only focused on retelling articles of exclusive WSJ articles that mention at least one ticker that we can map to CRSP. We narrow down the list of exclusive WSJ articles to those that do not contain any tickers or mentions of publically traded companies. We have 550 articles with no tickers mentioned that are retold. Moreover, we have 9,977 exclusive WSJ articles that do not mention any tickers that have not been retold.

In Table 10, we regress market returns on the average of story characteristics on a given day. We run the following regressions:

$$VWRETD_{t-1,t} = \alpha + \beta_1 \times retold_t + \beta_2 \times retelling_t + \beta_3 \times STORY_t \tag{12}$$

$$VWRETD_{t+1,t+21} = \alpha + \beta_1 \times retold_t + \beta_2 \times retelling_t + \beta_3 \times STORY_t \tag{13}$$

where $VWRETD$ is the CRSP value-weighted market return.

In the first three specifications, we focus on contemporaneous market returns. Unlike the firm-specific news sample, we do not find a significant relation between *retelling* and contemporaneous market return for the sample of macro news. In specification (3), we disentangle *retelling* into four specific features of the retelling articles. We note a negative and significant (at the 10% level) coefficient on $log(NewsOrgFreq)$ suggesting that on days with more retelling macro articles by less specialized news outlets, market returns tend to be lower. This is the opposite direction we observed for the firm-specific news sample. Consistent with the results using the firm-specific news sample, we find a positive and significant (at the 5% level) relation between $log(RetellingsSameDay)$ and contemporaneous market returns, suggesting that on days with more competition for retelling macro articles, the market return tends to be higher.

In specifications (4) to (6), we focus on future market returns. Again, we do not document a significant relation between market returns and *retelling*. However, we find a positive and significant (at the 5% level) coefficients on $NegPer$ and $log(RetellingsSameDay)$. In other words, on days with retelling macro articles that contain high level of negative personalization and during days with high competition between retelling articles, future market returns are higher, on average.

3.8. Analyst forecast dispersion

We explore whether distortion from retelling articles is associated with disagreement among analysts. One would expect that if retelling articles add distortion and confusion to the information environment, disagreement among investors will increase. However, if retelling articles are simply a way to increase investor attention (e.g., Engelberg and Parsons, 2011; Hillert et al., 2014) or sentiment (e.g., Tetlock, 2007; Hribar and McInnis, 2012; Garcia, 2013; Jiang et al., 2019), one would expect that analyst disagreement would not be impacted. We test this prediction in Table 11.

We regress dispersion in analyst forecasts in the quarter after the article is published $AnalystDisper_{t+1,t+63}$ on indicators for whether the article is a retelling of an exclusive WSJ article (*retelling*).

$$log(AnalystDisper_{i,j,t+1,t+63}) = \alpha + \beta_1 \times retold_{i,j,t} + \beta_2 \times retelling_{i,t} + \beta_3 \times STORY_{i,t} + \beta_4 \times FIRM_{i,t-1} \tag{14}$$

where $AnalystDisper_{i,j,t+1,t+63}$ is the standard deviation of analyst forecasts made in the quarter post article publication (between $t + 1$ and $t + 63$).

In the first specification of Table 11, we note a positive and significant (at the 1% level) coefficient on *retelling*. In other words, retelling articles lead to increase in analyst disagreement. In the second specification, we add firm-level controls and continue to observe a similar result. Finally, in specification (3), we add article characteristics and our results hold.

Our findings in Table 11 are consistent with a distortion story as opposed to an information or sentiment story.

4. Conclusion

This study advances the literature by studying how financial news is distorted as it proliferates through the market, challenging the prevailing assumption that investor behavior is predominantly influenced by acting on outdated information due to cognitive biases such as overconfidence. First, We demonstrate that the spread of news stories goes beyond mere repetition, revealing systematic patterns of sensationalization prior to reaching investors. Second, we identify several factors driving the distortion of original stories, including the news outlet's specialization, the time-lapse between the initial story and its retelling article, and competitive pressures among media outlets. Third, we examine the implications of these distortions for asset prices, trading behavior, and the overall information landscape. Fourth, we employ the advanced capabilities of ChatGPT-4 to measure the distortion in retelling articles from their sources across various dimensions.

Appendix A. News outlets

News ID	N	News Outlet
INVD	589	Investor's Business Daily
NYPO	356	New York Post
WPCO	226	Washington Post.com
USAT	212	USA Today
NYDN	98	New York Daily News
BSTN	80	The Boston Globe
WP00	66	The Washington Post
NYTF	65	NYTimes.com Feed
CHSM	62	The Christian Science Monitor
AMB0	52	American Banker
PPGZ	44	Pittsburgh Post-Gazette
SLMO	34	St. Louis Post-Dispatch
STPT	17	Tampa Bay Times
MSP0	15	STAR TRIBUNE (Mpls.-St. Paul)
ATJC	15	The Atlanta Journal
PHLI	6	The Philadelphia Inquirer
BFNW	2	Buffalo News
NYTA	1	The New York Times Abstracts

Appendix B. Search phrases

Search phrases to find retelling articles in Factiva
According to the Wall Street Journal
appeared in The Wall Street Journal
article in The Wall Street Journal
by The Wall Street Journal
information of The Wall Street Journal
quoted in the Wall Street Journal
report from The Wall Street Journal
report in The Wall Street Journal
reported by the Wall Street Journal
reports the Wall Street Journal
Sources tell The Wall Street Journal
Speaking to the Wall Street Journal
story in The Wall Street Journal
The Wall Street Journal first reported
The Wall Street Journal had reported
The Wall Street Journal initially reported
The Wall Street Journal is reporting
The Wall Street Journal looks
The Wall Street Journal previously reported
The Wall Street Journal publishes
The Wall Street Journal say
The Wall Street Journal writes
told the Wall Street Journal
Wall Street Journal claimed
Wall Street Journal claims
Wall Street Journal described
Wall Street Journal describes
Wall Street Journal found
Wall Street Journal has reported
Wall Street Journal published
Wall Street Journal report
Wall Street Journal reported
Wall Street Journal reports
Wall Street Journal revealed
Wall Street Journal said
Wall Street Journal says

Appendix C. ChatGPT-4 prompt

You will be provided two related news articles. These two articles (delimited with XML tags, the first article has the ¡firstarticle¿ tag, and the second one has the ¡secondarticle¿) come from different newspapers but are about the same story. Your task is to read and rate the two news articles. First, read the first article and rate it (on a 1: "not at all" to 7: "very much so" scale) based on how much you agree with the following statements related to the text:

1. This text is detailed;
2. This text is general / vague;
3. This text is interesting;
4. This text is well-written;
5. This text is positive about the subject matter;
6. This text is negative about the subject matter;
7. This text conveys interest in the subject matter;
8. This text is emotional about the subject matter;
9. In this text, the writer expressed support for the subject matter at hand;
10. In this text, the writer took an opposing stance towards the subject matter at hand;
11. In this text, the writer expressed no opinion about the subject matter itself;
12. In this text, the writer agreed with other people's opinions about the subject;
13. In this text, the writer disagreed with other people's opinions about the subject;
14. In this text, the writer expressed no opinion about other people's opinions.

