

Large Language Models, Price Discovery, and Synchronized Trading

Jiawei Wang^{*†}

September 2026

Abstract

Large language models can improve investors’ understanding of corporate disclosures, but reliance on common models can also make their errors more correlated. I study this trade-off using the public launch of ChatGPT and two pre-launch properties of corporate filings, readability and uncertainty. After the launch, order flow in hard-to-read firms becomes more one-sided relative to comparable firms, in both retail and non-retail trading. Relative to firms in the easiest-to-read decile, one-sidedness in the hardest-to-read decile rises by 6.4 percent of the sample mean. Price discovery also improves: for a firm one standard deviation harder to read, the share of the price impact of order flow that survives the following quarter rises from about 60 to about 80 percent, and earnings news enters prices faster. Non-retail order flow in hard-to-read stocks also moves more with the rest of the market. Greater disclosure uncertainty predicts less synchronized trading relative to other firms and no improvement in price discovery. To interpret these patterns, I develop a Kyle-style model in which investors can delegate reading to a common model that decodes better but introduces a shared error. Because trading on information that others also hold is less profitable, delegation is a strategic substitute, and crowding limits adoption. Greater synchronization can therefore coexist with better price discovery.

Keywords: large language models, disclosure complexity, order flow, price discovery, commonality

JEL Classification: G10, D40, O30

^{*}Warwick Business School, University of Warwick. Email: Jiawei.Wang.2@warwick.ac.uk

[†]Previously circulated under the title *Large Language Models, Price Discovery, and the Limits to Adoption*. I would like to thank two anonymous reviewers for providing helpful feedback, and Constantinos Antoniou, Matteo Benetton, Hongyi Chen, Patrick Coen, Anantha Divakaruni, Gaofeng Han, Tony Hu, Ruggero Jappelli, Yan Ji, Ali Kakhbod, Olga Klein, Roman Kozhan, John Kuong, Noah Lyman, Kebin Ma, Henry Mak, Peter Maxted, Philippe Mueller, Terrance Odean, Christine Parlour, Vikas Raman, Boyi Shen, David Skeie, David Sraer, Ruslan Sverchkov, Giorgio Valente, Ganesh Viswanath-Natraj, Johan Walden, Junxuan Wang, Pengguo Wang, Zhengge Zhou, the seminar participants at UC Berkeley, AsianFA Doctoral Consortium, Edinburgh Financial Technology Conference, Durham Conference for Finance Job Market Paper, Hong Kong Institute for Monetary and Financial Research, CES China Annual Conference, FMARC Doctoral Consortium, FMA Asia Doctoral Student Consortium, Warwick Business School Brown Bag, and the Gillmore Center Fintech PhD Workshop. I gratefully acknowledge financial support from the Hong Kong Institute for Monetary and Financial Research. This paper represents the views of the author, which are not necessarily the views of the Hong Kong Monetary Authority, Hong Kong Academy of Finance Limited, or Hong Kong Institute for Monetary and Financial Research. The above-mentioned entities except the author take no responsibility for any inaccuracies or omissions contained in the paper.

1 Introduction

Prices aggregate investors’ interpretations of information. A common technology for interpreting information can improve each investor’s reading while making the remaining errors more correlated. Large language models (LLMs) bring this trade-off to corporate disclosure. They summarize filings and translate technical language at low cost (Kim et al., 2023; Blankespoor et al., 2026). Nearly half of surveyed retail investors use generative AI, mainly to gather and interpret financial information (Blankespoor et al., 2026), and three quarters of financial institutions surveyed in Hong Kong have implemented, are piloting, or are exploring it (Hong Kong Institute for Monetary and Financial Research, 2025).

Investors have long turned to common sources to interpret financial information. Analysts and the business press interpret and disseminate filings for a broad audience (Healy and Palepu, 2001; Bushee et al., 2010; Leavy et al., 2011), and pricing models and trading strategies give their users a common lens. A language model differs from these sources in two ways. It helps investors decode the text itself, so its effect should be larger where disclosures are harder to read. And the same models serve retail and institutional investors across many firms, so a common interpretation can extend across investor classes and stocks.

