Generative AI and Investor Processing of Financial Media*

Tony Cho
The Hong Kong University of Science and Technology, accho@ust.hk

Allen H. Huang
The Hong Kong University of Science and Technology, allen.huang@ust.hk

Joseph Pacelli
Harvard Business School, jpacelli@hbs.edu

Kristina Rennekamp†
Cornell University, kmr52@cornell.edu

Abstract

We examine how AI-generated summaries in financial news influence investor information processing. Exploiting The Wall Street Journal's rollout of AI summaries, we find that articles featuring these summaries elicit stronger immediate market reactions, including higher trading volume and amplified price responses. Additional analyses show that the increased trading is concentrated in the content summarized by AI, but not in the content not summarized by AI. We investigate the underlying mechanism using a controlled experiment. Contrary to the concern that AI summaries encourage skimming, we find they increase both perceived confidence and objective recall of the full article text. These findings suggest that AI summaries do not substitute for reading but rather serve as interpretive cues, enhancing overall comprehension while directing investor attention toward highlighted elements. Cross-sectional analyses show that the effects of AI-generated summaries are strongest in settings with greater information frictions. Finally, articles with AI-generated summaries are associated with significantly greater subsequent information search by both retail and institutional investors. Overall, our study documents how the integration of AI in financial media reshapes the consumption of firm news.

JEL Codes: D83; G12; G14; G41; M41;

Keywords: Generative Artificial Intelligence; Financial Media; Information Processing; Investor Attention; Market Reaction

---

* We appreciate and acknowledge comments from Ed deHaan, Travis Dyer, Nick Guest, and Brady Twedt, as well as workshop participants at Central University of Finance and Economics, Fudan University, HKUST, and Renmin University. We used AI for limited editorial assistance, including proofreading and copyediting of the manuscript, and to help draft certain materials used in the experimental instrument. All AI-generated content was reviewed, edited, and verified by the authors, who take full responsibility for the final manuscript and study materials. The authors acknowledge the financial support provided by HKUST, HBS and Cornell. † Corresponding author.

1. Introduction

Recent advancements in artificial intelligence are reshaping how investors access information in capital markets. Financial news outlets, in particular, have begun experimenting with new ways to incorporate AI. One approach that has gained traction is the use of AI-generated summaries (AI summaries). Prominent news outlets, including The Wall Street Journal (WSJ), Bloomberg, and Yahoo Finance, have introduced summaries that algorithmically extract and highlight a short set of key points from selected articles before the full contents. While AI summaries are intended to help readers identify key points, it is unclear whether these summaries meaningfully amplify investor attention and trading activity, or merely repackage information that investors would have processed anyway from headlines and articles, and whether investors benefit from them.[1] In this study, we examine whether and how AI summaries in the business press influence investors' processing and responses to news.

We examine this question using two complementary approaches (Bloomfield, Nelson, and Soltes [2016]). First, we use archival data to assess whether the presence of AI summaries is associated with investors' market responses and information-seeking behaviors. We exploit the WSJ's rollout of AI summaries to selected articles and compare trading activities and market outcomes across similar articles with and without summaries. These tests allow us to examine real-world market reactions at scale and also assess whether our effects vary systematically with article characteristics, such as complexity and readability. The archival evidence, however, cannot reveal *how* investors process information or why they respond differently to articles with summaries. We

[1] Recent evidence raises concerns about AI summary accuracy and completeness. A European Broadcasting Union study found significant issues in 45% of AI-generated news summaries (EBU [2025]), and Bloomberg corrected at least 36 AI summaries in its first months of deployment (Robertson [2025]). Despite these concerns, major outlets continue expanding their use of AI summaries, creating an important setting to examine their effects on investor behavior.

thus also conduct a controlled experiment in which participants are randomly assigned to read a news article either with or without an AI-generated summary. This experiment allows us to directly measure investor comprehension through information recall, thereby identifying how summaries affect investors.

Ex ante, it is not clear whether AI summaries alter how investors react to firm news. On the one hand, summaries may enhance information processing and amplify trading responses. Theories of limited attention suggest that investors face cognitive constraints when processing large volumes of financial news (e.g., Kahneman [1973], Sims [2003], Hirshleifer, Lim, and Teoh [2009]). By previewing article content, summaries help investors identify which articles warrant attention, encouraging engagement with news they might otherwise overlook and thereby increasing trading (e.g., Engelberg and Parsons [2011]). Beyond expanding the audience, summaries may also enhance comprehension and trading confidence among readers by reducing the cognitive effort required to form an initial understanding (Navon [1997], Bargh and Chartrand [1999]). In addition, because the summaries are labeled as being algorithmically generated, investors may trust them more than content written by a human editor, who has incentives to selectively highlight information favorable to companies that generate advertising revenue (e.g., Reuter and Zitzewitz [2006], Gurun and Butler [2012]).

On the other hand, AI summaries repackage information already contained in the article, potentially offering little beyond what investors already glean from headlines and opening paragraphs (e.g., Huang, Nekrasov, and Teoh [2018]). Investors may also be uncertain whether the AI reliably identifies value-relevant content, given that the technology is still nascent, thus leading to "algorithm aversion" (Coleman, Merkley, and Pacelli [2022], Commerford, Dennis, and Joe [2022]). Moreover, sophisticated investors and algorithmic traders who rapidly process full articles

or generate their own summaries using AI tools may find little incremental value in publishers' AI summaries (Blankespoor, Croom, and Grant [2026], Ecker et al.[2026]). Under these conditions, AI summaries should not meaningfully alter investors' responses to firm news.

Our archival analyses draw on all WSJ articles covering publicly traded firms published in the one year after July 2024, when the first WSJ article with AI summaries appeared. We collect the full article, including headline and subheading (for simplicity, we refer to both as headline hereafter), and, when available, the associated AI summaries. After applying filters to ensure the availability of return and trading data, our final sample consists of 1,734 articles about 158 unique firms, of which approximately one-third include an AI summary.

We begin by testing whether the presence of an AI summary is systematically related to the characteristics of covered firms or to article features. Across a series of regressions and variance decompositions, we find little evidence that AI summary usage is associated with firm fundamentals, including performance, market-to-book, or leverage. Instead, most of the explained variation is attributable to time fixed effects and other media coverage. These patterns are inconsistent with AI being selectively deployed to certain topics or firms.

Our main market reaction tests proceed in two stages. First, we conduct across-article analyses to assess whether articles with AI summaries elicit differential market responses. To the extent that AI summaries highlight relevant information and help investors process firm news, we expect articles containing such summaries to generate more trading activity. Using a propensity-score-matched sample, we find that AI summaries are associated with significantly higher trading volume (3.5%) in the 30 minutes following publication, suggesting that investors react more strongly to articles with AI summaries. Absolute short-window returns over the same short window also increase, indicating that the summary amplifies price reactions. Additional tests using entropy

balancing and two different control groups (matched and same-firm pre-AI-rollout samples) mitigate endogeneity concerns and confirm the robustness of the results.

Next, we explore differential market reactions to content summarized by AI versus content not summarized by AI, providing further evidence that the presence of AI summaries drives these results. We construct our own GPT (hereafter GPT) to summarize all WSJ articles and confirm that its outcomes align with the WSJ summaries. We then identify paragraphs within each article that have been summarized by the GPT and classify the remaining paragraphs as those not summarized by GPT. We find that investor reactions to the tone of GPT-summarized paragraphs are stronger when an actual WSJ AI summary accompanies the article. In contrast, their reactions to the tone of non-GPT-summarized paragraphs do not depend on the presence of a WSJ AI summary. These results indicate that AI summaries are associated with stronger investor reactions to the news they summarize, consistent with the idea that summaries help investors better process and incorporate newly released information into trading.

Our findings in the market reaction tests are consistent with two competing mechanisms. First, AI summaries may lead investors to rely solely on the summarized points, as they read the rest of the article less thoroughly. A second possibility is that investors read and understand the whole article, but weigh the information that aligns with the summary more heavily because they interpret it as a cue about which elements are most important or most reliable. To shed light on which mechanism drives our findings, we conduct a controlled experiment that allows us to measure attention to news content. In the experiment, participants are randomly assigned to read a WSJ news article either with or without an accompanying AI summary. After reading the article, participants complete a series of measures capturing their perceived understanding of the firm's performance and their actual comprehension of information included and excluded from the

summary. This design allows us to address endogeneity concerns in the archival tests and compare how the presence of a summary affects both subjective and objective comprehension.

The results show that the presence of an AI summary does not reduce participants' effort or comprehension. Participants exposed to the summary condition had significantly higher perceived understanding of the firm's performance, consistent with summaries enhancing confidence in grasping the key points. More importantly, they also demonstrate higher factual recall of the article's content, including information featured in the summary as well as information appearing only in the body. These findings indicate that participants continue to read and process the whole article even when a summary is present, and that the summary enhances, rather than substitutes for, overall comprehension.

Limited-attention theories suggest that the effects we document thus far should be stronger when articles are more difficult to process or when readers are cognitively constrained. We therefore conduct additional archival analyses to test whether the impact of AI summaries varies with article complexity and the extent of cognitive constraints. Consistent with our intuition, the results indicate that AI summaries elicit a significantly stronger immediate market reaction to the articles when they are less readable and longer, when they contain a greater concentration of numerical information, and on days with many earnings announcements and industry news.

Beyond establishing that AI summaries increase trading activity, an important question is whether such effects benefit investors. Higher trading volume may be driven by informed responses that improve price discovery or attention-driven noise trading that degrades market quality. We next examine how AI summaries relate to several market outcomes, including liquidity (bid-ask spreads and price impact) and pricing efficiency (intra-period pricing error and post-article return drifts). If summaries help investors process firm news more effectively, we expect

articles with AI summaries to exhibit tighter spreads, lower price impact, and reduced pricing errors. Our results confirm these expectations.

Our final analyses examine whether AI-generated summaries affect the extent to which investors seek out additional information. On the one hand, summaries may lower the cost of forming an initial understanding and thus lead investors to seek more information because they can more easily identify which issues warrant further investigation. On the other hand, summaries may provide a "false" sense of completeness, thereby discouraging further information gathering. We test these competing predictions in both the archival and experimental settings.

In the archival setting, we examine how AI summaries relate to Google Search Volume and Bloomberg readership measures and find that articles with AI summaries are associated with significantly greater subsequent information search across both platforms, suggesting summaries prompt additional inquiry by both individual and professional investors. In the experimental setting, participants rate their interest in learning more about the firm and indicate which of nine pre-specified sources (e.g., the 10-K, equity analyst reports) they would consult before making an investment decision. In contrast to the archival results, we find little evidence that AI summaries meaningfully change participants' stated intentions to seek further information, suggesting that summary effects on information acquisition are stronger in real-world settings, where investors are likely already choosing to engage with a firm.[2]

Our study contributes to the media literature, particularly to the growing interest in technology's role in how the media processes and disseminates information. A large body of

[2] In the archival setting, investors who read firm-specific news are likely already interested in the firm, making subsequent information acquisition a natural margin of adjustment. In contrast, the experimental setting presents participants with a firm that may not be salient to them ex ante, which can limit variation in stated information-seeking intentions regardless of whether a summary is present. Importantly, this difference does not affect the experimental tests of attention and comprehension, which rely on participants' processing of the article content itself rather than on their motivation to pursue additional external information.

research shows that the media influences how investors incorporate firm news by helping to uncover and disseminate information (Drake, Guest, and Twedt [2014], Rogers, Skinner, and Zechman [2016], Twedt [2016], Lawrence et al. [2018], Guest [2021], Call et al. [2022]). Recent research has begun to examine the implications of newsrooms adopting new technologies to assist in content production. For example, Blankespoor, deHaan, and Zhu [2018] show that algorithmically generated earnings articles produced by the Associated Press increase trading volume and liquidity, with effects concentrated among retail traders.[3] Whereas their setting shows how automated technology synthesizes firm disclosures for previously under-covered firms, we examine how AI condenses journalist-produced content investors already access. Our setting yields two findings absent from prior work. First, pricing efficiency improves, consistent with the WSJ reaching a broader investor audience than Associated Press. Second, recall improves even for content not in the summary, indicating that summaries deepen engagement with the underlying article rather than substitute for reading it.