Next, read the second article and rate it like above. However, this time when rating the second text, please think about how it compares to the first text (e.g., relative to the first article, the second article is more... or relative to the first article, in the second article the writer...:).

Your response should be in this format [firstarticle = X, X, X, X, X, X, X, X, X, X, X, X, X, X, secondarticle = X, X, X, X, X, X, X, X, X, X, X, X, X, X], where X is a number between 1 and 7 that corresponds to the rating on the 14 elements above.

Appendix D. Definitions of variables

Variable	Definition
retold	An indicator variable for whether the exclusive WSJ article has been retold by other news media.
retelling	An indicator variable for whether the article is a retelling of a exclusive WSJ article.
$NegPer$	The difference between the negative personalization scores between the original story and the retelling. The negative personalization score for each article (the exclusive WSJ article and the corresponding retelling articles) is calculated as: $Neg - Pos + Opinion - Fact - Appeal$.
$log(NewsOrgFreq)$	The natural logarithm of one plus the number of retelling articles of exclusive WSJ articles by a news media organization in the past. The non-retelling articles have 0 for this variable.
$log(DaysBetween)$	The natural logarithm of one plus the days between the exclusive WSJ article and the retelling article. The non-retelling articles have 0 for this variable.
$log(RetellingsSameDay)$	The natural logarithm of one plus the number of retelling articles of the same exclusive WSJ article on the same day. The non-retelling articles have 0 for this variable.
$Opinion - Fact$	The difference between $Opinion$ and $Fact$. We sum ratings on elements 8, 11 (reverse-coded), and 14 (reverse-coded) in the prompt as our $Opinion$ variable. The sum of the ratings on the first two elements in the prompt is our $Fact$ variable. This variable captures the level of opinions less level of factual details in the text. The non-retelling articles have 0 for this variable.
$Neg - Pos$	The difference between Neg and Pos. Neg is the sum of elements 6, 10, and 13 in the prompt. Pos is the sum of elements 5, 9, and 12 in the prompt. Negativity or Neg (positivity or Pos) is the degree of opposition (support) or disagreement (agreement) conveyed in the text. This variable captures the net pessimism in the text. The non-retelling articles have 0 for this variable.
$log(Appeal)$	The natural logarithm of $Appeal$ which is the sum of elements 4, 5, and 7 in the prompt. This variable captures how appealing the text is perceived. The non-retelling articles have 0 for this variable.
$log(abret_{t+1,t+21})$	The natural logarithm of compounded abnormal returns between t and t+21 adjusted using the Fama-French three-factor model plus momentum (Fama and French, 1993; Carhart, 1997). We estimate expected returns using a 100-day trading window with a 50-day gap (70 day minimum window).

Variable	Definition
$log(abret_{t-1,t})$	The natural logarithm of compounded abnormal returns between $t-1$ and t adjusted using the Fama-French three-factor model plus momentum (Fama and French, 1993; Carhart, 1997). We estimate expected returns using a 100-day trading window with a 50-day gap (70 day minimum window).
$abret_{t-3,t-2}$	The compounded abnormal returns between $t-3$ and $t-2$ adjusted using the Fama-French three-factor model plus momentum (Fama and French, 1993; Carhart, 1997). We estimate expected returns using a 100-day trading window with a 50-day gap (70 day minimum window).
$abret_{t-5,t-4}$	The compounded abnormal returns between $t-5$ and $t-4$ adjusted using the Fama-French three-factor model plus momentum (Fama and French, 1993; Carhart, 1997). We estimate expected returns using a 100-day trading window with a 50-day gap (70 day minimum window).
$abret_{t-7,t-6}$	The compounded abnormal returns between $t-7$ and $t-6$ adjusted using the Fama-French three-factor model plus momentum (Fama and French, 1993; Carhart, 1997). We estimate expected returns using a 100-day trading window with a 50-day gap (70 day minimum window).
$abret_{t-19,t}$	The compounded abnormal returns between $t-19$ and t adjusted using the Fama-French three-factor model plus momentum (Fama and French, 1993; Carhart, 1997). We estimate expected returns using a 100-day trading window with a 50-day gap (70 day minimum window).
$AbLogTurnover_{t+1,t+63}$	The difference between average daily log turnover between $t+1$ and $t+63$ and the average daily log turnover between t–252 and t–63 (9-month period, skipping the most recent quarter). Turnover is calculated as the daily volume scaled by shares outstanding.
$AbLogTurnover_{t+1,t+21}$	The difference between average daily log turnover between $t-62$ and t and the average daily log turnover between t–140 and t–20 (6-month period, skipping the most recent month). Turnover is calculated as the daily volume scaled by shares outstanding.
$AbLogTurnover_{t-1,t}$	The difference between average daily log turnover between $t-62$ and t and the average daily log turnover between t–140 and t–20 (6-month period, skipping the most recent month). Turnover is calculated as the daily volume scaled by shares outstanding.
$AbLogTurnover_{t-62,t}$	The difference between average daily log turnover between $t-62$ and t and the average daily log turnover between t–252 and t–63 (9-month period, skipping the most recent quarter). Turnover is calculated as the daily volume scaled by shares outstanding.
$log(AnalystDisper_{t+1,t+63})$	The natural logarithm of standard deviation of analyst forecasts made in the quarter post article publication (between $t+1$ and $t+63$).
$log(TickerCount)$	The natural logarithm of the number of tickers mentioned in the exclusive WSJ article.
$readability$	The Flesch Reading Ease of the text to capture linguistic complexity (Kincaid et al., 1975).
$neg-pos_{lm}$	Is the net of the proportion of negative words in the text and positive words in the text according to the 2018 version of the Loughran and McDonald (2011) lexicon.