I ask three questions about how a common reading technology affects markets. First, does trading within a stock become more synchronized after LLMs become widely available? Investors may trade more similarly because they understand the disclosure better or because they inherit the same mistake (Kim et al., 2025; Huang et al., 2026). Better understanding improves price discovery, a shared mistake can distort prices, and synchronization alone cannot tell the two apart. The second question is therefore whether prices become more informative. Third, does order flow become more correlated across stocks? If investors use the same models to interpret many firms’ disclosures, the common component of their readings may extend across firms.

To answer these questions, I distinguish difficulty in reading a disclosure from uncertainty about what it describes. An LLM can help investors parse difficult prose (Kim et al., 2023; Blankespoor et al., 2026), but understanding the text does not resolve uncertainty about information that the firm does not have or does not disclose, and the model cannot supply that information: language models are unreliable on facts that are rare or absent in their training data (Kandpal et al., 2023; Kalai and Vempala, 2024). I therefore measure two dimensions of disclosure difficulty drawn from the disclosure literature (Li, 2008; Loughran and McDonald, 2011). Readability, proxied by the Flesch–Kincaid grade level and the Gunning Fog index, which rise with sentence length and with the share of complex words, measures how difficult the language is to understand. Uncertainty, proxied by the

frequency of words in the Loughran–McDonald uncertainty list such as “may” and “uncertain,” measures how far the outcomes a filing describes remain unresolved once the text is understood. Both measures come from each firm’s 10-K and 10-Q filings between January 1, 2021 and November 29, 2022, the day before the public launch of ChatGPT, and are held fixed throughout the sample.¹ Using U.S. common stocks from 2021 through 2025, I compare the changes after the launch across firms whose pre-launch disclosures differ along the two dimensions.

For firms with hard-to-read disclosures, the answer to all three questions is yes. After the launch, trading in these firms becomes more one-sided relative to comparable firms, by 2.2 percent of the sample mean per standard deviation of readability difficulty, in both retail and non-retail flow, and the difference more than doubles from 2023 to 2025 as frontier models accumulate. Between the easiest and hardest deciles of pre-launch readability, the difference is 6.4 percent of the sample mean. The added synchronization carries information. For a firm one standard deviation harder to read, the share of the same-day price impact of order flow that survives the following quarter rises from about 60 to about 80 percent, and the post-announcement return spread between good- and bad-news announcements narrows by 1.4 percentage points. Non-retail order flow in these firms also moves more with the rest of the market, by 11 percent of its pre-launch comovement per standard deviation. The increase in one-sidedness is larger among retail-heavy, small, and thinly covered firms, and the gain in price discovery is larger among retail-heavy and small firms, where investors previously had less help interpreting disclosures. Where filings are uncertain, the differential runs the other way: trading becomes less one-sided relative to other firms, the surviving share of the price impact falls from about 90 to about 60 percent for a firm one standard deviation more uncertain, and earnings news is incorporated no faster.

I measure synchronization by the firm-month average of absolute daily order imbalance, which is high when daily trading is predominantly buyer- or seller-initiated, even if the direction varies across days. For the market as a whole, dollar one-sidedness falls from 13.8 percent before the launch to 12.5 percent afterward, but it falls less in firms with harder-to-read disclosures. For a firm one standard deviation harder to read, equivalent to about three additional words per sentence, the change in dollar one-sidedness is 0.29 percentage points larger than for a comparable firm, or 2.2 percent of the sample mean.² The result

¹Firms can adapt their disclosures to new readers: filings became more machine-friendly as algorithmic readership expanded (Cao et al., 2023), and by 2024 generative AI accounted for up to 4.5 percent of new disclosure text (Blankespoor et al., 2026). Firm rankings on both dimensions are nevertheless persistent, with correlations of about 0.8 between pre- and post-launch firm means.

²The raw range of readability is narrow because filings share a legal and financial vocabulary, so most cross-firm variation comes from sentence structure. Both the sentence-length and vocabulary components

holds for dollar and trade-count imbalance and for retail and non-retail flow, and retail and non-retail flow more often trade in the same direction. Signed imbalance does not change, so trading becomes relatively more one-sided without a systematic shift toward buying or selling. The convergence also appears outside the trading data: analysts' earnings forecast dispersion falls by five percent per standard deviation of readability difficulty.