More broadly, our study also contributes to the growing literature examining how generative AI shapes information production and processing in the capital markets, outside of the media setting (Bertomeu et al. [2025], Blankespoor, deHaan, and Li [2025], Choi and Xie [2025], Cheng, Lin, and Zhao [2025], Croom [2025], Jia et al. [2025], Blankespoor, Croom, and Grant [2026], Bradshaw et al. [2026], Chang et al. [2026], Ecker et al. [2026], Hirshleifer et al. [2026], Ramella [2026]). Although this nascent stream of research emphasizes AI's generative capabilities and the associated benefits for investors, little is known about how AI-mediated summaries affect how investors attend to and interpret information. This represents an important gap in the literature,

[3] Recent studies in communication and journalism literature have explored how generative AI changes the production and consumption of news (Hong, Chang, and Tewksbury [2025], Leuppert, Weinmann, and Eiden [2025], Ashuri, Zalmanson, and Goldstein [2026]).

as investors are increasingly relying on AI not only to generate new information but also to filter and condense the information they already face.[4] Understanding how summaries influence investors is thus an important extension to the existing literature.

Finally, our study contributes to the literature on how presentation formats influence investors' information processing. Prior research shows that the way information is displayed, including its placement and sequencing, can materially affect investor processing and decision-making (e.g., Elliott [2006], Files, Swanson, and Tse [2009], Rennekamp [2012], Cardinaels, Hollander, and White [2019], Chen, Gee, and Neilson [2021]). Much of this work examines settings in which firms may strategically choose which information to summarize or emphasize, such as in earnings press releases, earnings call highlights, and narrative disclosures (Huang, Nekrasov, and Teoh [2018], Chen, Gee, and Neilson [2021], Asay, Hinds, and Rennekamp [2025]). We extend this literature by examining a technology-driven presentation format.[5] Unlike traditional formatting choices, the summaries in our setting are produced through an automated process that operates independently of firm discretion, introducing a non-strategic layering of information that elevates content based on importance rather than intent. Our findings show that this format facilitates information processing, particularly when investors are more cognitively constrained, and that AI summaries enhance rather than substitute for deep processing of news.

2. Institutional Background

[4] Several recent studies leverage the availability of generative AI tools to document that investors benefit from these tools (Bertomeu et al. [2025], Cheng, Lin, and Zhao [2025], Even-Tov et al. [2025], Chang et al. [2026]). Nevertheless, it is not clear how investors use the tool, e.g., search for, retrieve or summarize information, or evaluate, interpret or analyze information (Croom, Gale, and Grant [2026], Ecker et al. [2026]).

[5] Our study complements prior work by Cardinaels, Hollander, and White [2019], who show that management summaries in earnings press releases are more optimistically biased than summarizes generated automatically using technology to process an earnings press release. While Cardinaels, Hollander, and White [2019] primarily focus on whether automated summaries can be a useful tool for "debiasing" disclosures prepared by firms, our focus is instead on whether information intermediaries enhance the usefulness of their information when they include AI-generated summaries for users.

Over the past several years, news organizations have increasingly experimented with artificial intelligence to support the production and distribution of journalism (Blankespoor, deHaan, and Zhu [2018], Blankespoor et al. [2019], Ashuri, Zalmanson, and Goldstein [2026]). These efforts reflect broader changes in how readers consume news, including rising information volume and shorter attention spans (Hong et al. [2025], Scire [2025]). In financial journalism, where articles often contain dense quantitative information and time-sensitive disclosures, publishers have begun exploring AI-generated summaries to help readers quickly identify the main points of an article without replacing the underlying reporting.

Major financial news outlets such as Bloomberg, The Wall Street Journal, and Yahoo News have publicly described AI summaries as a reader-facing convenience feature designed to complement, rather than substitute for, full-length articles (Micklethwait [2025], Scire [2025]). These summaries typically appear at the top of an article as short bullet points that condense the story's core factual elements. Publishers emphasize that the summaries are generated from the article text itself and are subject to editorial oversight, reflecting concerns about accuracy, trust, and transparency (Gallardo [2025], Robertson [2025]).

The WSJ was among the earliest major U.S. financial newspapers to deploy AI-generated summaries at scale. According to newsroom accounts, the Journal began developing the feature in early 2024 and rollout the summaries to its core news articles following internal testing (Scire [2025]). The summaries, labeled as "Quick Summary," are generated using an AI model integrated directly into the Journal's content management system and are accompanied by explicit disclosures informing readers that an artificial intelligence tool created the summary and that an editor reviewed it. Journal editors retain discretion over whether summaries appear on individual articles (Gallardo [2025], Scire [2025]).

Public reporting suggests that the Journal's primary motivation for introducing AI-generated summaries was to improve reader orientation and engagement, particularly for longer or more complex articles, rather than to alter the substance of its reporting (Scire [2025]). Internal presentations emphasize metrics such as time spent on articles, articles per session, and subscriber engagement, reflecting a focus on how presentation affects consumption rather than information production (Gallardo [2025]). At the same time, experiences at other outlets highlight potential concerns associated with AI summaries, including occasional inaccuracies and the risk that readers may rely too heavily on summaries at the expense of full articles, underscoring the importance of understanding how such features shape information processing (Robertson [2025], Sato [2024]).

Taken together, these institutional features make the WSJ a particularly useful setting for studying the effects of AI-generated summaries on investor information processing. The summaries are prominently displayed and explicitly labeled. Moreover, as a leading source of financial news for both retail and institutional investors, changes in how WSJ articles are presented can influence capital market behavior in economically meaningful ways.

3. Sample

3.1. Data

We begin by identifying all firms with valid tickers in the CRSP/Compustat Merged database. For these firms, we use their corresponding Factiva firm identifiers to collect WSJ online articles published between July 2024 and June 2025. For each article, we extract the headline and subheading, full article text, publication timestamp, and WSJ section. This initial step yields 2,877 articles covering 502 unique firms. We further impose several filters. First, we delete 209 articles from WSJ's "Opinion," "Arts," "Style News," and "Science" sections because all articles in these sections do not use AI summaries. Second, we require intraday TAQ trading data and daily CRSP

return data. This requirement removes 632 observations. Third, we delete singleton observations to ensure within-firm variation in article coverage, consistent with our use of firm fixed effects in the main analyses. We also remove a small number of articles for which GPT-based text analysis cannot be applied due to content restrictions.[6] The final sample comprises 1,734 articles from 158 unique firms.[7] Table 1, Panel A, summarizes the sample selection procedure.

<INSERT TABLE 1 ABOUT HERE>

To identify articles with AI summaries, we search for the labels "Quick Summary" and "An artificial-intelligence tool created this summary" on the WSJ webpage of each article. Approximately 37 percent of articles in the final sample include an AI summary.[8] We extract the AI summaries, presented in bullet-point format between these two labels. Untabulated statistics show that the vast majority of summaries have three bullet points (97.3%), with the remaining articles having two or four bullet points (2.1% and 0.6%, respectively).[9]

Figure 1 presents the monthly time-series distribution of articles with AI summaries in the sample. The blue bars represent the number of articles with AI summaries, while the red line depicts the percentage of articles with AI summaries. The first article with AI summaries appears in July 2024. The number jumps to about 40 articles in September 2024 and surges to over 110 by March 2025, accounting for about 70% of the articles that month. The trend is consistent with the WSJ testing the waters with AI summaries early in the sample period and ramping up the use later.

[6] OpenAI prohibits using its services for violence, ilicit activites, self-harm, and more. An example article with content of such nature is "Violent Attacks on Tesla Are Surging. Is It Domestic Terrorism?".
[7] On average, each article in this final sample concerns 1.62 firms. An untabulated robustness test shows that limiting to articles that only concern one firm produces similar results.
[8] Around late July and early August of 2025, WSJ changed the presentation format of AI summaries. It changed "Quick Summary" to "Key Points" and presented the AI summaries in entirety. It soon reverted back to the original format in about a week.
[9] The results are robust to removing articles with two or four bullet points.

Table 1, Panel B, presents the Fama-French 12 industry distribution of the articles. We observe that articles cover a broad range of industries, with most industries receiving a fairly even number of articles. More importantly, AI summaries are not concentrated in specific industries. Across industries, about 30% to 44% of articles have AI summaries, except in Utilities (53%), which comprises only 15 articles. Panel C provides a breakdown of sample composition by WSJ sections. Not surprisingly, most articles appear in the "Business" and "Markets" sections, consistent with their focus on specific firms. About 30% to 40% of articles in these two sections are summarized by AI.

3.2. Determinants of AI Summary

A potential concern in our analyses is that AI summary usage may be systematically related to firm characteristics or the nature of the news being covered. If the WSJ selectively deploys AI summaries on articles about firms with greater investor interest or better performance, then any observed differences in market reactions could partly reflect these underlying differences rather than the presence of the summary itself. To assess this concern, we examine the determinants of AI summary usage in two ways.

We begin with a variance decomposition of AI summary usage into different factors. We examine time (year-month), industry, firm, and section. The results, tabulated in Table 2, Panel A, show that AI usage is not well explained by industry, firm, or sections (0.63%, 4.91%, and 3.78%, respectively). Instead, most of the variation is explained by time fixed effects (29.62%), consistent with WSJ not deploying it for select topics or firms.

<INSERT TABLE 2 ABOUT HERE>

We further estimate a determinants model of AI summary use at the article level. The model includes time-varying firm characteristics such as firm size, profitability, market-to-book ratio,

and leverage. We also add article details, including the number of other media articles, the overall tone of those articles, and whether the firm issued a press release on the day the article was published. Table 2, Panel B, compares articles with and without AI summaries. Articles with AI summaries tend to be shorter and are accompanied by more articles from other media and by firm-initiated press releases. By contrast, firm size and media sentiment do not differ meaningfully between articles with and without AI summaries.

Table 2, Panel C, presents regressions of an indicator for AI summary usage on article characteristics and firm fundamentals, estimated with year-month and section fixed effects and, in some specifications, firm fixed effects. The results reveal three noteworthy patterns. First, article length is a robust predictor of AI summary usage. The coefficient on *WordCount* is negative and statistically significant across specifications, indicating that a one-standard-deviation increase in article length is associated with a roughly 2 to 4 percentage-point decrease in the likelihood that an article includes an AI summary. Second, contemporaneous media coverage is positively associated with AI usage. A one-standard-deviation increase in same-day media coverage corresponds to a 5 to 10 percentage point increase in the probability of an AI summary. By contrast, firm fundamentals are largely unrelated to AI summary usage. Firm size, profitability, growth opportunities, leverage, and the presence of a same-day press release are all weak predictors of AI summary usage, and their significance disappears once firm fixed effects are included. Third, firm and article characteristics explain minimal variation in AI summaries. The adjusted R-squareds of the specifications range from 32.40% to 34.04%, suggesting that these additional variables add little beyond the time and section fixed effects (which had an adjusted R-squared of 32.55% from an untabulated analysis).

Overall, the results indicate that AI summaries are more likely to appear in shorter articles and in periods of heightened media attention, but do not appear to be selectively deployed based on economically informative firm characteristics. This pattern is consistent with AI summaries being introduced through internal testing and rollout processes, rather than targeted placement with particular firms or types of news.

4. AI Summary and Investor Trading – Archival Results

4.1. AI Summary and Investor Trading – Across-Article Tests on Average Investor Responses

We begin with *across-article* analyses to assess whether and to what extent articles with AI summaries elicit differential market responses compared with those without. We estimate the following OLS regression model:

$$y_{i,s,t} = \beta AI_{i,s,t} + Controls_{i,s,t} + \gamma_i + \eta_s + \delta_t + \epsilon_{i,s,t}$$

where i indexes a firm, s indexes a section, and t indexes time. The dependent variable is either abnormal trading volume ($AbnTradeVol$) or absolute abnormal return ($|Return|$). Abnormal trading volume is the natural logarithm of one plus the total number of shares traded during the 30-minute window following the article publication, minus the natural logarithm of one plus the average total number of shares traded during the same 30-minute window on the same day of the week during the past month. Absolute abnormal return is the absolute value of market-adjusted return in percentage during the 30-minute window following the article publication (Tetlock [2008], Chen, Nagar, and Schoenfeld [2018]). The independent variable of interest is AI, an indicator variable that takes the value of one if an article is summarized by AI, and zero otherwise. If AI summaries meaningfully shape investor processing, we expect β to be positive.

We control for firm and article characteristics that may influence investor reaction. We control for firm size, word count, media coverage, media sentiment, and whether a press release is issued on the day of article publication. Appendix A includes detailed definitions of all variables. Firm fixed effects (γ_i) and section fixed effects (η_s) control for time-invariant heterogeneity across firms and sections, respectively, and time fixed effects (δ_t) control for time-varying heterogeneity.

Table 3, Panel A, presents the descriptive statistics for the full sample. About 37% of articles have AI summaries. The articles, on average, have 735 words and cover firms with an average market capitalization of $142 billion. About 144 news articles are published on the same day as the WSJ article, with media sentiment slightly positive. Roughly 66% of the articles are published on the same day as a press release.