Variable	Definition
numbers	The total instances of digits or numbers identified in the text.
complex	The percentage of words in the text that are complex. We define complex words as those with two more syllables.
beta	Market beta is estimated using the Fama and MacBeth (1973) regressions using weekly returns and equal weighted market returns for three years.
idiovol	Idiosyncratic volatility is calculated following Ali et al. (2003) by taking the standard deviation of residuals of weekly returns on weekly equal weighted market returns for three year window.
ill	Amihud (2002) illiquidity measure is calculated by taking the average of daily absolute returns divided by dollar volume.
lev	Leverage ratio is calculated following Bhandari (1988) by taking total liabilities and divide by the market capitalization.
roic	The profitability of companies is calculated using the definition in Brown and Rowe (2007) by taking the annual earnings before interest and taxes net of non-operating income and scale by non-cash enterprise value.
bm	The book-to-market ratio is calculated following Rosenberg et al. (1985) as the book value of equity divided by market capitalization.
$log(mv)$	The natural logarithm of market capitalization of the stock calculated by multiplying the price of the share by the shares outstanding.

References

Ali, A., Hwang, L.-S., Trombley, M. A., 2003. Arbitrage risk and the book-to-market anomaly. Journal of Financial Economics 69, 355–373.

Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5, 31–56.

Barasch, A., Berger, J., 2014. Broadcasting and narrowcasting: How audience size affects what people share. Journal of Marketing Research 51, 286–299.

Barber, B. M., Lee, Y.-T., Liu, Y.-J., Odean, T., 2009. Just how much do individual investors lose by trading? Review of Financial Studies 22, 609–632.

Barber, B. M., Odean, T., 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. Journal of Finance 55, 773–806.

Barber, B. M., Odean, T., 2008. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. Review of Financial Studies 21, 785–818.

Barber, B. M., Odean, T., Zhu, N., 2008. Do retail trades move markets? Review of Financial Studies 22, 151–186.

Baron, D. P., 2005. Competing for the public through the news media. Journal of Economics & Management Strategy 14, 339–376.

Barrett, J., Nyhof, M., 2001. Spreading non-natural concepts: The role of intuitive conceptual structures in memory and transmission of cultural materials. Journal of Cognition and Culture 1, 69–100.

Besley, T., Prat, A., 2006. Handcuffs for the grabbing hand? media capture and government accountability. American Economic Review 96, 720–736.

Bhandari, L. C., 1988. Debt/equity ratio and expected common stock returns: Empirical evidence. Journal of Finance 43, 507–528.

Boehmer, E., Jones, C. M., Zhang, X., Zhang, X., 2021. Tracking retail investor activity. Journal of Finance 76, 2249–2305.

Brown, D. P., Rowe, B., 2007. The productivity premium in equity returns. Available at SSRN 993467 .

Carhart, M. M., 1997. On persistence in mutual fund performance. Journal of Finance 52, 57–82.

Cookson, J. A., Engelberg, J. E., Mullins, W., 2023. Echo chambers. Review of Financial Studies 36, 450–500.

Devlin, J., Chang, M.-W., Lee, K., Toutanova, K., 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 .

Diether, K. B., Malloy, C. J., Scherbina, A., 2002. Differences of opinion and the cross section of stock returns. Journal of Finance 57, 2113–2141.

Dougal, C., Engelberg, J., Garcia, D., Parsons, C. A., 2012. Journalists and the stock market. Review of Financial Studies 25, 639–679.

Engelberg, J., Sasseville, C., Williams, J., 2012. Market madness? the case of mad money. Management Science 58, 351–364.

Engelberg, J. E., Parsons, C. A., 2011. The causal impact of media in financial markets. Journal of Finance 66, 67–97.

Fama, E. F., French, K. R., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3–56.

Fama, E. F., MacBeth, J. D., 1973. Risk, return, and equilibrium: Empirical tests. Journal of Political Economy 81, 607–636.

Fang, L., Peress, J., 2009. Media coverage and the cross-section of stock returns. Journal of Finance 64, 2023–2052.

Garcia, D., 2013. Sentiment during recessions. Journal of Finance 68, 1267–1300.

Gentzkow, M., Shapiro, J. M., 2006. Media bias and reputation. Journal of Political Economy 114, 280–316.

Gentzkow, M., Shapiro, J. M., 2008. Competition and truth in the market for news. Journal of Economic perspectives 22, 133–154.

Gentzkow, M., Shapiro, J. M., 2010. What drives media slant? evidence from us daily newspapers. Econometrica 78, 35–71.

Groseclose, T., Milyo, J., 2005. A measure of media bias. The quarterly journal of economics 120, 1191–1237.

Gurun, U. G., Butler, A. W., 2012. Don't believe the hype: Local media slant, local advertising, and firm value. Journal of Finance 67, 561–598.

Hansen, A. L., Kazinnik, S., 2023. Can chatgpt decipher fedspeak? Working Paper .

Hillert, A., Jacobs, H., Müller, S., 2014. Media makes momentum. Review of Financial Studies 27, 3467–3501.

Hribar, P., McInnis, J., 2012. Investor sentiment and analysts' earnings forecast errors. Management science 58, 293–307.

Huberman, G., Regev, T., 2001. Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. Journal of Finance 56, 387–396.

Hyman, I. E., 1994. Conversational remembering: Story recall with a peer versus for an experimenter. Applied Cognitive Psychology 8, 49–66.

Jaccard, P., 1901. Distribution de la flore alpine dans le bassin des dranses et dans quelques régions voisines. Bull Soc Vaudoise Sci Nat 37, 241–272.

Jiang, F., Lee, J., Martin, X., Zhou, G., 2019. Manager sentiment and stock returns. Journal of Financial Economics 132, 126–149.

Jones, C. M., Lamont, O. A., 2002. Short-sale constraints and stock returns. Journal of Financial Economics 66, 207–239, limits on Arbitrage.

Karpoff, J. M., 1986. A theory of trading volume. The Journal of Finance 41, 1069–1087.

Kincaid, J. P., Fishburne Jr, R. P., Rogers, R. L., Chissom, B. S., 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel .

Lopez-Lira, A., Tang, Y., 2023. Can chatgpt forecast stock price movements? return predictability and large language models. Working Paper .

Loughran, T., McDonald, B., 2011. When is a liability not a liability? textual analysis, dictionaries, and 10-ks. Journal of Finance 66, 35–65.