The response has a threshold. With difficulty deciles in place of the continuous exposure, the two easiest readability deciles show little response, the effect emerges sharply in the third decile, and it rises with difficulty from there, as if investors turn to model-assisted reading only when the text is hard enough to make assistance worthwhile. An event study shows a differential that is flat before December 2022 and positive in every quarter afterward. It also grows over time, from 0.19 percentage points in 2023 to 0.28 in 2024 and 0.42 in 2025.

Uncertainty produces the opposite pattern. For a firm one standard deviation more uncertain, equivalent to about four additional uncertainty words per thousand, the post-launch change in dollar one-sidedness is 0.09 percentage points smaller than for a comparable firm, and analysts' forecast dispersion rises by six percent relative to comparable firms. A common reader can clarify what a disclosure says, but it cannot settle outcomes that the disclosure leaves unresolved. When dense language obscures a definite message, a model can help investors recover it. When the outcome itself is uncertain, as with a pending lawsuit or an unproven product, a model helps less and can even widen disagreement: it points investors to scenarios they had not fully considered, and investors who weigh those scenarios with their own judgment can end up further apart.

I then ask whether the synchronization carries information, with two tests of whether prices move by the right amount: a price that moves too far is later reversed, and a price that moves too little keeps drifting in the same direction. The first separates the price impact of one-sided order flow into a permanent part and a transitory part that is later reversed. For a firm one standard deviation harder to read, the part of the same-day price impact that reverses within a quarter roughly halves after the launch, while the impact itself rises. The second examines how quickly prices incorporate earnings news. Per standard deviation of readability difficulty, abnormal announcement volume rises by about 1.9 percent and the share of the price response on the announcement days by 0.8 percentage points, and the drift that follows narrows. In firms with more uncertain disclosures, more of the price impact of one-sided flow reverses, and announcement volume rises without faster incorporation or less drift, as expected when investors trade on different assessments of the same scenarios. The greater synchronization in hard-to-read firms thus reflects, at least in

of the measure predict the effect (Section [4.1.1](#)).

part, a better understanding of disclosures, and the gain in price discovery does not appear where the difficulty lies in unresolved outcomes.

The common reader can also link trading across stocks. Better reading of one firm’s filing gives trading in other firms no reason to move with it, but a model that reads many firms’ filings can. I measure commonality by regressing each firm’s daily signed order imbalance on that of all other sample firms. Before the launch, the average firm’s imbalance moves by 0.83 percentage points for each percentage point in the rest of the market’s. After the launch, this response rises in hard-to-read firms by 0.105 for total flow and 0.091 for non-retail flow relative to other firms, about 13 and 11 percent of the pre-launch average. The cross-stock effect is present in 2023 and 2024 but fades in 2025, while within-stock synchronization keeps strengthening. Errors are correlated across language models, and more so among models that share an architecture or provider (Kim et al., 2025), but a shared error links trading across firms only if it also carries from one firm’s disclosure to another, as a model family’s reading habits might. As investors spread across model families, this cross-firm link can weaken even as better readings continue to strengthen synchronization within stocks.

To interpret the results, I develop a model in the spirit of Kyle (1985). Each speculator forms a signal by reading the disclosure, either independently or through a common model. An independent reading contains decoding error, which a better reader can reduce, and irreducible uncertainty about outcomes the disclosure leaves unresolved. Delegation reduces decoding error but makes part of the remaining error common across model users. Adoption begins when an investor’s own decoding error exceeds the model’s, but as more investors delegate, their information overlaps and the private value of joining them falls. Delegation is therefore a strategic substitute, and harder disclosures attract more model use because the precision gain is larger.

Endogenous adoption limits the cost of shared errors. If all investors had to use the model, the common error could outweigh the gain in precision and make prices less informative. When adoption is a choice, independent readers remain once delegation becomes crowded, and adoption expands only as far as the precision gain justifies. The equilibrium therefore combines more synchronized trading with better price discovery, and shared errors harm price discovery only when adoption is forced. Across firms, greater decoding difficulty raises delegation and synchronization, while uncertainty limits the benefits of better reading and reduces the share of reading error common across model users.