<INSERT TABLE 3 ABOUT HERE>

Results from Section 3.2 show that articles with AI summaries differ from non-summarized articles in certain characteristics. For instance, articles with greater media coverage on the publication date are more likely to be summarized by AI. To account for any such differences and mitigate endogeneity, we control for firm and article characteristics in the regressions and use a propensity-score-matched sample (Rosenbaum and Rubin [1983], Shipman, Swanquist, and Whited [2017]).[10] Specifically, each article with AI summaries is matched to an article without AI summaries on $Size$, ROA, MTB, $Leverage$, $MediaCoverage$, $MediaTone$, and $PressRelease$.[11] We use the nearest-neighbor method with a caliper of 0.2, without replacement. The process yields a match for each of 574 articles with AI summaries. Table 3, Panel B, provides

[10] Untabulated analyses using the full sample yield qualitatively similar results.

[11] An indicator variable for whether the articles are published during trading hours weakly positively predicts the use of AI summaries. The tone of WSJ article positively predicts the use of AI summaries. The Fog Index of WSJ article insignificantly positively predicts the use of AI summaries. Untabulated tests find that all results are robust to matching on and including these variables as additional control variables. Furthermore, untabulated tests show that the results are robust to matching on sections.

covariate balance after propensity score matching, showing no significant differences in the matched sample.

Table 4 presents the regression results across different specifications. Columns (1) and (3) do not include any controls, while Columns (2) and (4) control for firm and article characteristics. In Column (1), a significantly positive coefficient on AI indicates that abnormal trading volume is higher for articles with AI summaries. Controlling for firm and article characteristics, Column (2) retains a significantly positive coefficient on AI, indicating about 3.5% higher abnormal trading volume for articles with AI summaries.

Turning to absolute abnormal returns, Column (3) shows that articles with AI summaries have, on average, higher absolute abnormal returns. Including firm and article characteristics as controls in Column (4) yields a similar result. A significantly positive coefficient on $|Return|$ suggests that articles with AI summaries have a modest increase in price reactions of 0.209%. Taken together, the results are consistent with the idea that AI summaries amplify investor reactions to news articles.

<INSERT TABLE 4 ABOUT HERE>

In addition to the propensity score-matched sample, we consider three alternative comparison groups to address the concern that unobserved contemporaneous factors drive our results. Table 5 presents the results from using the alternative control groups. First, we use an entropy-balanced sample matched on the means and variances of the covariates (Hainmueller [2012], McMullin and Schonberger [2020]). Panel A presents the covariate balance of the samples with respect to their means and variances, indicating no significant differences. Panel B presents the results from the regressions using the entropy-balanced sample. We find significantly positive coefficients on AI, suggesting 10.2% higher abnormal trading volume and 3.1% higher price

reactions for articles with AI summaries.

The remaining two groups use pre-rollout articles as controls. The first group matches each post-rollout article with an AI summary to a pre-rollout article that would have been similarly likely to receive one. To construct this matched sample, we apply the propensity score model estimated in Table 3, Panel B, to articles published between July 2023 and June 2024, and match each AI-summarized article to a pre-rollout article with the nearest propensity score, using a caliper of 0.2 without replacement. The second group relaxes the matching requirement while maintaining firm-level comparability. Rather than matched pairs, we pool all pre-rollout articles for the firms that subsequently appear in our (post-rollout) AI-summary sample and use them as the baseline. Panel C presents the covariate balance between articles with AI summaries and pre-rollout-matched articles, showing no significant differences. Panel D presents the results from the regressions using the pre-rollout matched sample. Across different specifications, we find positively significant coefficients on AI, indicating about 9.2% higher abnormal trading and 0.591% higher absolute abnormal returns. Panel E presents the results using the same-firm pooled sample. Similarly, we find positively significant coefficients on AI, suggesting that articles with AI summaries have higher abnormal trading volume and price reactions of 6.1% and 0.345%, respectively.

<INSERT TABLE 5 ABOUT HERE>

4.2. AI Summary and Investor Trading – Across-Article Tests on Investor Responses to Summarized Contents

Next, we explore whether the differential market reactions to articles with and without AI summaries are due to investors' processing of content summarized by AI. To do so, we examine

whether investor reactions to a given portion of an article's content differ depending on whether the article carries an AI summary. We estimate the following regression using all WSJ articles:

$$y_{i,s,t} = \beta_1 |Tone_{ParaSum_{i,s,t}}| \times AI_{i,s,t} + \beta_2 |Tone_{ParaNotSum_{i,s,t}}| \times AI_{i,s,t} + \beta_3 |Tone_{ParaSum_{i,s,t}}| + \beta_4 |Tone_{ParaNotSum_{i,s,t}}| + \beta_5 AI_{i,s,t} + Controls_{i,s,t} + \gamma_i + \eta_s + \delta_t + \epsilon_{i,s,t},$$

where the dependent variable is either $AbnTradeVol$ or $|Return|$. The main independent variables include the absolute tone of the contents that AI would have summarized ($|Tone_{ParaSum}|$), the AI summary indicator (AI), and their interactions. We also include the absolute tone of other contents in the article ($|Tone_{ParaNotSum}|$), i.e., contents that would not have been summarized by AI, as placebo tests.

To identify which content AI would have summarized across all WSJ articles, not just the subset that carries actual WSJ summaries, we use ChatGPT to generate summaries for the full sample.[12] We prompt GPT with examples of WSJ articles and their associated AI summaries, instructing it to produce three bullet points per article, consistent with the format used by WSJ. To validate this approach, we compare GPT-generated summaries against actual WSJ summaries for the subset of articles where both are available, using BERTScore to measure semantic similarity. We find a close one-to-one correspondence between the two, with an average pairwise BERTScore of 0.67, confirming that GPT reliably approximates the content the WSJ AI would have highlighted.

We then identify paragraphs in each article that GPT has summarized. For each bullet point, we label the paragraph with the highest BERTScore as the one it summarizes. We measure the tone of each sentence in the summarized paragraphs using FinBERT and average across sentences in

[12] We use GPT-4o accessed through API, with a one-shot prompt. That is, we give GPT-4o one example from WSJ and ask it to summarize other articles. See Appendix B for details. In untabulated robustness tests, we find similar results using various other models including GPT-5 and Gemini 2.5 Pro.

each paragraph, and define $Tone_{ParaSum}$ as the average tone of the three summarized paragraphs (Huang, Wang, and Yang [2023]). We label remaining paragraphs as non-summarized paragraphs and calculate their tone ($Tone_{ParaNotSum}$) in a similar manner.

Table 6 presents the results. In Columns (1) and (3), we find significantly positive coefficients on $|Tone_{ParaSum}| \times AI$, indicating that investor reactions to the tone of these paragraphs are stronger when the article is accompanied by an actual WSJ AI summary. On the contrary, we do not observe a similar pattern in non-summarized paragraphs.

As an alternative to $Tone_{ParaSum}$, we measure the tone of GPT summaries and define $Tone_{GPTSum}$ as the average tone across the three bullet points. Using the tone of GPT summaries instead in Columns (2) and (4) yields similar results. Investor reactions to the summarized contents are amplified if the article is released with an AI summary. Across all specifications, there is no significant difference in investor reactions to non-summarized paragraphs between articles with and without an actual WSJ AI summary. Taken together, these results provide evidence that AI summaries are associated with stronger investor reactions to firm news, consistent with the idea that summaries help investors better process and incorporate newly released information into trading.[13]

<INSERT TABLE 6 ABOUT HERE>

4.3. AI Summary and Investor Trading – Cross-Sectional Tests

The results thus far show that AI summaries amplify investor reactions to firm news, with this amplification concentrated in the content the summary highlights. A natural follow-up question is whether these effects vary systematically with the difficulty of processing the underlying information. Limited attention theories suggest that salient cues are most valuable when

[13] The results presented above and below remain qualitatively similar when the alternative comparison groups, as outlined in Section 4.1, are used.

cognitive demands are highest, as investors are more likely to rely on interpretive shortcuts when articles are harder to parse or when they are simultaneously processing a larger volume of competing news. If AI-generated summaries mitigate information friction and serve as a salient cue, we expect their effect to be greater for more complex articles and on busy days. We use three different proxies for article complexity: readability measured using the Gunning Fog Index, word count, and proportion of numerical content. We use two proxies for the degree of cognitive constraint among investors: the number of same-day earnings announcements and the number of same-day same-industry news articles.

Table 7 presents the results from regressing abnormal trading volume and absolute abnormal return on AI usage, estimated separately for subsamples partitioned on each proxy. Panel A shows that the effect of AI summaries on both trading volume and price reactions is significantly stronger for articles with above-median Fog Index, indicating that summaries are especially valuable when articles are harder to read. Panel B reveals a similar pattern for word count. The effect of AI summaries is concentrated among longer articles, consistent with the idea that summaries help investors navigate more information-dense coverage. Panel C shows that articles with a higher proportion of numerical content also exhibit stronger AI summary effects, suggesting that summaries are particularly useful when articles require investors to process and interpret quantitative information.

Turning to cognitive constraints, Panel D shows that the effect of AI summaries is stronger on days with more concurrent earnings announcements, consistent with investors facing higher cognitive load on those days. Finally, Panel E shows a similar pattern for same-day industry news, though the difference across subsamples is less precisely estimated. Taken together, the evidence across all five proxies collectively suggests that AI summaries have a greater impact when

investors face higher information frictions, supporting the view that summaries function as salient cues whose value is highest when processing demands are most acute.

<INSERT TABLE 7 ABOUT HERE>

4.4. AI Summary and Investor Trading – Market Outcomes

Having established that AI summaries amplify investor reactions to firm news, we next examine whether these heightened responses reflect informed trading that improves market quality or attention-driven noise trading that degrades it. If AI summaries facilitate more effective information processing, we expect articles with AI summaries to exhibit improved liquidity and pricing efficiency. Conversely, if summaries merely attract uninformed attention without enhancing comprehension, we would expect to see worse market outcomes.

We assess liquidity using two measures. First, we evaluate liquidity using abnormal bid-ask spreads ($AbnSpread$), i.e., the natural logarithm of the weighted average daily spread over trading days [0, 1] divided by the weighted average daily spread over days [-41, -11], where lower values indicate lower transaction costs and better market liquidity (Hasbrouck [2009]). Second, we measure abnormal price impact ($AbnImpact$), i.e., the natural logarithm of the weighted average daily impact over trading days [0, 1] divided by the weighted average daily impact over days [-41, -11], which reflects the degree to which trades move prices; lower price impact suggests greater market depth and resilience (Amihud [2002]).

Panel A of Table 8 presents the results. Column (1) shows that articles with AI summaries exhibit significantly tighter bid-ask spreads. The coefficient on AI is -0.090 (t = -5.12), corresponding to a 9.0% reduction in abnormal spreads. This narrowing of spreads indicates that AI summaries reduce trading costs and improve market liquidity. Similarly, Column (2) documents a significant reduction in price impact. The coefficient on AI is -0.095 (t = -2.25),

representing approximately a 9.5% decrease in abnormal price impact. This finding suggests that trades following articles with AI summaries encounter less price resistance, consistent with greater market depth.

<INSERT TABLE 8 ABOUT HERE>

Next, we examine pricing efficiency using two approaches. The first one is intraperiod pricing efficiency (IPE), which captures the extent to which prices reflect information efficiently within the trading period of the week after the article day (Blankespoor, deHaan, and Zhu [2018]). Higher values indicate more efficient price discovery. We also supplement this analysis by testing whether AI summaries affect the extent of post-publication price drift. If AI summaries enable more complete, immediate incorporation of news, we expect to observe reduced drift following article publication.

Panel B of Table 8 tests these predictions. Column (1) shows that articles with AI summaries are associated with significantly higher intraperiod pricing efficiency. The coefficient on AI is 0.185 (t = 4.01), indicating that AI summaries improve pricing efficiency by approximately 18.5 percentage points relative to the sample mean. This economically meaningful improvement suggests that AI summaries help the market incorporate news more efficiently.

Column (2) establishes a baseline relationship between article tone and two-day returns. The positive and significant coefficient on $Tone$ (t = 2.68), confirming that more positive articles are associated with higher returns on the publication date. More importantly, the interaction term $AI \times Tone$ is positive and significant (t = 1.67), indicating that articles with AI summaries elicit stronger (around 81% = 0.025 / 0.031) immediate price reactions to news tone.

Column (3) examines returns over the subsequent four days ([2,5]) and includes both initial return response ($Return_{[0,1]}$) and its interaction with AI. We find a significantly positive

coefficient on $Return_{[0,1]}$ (0.150 with a *t*-stat of 3.200) and a significantly negative coefficient on $AI \times Return_{[0,1]}$ (-0.143 with a *t*-stat of -2.150). This pattern indicates that while there is a typical post-article drift, this drift does not apply to articles with AI summaries.[14] These results suggest that AI summaries facilitate more complete immediate price adjustment, leaving less residual information to be incorporated gradually over subsequent days.[15]

Taken together, the evidence in Table 8 demonstrates that the amplified trading responses to articles with AI summaries reflect informed investor activity rather than noise. Articles with AI summaries exhibit tighter spreads, lower price impact, and better pricing efficiency. These findings indicate that AI summaries enhance investors' ability to process firm news effectively, with beneficial implications for market quality and price discovery.