Melumad, S., Meyer, R., Kim, Y. D., 2021. The dynamics of distortion: How successive summarization alters the retelling of news. Journal of Marketing Research 58, 1058–1078.

Mullainathan, S., Shleifer, A., 2005. The market for news. American economic review 95, 1031–1053.

Obaid, K., Pukthuanthong, K., 2022. A picture is worth a thousand words: Measuring investor sentiment by combining machine learning and photos from news. Journal of Financial Economics 144, 273–297.

Reuter, J., Zitzewitz, E., 2006. Do ads influence editors? advertising and bias in the financial media. The Quarterly Journal of Economics 121, 197–227.

Rosenberg, B., Reid, K., Lanstein, R., 1985. Persuasive evidence of market inefficiency. Journal of Portfolio Management 11, 9–16.

Solomon, D. H., Soltes, E., Sosyura, D., 2014. Winners in the spotlight: Media coverage of fund holdings as a driver of flows. Journal of Financial Economics 113, 53–72.

Tetlock, P. C., 2007. Giving content to investor sentiment: The role of media in the stock market. Journal of Finance 62, 1139–1168.

Tetlock, P. C., 2010. Does public financial news resolve asymmetric information? Review of Financial Studies 23, 3520–3557.

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

Wiggers, K., 2023. Openai releases gpt-4, a multimodal ai that it claims is state-of-the-art. Accessed on April 2, 2024.

Yang, S., 2023. Predictive patentomics: Forecasting innovation success and valuation with chatgpt. Working Paper .

Table 1: Exclusive WSJ articles sample

This table displays descriptive information for the exclusive WSJ articles sample. The sample consists of all the WSJ articles labelled as exclusive ('PMDM') directly by Dow Jones. The sample period is between 2013 and 2022.

	N				
# of retold articles-match observations	1,940				
# of unique retold articles	1,351				
# of non-retold articles	27,524				
# of retold articles by year	2013: 100	2014: 174	2015: 174	2016: 121	2017: 143
	2018: 150	2019: 132	2020: 137	2021: 113	2022: 107
# of non-retold articles by year	2013: 4957	2014: 6485	2015: 4396	2016: 1852	2017: 1448
	2018: 1484	2019: 1772	2020: 1741	2021: 1753	2022: 1636
# of retold articles with tickers	359				
# of non-retold articles with tickers	9,388				
# of news organizations	18				
News organizations with most retellings	INVD': 589	'NYPO': 356	'WPCO': 226	'USAT': 212	'NYDN': 98
	mean	median	std	max	min
# of tickers reflected in retold articles	2.702	2	3.531	42	1
# of tickers reflected in non-retold articles	1.661	1	1.38	30	1
# of retellings per article	1.404	1	0.818	5	1
# of days between and retellings	2.635	1	3.38	14	0

Table 2: ChatGPT-4 distortion variables

Panel A reports the descriptive statistics for the distortion measures derived from the exclusive WSJ articles that have been retold (A) and the respective retelling articles ($Match$) utilizing ChatGPT. Panel B presents the mean difference for each of the distortion measures across the exclusive WSJ article sample and the respective retelling articles sample. The sample is between 2013 and 2022. All variables are defined in Appendix D.

Panel A: Distortion Variables (ChatGPT)						
Variable	Number of Obs	Mean	Median	Std Dev	Max	Min
$Fact_A$	1,940	11.915	12.000	1.141	14.000	3.000
$Fact_{Match}$	1,940	10.458	11.000	2.697	14.000	2.000
$Opinion_A$	1,940	8.728	8.000	3.129	20.000	3.000
$Opinion_{Match}$	1,940	9.159	9.000	3.204	21.000	3.000
Neg_A	1,940	6.439	6.000	2.465	21.000	3.000
Neg_{Match}	1,940	7.103	7.000	3.097	21.000	3.000
Pos_A	1,940	6.096	6.000	2.064	15.000	3.000
Pos_{Match}	1,940	5.978	6.000	2.107	17.000	3.000
$Appeal_A$	1,940	15.677	16.000	2.026	21.000	4.000
$Appeal_{Match}$	1,940	14.809	15.000	2.990	21.000	4.000
Panel B: Means Difference						
Variable	Diff (A-Match)	t-statistic	p-value			
Fact	1.457	21.910	0.000			
Opinion	-0.431	-4.243	0.000			
Neg	-0.664	-7.388	0.000			
Pos	0.119	1.771	0.077			
Appeal	0.868	10.585	0.000			

Table 3: Textual variables

This table displays the descriptive statistics for various textual variables derived from the exclusive WSJ articles that have been retold (A) and the respective retelling articles ($Match$). *readability* is the Flesch Reading Ease designed to assess the clarity of English writing. *numbers* is the count of numbers in the text. *complex* is the percentage of words in the text that have two or more syllables. $neg - pos_{lm}$ is the percentage of negative words minus positive words using the 2018 version of the Loughran and McDonald (2011) lexicon. The sample is between 2013 and 2022. All variables are defined in Appendix D.

Automated Text Variables						
Variable	Number of Obs	Mean	Median	Std Dev	Max	Min
$readability_A$	1,940	57.225	57.060	8.301	84.170	15.650
$readability_{Match}$	1,940	51.108	51.565	12.872	98.410	-9.820
$numbers_A$	1,940	14.814	13.000	8.256	69.000	6.000
$numbers_{Match}$	1,940	9.880	6.000	11.548	187.000	0.000
$complex_A$	1,940	0.170	0.169	0.035	0.312	0.050
$complex_{Match}$	1,940	0.162	0.162	0.037	0.321	0.039
$neg - pos_A$	1,940	1.844	1.373	2.088	19.231	-2.308
$neg - pos_{Match}$	1,940	1.863	1.404	2.286	14.894	-4.930
$cosine\_similarity$	1,940	0.704	0.721	0.114	0.929	0.178

Table 4: Firm characteristics and market variables

Panel A reports descriptive statistics for the characteristics of firms in the exclusive WSJ articles. Panel B presents the descriptive statistics for the market variables of firms in the exclusive WSJ articles. Panels C and D focus on comparing the group of variables in Panels A and B, respectively, for exclusive WSJ articles that are *retold* (*retold*1) and those that are not retold (*retold*0). The sample is between 2013 and 2022. All variables are defined in Appendix D.