This paper contributes to the literature on generative AI in financial markets. Existing work shows that LLMs help investors process financial information, affect trading and market quality, and predict market outcomes (Bertomeu et al., 2025; Cheng et al., 2025; Stanco

and Chung, 2025; Blankespoor et al., 2026; Lopez-Lira and Tang, 2026). I study a different margin: how reliance on a common reading technology changes the correlation structure of investors’ information. Prices respond to that structure, because idiosyncratic errors diversify away in aggregate order flow and common errors do not. By examining synchronization together with price discovery and cross-stock commonality, I provide evidence on both the gains from better reading and the shared component of investors’ interpretations.

A second contribution is to show that the effects of LLMs depend on what makes a disclosure difficult to process. Research on disclosure complexity measures processing costs through features such as the Fog index used by Li (2008), the file-size measure of Loughran and McDonald (2014), and the bloat that Kim et al. (2023) use ChatGPT to remove. I separate difficulty in decoding language, which a better reader can reduce, from uncertainty about underlying outcomes, which remains once the text is understood. After the launch, harder-to-read disclosures are associated with more synchronization and better price discovery, and more uncertain disclosures with less synchronization and no gain in price discovery. The benefits of a common reading technology therefore depend on whether the obstacle lies in understanding the text or in what the text leaves unresolved.

The paper also connects the literature on information aggregation with work on commonality in trading. Prices aggregate dispersed information, and the quality of that aggregation depends on who holds the information and how correlated it is (Grossman and Stiglitz, 1980; Stein, 2009; Goldstein et al., 2025); common components in order flow and liquidity are well documented (Chordia et al., 2000; Hasbrouck and Seppi, 2001). LLMs add a new channel for such commonality, because investors can use the same technology to interpret the same disclosure and the disclosures of many firms.

Section 2 discusses related work, Section 3 describes the data and the two disclosure dimensions, Section 4 presents the evidence, Section 5 develops the model, and Section 6 concludes.

2 Related literature

The first related literature studies how generative AI changes the production and processing of financial information. LLMs can analyze disclosures and other financial text at low cost (Kim et al., 2023; Shaffer and Wang, 2024; Blankespoor et al., 2026), and the applications now extend to earnings calls and corporate policies (Beckmann et al., 2024; Jha et al., 2024), central-bank language (Hansen and Kazinnik, 2024), investor chatter on social media (Huang et al., 2024; Chen et al., 2025), corporate culture (Li et al., 2026), and forecasts of

returns and the macroeconomy (Lopez-Lira and Tang, 2026; Chen et al., 2025). Evidence from the Italian ban of ChatGPT studies the role of access to the technology in investor information processing and in the value of exposed firms (Bertomeu et al., 2023, 2025), and account-level evidence from the same ban shows that losing access changes the assets retail investors choose to trade (Even-Tov et al., 2025). Outages at ChatGPT and other LLM providers supply short disruptions to AI-assisted trading: trading volume falls while access is unavailable (Cheng et al., 2025), more post-news drift persists into the following day (Lu et al., 2026), and retail participation in earnings-news trading falls (Campbell et al., 2026). Other work evaluates LLM forecasts, investment advice and asset management (Lopez-Lira and Tang, 2026; Lu et al., 2023; Sheng et al., 2026), and studies whether broader access to AI narrows the gap between retail and institutional investors (Chang et al., 2026). Stanco and Chung (2025) use the same ChatGPT launch and find that trading volume rises, and liquidity, price efficiency, and analyst accuracy improve, more for firms with long 10-K filings and high investor attention. Those results concern the amount of trading and market quality. I examine the direction of trading: whether it becomes more one-sided, whether this convergence accompanies better price discovery, and whether it extends across stocks. These tests help separate better understanding, which can synchronize trading within a stock, from shared reading errors, which can also link trading across firms.