4.5. AI Summary and Investor Trading – Information Search

Having established that AI summaries amplify investor reactions and improve market quality, we turn to a final question: do summaries affect the extent to which investors seek out additional information after reading an article? As highlighted above, summaries could influence information acquisition in competing ways. On the one hand, summaries can help investors quickly identify the key aspects of a story, thus lowering the cost of forming an initial understanding and making it easier to recognize which issues warrant further investigation, thereby encouraging additional search. On the other hand, summaries may create a false sense of completeness, leading investors to feel sufficiently informed without consulting other sources.

[14] An untabulated test shows that the sum of the coefficients of $AI \times Return_{[0,1]}$ and $Return_{[0,1]}$ is not significantly different from zero.

[15] An untabulated test using a longer return window ([2,20]) shows that $Return_{[2,20]}$ is not significantly associated with either $Return_{[0,1]}$ or $AI \times Return_{[0,1]}$.

We test these competing predictions using two measures of post-article information search that capture the behavior of different investor types. Columns (1) and (2) in Table 9 use the Google Search Volume Index as a measure of retail investor search activity, while Columns (3) and (4) use an indicator for high Bloomberg readership to capture information search by institutional investors (Da, Engelberg, and Gao [2011], Ben-Rephael, Da, and Israelsen [2017]). The results indicate that AI usage is significantly positively associated with information search activity. Search activity increases by 3.4% if articles have AI summaries in the Google Search Volume setting. Replacing the dependent variable with an indicator variable for high Bloomberg readership in Columns (3)-(4), we find significantly positive coefficients in both specifications. Articles with AI summaries are 1.5% more likely to be highly searched in Bloomberg.[16] Taken together, the evidence is consistent with both retail and institutional investors seeking additional information after reading articles with AI summaries.

<INSERT TABLE 9 ABOUT HERE>

5. AI Summary and Investor Processing – Experiment

The archival analyses in Section 4 establish that AI summaries are associated with amplified investor trading reactions, improved market quality, and stimulated subsequent information search. However, archival data alone cannot reveal why investors respond differently to articles with summaries. In particular, two competing mechanisms could explain the patterns we observe. First, summaries may lead investors to skim the rest of the article, concentrating attention on highlighted content at the expense of broader comprehension. Second, summaries may enhance overall understanding by providing an interpretive structure that helps investors process the full article more effectively. Distinguishing between these mechanisms requires direct

[16] Excluding the day of publication from the two measures of search activity yields qualitatively similar results but with smaller economic magnitude.

measurement of investor comprehension. We therefore conduct a controlled experiment in which participants are randomly assigned to read a news article with or without an AI summary, allowing us to measure both perceived and objective understanding of the article's content.

5.1. Method

Participants

Participants are 124 graduate students recruited from the behavioral lab of the business school at a private university in exchange for a $15 Amazon Gift Card.[17] The study takes an average of 13.51 minutes to complete. Consistent with prior work using student investor populations (e.g., Elliott et al. [2007], Asay, Elliott, and Rennekamp [2017]), these participants possess sufficient familiarity with basic financial concepts to complete the task. On average, participants have completed 1.64 accounting courses and 1.85 finance courses, and 52% percent report having previously invested in individual securities. In addition, 79% percent indicate that they expect to invest in individual securities in the future. The sample includes 63% percent female and 34% percent male participants, with the remaining identifying as non-binary or preferring not to respond.

Design and Manipulations

To test whether and how AI-generated news summaries affect investors' information processing, we employ a single-factor between-subjects design that manipulates whether participants are exposed to an AI-generated summary embedded at the top of a news article about a retailer, Ridgeway Retail. Participants are randomly assigned to one of two conditions, where an AI summary is either present (AI-Present) or absent (AI-Absent).

[17] This study received the required approval from the Institutional Review Board (IRB) of the university where it was conducted.

All participants read the same Wall Street Journal–style article describing Ridgeway's recent quarterly performance. In the AI-Present condition, the article begins with a brief "Quick Summary" produced by an AI system and labeled as such. This summary highlights three key points about the firm's performance and outlook. In the AI-Absent condition, the article is identical except that the AI-generated summary is omitted. Consistent with our archival setting, the full underlying article remains constant across conditions, ensuring that any observed differences result from the presence of the AI-generated summary rather than from variation in the article's content.

Task and Procedure

Participants first receive general instructions explaining that they will evaluate a publicly traded company and answer questions about the materials they review. They are then presented with a news article about Ridgeway Retail.[18] The article provides mixed but realistic information, including stronger-than-expected revenue, weather-related demand softness, management commentary, and forward-looking sales guidance.

After reading the article, participants respond to two questions in randomized order. Specifically, they rate (1) how well they feel they understand Ridgeway's recent performance and (2) how interested they are in learning more about the company. These measures are collected immediately after exposure to the article and before any additional tasks.

Participants then answer seven comprehension questions, which assess recall of the information provided in the article. Three questions relate to items explicitly covered in the AI-generated summary (AI items), while four questions relate to information contained only in the body of the article (Article items). All seven questions appear in randomized order, and all

[18] The article used in the experiment is based on an actual article from the Wall Street Journal about a real discount retailer (Burlington Coat Factory). However, we modify the article to remove any identifying information about the company or the individuals quoted in the article, in order to prevent participants from responding based on their actual knowledge or opinions about the company.

information is available in both conditions, ensuring that performance differences reflect the cognitive influence of the summary rather than differences in access to information.

After completing the comprehension questions, participants respond to a question that captures their information search intentions. Specifically, they view a list of nine possible external information sources (e.g., the most recent annual or quarterly report, the company website, news outlets) and indicate which sources they would consult before making an investment decision. Following Asay, Elliott, and Rennekamp [2017], instructions emphasize that there are no correct answers, mitigating artificial inflation of selections due to self-presentation concerns. Consistent with this, participants, on average, indicate that they would consult 3.65 of the nine possible sources of information. Finally, participants answer manipulation check and demographic questions.

Primary Dependent Variables

To evaluate how AI-generated summaries influence investors' processing of financial information, we use two different types of measures.

Perceptions of Understanding. Participants rate their agreement with the statement, "I feel like I understand the recent performance of Ridgeway Retail" (1 = strongly disagree, 7 = strongly agree). This measure captures subjective feelings about their understanding of the firm's performance after reading the article.

Actual Understanding. Participants' responses to the comprehension questions form two separate measures: an AI Score (0–3) based on questions tied to content that is included in the AI summary, and an Article Score (0–4) based on questions tied exclusively to unsummarized content. Again, even in the condition where there is no AI-generated summary, the *content* is still present in the news article. Thus, the AI Score allows us to test whether, for a given piece of information,

recall is greater when that information is presented in an AI summary than when it is not. These two measures allow us to objectively assess how the presence of an AI-generated summary affects recall of information presented in different parts of the article.

Supplemental Measures on Investors' Information-Seeking

Interest in Accessing Additional Information. Participants rate their agreement with the statement, "I am interested in learning more about Ridgeway Retail as a Potential Investment" (1 = strongly disagree, 7 = strongly agree). This measure captures whether the inclusion of an AI summary affects investors' interest in seeking out more information about the firm discussed in an article.

Information Acquisition Intentions. Participants select which of nine additional sources of information they would seek out before making an investment decision, if any, where the options are an annual report, a quarterly report, a conference call recording, the company website, equity analyst reports, credit rating agency reports, industry trade publications, online discussion forms, or other financial news outlets. This measure captures whether the inclusion of an AI summary in a news article affects investors' intentions to search for subsequent information and which sources they would choose to access.

5.2. Results

Table 10 reports how the presence of an AI-generated summary affects participants' recall of article information, and Table 11 reports their *perceptions* of their own understanding. Panel A of Table 10 compares participants' accuracy on questions tied directly to the topics emphasized in the summary. Participants in the AI-Present condition answer significantly more of these questions correctly than participants in the AI-Absent condition (2.23 versus 1.90; $p = 0.004$).[19] This pattern

[19] All reported *p*-values are two-tailed.

suggests that the summary effectively highlights and reinforces the article's most salient points, helping participants encode and process the information it presents more comprehensively.

In Panel B, we examine questions about content that does not appear in the summary. In this setting, we again find that participants in the AI-Present condition perform better, scoring 2.49 out of 4, compared with 2.05 in the AI-Absent condition ($p = 0.037$). The fact that recall improves even for non-highlighted information is informative, as it suggests that the summary does not crowd out attention to the rest of the article. One possible explanation is that the summary provides a structure that helps participants integrate and process the broader set of facts. In this sense, the summary operates not only as a cueing device but also as a means for supporting broader comprehension of the full article.

Panel C evaluates total recall across all seven questions. Participants in the AI-Present condition answer approximately one additional question correctly, on average, compared to participants in the AI-Absent condition (4.72 versus 3.95; $p = 0.014$). The magnitude of this difference, combined with the patterns in Panels A and B, suggests that the summary enhances overall comprehension rather than redistributing effort across different pieces of information. In short, providing a brief AI-generated summary meaningfully improves what readers recall from the article as a whole.

<INSERT TABLE 10 ABOUT HERE>

Having provided objective evidence on how readers recall information, we next explore participants' subjective perceptions of their understanding after reading the news article. Table 11 presents participants' reported understanding of the firm's recent performance across both conditions. Those in the AI-Present condition report greater understanding of the firm's recent performance than those in the AI-Absent condition (5.57 versus 5.13 on a 7-point scale; $p = 0.036$),

mirroring the improvements in objective recall. This alignment suggests that participants are at least partially aware of the benefits the summary provides.

<INSERT TABLE 11 ABOUT HERE>

Finally, Table 12 presents our supplementary evidence on how the presence of AI summaries affects participants' subsequent information-seeking behavior. First, Panel A examines participants' interest in learning more about the firm as a potential investment. Here, we do not find a significant difference between the AI-Present and AI-Absent conditions (mean responses are 5.31 and 5.14, respectively; $p = 0.452$).

Panel B of Table 12 presents results examining whether the presence of AI summaries influences participants' stated inclination to seek additional information sources before making an investment decision. Across nine potential sources, selection rates are generally similar across conditions, and no difference is statistically significant at conventional levels. Overall, the results do not support the idea that AI summaries meaningfully crowd out or replace other forms of due diligence. These findings suggest that, while summaries help participants feel more informed, they do not necessarily increase or decrease their motivation to learn more about the investment.

<INSERT TABLE 12 ABOUT HERE>

Taken together, the experimental evidence shows that the presence of an AI summary enhances, rather than diminishes, participants' comprehension of the article. Participants in the AI-Present condition report higher perceived understanding and demonstrate significantly greater factual recall, not only for the information emphasized in the summary but also for information that appears exclusively in the body of the article. These patterns indicate that participants continue to read and process the full article even when a summary is available, and that the summary reinforces rather than replaces their engagement with the underlying content. Consistent with this

view, participants' stated information-seeking intentions remain largely unchanged across conditions. This evidence is inconsistent with a reduced-effort explanation and instead supports the interpretation that AI summaries shape how readers interpret and weigh the information they encounter.

6. Conclusion

This study examines whether and how AI-generated summaries in financial news shape investor information processing. Exploiting the Wall Street Journal's rollout of AI "Quick Summary," we provide evidence that AI summaries affect how investors react to firm news, even though such summaries do not introduce new information. In our archival analyses, articles featuring AI summaries elicit stronger immediate market responses. Trading volume in the short window after publication is higher when an AI summary is present, and price reactions are also amplified. We also show that the effect operates through differential weighting of information within the same article. Specifically, investor reactions to the tone of paragraphs that AI would have summarized are stronger when an actual WSJ AI summary accompanies the article. In contrast, we observe no such difference in investor reactions to the tone of non-summarized paragraphs.

Beyond establishing that AI summaries affect trading activity, we demonstrate that these effects benefit market participants. Articles with AI summaries exhibit tighter bid-ask spreads, lower price impact, and improved pricing efficiency. These findings indicate that the amplified trading responses reflect informed investor activity rather than attention-driven noise trading, with beneficial implications for market quality and price discovery.

To better understand why AI summaries have these effects and to alleviate concerns that archival tests suffer from endogeneity, we conduct a controlled experiment that directly measures

readers' comprehension. The experimental evidence is inconsistent with the concern that summaries encourage skimming or reduced effort. Participants who view an AI summary report greater perceived understanding of the firm's performance. More importantly, they demonstrate higher factual recall, not only for information included in the summary but also for information that appears only in the body of the article. This pattern indicates that the summary does not substitute for reading. Instead, it appears to provide an initial structure that improves comprehension while simultaneously cueing which elements of the article are most important. In this sense, AI summaries serve as interpretive cues that shape how readers process the news they encounter.