Panel A: Firm Characteristics

Variable	Number of Obs	Mean	Median	Std Dev	Max	Min
retelling	16,863	0.071	0.000	0.257	1.000	0.000
retold	16,863	0.061	0.000	0.240	1.000	0.000
beta	16,819	1.118	1.072	0.454	3.325	-0.124
idiovol	16,819	0.035	0.029	0.019	0.260	0.010
ill	16,863	0.002	0.000	0.037	2.205	0.000
lev	16,863	2.807	0.706	4.723	59.307	0.000
roic	16,819	0.095	0.081	0.296	1.090	-12.468
bm	16,863	0.481	0.333	0.474	4.008	-1.212
mve_m	16,863	139028.347	60530.352	248767.872	2902368.097	7.877

Panel B: Returns

Variable	Number of Obs	Mean	Median	Std Dev	Max	Min
$ret_{t-1,t}$	16,863	0.006	0.002	0.096	7.816	-0.609
$ret_{t+1,t+21}$	16,862	0.012	0.012	0.111	3.392	-0.850
$abret_{t-1,t}$	16,863	0.005	0.000	0.094	7.808	-0.643
$abret_{t+1,t+21}$	16,862	-0.002	-0.003	0.105	2.860	-0.942
$AbLogTurnover_{t-1,t}$	16,857	0.005	0.000	0.042	3.563	-0.208
$AbLogTurnover_{t+1,t+21}$	16,855	0.001	0.000	0.015	0.587	-0.214

Panel C: Firm Characteristics by Retelling						
Variable	Number of Obs	Mean	Median	Std Dev	Max	Min
beta_retold0	15,791	1.125	1.080	0.455	3.325	-0.124
*beta_retold*1	1,028	1.000	0.950	0.414	3.140	-0.109
*idiovol_retold*0	15,791	0.035	0.029	0.019	0.260	0.010
*idiovol_retold*1	1,028	0.034	0.029	0.020	0.260	0.014
*ill_retold*0	15,830	0.002	0.000	0.038	2.205	0.000
*ill_retold*1	1,033	0.001	0.000	0.015	0.348	0.000
*lev_retold*0	15,830	2.890	0.723	4.812	59.307	0.000
*lev_retold*1	1,033	1.535	0.558	2.769	22.025	0.000
*roic_retold*0	15,787	0.095	0.079	0.300	1.090	-12.468
*roic_retold*1	1,032	0.104	0.098	0.227	0.767	-2.547
*bm_retold*0	15,830	0.490	0.340	0.481	4.008	-1.212
*bm_retold*1	1,033	0.342	0.276	0.325	2.477	-0.879
*mve_m_retold*0	15,830	136427.678	59605.010	246043.314	2902368.097	7.877
*mve_m_retold*1	1,033	178881.769	84148.665	284481.774	2509775.134	27.040

Panel D: Returns by Retelling						
Variable	Number of Obs	Mean	Median	Std Dev	Max	Min
$ret\_retold0_{t-1\_t}$	15,830	0.006	0.002	0.097	7.816	-0.609
$ret\_retold1_{t-1,t}$	1,033	0.010	0.001	0.072	1.242	-0.341
$ret\_retold0_{t+1,t+21}$	15,829	0.012	0.012	0.110	3.392	-0.780
$ret\_retold1_{t+1,t+21}$	1,033	0.004	0.007	0.126	1.170	-0.850
$abret\_retold0_{t-1,t}$	15,830	0.004	0.000	0.096	7.808	-0.643
$abret\_retold1_{t-1,t}$	1,033	0.010	0.001	0.069	1.197	-0.353
$abret\_retold0_{t+1,t+21}$	15,829	-0.002	-0.002	0.104	2.860	-0.901
$abret\_retold1_{t+1,t+21}$	1,033	-0.012	-0.011	0.123	1.096	-0.942
$AbLogTurnover\_retold0_{t-1,t}$	15,827	0.005	0.000	0.042	3.563	-0.208
$AbLogTurnover\_retold1_{t-1,t}$	1,030	0.008	0.000	0.040	0.818	-0.064
$AbLogTurnover\_retold0_{t+1,t+21}$	15,825	0.001	0.000	0.015	0.587	-0.214
$AbLogTurnover\_retold1_{t+1,t+21}$	1,030	0.004	0.000	0.022	0.335	-0.168

Panel E: Means Difference			
Variable	Diff (retold1-retold0)	*t*-statistic	p-value
beta	-0.125	9.334	0.000
idiovol	0.000	0.048	0.962
ill	-0.001	1.249	0.212
lev	-1.354	14.366	0.000
roic	0.009	-1.238	0.216
bm	-0.148	13.701	0.000
mve_m	42454.091	-4.683	0.000
$ret_{t-1,t}$	0.005	-2.095	0.036
$ret_{t+1,t+21}$	-0.008	1.994	0.046
$abret_{t-1,t}$	0.006	-2.508	0.012
$abret_{t+1,t+21}$	-0.011	2.747	0.006
$AbLogTurnover_{t-1,t}$	0.003	-2.274	0.023
$AbLogTurnover_{t+1,t+21}$	0.003	-3.745	0.000

Table 5: Retelling articles and distortion

This table presents coefficients from article level ordinary least squares regressions of ChatGPT-4 distortion measures on article characteristics. The sample is between 2013 and 2022. All variables are defined in Appendix D. In all regressions, we include month-year fixed effects and the t-statistics are based on standard errors clustered by article. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Dept Variable	(1) $Opinion - Fact$	(2) $Opinion - Fact$	(3) $Neg - Pos$	(4) $Neg - Pos$	(5) $log(Appeal)$	(6) $log(Appeal)$
$retelling$	1.9146***	1.9514***	0.7986***	0.8488***	-0.0656***	-0.0533***
	(28.82)	(20.77)	(10.48)	(8.73)	(-10.97)	(-6.33)
$log(TickerCount)$		-0.7566***		-0.6240***		-0.0152***
		(-5.79)		(-5.22)		(-2.61)
$readability$		0.0252***		0.0013		-0.0010**
		(3.76)		(0.22)		(-2.25)
$neg - pos_{lm}$		0.1301***		0.5099***		-0.0093***
		(4.20)		(15.92)		(-5.18)
$numbers$		-0.0284***		0.0030		0.0036***
		(-4.52)		(0.56)		(5.92)
$complex$		3.2238		4.2538**		0.0267
		(1.46)		(2.16)		(0.20)
_cons	-3.2182***	-4.4181***	0.2820***	-0.9951*	2.8045***	2.8323***
	(-44.37)	(-6.24)	(3.61)	(-1.66)	(817.17)	(60.43)
N	3794	3794	3794	3794	3794	3794
adj. R-sq	0.111	0.150	0.094	0.211	0.049	0.107