A second strand of research examines the risks that arise when market participants rely on AI systems. LLMs can produce hallucinations and misinformation in financial tasks (Kang and Liu, 2023; Xu et al., 2025), and recent work studies correlated AI errors, AI-generated bubbles, opinion ecosystems, and strategic interaction among algorithmic traders (Dou et al., 2025; Cartea et al., 2026; Hirshleifer et al., 2026). Evidence on model errors supports the premise that a common reader can leave shared mistakes. Huang et al. (2026) show that models concentrate on a narrow set of analytical paths and produce substantially less dispersion than human teams. In an evaluation of more than three hundred models, Kim et al. (2025) find substantial correlation in errors: on one leaderboard dataset, two models agree about sixty percent of the time when both are wrong, and larger and more accurate models have highly correlated errors even across architectures and providers. These concerns feature in policy discussions: the International Monetary Fund examines the financial-stability implications of widespread AI use (International Monetary Fund, 2024), and the Hong Kong Institute for Monetary and Financial Research highlights model accuracy and reliance on external providers in institutional adoption (Hong Kong Institute for Monetary and Financial Research, 2025). My model adds investors' choice to use the technology. Because the private return to delegation falls as adoption spreads, endogenous adoption limits crowding, and shared errors need not overturn the price-discovery benefits

of more accurate interpretation.

The mechanism also relates to the literature on information aggregation. In [Kyle \(1985\)](#), informed orders reveal information to a competitive market maker, and in [Grossman and Stiglitz \(1980\)](#), information acquisition and price informativeness are jointly determined. [Stein \(2009\)](#) emphasizes that adding sophisticated investors need not improve efficiency when their information or strategies overlap. I add a choice between reading a disclosure independently and relying on a common reading technology, in which investors weigh the gain in precision against the cost of trading on information shared with other users. Through this choice, disclosure difficulty and model quality jointly determine the precision and overlap of the information investors trade on.

The paper also draws on research linking disclosure attributes to trading on public news. My readability measure follows [Li \(2008\)](#), and the uncertainty measure uses the word list of [Loughran and McDonald \(2011\)](#); filing length, which [Loughran and McDonald \(2014\)](#) propose in place of the Fog index, enters every specification as a control. [Kim et al. \(2023\)](#) show that LLM summaries remove substantial bloat from long filings, and [Cao et al. \(2023\)](#) show that firms adapt their disclosures to machine readers, which is why I measure disclosure exposures before the launch. Earlier changes in disclosure technology worked through access: the move to electronic filing lowered the cost of obtaining a filing and raised information production by corporate outsiders ([Gao and Huang, 2020](#)). The launch studied here lowers the cost of interpreting one. The finding that more uncertain disclosures generate additional announcement trading without faster price discovery echoes [Kandel and Pearson \(1995\)](#), in which different interpretations of the same public signal increase trading volume without necessarily reflecting superior information.

Finally, the paper relates to evidence on common factors in trading and prices. [Chordia et al. \(2000\)](#) document commonality in liquidity, and [Hasbrouck and Seppi \(2001\)](#) identify common factors in order flow, prices, and liquidity. Related work uses return comovement to study the incorporation of firm-specific information into prices ([Morck et al., 2000](#)). I measure trading commonality with a leave-one-out order-flow beta, the response of a firm’s signed order imbalance to that of all other sample firms, and ask whether it strengthens after LLMs become widely available, particularly among firms whose pre-launch disclosures are harder to read.

3 Data and measurement

The sample contains U.S. common stocks (CRSP share codes 10 and 11) from January 2021 through December 2025. I require a lagged price of at least five dollars, which excludes microcap stocks whose imbalance measures are dominated by microstructure noise, and at least fifteen trading days in a firm-month so that monthly averages are not driven by a small number of trading days.

Two persistent dimensions of disclosure difficulty. Readability is the standardized sum of the Flesch–Kincaid grade level and the Gunning Fog index, computed from the sentence-level parse of the WRDS SEC Analytics Suite. Both indices combine sentence length with syllable-based measures of word complexity, and higher values indicate greater difficulty.³ The premise is that these features slow a human reader more than a language model, which can restate long sentences and technical terms in plain language. Uncertainty is the frequency of words in the Loughran–McDonald uncertainty list, obtained from the Loughran–McDonald 10-X Summaries. It counts language that presents outcomes as contingent or imprecise, and I use it as a linguistic proxy for how far the firm’s prospects remain unresolved.⁴