Consistent with this interpretation, we find that the effects of AI summaries are stronger in settings where information is more difficult to process. The market response to AI summaries increases when articles are longer, less readable, and more numerically dense, and when investors face greater cognitive constraints due to high volumes of concurrent news or earnings announcements. These cross-sectional patterns suggest that the incremental value of AI summaries is greatest when information frictions are greater. We also examine downstream information acquisition. In the archival setting, articles with AI summaries are followed by greater subsequent information search by both retail and institutional investors, as captured by Google search activity and Bloomberg readership. While we find little evidence that summaries shift stated search intentions in the experiment, the combined evidence suggests that, in real-world settings, summaries can facilitate deeper engagement by helping investors identify which issues warrant additional attention.

Our findings contribute to several literatures. First, we add to the research on financial media by showing that the press's intermediary role depends not only on what journalists write,

but also on the tools used to package and present reporting to readers. Second, we contribute to the emerging literature on generative AI in capital markets by documenting that AI can shape investor behavior through presentation and salience, even when it does not generate new economic content. Third, we extend work on disclosure and presentation formats by examining a setting in which summary and emphasis are introduced through a largely automated, non-strategic process, rather than through firm discretion. Across approaches, the evidence indicates that AI summaries enhance comprehension while shifting relative attention toward the information they highlight, with economically meaningful implications for trading and price formation.

An important limitation of our setting is that it reflects the early adoption of a specific summary format at a single major outlet. Future research could examine whether similar effects arise when summaries are personalized, deployed across different platforms, or under other oversight regimes. More broadly, as AI becomes increasingly embedded in the production and delivery of financial information, understanding how these tools shape what investors attend to, remember, and trade on will be central to assessing how information intermediaries influence capital markets.

References

Amihud, Y. "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects." *Journal of Financial Markets* 5 (2002): 31-56. https://doi.org/10.1016/S1386-4181(01)00024-6.

Asay, H. S.; W. B. Elliott; and K. Rennekamp. "Disclosure Readability and the Sensitivity of Investors' Valuation Judgments to Outside Information." *The Accounting Review* 92 (2017): 1-25. https://doi.org/10.2308/accr-51570.

Asay, H. S.; C. Hinds; and K. Rennekamp. "Strategic Formatting in Firm Disclosures." Working Paper, 2025. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4089745.

Ashuri, T.; L. Zalmanson; and D. Goldstein. "Journalistic Epistemic Authority in the Age of AI: The Incorporation of Generative AI in Newsrooms." *Digital Journalism* (2026): 1-20. https://doi.org/10.1080/21670811.2026.2640421.

Bargh, J. A., and T. L. Chartrand. "The Unbearable Automaticity of Being." *American Psychologist* 54 (1999): 462–479. https://doi.org/10.1037/0003-066X.54.7.462.

Ben-Rephael, A.; Z. Da; and R. D. Israelsen. "It Depends on Where You Search: Institutional Investor Attention and Underreaction to News." *The Review of Financial Studies* 30 (2017): 3009-3047. https://doi.org/10.1093/rfs/hhx031.

Bertomeu, J.; Y. Lin; Y. Liu; and Z. Ni. "The Impact of Generative AI on Information Processing: Evidence from the Ban of ChatGPT in Italy." *Journal of Accounting and Economics* 80 (2025). https://doi.org/10.1016/j.jacceco.2025.101782.

Blankespoor, E.; J. Croom; and S. Grant. "Generative AI and Investor Processing of Financial Information." Working Paper, 2026. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5053905.

Blankespoor, E.; E. deHaan; and Q. Li. "Generative AI in Financial Reporting." Working Paper, 2025. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4986017.

Blankespoor, E.; E. deHaan; J. Wertz; and C. Zhu. "Why Do Individual Investors Disregard Accounting Information? The Roles of Information Awareness and Acquisition Costs." *Journal of Accounting Research* 57 (2019): 53-84. https://doi.org/10.1111/1475-679X.12248.

Blankespoor, E.; E. deHaan; and C. Zhu. "Capital Market Effects of Media Synthesis and Dissemination: Evidence from Robo-Journalism." *Review of Accounting Studies* 23 (2018): 1-36. https://doi.org/10.1007/s11142-017-9422-2.

Bloomfield, R.; M. W. Nelson; and E. Soltes. "Gathering Data for Archival, Field, Survey, and Experimental Accounting Research." *Journal of Accounting Research* 54 (2016), 341-395. https://doi.org/10.1111/1475-679X.12104.

Bradshaw, M. T.; C. Ma; B. P. Yost; and Y. Zou. "Generative AI Use by Capital Market Information Intermediaries: Evidence from Seeking Alpha." Working Paper, 2026. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5226562.

Call, A. C.; S. A. Emett; E. Maksymov; and N. Y. Sharp. "Meet the Press: Survey Evidence on Financial Journalists as Information Intermediaries." *Journal of Accounting and Economics* 73 (2022). https://doi.org/10.1016/j.jacceco.2021.101455.

Cardinaels, E.; S. Hollander; and B. J. White. "Automatic Summarization of Earnings Releases, Attributes, and Effects on Investors' Judgments." *Review of Accounting Studies* 24 (2019): 860-890. https://doi.org/10.1007/s11142-019-9488-0.

Chang, A.; X. Dong; X. Martin; and C. Zhou. "AI (ChatGPT) Democratization and Trading Inequality." Working Paper, 2026. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543999.
Chen, J. V.; K. H. Gee; and J. J. Neilson. "Disclosure Prominence and the Quality of Non-GAAP Earnings." *Journal of Accounting Research* 59 (2021): 163-213. https://doi.org/10.1111/1475-679X.12344.
Cheng, Q.; P. Lin; and Y. Zhao. "Does Generative AI Facilitate Investor Trading? Early Evidence from ChatGPT Outages." *Journal of Accounting and Economics* 80 (2025). https://doi.org/10.1016/j.jacceco.2025.101821.
Cheng, L.; D. T. Roulstone; and A. Van Buskirk. "Are Investors Influenced by the Order of Information in Earnings Press Releases?" *The Accounting Review* 96 (2021): 413-433. https://doi.org/10.2308/TAR-2017-0478.
Choi, J. H., and C. Xie. "Human + AI in Accounting: Early Evidence from the Field." Working Paper, 2025. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5240924.
Coleman, B., K. Merkley, and J. Pacelli. "Human versus Machine: A Comparison of Robo-Analyst and Traditional Research Analyst Investment Recommendations." *The Accounting Review* 97 (2022): 221–244. https://doi.org/10.2308/TAR-2020-0096.
Commerford, B. P.; S. A. Dennis; J. R. Joe; and J. W. Ulla. "Man Versus Machine: Complex Estimates and Auditor Reliance on Artificial Intelligence." *Journal of Accounting Review* 60 (2022): 171-201. https://doi.org/10.1111/1475-679X.12407.
Croom, J. "Interactivity and Illusions of Ability: How Using Generative AI Affects Investor Judgments." *Journal of Accounting Research* (2025) Forthcoming. https://doi.org/10.1111/1475-679X.70017.
Croom, J.; B. Gale; S. M. Grant. "Generative AI Prompting and Investor Reactions to Disclosure Tone." Working Paper, 2026. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5040309.
Da, Z.; J. Engelberg; and P. Gao. "In Search of Attention." *The Journal of Finance* 66 (2011): 1461-1499. https://doi.org/10.1111/j.1540-6261.2011.01679.x.
deHaan, E.; A. Lawrence, A.; and R. Litjens. "Measuring Investor Attention Using Google Search." *Management Science* 71 (2025): 6275-6297. https://doi.org/10.1287/mnsc.2022.02174.
Drake, M. S.; N. M. Guest; and B. J. Twedt. "The Media and Mispricing: The Role of the Business Press in the Pricing of Accounting Information." *The Accounting Review* 89 (2014): 1673-1701. https://doi.org/10.2308/accr-50757.
EBU. "News Integrity in AI Assistants." *BBC*, October 22, 2025. Available at https://www.bbc.co.uk/mediacentre/2025/new-ebu-research-ai-assistants-news-content
Ecker, F.; X. Li; Y. Li; and F. Wu. "How Stock Market Participants Use Generative Artificial Intelligence: Evidence from User-Platform Interaction Data." Working Paper, 2026. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5224596.
Elliott, W. B. "Are Investors Influenced by Pro Forma Emphasis and Reconciliations in Earnings Announcements?" *The Accounting Review* 81 (2006): 113-133. https://doi.org/10.2308/accr.2006.81.1.113.
Elliott, W. B.; F. D. Hodge; J. J. Kennedy; and M. Pronk. "Are M.B.A. students a good proxy for nonprofessional investors?" *The Accounting Review* 82 (2007): 139-168. https://doi.org/10.2308/accr.2007.82.1.139.

Engelberg, J. E., and C. A. Parsons. "The Causal Impact of Media in Financial Markets." *The Journal of Finance* 66 (2011): 67-97. https://doi.org/10.1111/j.1540-6261.2010.01626.x.

Even-Tov, O.; B. Lourie; K. Munevar; and A. Nekrasov. "The Effect of AI on Retail Investor Behavior: Evidence from Account-Level Trading Data." Working Paper, 2025. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5778246.

Files, R.; E. P. Swanson; and S. Tse. "Stealth Disclosure of Accounting Restatements." *The Accounting Review* 84 (2009): 1495-1520. https://doi.org/10.2308/accr.2009.84.5.1495.

Gallardo, R. D. "SVP of Product, Consumer at Dow Jones | How to Build High-Quality AI Features." *YouTube*, August 7, 2025. Available at https://youtu.be/C3DO6lKtHOY?si=BY2w6syBGHfiCtPi

Guest, N. M. "The Information Role of the Media in Earnings News." *Journal of Accounting Research* 59 (2021): 1021-1076. https://doi.org/10.1111/1475-679X.12349.

Gurun, U.G., and A. W. Butler. "Don't Believe the Hype: Local Media Slant, Local Advertising, and Firm Value." *The Journal of Finance* 67 (2012): 561-598. https://doi.org/10.1111/j.1540-6261.2012.01725.x.

Hainmueller, J. "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies." *Political Analysis* 20 (2012): 25-46. https://doi.org/10.1093/pan/mpr025.

Hasbrouck, J. "Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data." *The Journal of Finance* 64 (2009): 1445-1477. https://doi.org/10.1111/j.1540-6261.2009.01469.x.

Hirshleifer, D.; S. S. Lim; and S. H. Teoh. "Driven to Distraction: Extraneous Events and Underreaction to Earnings News." *The Journal of Finance* 64 (2009): 2289-2325. https://doi.org/10.1111/j.1540-6261.2009.01501.x.

Hirshleifer, D.; L. Peng; Q. Wang; W. Zhang; and X. Zhang. "AI, Opinion Ecosystems, and Finance." Working Paper, 2026. Available at https://www.nber.org/system/files/working_papers/w34807/w34807.pdf.

Hong, J. W.; H. C. H. Chang; and D. Tewksbury. "Can AI Become Walter Cronkite? Testing the Machine Heuristic, the Hostile Media Effect, and Political News Written by Artificial Intelligence." *Digital Journalism* 13 (2025): 845-868. https://doi.org/10.1080/21670811.2024.2323000.

Huang, X.; A. Nekrasov; and S. H. Teoh. "Headline Salience, Managerial Opportunism, and Over- and-Underreaction to Earnings." *The Accounting Review* 93 (2018): 231-255. https://doi.org/10.2308/accr-52010.

Huang, A. H.; H. Wang; and Y. Yang. "FinBERT: A Large Language Model for Extracting Information from Financial Text." *Contemporary Accounting Research* 40 (2023): 806-841. https://doi.org/10.1111/1911-3846.12832.

Jia, N.; N. Li; D. Xu; and Y. Zhang. "When IR Meets GenAI." Working Paper, 2025. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5317057.

Kahneman, D. "Attention and Effort." *Prentice-Hall, Inc.* (1973).

Lawrence, A.; J. Ryans; E. Sun; and N. Laptev. "Earnings Announcement Promotions: A Yahoo Finance Field Experiment." *Journal of Accounting and Economics* 66 (2018): 399-414. https://doi.org/10.1016/j.jacceco.2018.08.004.

Leuppert, R.; C. Weinmann; and J. Elden. "AI-Reporters as the Future of Journalism? Investigating Recipients' Credibility Evaluations of AI-Authored News." *Journalism* (2025). https://doi.org/10.1177/14648849251382484.