Table 6: The role of specialization, memory, and motivation on distortion

This table presents coefficients from article level ordinary least squares regressions of ChatGPT-4 distortion measures on article characteristics. The sample is between 2013 and 2022. All variables are defined in Appendix D. In all regressions, we include month-year fixed effects and the t-statistics are based on standard errors clustered by article. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Dept Variable	(1) $Opinion - Fact$	(2) $Neg - Pos$	(3) $log(Appeal)$
$log(NewsOrgFreq)$	0.2926***	-0.0750*	-0.0189***
	(6.81)	(-1.70)	(-6.65)
$log(DaysBetween)$	0.3454***	0.5110***	0.0070
	(3.54)	(4.44)	(1.25)
$log(RetellingsSameDay)$	0.1570	0.6224***	0.0241**
	(0.84)	(3.37)	(2.19)
$log(TickerCount)$	-0.7631***	-0.6144***	-0.0146**
	(-5.87)	(-5.18)	(-2.56)
$readability$	0.0257***	-0.0046	-0.0014***
	(3.65)	(-0.72)	(-2.93)
$neg - pos_{lm}$	0.1310***	0.5066***	-0.0094***
	(4.19)	(15.77)	(-5.31)
$numbers$	-0.0329***	0.0030	0.0038***
	(-5.02)	(0.57)	(6.22)
$complex$	3.1124	3.6203*	-0.0236
	(1.38)	(1.80)	(-0.18)
_cons	-4.2814***	-0.5389	2.8610***
	(-5.81)	(-0.85)	(57.45)
N	3794	3794	3794
adj. R-sq	0.143	0.218	0.114

Table 7: Abnormal returns and distortion

This table presents coefficients from a stock-article level ordinary least squares regressions of contemporaneous (Panel A) and future (Panel B) abnormal performance on article characteristics including the ChatGPT-4 distortion variables and controls. The sample is between 2013 and 2022. All variables are defined in Appendix D. In all regressions, we include industry fixed effects and the t-statistics are based on standard errors clustered by firm. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A

	(1)	(2)	(3)	(4)
Dept Variable		$log(abret_{t-1,t})$		
$retold$	0.0066***	0.0073***	0.0075***	0.0074***
	(2.76)	(3.02)	(3.07)	(3.04)
$retelling$	0.0065**	0.0071***	0.0073***	
	(2.45)	(2.62)	(2.59)	
$NegPer$				0.0007***
				(2.69)
$log(NewsOrgFreq)$				0.0028***
				(2.74)
$log(DaysBetween)$				-0.0106***
				(-4.80)
$log(RetellingsSameDay)$				0.0071*
				(1.71)
$log(TickerCount)$			-0.0025**	-0.0020**
			(-2.48)	(-2.01)
$readability$			0.0000	0.0001
			(0.48)	(1.45)
$neg-pos_{lm}$			-0.0018***	-0.0017***
			(-5.64)	(-5.46)
$numbers$			-0.0001*	-0.0001***
			(-1.85)	(-2.67)
$complex$			-0.0123	-0.0029
			(-0.59)	(-0.14)
$beta$		0.0013	0.0014	0.0015
		(0.36)	(0.40)	(0.43)
$idiovol$		-0.2680**	-0.2676**	-0.2687**
		(-2.14)	(-2.16)	(-2.17)
ill		0.1289**	0.1294**	0.1292**
		(2.21)	(2.24)	(2.24)
lev		0.0003	0.0004	0.0004
		(0.94)	(1.05)	(1.08)
$roic$		-0.0390**	-0.0387**	-0.0386**
		(-2.15)	(-2.13)	(-2.13)
bm		-0.0075*	-0.0073*	-0.0075*
		(-1.85)	(-1.83)	(-1.88)
$log(mv)$		-0.0038***	-0.0038***	-0.0038***
		(-5.41)	(-5.30)	(-5.31)
$abret_{t-3,t-2}$	-0.0410	-0.0422	-0.0436	-0.0434
	(-1.21)	(-1.34)	(-1.40)	(-1.39)
$abret_{t-5,t-4}$	-0.0077	-0.0114	-0.0113	-0.0115
	(-0.24)	(-0.38)	(-0.38)	(-0.39)
$abret_{t-7,t-6}$	-0.0889	-0.0862	-0.0867	-0.0872
	(-1.05)	(-1.04)	(-1.04)	(-1.05)
_cons	0.0020***	0.0831***	0.0883***	0.0818***
	(3.44)	(5.42)	(4.43)	(4.17)
N	16856	16768	16768	16768
adj. R-sq	0.019	0.074	0.077	0.080

Panel B

Dept Variable	(1)	(2)	(3)	(4)
		$log(abret_{t+1,t+21})$		
retold	-0.0141***	-0.0144***	-0.0144***	-0.0144***
	(-2.59)	(-2.65)	(-2.67)	(-2.67)
retelling	-0.0099**	-0.0102***	-0.0107***	
	(-2.48)	(-2.60)	(-2.58)	
$NegPer$				-0.0013**
				(-2.42)
$log(NewsOrgFreq)$				-0.0017
				(-1.09)
$log(DaysBetween)$				-0.0077
				(-1.60)
$log(RetellingsSameDay)$				-0.0005
				(-0.06)
$log(TickerCount)$			-0.0010	-0.0011
			(-0.48)	(-0.54)
readability			-0.0001	-0.0001
			(-0.84)	(-1.11)
$neg-pos_{lm}$			-0.0004	-0.0004
			(-0.60)	(-0.66)
numbers			-0.0000	0.0000
			(-0.10)	(0.15)
complex			0.0128	0.0070
			(0.40)	(0.21)
beta		-0.0051	-0.0050	-0.0052
		(-0.68)	(-0.66)	(-0.69)
idiovol		-0.9492***	-0.9536***	-0.9540***
		(-4.74)	(-4.77)	(-4.78)
ill		-0.0306	-0.0310	-0.0310
		(-0.42)	(-0.42)	(-0.42)
lev		-0.0005	-0.0005	-0.0005
		(-0.90)	(-0.87)	(-0.86)
roic		-0.0078	-0.0077	-0.0078
		(-0.90)	(-0.89)	(-0.90)
bm		0.0042	0.0043	0.0042
		(0.68)	(0.71)	(0.68)
$log(mv)$		-0.0037***	-0.0038***	-0.0039***
		(-3.43)	(-3.48)	(-3.53)
$abret_{t-19,t}$	-0.0290	-0.0268	-0.0269	-0.0269
	(-0.97)	(-0.89)	(-0.89)	(-0.89)
_cons	-0.0059***	0.0987***	0.1055***	0.1101***
	(-5.16)	(4.45)	(4.13)	(4.26)
N	16855	16767	16767	16767
adj. R-sq	0.073	0.094	0.094	0.094