Each measure is the firm’s average across its 10-K and 10-Q filings from January 1, 2021 through November 29, 2022, standardized across firms. I hold these exposures fixed so that they capture disclosure characteristics before the launch and exclude later changes in disclosure practice, which may respond to the technology itself: firms adapt their disclosures to anticipated machine readership (Cao et al., 2023), and generative AI is increasingly used in disclosure production (Blankespoor et al., 2026). Firms’ relative disclosure characteristics are persistent. On consistently constructed measures of vocabulary complexity and uncertainty, pre- and post-launch firm means correlate at about 0.8, and adjacent-year rank correlations range from 0.76 to 0.89 (Appendix A).⁵

Other variables. CRSP provides returns, prices, shares outstanding, and market equity. Earnings-announcement dates come from IBES and Compustat, and factor returns

³Flesch–Kincaid grade = $0.39 \times (\text{words per sentence}) + 11.8 \times (\text{syllables per word}) - 15.59$; Gunning Fog = $0.4 \times [(\text{words per sentence}) + 100 \times (\text{complex-word share})]$, where complex words have three or more syllables. The composite correlates at 0.91 with words per sentence and 0.43 with the complex-word share.

⁴Disclosure length, defined below, enters every specification as a control; Appendix A treats it as a third dimension.

⁵These comparisons use the Loughran–McDonald 10-X Summaries, whose word counts are constructed consistently throughout the sample; the WRDS sentence parser is revised in 2022 and 2023, so its scores are not compared across years (Appendix A).

come from the [Fama and French \(2015\)](#) five-factor series. Every specification allows the post-launch change to vary with filing length and nine firm characteristics. Length is the standardized sum of log net file size ([Loughran and McDonald, 2014](#)) and log word count, both from the Loughran–McDonald 10-X Summaries and averaged over the same filings as the two disclosure dimensions. The nine characteristics are computed from daily data over the characteristic window, January 1, 2022 through November 29, 2022, and standardized across firms. Size is log market equity. Idiosyncratic volatility is the average daily volatility of residual returns from the factor model, and effective spread is the average daily trade-based spread from TAQ. Momentum is the firm’s cumulative return over the characteristic window. Turnover is average daily dollar volume divided by average market equity, and retail share is retail dollar volume, identified with the [Boehmer et al. \(2021\)](#) algorithm, divided by total dollar volume. The remaining controls are price, market beta, and AI-sector beta, which captures a stock’s exposure to the software and semiconductor industries beyond its exposure to the market: I construct the value-weighted return of sample firms in these industries, remove its market component, and estimate each firm’s beta with respect to the residual.⁶

Some tests use a cumulative stock of frontier language-model releases, denoted $N(m)$. A release enters the calendar if it is a usable, general-purpose flagship model from a frontier developer and its date is confirmed against Epoch AI’s Notable AI Models database ([Epoch AI, 2026](#)). Table 1 lists the twenty-three dates: seven from OpenAI, six from Anthropic, five from Google DeepMind, two each from Meta and xAI, and one from DeepSeek.

The sample. Table 2 describes 151,022 firm-months for 3,876 firms.⁷ Continuous variables are winsorized at the 1st and 99th percentiles. The average firm-month has an absolute dollar imbalance of 13.1 percentage points. On an average day, the larger side accounts for about 56.6 percent of classified dollar volume. Trade-count imbalance is lower, at 10.1 percentage points, because a small number of large trades can affect dollar imbalance more than trade-count imbalance. Retail imbalance is much higher, at 24.4 percentage points. Retail trading in an individual stock is thinner, so buy and sell orders average out less

⁶The launch coincided with a repricing of software and semiconductor stocks, so firms that comove with these industries could attract more one-sided trading afterward for reasons unrelated to their disclosures. The sector portfolio contains firms in SIC 7372 and 3674. Replacing it with the 20 largest firms in these industries, a familiar mega-cap portfolio, or omitting the AI-sector beta leaves the signs and statistical conclusions unchanged.