McMullin, J. L., and B. Schonberger. "Entropy-Balanced Accruals." *Review of Accounting Studies* 25 (2020): 84-119. https://doi.org/10.1007/s11142-019-09525-9.

Micklethwait, J. "How Journalism Will Adapt in the Age of AI." *Bloomberg*, January 10, 2025. Available at https://www.bloomberg.com/news/articles/2025-01-10/8-ways-ai-will-transform-journalism/.

Navon, D. "Forest Before Trees: The Precedence of Global Features in Visual Perception." *Cognitive Psychology* 9 (1977): 353-383. https://doi.org/10.1016/0010-0285(77)90012-3.

Ramella, P. "Better AI, Worse Disclosures?" Working Paper, 2026. Available at https://pietroramella.com/material/JMP.pdf.

Rennekamp, K. "Processing Fluency and Investors' Reactions to Disclosure Readability." *Journal of Accounting Research* 50 (2012):1319-1354. https://doi.org/10.1111/j.1475-679X.2012.00460.x.

Reuter, J., and E. Zitzewitz. "Do Ads Influence Editors? Advertising and Bias in the Financial Media." *The Quarterly Journal of Economics* 121 (2006): 197-227. https://doi.org/10.1093/qje/121.1.197.

Robertson, K. "Bloomberg Has a Rocky Start With A.I. Summaries." *The New York Times*, March 29, 2025. Available at https://www.nytimes.com/2025/03/29/business/media/bloomberg-ai-summaries.html.

Rogers, J. L.; D. J. Skinner; and S. L. C. Zechman. "The Role of the Media in Disseminating Insider-Trading News." *Review of Accounting Studies* 21 (2016): 711-739. https://doi.org/10.1007/s11142-016-9354-2.

Rosenbaum, P. R., and D. B. Rubin. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70 (1983): 41-55. https://doi.org/10.1093/biomet/70.1.41.

Sato, M. "Newspaper conglomerate Gannett is adding AI-generated summaries to the top of its articles / Journalists participating in the pilot program will use AI to produce bulleted 'key points' of their story." *The Verge*, May 17, 2024. Available at https://www.theverge.com/2024/5/16/24158531/gannett-ai-generated-overviews-usa-today-memo.

Scire, S. "Let's get to the point: Three Newsrooms on Generating AI Summaries for News." *Nieman Lab*, June 5, 2025. Available at https://www.niemanlab.org/2025/06/lets-get-to-the-point-three-newsrooms-on-generating-ai-summaries-for-news/.

Shipman, J. E.; Q. T. Swanquist; and R. L. Whited. "Propensity Score Matching in Accounting Research." *The Accounting Review* 92 (2017): 213-244. https://doi.org/10.2308/accr-51449.

Sims, C. A. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (2003): 665-690. https://doi.org/10.1016/S0304-3932(03)00029-1.

Twedt, B. "Spreading the Word: Price Discovery and Newswire Dissemination of Management Earnings Guidance." *The Accounting Review* 91 (2016): 317-346. https://doi.org/10.2308/accr-51129.

Figure 1: Frequency of Articles with AI Summaries Over Time

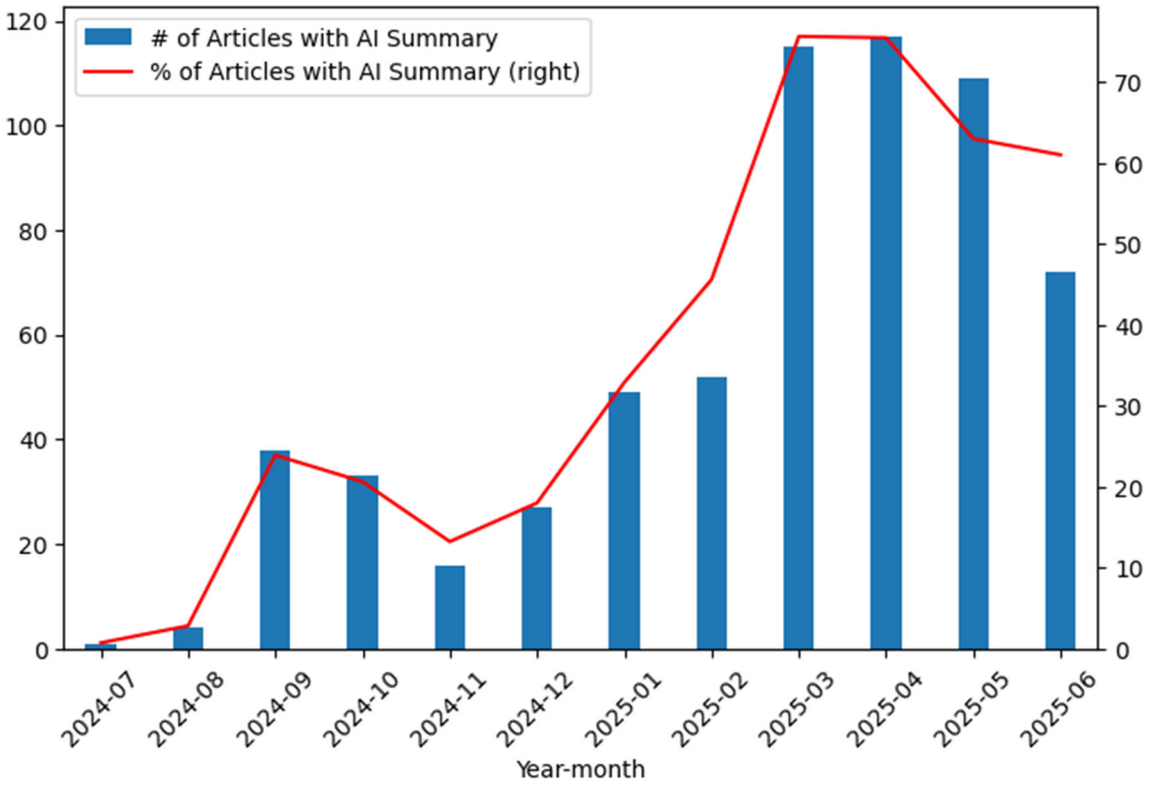

This figure depicts the monthly trend in the use of AI-generated summaries in articles published between July 2024 and June 2025. The blue bars represent the number of articles with AI summaries, measured against the left y-axis. The red line shows the percentage of total articles with AI summaries, aligned with the right y-axis.

Appendix A: Variable Definition

Variables	Definition
AI	An indicator variable that equals one if the article includes an AI summary, and zero otherwise.
$WordCount$	Natural logarithm of word counts of the paragraph or the article
$Size$	Natural logarithm of the firm's market capitalization
ROA	Income before extraordinary items divided by total assets
MTB	Market value of equity divided by book value of equity
$Leverage$	Sum of long-term debts and debt in current liabilities divided by total assets
$MediaCoverage$	Natural logarithm of one plus the number of articles published excluding Wall Street Journal on the article publication date
$MediaTone$	Mean tone of all unique published articles excluding Wall Street Journal on the article publication date. Unique articles are identified as articles with RavenPack's Event Similarity Days greater than one.
$PressRelease$	An indicator variable that equals one if there is a press release on the article publication date, and zero otherwise.
$AbnTradeVol$	Natural logarithm of one plus the total number of shares traded during the 30-minute window following the article publication minus natural logarithm of one plus the average total number of shares traded during the same day and time for the past month.
$Return$	Market-adjusted return in percentage during the 30-minute window following the article publication
$Tone_{ParaSum}$	Tone of the paragraphs that are summarized by GPT. For each GPT summary, a paragraph with the highest BERTScore is selected as having been summarized.
$Tone_{ParaNotSum}$	Tone of the paragraphs that are not summarized by GPT
$Tone_{GPTSum}$	Average tone of the three GPT summaries
$HighFOG$	An indicator variable that equals one if the article's Gunning-Fog Index exceeds the sample median, and zero otherwise.
$HighWC$	An indicator variable that equals one if the article's word count exceeds the sample median, and zero otherwise.
$HighNum$	An indicator variable that equals one if the article's proportion of numbers to word count exceeds the sample median, and zero otherwise.
$HighEA$	An indicator variable that equals one if the number of same-day earnings announcements exceeds the sample median, and zero otherwise.
$HighNews$	An indicator variable that equals one if the number of same-day same-Fama-French 12 industry WSJ articles exceeds the sample median, and zero otherwise.
$AbnSpread$	Natural logarithm of the weighted average daily spread over trading days [0, 1] divided by the weighted average daily spread over days [-41, -11]. Daily spread is average percent effective spread weighted by total number of trades during market hours.
$AbnImpact$	Natural logarithm of the weighted average daily impact over trading days [0, 1] divided by the weighted average daily impact over days [-41, -11]. Daily price impact is average percent price impact of each trade over a 5-minute window, weighted by total number of trades during market hours.
IPE	Intra-Period Efficiency measured over the five trading days following the article publication: average of $[1 - (\lvert AR_5 - AR_t \rvert)/\lvert AR_5 \rvert]$ where AR_t is the buy-and-hold market-adjusted return over $[0, t]$ (Blankespoor, deHaan, and Zhu [2018])
$Return_{[t,t+i]}$	Buy-and-hold market-adjusted returns over trading days $[t, t+i]$
$Tone$	Tone of the WSJ article

$AbnGoogleSV$	Mean daily Google Search Volume Index (SVI) of the stock on the article publication date and the following day minus mean daily SVI for five days ending one day before the publication. Daily SVI is computed following deHaan, Lawrence, and Litjens [2025].
$HighBlmbgSV$	An indicator variable that equals one if Bloomberg News Heat - Daily Max Readership on the article publication date is either 3 or 4, and zero if it is either 1 or 2.

Appendix B: One-Shot Prompt

System Prompt:
For a news article identified by accession number, you are given its article content. You need to summarize the given news article into three bullet points. Below is an example article with its summary in three bullet points: ### [an actual article from WSJ] Summary #1: [the first bullet point summarized by WSJ] Summary #2: [the second bullet point summarized by WSJ] Summary #3: [the third bullet point summarized by WSJ] ###

User Prompt:
You are given below an article with accession number {Article ID}: ### {Article Text} ### Return a JSON object with the following structure: AccessionNo: Article ID GPT_summary: [summary #1, summary #2, summary #3]

This appendix includes the one-shot prompt we used to summarize WSJ articles. We access GPT-4o-2024-08-06 via the API, with the random seed set to 1921.

Table 1: Sample Selection and Composition

Panel A: Sample Selection		
	# of firm-article observations	# of unique firms
Wall Street Journal articles from July 2024 to June 2025 for firms with tickers in CRSP/Compustat Merged database	2,877	502
Delete articles from sections that do not use AI summaries	(209)	(11)
Require TAQ and CRSP data	(632)	(128)
Delete singleton observations	(300)	(205)
Delete articles with content for which GPT cannot be used	(2)	
Final sample	1,734	158

Panel B: Sample Composition by Industry			
FF-12 Industry	# of articles	# of articles with AI summary	% of articles with AI summary
Business Equipment	297	111	37.37
Chemicals and Allied Products	6	2	33.33
Consumer Durables	146	55	37.67
Consumer Nondurables	80	34	42.50
Finance	282	84	29.79
Healthcare, Medical Equipment, and Drugs	86	26	30.23
Manufacturing	153	62	40.52
Oil, Gas, and Coal Extraction and Products	71	31	43.66
Telephone and Television Transmission	93	31	33.33
Utilities	15	8	53.33
Wholesale, Retail, and Some Services	317	114	35.96
Other	188	75	39.89
Total	1,734	633	36.51

Panel C: Sample Composition by Section			
Section	# of articles	# of articles with AI summary	% of articles with AI summary
Business	801	346	43.20
Markets	323	100	30.96
Tech	149	63	42.28
C Suite	96	36	37.50
WSJ Pro	65	1	1.54
Politics	60	24	40.00
US	59	9	15.25
Economy	13	9	69.23
World	17	9	52.94
Health	41	14	34.15
Life	34	10	29.41
Real Estate	30	7	23.33
Page One	22	3	13.64
NewsPlus	21	1	4.76
Management	3	1	33.33
Total	1,734	633	36.51

This table presents the sample selection procedure (Panel A), sample distribution by Fama-French industry (Panel B), and sample distribution by the Wall Street Journal section (Panel C).