Table 8: Abnormal turnover and distortion

This table presents coefficients from a stock-article level ordinary least squares regressions of contemporaneous (Panel A) and future (Panel B) abnormal turnover on article characteristics including the ChatGPT-4 distortion variables and controls. The sample is between 2013 and 2022. All variables are defined in Appendix D. In all regressions, we include industry fixed effects and the t-statistics are based on standard errors clustered by firm. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A

Dept Variable	(1)	(2) $AbLogTurnover_{t-1,t}$	(3)	(4)
$retold$	0.0015*	0.0025***	0.0024***	0.0023***
	(1.88)	(3.09)	(2.87)	(2.86)
$retelling$	0.0027***	0.0037***	0.0039***	
	(3.01)	(4.33)	(3.51)	
$NegPer$				0.0004***
				(4.04)
$log(NewsOrgFreq)$				0.0008**
				(2.01)
$log(DaysBetween)$				0.0004
				(0.41)
$log(RetellingsSameDay)$				0.0014
				(0.84)
$log(TickerCount)$			-0.0033***	-0.0033***
			(-4.08)	(-3.98)
$readability$			0.0001	0.0001
			(0.85)	(1.16)
$neg - pos_{lm}$			0.0002	0.0002*
			(1.58)	(1.74)
$numbers$			0.0000	0.0000
			(0.53)	(0.30)
$complex$			-0.0316***	-0.0289**
			(-2.58)	(-2.35)
$beta$		-0.0033	-0.0034	-0.0033
		(-0.71)	(-0.72)	(-0.71)
$idiovol$		0.2157	0.2115	0.2115
		(1.20)	(1.18)	(1.18)
ill		0.0570	0.0573	0.0573
		(1.00)	(1.01)	(1.01)
lev		0.0004	0.0004	0.0004
		(1.49)	(1.51)	(1.50)
$roic$		-0.0170	-0.0170	-0.0169
		(-1.60)	(-1.60)	(-1.60)
bm		-0.0022	-0.0024	-0.0024
		(-0.97)	(-1.03)	(-1.03)
$log(mv)$		-0.0035***	-0.0035***	-0.0035***
		(-5.62)	(-5.70)	(-5.68)
$abret_{t-3,t-2}$	-0.0105	-0.0117	-0.0118	-0.0118
	(-0.17)	(-0.19)	(-0.19)	(-0.19)
$abret_{t-5,t-4}$	0.0612***	0.0554***	0.0554***	0.0554***
	(3.78)	(3.46)	(3.47)	(3.46)
$abret_{t-7,t-6}$	0.0447	0.0393	0.0389	0.0388
	(0.57)	(0.50)	(0.49)	(0.49)
_cons	0.0044***	0.0639***	0.0691***	0.0671***
	(7.82)	(5.04)	(5.13)	(4.96)
N	16850	16762	16762	16762
adj. R-sq	0.048	0.107	0.109	0.110

Panel B

Dept Variable	(1)	(2)	(3)	(4)
		$AbLogTurnover_{t+1,t+21}$		
$retold$	0.0021***	0.0024***	0.0023***	0.0023***
	(3.63)	(4.13)	(3.95)	(3.95)
$retelling$	0.0012***	0.0015***	0.0014***	
	(2.60)	(3.16)	(2.80)	
$NegPer$				0.0002***
				(4.16)
$log(NewsOrgFreq)$				0.0006***
				(2.75)
$log(DaysBetween)$				-0.0001
				(-0.12)
$log(RetellingsSameDay)$				-0.0000
				(-0.04)
$log(TickerCount)$			-0.0006***	-0.0006**
			(-2.59)	(-2.28)
$readability$			0.0000	0.0000
			(0.17)	(1.08)
$neg-pos\_lm$			0.0002***	0.0002***
			(3.27)	(3.44)
$numbers$			-0.0000	-0.0000
			(-0.76)	(-1.38)
$complex$			0.0007	0.0028
			(0.16)	(0.63)
$beta$		0.0005	0.0005	0.0005
		(0.64)	(0.65)	(0.69)
$idiovol$		-0.0016	-0.0034	-0.0036
		(-0.03)	(-0.07)	(-0.08)
ill		0.0059	0.0058	0.0057
		(0.83)	(0.82)	(0.82)
lev		0.0003**	0.0003**	0.0003**
		(2.16)	(2.18)	(2.17)
$roic$		-0.0001	-0.0001	-0.0001
		(-0.05)	(-0.06)	(-0.05)
bm		-0.0016	-0.0016	-0.0016
		(-1.31)	(-1.34)	(-1.34)
$log(mv)$		-0.0009***	-0.0009***	-0.0009***
		(-5.11)	(-5.06)	(-5.03)
$abret_{t-19,t}$	0.0214***	0.0208***	0.0208***	0.0208***
	(2.99)	(2.95)	(2.96)	(2.95)
_cons	0.0008***	0.0163***	0.0168***	0.0153***
	(5.94)	(4.18)	(3.29)	(3.01)
N	16848	16760	16760	16760
adj. R-sq	0.159	0.171	0.172	0.173