⁷The firm count differs across tests with the variables each requires. Table 2 imposes only the price and trading-day filters and the three text measures, readability, uncertainty, and length (3,876 firms). The order-flow regressions also require the nine firm characteristics and the activity controls (3,670), and the daily flow-return panel requires only the price and volume filters (4,034). On the 3,669 firms present in both estimation panels, the estimates are nearly identical (Table A1).

within a month.

Each standard deviation of the two standardized measures has a natural interpretation in the underlying units. For readability, one standard deviation corresponds to 1.7 Flesch–Kincaid grade points, about three additional words per average sentence, and 0.4 percentage points more complex words on a mean share of 28.7 percent.⁸ A firm at the 10th percentile of readability writes sentences averaging about 32 words, of which 28.1 percent are complex, compared with 41 words and 29.1 percent at the 90th percentile; the complex-word share itself has a standard deviation of 0.9 percentage points across firms, about one complex word per hundred. For the uncertainty measure, one standard deviation corresponds to a 0.41 percentage-point increase in the share of uncertainty words, relative to a mean of 1.38 percent. This is about four additional uncertainty words, such as “may” or “uncertain,” per thousand words, or roughly 120 additional uncertainty words in a median-length filing. Table 3 describes the firms at each tercile of the two disclosure dimensions. Firms with harder-to-read disclosures are larger, more liquid, and more widely covered by analysts, with less retail participation. Firms with more uncertain disclosures tend to be smaller and more volatile, with higher market betas, greater retail participation, and wider analyst coverage. The uncertainty measure is also positively correlated with pre-launch option-implied volatility and with the volatility of return on assets, with rank correlations of 0.27 and 0.19 that change little when size is removed, whereas readability is negatively correlated with both (Table A6). The two disclosure measures have a cross-sectional correlation of -0.18 .

4 Empirical results

The empirical analysis follows the three questions. If many investors read the same disclosure through the same model, their conclusions should become more similar and trading within a stock more one-sided; Section 4.1 examines this. Similar trading alone does not show whether the shared interpretation is useful, since investors may converge because they read the disclosure better or because they share the same error, so Section 4.2 asks whether prices become more informative. Section 4.3 then tests whether order flow becomes more correlated across firms, as it would if a general-purpose model applied similar reading patterns to many firms.

⁸All filings score well above conventional schooling levels, so I use the grade measure as an index of syntactic difficulty rather than reading the grade level literally.

4.1 Synchronized trading

Order flow. I sign trades in the NYSE Daily TAQ millisecond files using the [Lee and Ready \(1991\)](#) algorithm. Retail trades follow [Boehmer et al. \(2021\)](#), and non-retail flow is total flow minus retail flow. Daily signed order imbalance is

$$OIB_{id} = \frac{Buy_{id} - Sell_{id}}{Buy_{id} + Sell_{id}}.$$

I construct the measure using dollar volume and, separately, the number of trades. Dollar imbalance places more weight on large trades, while trade-count imbalance captures how broadly trades line up on one side. Order imbalance is widely used in market microstructure because signed imbalances contain information about individual stock returns ([Chordia and Subrahmanyam, 2004](#)). Here, the main dependent variable is the firm-month mean of $|OIB_{id}|$. A high value indicates that daily trading repeatedly falls on one side of the market, whether toward buying or selling.

4.1.1 Main specification

The baseline regression enters the two disclosure dimensions jointly:

$$y_{im} = \alpha_i + \gamma_m + \sum_{d \in \{R, U\}} \beta_d (Post_m \times D_{di}) + \psi' (Post_m \times Z_i) + \theta' Act_{im} + e_{im}. \quad (1)$$

Here y_{im} is a synchronization measure for firm i in month m ; D_{di} is firm i 's pre-launch score on dimension d , readability (R) or uncertainty (U), standardized across firms and held fixed over the sample; and $Post_m$ equals one from December 2022 onward. Firm fixed effects absorb differences across stocks, and month fixed effects absorb market-wide changes, including the average post-launch change in y_{im} . The coefficient β_d is therefore a difference across firms: how much larger the post-launch change is for a firm one standard deviation higher on dimension d than for an otherwise comparable firm at the mean, holding the other dimension fixed. For the market as a whole, dollar one-sidedness fell after the launch, from 13.8 to 12.5 percentage points, so a positive readability coefficient means that hard-to-read firms fell by less than comparable easier-to-read firms (Table A2).⁹ The vector Z_i contains ten predetermined controls, filing length and the nine firm characteristics, each interacted with $Post_m$, so that the estimate compares firms that differ in disclosure difficulty but were comparable in size, liquidity, volatility, retail participation, and sector exposure before