Table 2: Determinants of AI Summary

Panel A: Variance Decomposition	
Fixed Effects	% Explained
Year-Month	29.62%
FF-12 Industry	0.63%
Firm	4.91%
Section	3.78%
Total R^2	38.94%
Adj. R^2	31.77%

Panel B: Comparison by AI Summary						
	Mean			Median		
Variables	$AI = 1$	$AI = 0$	Diff.	$AI = 1$	$AI = 0$	Diff.
$WordCount$	6.520	6.646	-0.126***	6.581	6.712	-0.131***
$Size$	18.679	18.829	-0.151	18.788	18.852	-0.064
ROA	0.083	0.081	0.002	0.051	0.058	-0.007
MTB	7.633	6.943	0.690	3.344	3.524	-0.179
$Leverage$	0.312	0.321	-0.009	0.283	0.283	-0.000
$MediaCoverage$	5.112	4.900	0.212**	5.447	5.198	0.248***
$MediaTone$	0.103	0.092	0.011	0.121	0.095	0.026
$PressRelease$	0.720	0.631	0.089***	1.000	1.000	0.000***

Panel C: OLS Regression Estimates						
	(1)	(2)	(3)	(4)	(5)	(6)
Variables			AI			
$WordCount$	-0.074***	-0.082***			-0.042**	-0.051**
	(-3.924)	(-3.952)			(-2.026)	(-2.211)
$Size$			-0.023***	-0.050	-0.021***	-0.048
			(-2.928)	(-0.556)	(-2.693)	(-0.532)
ROA			-0.127	-0.328	-0.134	-0.320
			(-1.242)	(-0.803)	(-1.310)	(-0.785)
MTB			0.0004	-0.004	0.0004	-0.004
			(0.625)	(-1.120)	(0.558)	(-1.109)
$Leverage$			-0.094	-0.162	-0.092	-0.123
			(-1.536)	(-0.259)	(-1.492)	(-0.197)
$MediaCoverage$			0.031***	0.058***	0.031***	0.055***
			(3.186)	(4.497)	(3.184)	(4.235)
$MediaTone$			0.096*	0.088	0.085	0.078
			(1.825)	(1.445)	(1.624)	(1.278)
$PressRelease$			-0.008	-0.035	-0.012	-0.038
			(-0.281)	(-1.124)	(-0.448)	(-1.231)
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	No	Yes
N	1,734	1,734	1,734	1,734	1,734	1,734
$Adj.\ R^2$	33.12%	32.40%	33.90%	33.18%	34.04%	33.38%

This table examines the determinants of AI summary usage. Panel A presents the variance decomposition of AI usage. % Explained is the variation explained by a factor incremental to other factors. Panel B provides the univariate comparisons of articles with and without AI summaries. Panel C presents the results from estimating OLS regressions on whether articles are summarized by AI, using article and firm characteristics. The *t*-statistics based on standard errors clustered by firm are reported in parentheses below the coefficients. *, **, and *** indicate two-tailed statistical significance at the 0.10, 0.05, and 0.01 levels, respectively. See Appendix A for variable definitions.

Table 3: Summary Statistics

Panel A: Descriptive Statistics

Variables	*N*	Mean	SD	25th	Median	75th
AI	1,734	0.365	0.482	0.000	0.000	1.000
$WordCount$	1,734	6.600	0.531	6.252	6.662	6.973
$Size$	1,734	18.774	1.991	17.598	18.825	20.147
ROA	1,734	0.082	0.135	0.016	0.057	0.105
MTB	1,734	7.195	17.55	1.281	3.400	9.296
$Leverage$	1,734	0.318	0.189	0.191	0.283	0.414
$MediaCoverage$	1,734	4.977	1.664	3.917	5.296	6.184
$MediaTone$	1,734	0.096	0.206	-0.016	0.106	0.241
$PressRelease$	1,734	0.663	0.478	0.000	1.000	1.000

Panel B: Covariate Balance After Propensity Score Matching

	Mean			Median		
Variables	$AI = 1$	$AI = 0$	Diff.	$AI = 1$	$AI = 0$	Diff.
$WordCount$	6.503	6.490	0.013	6.560	6.533	0.027
$Size$	18.681	18.783	-0.102	18.788	18.722	0.067
ROA	0.082	0.088	-0.006	0.050	0.058	-0.008
MTB	7.115	7.574	-0.459	3.344	3.513	-0.168
$Leverage$	0.314	0.319	-0.005	0.292	0.289	0.003
$MediaCoverage$	5.260	5.289	-0.028	5.535	5.535	0.000
$MediaTone$	0.104	0.101	0.003	0.125	0.113	0.013
$PressRelease$	0.751	0.742	0.009	1.000	1.000	0.000

Panel C: Descriptive Statistics of the Propensity Score-Matched Sample

Variables	*N*	Mean	SD	25th	Median	75th
AI	1,148	0.500	0.500	0.000	0.500	1.000
$WordCount$	1,148	6.496	0.523	6.157	6.545	6.881
$Size$	1,148	18.732	2.027	17.543	18.767	20.186
ROA	1,148	0.085	0.146	0.016	0.051	0.097
MTB	1,148	7.344	17.206	1.301	3.400	8.895
$Leverage$	1,148	0.317	0.186	0.191	0.289	0.420
$MediaCoverage$	1,148	5.275	1.525	4.317	5.535	6.337
$MediaTone$	1,148	0.102	0.203	-0.009	0.118	0.248
$PressRelease$	1,148	0.747	0.435	0.000	1.000	1.000
$AbnTradeVol$	1,148	0.243	0.860	-0.308	0.096	0.709
$Return$	1,148	-0.085	2.455	-0.811	0.026	0.777
$\lvert Return \rvert$	1,148	1.611	2.771	0.300	0.793	1.753
$Tone_{ParaSum}$	1,148	-0.095	0.352	-0.332	-0.006	0.111
$Tone_{ParaNotSum}$	1,148	-0.026	0.193	-0.145	-0.028	0.070
$Tone_{GPTSum}$	1,148	-0.146	0.516	-0.576	-0.182	0.330
$AbnSpread$	1,140	-0.018	0.279	-0.180	-0.008	0.159
$AbnImpact$	1,148	-0.229	0.674	-0.630	-0.157	0.265
IPE	894	-0.400	2.406	-0.135	0.358	0.586
$Return_{[0,1]}$	1,148	-0.001	0.052	-0.018	0.002	0.021
$Return_{[2,5]}$	1,148	0.002	0.057	-0.021	0.003	0.029

$Tone$	1,148	-0.044	0.180	-0.157	-0.044	0.057
$AbnGoogleSV$	1,148	0.112	0.189	0.000	0.060	0.186
$HighBlmbgSV$	1,148	0.555	0.497	0.000	1.000	1.000

This table reports the summary statistics. Panel A provides the descriptive statistics of the full sample. Panel B reports the covariate balance after propensity score matching. Panel C reports the descriptive statistics of the propensity score-matched sample. Each article with AI summaries is matched to an article without AI summaries on $WordCount$, $Size$, ROA, MTB, $Leverage$, $MediaCoverage$, $MediaTone$, and $PressRelease$. The nearest-neighbour method with a caliper of 0.2 and without replacement is used. Of 633 articles with AI summaries, 574 articles with AI summaries are each matched to an article without AI summaries. See Appendix A for variable definitions.

Table 4: AI Summary and Investor Trading

	(1)	(2)	(3)	(4)
Variables	$AbnTradeVol$		$\lvert Return \rvert$	
AI	0.041**	0.035**	1.711*	0.209*
	(2.509)	(2.044)	(1.711)	(1.748)
$Size$		-0.238		-0.570
		(-1.268)		(-0.967)
$WordCount$		-0.0002		-0.157
		(-0.004)		(-1.001)
$MediaCoverage$		0.328***		0.438***
		(12.773)		(5.436)
$MediaTone$		0.113		-0.502
		(0.866)		(-1.221)
$PressRelease$		0.049		-0.008
		(0.735)		(-0.037)
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	1,148	1,148	1,148	1,148
$Adj.\ R^2$	14.96%	29.97%	30.86%	33.47%

This table presents the associations between investor trading and AI summary usage. We estimate OLS regressions of the market reaction on whether articles are summarized by AI using a propensity-score-matched sample. The *t*-statistics based on standard errors clustered by firm are reported in parentheses below the coefficients. *, **, and *** indicate two-tailed statistical significance at the 0.10, 0.05, and 0.01 levels, respectively. See Appendix A for variable definitions.

Table 5: AI Summary and Investor Trading – Alternative Comparison Groups

Panel A: Covariate Balance After Entropy Balancing

	Mean			Variance		
Variables	$AI = 1$	$AI = 0$	Diff.	$AI = 1$	$AI = 0$	Diff.
$WordCount$	6.520	6.520	0.000	0.250	0.250	0.000
$Size$	18.679	18.679	0.000	4.319	4.311	0.007
ROA	0.083	0.083	-0.000	0.024	0.024	0.000
MTB	7.633	7.633	-0.000	338.258	337.671	0.586
$Leverage$	0.312	0.312	0.000	0.034	0.034	0.000
$MediaCoverage$	5.112	5.112	0.000	2.433	2.429	0.004
$MediaTone$	0.103	0.103	0.000	0.038	0.037	0.000
$PressRelease$	0.720	0.720	0.000	0.187	0.186	0.000

Panel B: Entropy Balanced Sample

	(1)	(2)	(3)	(4)
Variables	$AbnTradeVol$		$\lvert Return \rvert$	
AI	0.099*	0.102**	0.304**	0.310**
	(1.868)	(2.097)	(1.979)	(2.033)
Controls	No	Yes	No	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	1,734	1,734	1,734	1,734
$Adj.\ R^2$	17.07%	30.33%	14.31%	16.70%

Panel C: Covariate Balance After Pre-Rollout Matching

	Mean			Median		
Variables	$AI = 1$	$AI = 0$	Diff.	$AI = 1$	$AI = 0$	Diff.
$WordCount$	6.501	6.507	-0.006	6.558	6.555	0.003
$Size$	18.679	18.699	-0.020	18.788	18.660	0.128
ROA	0.083	0.077	0.005	0.050	0.044	0.006
MTB	7.148	7.814	-0.666	3.344	2.631	0.714
$Leverage$	0.314	0.330	-0.016	0.289	0.306	-0.018
$MediaCoverage$	5.269	5.306	-0.038	5.541	5.613	-0.072
$MediaTone$	0.104	0.108	-0.004	0.125	0.115	0.010
$PressRelease$	0.752	0.745	0.007	1.000	1.000	0.000

Panel D: Pre-Rollout Matched Sample

	(1)	(2)	(3)	(4)
Variables	$AbnTradeVol$		$\lvert Return \rvert$	
AI	0.093*	0.092*	0.507***	0.591***
	(1.746)	(1.770)	(3.944)	(4.409)
Controls	No	Yes	No	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	1,154	1,154	1,154	1,154
$Adj.\ R^2$	15.31%	28.07%	29.86%	32.68%

Panel E: Pre-Rollout Same-Firm Pooled Sample				
	(1)	(2)	(3)	(4)
Variables	$AbnTradeVol$		$\lvert Return \rvert$	
AI	0.063**	0.061*	0.309***	0.345***
	(2.092)	(1.941)	(3.742)	(3.841)
Controls	No	Yes	No	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	2,088	2,088	2,088	2,088
$Adj.\ R^2$	4.98%	11.26%	14.18%	15.14%

This table presents the associations between investor trading and AI summary usage using three alternative comparison groups. Panel A reports the covariate balance for the entropy-balanced sample, matched on means and variances. Panel B presents the results from OLS regressions estimated using the entropy-balanced sample. Panel C reports the covariate balance for a sample of articles with AI summaries, each matched to an article before the rollout of AI summaries. The estimated parameters of the propensity score model from Table 3, Panel B, are used to match each article with AI summaries to an article published between July 2023 and June 2024, with a caliper of 0.2 and without replacement. Panel D presents the results from estimating OLS regressions using the pre-rollout matched sample. Panel E presents results using all pre-rollout articles for firms that subsequently appear in our post-rollout AI summary sample as the baseline. The *t*-statistics based on standard errors clustered by firm are reported in parentheses below the coefficients. *, **, and *** indicate two-tailed statistical significance at the 0.10, 0.05, and 0.01 levels, respectively. See Appendix A for variable definitions.

Table 6: Investor Responses to Summarized and Non-Summarized Contents

Variables	(1)	(2)	(3)	(4)
	$AbnTradeVol$		$\|Return\|$	
$\|Tone_{ParaSum}\| \times AI$	0.173**		0.425**	
	(2.469)		(2.314)	
$\|Tone_{GPTSum}\| \times AI$		0.184**		0.401**
		(2.434)		(2.125)
$\|Tone_{ParaNotSum}\| \times AI$	0.375	0.345	0.166	0.165
	(1.351)	(1.240)	(1.393)	(1.380)
$\|Tone_{ParaSum}\|$	0.020**		0.561**	
	(2.092)		(2.180)	
$\|Tone_{GPTSum}\|$		0.020**		0.511**
		(2.411)		(1.974)
$\|Tone_{ParaNotSum}\|$	0.142*	0.134*	0.313*	0.325*
	(1.946)	(1.843)	(1.662)	(1.708)
AI	0.027**	0.030**	0.174**	0.171**
	(1.970)	(2.099)	(2.252)	(2.481)
Controls	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	1,148	1,148	1,148	1,148
$Adj.\ R^2$	30.56%	30.65%	33.58%	33.48%

This table reports the associations between market reaction and summarized and non-summarized contents. It presents results from estimating OLS regressions of the market reaction on whether articles are summarized by AI, the absolute tones of summarized and non-summarized paragraphs, and their interactions. Summarized (non-summarized) paragraphs are those that would (not) have been summarized by AI, identified using ChatGPT. The *t*-statistics based on standard errors clustered by firm are reported in parentheses below the coefficients. *, **, and *** indicate two-tailed statistical significance at the 0.10, 0.05, and 0.01 levels, respectively. See Appendix A for variable definitions.