Table 9: Abnormal turnover by retail and institutional investors

This table presents coefficients from a stock-article level ordinary least squares regressions of contemporaneous and future abnormal turnover for retail and institutional traders on article characteristics and controls. The sample is between 2013 and 2022. All variables are defined in Appendix D. In all regressions, we include industry fixed effects and the t-statistics are based on standard errors clustered by firm. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)
	$AbLogTurnover_{t-62,t}$		$AbLogTurnover_{t+1,t+63}$	
Dept Variable	Retail	Institutional	Retail	Institutional
$retold$	0.0515*	-0.0922	0.0667***	-0.1994
	(1.94)	(-0.39)	(2.59)	(-1.01)
$retelling$	0.0488**	0.1216	0.0505**	0.1125
	(1.99)	(0.46)	(2.08)	(0.54)
$log(TickerCount)$	-0.0249***	-0.0214	-0.0191**	0.0519
	(-2.89)	(-0.15)	(-2.29)	(0.41)
$readability$	-0.0006	0.0150*	-0.0005	0.0135*
	(-0.74)	(1.73)	(-0.72)	(1.68)
$neg-pos_{lm}$	0.0116***	-0.0632**	0.0131***	-0.0832***
	(3.18)	(-2.02)	(3.95)	(-2.70)
$numbers$	0.0004	0.0020	0.0001	0.0011
	(0.93)	(0.54)	(0.30)	(0.27)
$complex$	-0.1906	2.5663	0.0461	4.9912**
	(-1.02)	(1.34)	(0.29)	(2.08)
$beta$	-0.0168	-0.5408	0.0255	-0.7940**
	(-0.37)	(-1.21)	(0.68)	(-2.35)
$idiovol$	1.7842	-12.4441	0.5068	1.9009
	(1.07)	(-0.82)	(0.31)	(0.24)
ill	-0.0386	1.6522	0.3344	1.2184
	(-0.41)	(1.37)	(1.19)	(1.59)
lev	0.0032	0.0538*	0.0085**	0.0339
	(1.47)	(1.72)	(2.00)	(1.13)
$roic$	0.0116	-0.4558	0.0297	0.0146
	(0.35)	(-0.89)	(0.64)	(0.05)
bm	0.0100	-0.5803	-0.0294	-0.0776
	(0.22)	(-1.23)	(-0.73)	(-0.27)
$log_m v$	-0.0164**	0.0386	-0.0194***	0.0608
	(-2.43)	(0.50)	(-2.77)	(0.94)
$abret_{t-19,t}$	0.6889***	1.0708***	0.7104***	0.0120
	(4.13)	(2.70)	(3.37)	(0.01)
_cons	0.3187*	-0.7258	0.3194*	-1.9684
	(1.79)	(-0.36)	(1.70)	(-1.22)
N	15378	14242	15357	14431
adj. R-sq	0.177	0.013	0.252	0.012

Table 10: Macro articles and market returns

This table presents coefficients from a time-series regressions of contemporaneous and future market returns (CRSP value-weighted index, $VWRETD$) on average article characteristics including the ChatGPT-4 distortion variables and controls. We only include articles that do not have tickers referenced when averaging article characteristics on a given day. The sample is between 2013 and 2022. All variables are defined in Appendix D. In all regressions, the t-statistics are based on robust standard errors. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dept Variable	$VWRETD_{t-1,t}$			$VWRETD_{t+1,t+21}$		
$retold$	0.0052**	0.0055**	0.0055**	-0.0021	-0.0034	-0.0034
	(2.01)	(2.10)	(2.11)	(-0.25)	(-0.37)	(-0.37)
$retelling$	-0.0016	-0.0014		-0.0054	-0.0052	
	(-0.59)	(-0.55)		(-0.74)	(-0.70)	
$NegPer$			0.0002			0.0030**
			(0.31)			(2.20)
$log(NewsOrgFreq)$			-0.0034*			-0.0068
			(-1.94)			(-1.61)
$log(DaysBetween)$			-0.0028			-0.0017
			(-1.13)			(-0.19)
$log(RetellingsSameDay)$			0.0155**			0.0361**
			(2.48)			(2.07)
$readability$		0.0000	0.0000		0.0000	0.0001
		(0.34)	(0.45)		(0.36)	(0.53)
$neg - pos_{lm}$		-0.0002	-0.0002		0.0004	0.0004
		(-0.93)	(-1.05)		(0.71)	(0.66)
$numbers$		-0.0000	-0.0000		0.0001	0.0000
		(-0.55)	(-0.62)		(0.67)	(0.48)
$complex$		-0.0006	0.0004		0.0612*	0.0659**
		(-0.05)	(0.04)		(1.85)	(1.99)
$abret_{t-3,t-2}$	-0.0101	-0.0099	-0.0093			
	(-0.26)	(-0.26)	(-0.24)			
$abret_{t-5,t-4}$	-0.0471	-0.0469	-0.0456			
	(-1.33)	(-1.32)	(-1.28)			
$abret_{t-7,t-6}$	-0.0254	-0.0251	-0.0263			
	(-0.68)	(-0.67)	(-0.70)			
$abret_{t-19,t}$				-0.1775***	-0.1780***	-0.1795***
				(-5.25)	(-5.27)	(-5.31)
_cons	0.0011***	0.0009	0.0005	0.0130***	-0.0021	-0.0039
	(3.26)	(0.21)	(0.11)	(12.45)	(-0.17)	(-0.31)
N	2303	2303	2303	2303	2303	2303
adj. R-sq	0.002	0.001	0.002	0.030	0.031	0.032

Table 11: Analyst forecast dispersion

This table presents coefficients from a stock-article level ordinary least squares regressions of future analyst forecast dispersion on article characteristics and controls. The sample is between 2013 and 2022. All variables are defined in Appendix D. In all regressions, we include industry fixed effects and the t-statistics are based on standard errors clustered by firm. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)
Dept Variable	$log(AnalystDisper_{t+1,t+63})$		
$retold$	0.2535***	0.2745***	0.2777***
	(4.48)	(5.15)	(5.17)
$retelling$	0.2442***	0.2775***	0.2930***
	(4.36)	(5.12)	(5.14)
$log(TickerCount)$	0.2385***	0.1925***	0.1916***
	(5.80)	(2.68)	(2.68)
$readability$			-0.0358
			(-1.15)
$neg - pos_{lm}$			0.0000
			(0.03)
$numbers$			0.0029
			(0.24)
$complex$			0.0012
			(1.14)
$beta$			0.7147
			(1.41)
$idiovol$		0.2916***	0.2915***
		(4.69)	(4.75)
ill		6.5722***	6.5084***
		(3.10)	(3.08)
lev		-0.4791	-0.4882
		(-1.16)	(-1.17)
$roic$		-0.0292***	-0.0299***
		(-2.81)	(-2.88)
bm		-0.0061	-0.0048
		(-0.09)	(-0.07)
$log(mv)$		-0.1129	-0.1110
		(-1.40)	(-1.38)
$abret_{t-19,t}$		-0.0949***	-0.0962***
		(-5.41)	(-5.47)
_cons	-0.3011***	0.8968**	0.8132*
	(-9.49)	(2.37)	(1.75)
N	7287	7255	7255
adj. R-sq	0.071	0.140	0.140