Table 7: AI Summary and Investor Trading – Cross-Sectional Analyses

Panel A: Article Readability				
	(1)	(2)	(3)	(4)
	$AbnTradeVol$		$\lvert Return \rvert$	
Variable	$HighFOG = 1$	$HighFOG = 0$	$HighFOG = 1$	$HighFOG = 0$
AI	0.019**	-0.004	0.423**	0.027
	(2.369)	(0.464)	(2.046)	(0.240)
Controls	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	574	574	574	574
$Adj.\ R^2$	36.59%	33.94%	39.78%	38.19%
Difference b/w coefficients	0.023*		0.396*	

Panel B: Word Count				
	(1)	(2)	(3)	(4)
	$AbnTradeVol$		$\lvert Return \rvert$	
Variable	$HighWC = 1$	$HighWC = 0$	$HighWC = 1$	$HighWC = 0$
AI	0.016**	-0.004	0.348*	0.007
	(1.974)	(-0.480)	(1.718)	(0.027)
Controls	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	576	572	576	572
$Adj.\ R^2$	31.24%	36.01%	54.17%	43.88%
Difference b/w coefficients	0.012*		0.341	

Panel C: Proportion of Numbers				
	(1)	(2)	(3)	(4)
	$AbnTradeVol$		$\lvert Return \rvert$	
Variable	$HighNum = 1$	$HighNum = 0$	$HighNum = 1$	$HighNum = 0$
AI	0.041**	-0.011	0.305	-0.112
	(2.168)	(-0.604)	(1.644)	(-0.542)
Controls	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	574	574	574	574
$Adj.\ R^2$	39.11%	26.72%	35.70%	61.02%
Difference b/w coefficients	0.052**		0.417	

Panel D: Number of Earnings Announcements				
	(1)	(2)	(3)	(4)
	$AbnTradeVol$		$\|Return\|$	
Variable	$HighEA = 1$	$HighEA = 0$	$HighEA = 1$	$HighEA = 0$
AI	0.034*	-0.034*	0.346*	-0.071
	(1.833)	(-1.661)	(1.845)	(-0.367)
Controls	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	582	566	582	566
$Adj.\ R^2$	38.88%	29.50%	52.56%	24.64%
Difference b/w coefficients	0.068**		0.417	

Panel E: Number of Industry News				
	(1)	(2)	(3)	(4)
	$AbnTradeVol$		$\|Return\|$	
Variable	$HighNews = 1$	$HighNews = 0$	$HighNews = 1$	$HighNews = 0$
AI	0.031*	0.004	0.327*	-0.039
	(1.850)	(0.289)	(1.772)	(-0.110)
Controls	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	832	316	832	316
$Adj.\ R^2$	28.59%	39.43%	32.01%	59.07%
Difference b/w coefficients	0.027		0.166	

This table reports cross-sectional analyses of the association between investor trading and AI summary usage. We estimate OLS regressions of the market reaction on whether articles are summarized by AI, in subsamples partitioned by article readability, word count, numerical content, the number of same-day earnings announcements, or the number of same-day WSJ articles. The *t*-statistics based on standard errors clustered by firm are reported in parentheses below the coefficients. *, **, and *** indicate two-tailed statistical significance at the 0.10, 0.05, and 0.01 levels, respectively. See Appendix A for variable definitions.

Table 8: Market Outcomes of AI Summary

Panel A: AI Summary and Market Liquidity

Variables	(1) $AbnSpread$	(2) $AbnImpact$
AI	-0.090***	-0.095**
	(-5.115)	(-2.250)
$Size$	-0.251***	-0.409***
	(-3.894)	(-2.713)
$WordCount$	0.022	0.002
	(1.223)	(0.040)
$MediaCoverage$	0.026***	-0.010
	(2.806)	(-0.456)
$MediaTone$	-0.047	-0.346***
	(-0.973)	(-3.014)
$PressRelease$	-0.004	0.007
	(-0.161)	(0.123)
Year-Month FE	Yes	Yes
Section FE	Yes	Yes
Firm FE	Yes	Yes
N	1,140	1,148
$Adj.R^2$	39.08%	12.46%

Panel B: AI Summary and Post-Article Return Drift

Variables	(1) IPE	(2) $Return_{[0,1]}$	(3) $Return_{[2,5]}$
AI	0.185***	0.004	0.005
	(4.006)	(1.159)	(1.259)
$AI \times Tone$		0.025*	
		(1.672)	
$AI \times Return_{[0,1]}$			-0.143**
			(-2.150)
$Tone$		0.031***	
		(2.685)	
$Return_{[0,1]}$			0.150***
			(3.200)
$Size$	0.474***	-0.059***	-0.058***
	(3.156)	(-5.216)	(-4.697)
$WordCount$	-0.031	-0.001	0.002
	(-0.751)	(-0.249)	(0.530)
$MediaCoverage$	0.052**	-0.001	-0.001
	(2.306)	(-0.678)	(-0.293)
$MediaTone$	0.072	0.047***	-0.024***
	(0.678)	(5.818)	(-2.760)
$PressRelease$	-0.008	-0.008**	-0.007
	(-0.148)	(-1.990)	(-1.631)
Year-Month FE	Yes	Yes	Yes
Section FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes

N	894	1,148	1,148
$Adj.R^2$	14.81%	23.16%	12.13%

This table reports the market outcomes of AI summary usage. Panel A reports the associations between liquidity and AI summary usage. We estimate OLS regressions of bid-ask spread and price impact on whether articles are summarized by AI. Panel B reports the associations between pricing efficiency and AI summary usage. In Column (1), we estimate an OLS regression of intraperiod pricing efficiency on whether articles are summarized by AI. In Column (2), we estimate an OLS regression of two-day return on article tone, AI usage, and their interaction. In Column (3), we estimate an OLS regression of the subsequent four-day return on the initial two-day return, AI usage, and their interaction. For the test involving IPE, we exclude observations with an absolute five-day buy-and-hold return of less than 1%. The *t*-statistics based on standard errors clustered by firm are reported in parentheses below the coefficients. *, **, and *** indicate two-tailed statistical significance at the 0.10, 0.05, and 0.01 levels, respectively. See Appendix A for variable definitions.

Table 9: Effect of AI Summary on Information Search

Variables	(1)	(2)	(3)	(4)
	AbnGoogleSV		*HighBlmbgSV*	
AI	0.040***	0.034***	0.011**	0.015**
	(2.875)	(2.720)	(1.983)	(2.144)
Size		-0.021		-0.032
		(-0.507)		(-0.306)
WordCount		0.003		-0.039
		(0.314)		(-1.405)
MediaCoverage		0.077***		0.196***
		(13.795)		(13.825)
MediaTone		0.025		-0.103
		(0.889)		(-1.417)
PressRelease		0.018		0.003
		(1.252)		(0.087)
Year-Month FE	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	1,148	1,148	1,148	1,148
$Adj.R^2$	14.29%	31.73%	19.46%	35.76%

This table reports the associations between information search activity and AI summary usage. We estimate OLS regressions of information search activity on whether articles are summarized by AI. We use *AbnGoogleSV* and *HighBlmbgSV* as measures of information search activity in Columns (1)-(2) and (3)-(4), respectively. The *t*-statistics based on standard errors clustered by firm are reported in parentheses below the coefficients. *, **, and *** indicate two-tailed statistical significance at the 0.10, 0.05, and 0.01 levels, respectively. See Appendix A for variable definitions.

Table 10: Experimental Results – Correct Responses to Questions on Topics Covered in the News Article

Panel A: Correct Responses to Questions on Topics Included in AI Summary (AI Score)				
	AI Summary Absent Condition (*N*=63)	AI Summary Present Condition (*N*=61)	*t*-stat	*p*-value
Mean Correct Responses (out of three)	1.90	2.23	2.036	0.004
Std. Error	(0.11)	(0.11)		

Panel B: Correct Responses to Questions on Topics Not Included in AI Summary (Article Score)				
	AI Summary Absent Condition (*N*=63)	AI Summary Present Condition (*N*=61)	*t*-stat	*p*-value
Mean Correct Responses (out of four)	2.05	2.49	2.107	0.037
Std. Error	(0.15)	(0.15)		

Panel C: Total Correct Responses to Questions on Topics Covered in the News Article (Total Score)				
	AI Summary Absent Condition (*N*=63)	AI Summary Present Condition (*N*=61)	*t*-stat	*p*-value
Mean Correct Responses (out of seven)	3.95	4.72	2.495	0.014
Std. Error	(0.22)	(0.22)		

This table summarizes the number of questions that participants answer correctly, by condition, about the content of the news article in our experiment. All participants receive an article about the hypothetical firm, Ridgeway Retail. We manipulate whether or not the article provides an AI summary, which includes three bullet points. After reading the article, all participants answer seven multiple-choice questions based on the information in it. Three of the questions pertain to information included in the AI summary (when present), although the information needed to answer them is in the article even when not presented in the AI summary. Four additional questions relate to information contained in the article, but not in the AI summary in either condition. Panel A presents the mean number of questions answered correctly, by condition, out of the three total questions related to the information in the AI summary, as well as results of testing for differences across conditions. Panel B presents the mean number of questions answered correctly, by condition, out of the four total questions related to information not in the AI summary, as well as results of testing for differences across conditions. Panel C presents the mean number of questions answered correctly, by condition, out of the seven total questions, as well as results of testing for differences across conditions. All *p*-values are two-tailed.

Table 11: Experimental Results – Participant Ratings of their Understanding of Recent Firm Performance

	AI Summary Absent Condition (*N*=63)	AI Summary Present Condition (*N*=61)	*t*-stat	*p*-value
Mean Response	5.13	5.57	2.120	0.036
Std. Error	(0.15)	(0.15)		

This table summarizes participants' ratings, by condition, of how well they understand a firm's recent performance. All participants receive an article about the hypothetical firm, Ridgeway Retail. We manipulate whether or not the article includes an AI summary, which includes three bullet points. After reading the article, we ask participants to rate their agreement with the statement, "I feel like I understand the recent performance of Ridgeway Retail" (1 = strongly disagree, 7 = strongly agree). The table also presents results of testing for differences across conditions on this measure. The reported *p*-value is two-tailed.

Table 12: Experimental Results – Participants' Interest in Accessing More Information about the Firm

Panel A: Participant Ratings of their Interest in Learning More about the Firm as an Investment				
	AI Summary Absent Condition (*N*=63)	AI Summary Present Condition (*N*=61)	*t*-stat	*p*-value
Mean Response	5.14	5.31	0.755	0.452
Std. Error	(0.16)	(0.16)		

Panel B: Participant Choices about Additional Information to Access before Investing				
	AI Summary Absent Condition (*N*=63)	AI Summary Present Condition (*N*=61)	*t*-stat	*p*-value
Annual Report	0.73	0.67	-0.702	0.484
Quarterly Report	0.59	0.61	0.217	0.829
Conference Call Recording	0.21	0.26	0.731	0.466
Company Website	0.35	0.20	-1.916	0.058
Equity Analyst Reports	0.54	0.57	0.379	0.705
Credit Rating Agency Reports	0.35	0.26	-1.046	0.298
Industry Trade Publications	0.33	0.39	0.692	0.491
Online Discussion Forums	0.29	0.28	-0.086	0.932
Other Financial News Outlets	0.35	0.31	-0.443	0.658

This table summarizes participants' responses, by condition, to supplemental questions about their interest in accessing more information about a firm before investing. Panel A presents mean responses, by condition, where participants rate their agreement with the statement, "I am interested in learning more about Ridgeway Retail as a potential investment" (1 = strongly disagree, 7 = strongly agree). Panel A also presents results from a test of differences in this measure across conditions. Panel B reports results from a supplemental question asking them what other sources of information they would feel the need to consult before making an investment decision in Ridgeway Retail. Participants are given the choice of nine potential sources of information and indicate whether they would access each. Panel B reports, by condition, the proportion of individuals who select a given source of information, and also reports tests for differences across conditions. All *p*-values are two-tailed.