⁹Throughout the paper, statements such as “trading becomes more one-sided in hard-to-read firms” are shorthand for this comparison and do not describe the level of the change in any group of firms.

the launch. The vector Act_{im} contains two contemporaneous activity controls, log dollar volume and the log number of intraday intervals with trading activity. Absolute imbalance falls mechanically as trading intensifies, because buy and sell trades offset more completely when more trades arrive, and these controls separate changes in the direction of trading from changes in its amount. Reading β_d as an effect of language models requires that, given these controls, nothing else that changed around the launch moved trading differently in firms with different disclosures. The event study and the placebo launch dates below check the pre-launch period for such differences.

Table 4 reports the main synchronization results, with both dimensions entered jointly in the first row of Panel A. A one-standard-deviation increase in readability difficulty, about three additional words per average sentence, is associated with a post-launch change in dollar one-sidedness that is 0.29 percentage points larger, or 2.2 percent of the sample mean of 13.1 percentage points.¹⁰ A one-standard-deviation increase in uncertainty is associated with a change that is 0.09 percentage points smaller, or 0.7 percent of the sample mean.¹¹

If investors read through a common model, the two coefficients have a natural interpretation. A common reader brings investors to the same conclusion only about what the text settles. When a filing has a definite message that is merely hard to extract, investors who read it through the same model extract close to the same message, and their trades line up. When a filing says the outcome depends on a pending lawsuit, a regulator, or an unproven product, the model reads it just as accurately but cannot settle the outcome, and the negative uncertainty coefficient has two readings. In the first, the model simply helps less in these firms: because the coefficient compares firms, trading in uncertain firms need not become less aligned, only less aligned than in firms whose message the model can recover. In the second, the model widens disagreement: it points investors to scenarios they had not fully considered, and investors who weigh the same scenarios with their own

¹⁰Table A3 separates the two components of the measure, sentence length and vocabulary. Entered together with uncertainty and the ten controls, words per sentence carries 0.11 percentage points per standard deviation, three words per sentence, and the complex-word share 0.34 per standard deviation, one more complex word per hundred; the uncertainty coefficient is -0.16 . Sentence length accounts for most of the cross-firm variation in the composite, and the vocabulary component carries the larger coefficient per standard deviation, so the effect is present for both kinds of difficulty, long sentences and technical vocabulary.

¹¹Appendix A shows similar results when readability is constructed from within-quarter ranks, an earlier measurement window, or the independent Loughran–McDonald parse. The result also does not come from firms that changed their writing after the launch: excluding the ten percent of firms that move by at least two quintiles of the complex-word share between the pre- and post-launch periods leaves the readability coefficient at 0.26. The results still hold when the pre-launch period begins in July 2021, after the episode of concentrated retail trading in early 2021, or when the firms whose retail participation rose most in that episode are dropped (Appendix A). In unreported results, they also hold after dropping the activity controls or after controlling for each firm’s current complex-word share and its post-launch interaction.

judgment end up further apart than before.

The readability coefficient is 0.29 for dollar imbalance and 0.15 for trade-count imbalance. The larger dollar estimate means that large trades line up on one side more than small trades do, but the positive trade-count estimate shows that the shift involves many trades rather than a few large ones. The dollar coefficients are similar for non-retail and retail flow, at 0.26 and 0.24, so the shift appears in both groups. The uncertainty gradient again runs the other way but is statistically significant only for retail flow. Non-retail investors can supplement the public text with other sources, such as analyst access, management interaction, and in-house research, so uncertain language may matter less for their trading, while retail investors remain more exposed to unresolved outcomes. Signed imbalance, in the last row of Panel A, changes significantly with neither dimension, so trading becomes more one-sided without an average shift toward buying or selling.

Greater one-sidedness could also come from a few informed traders trading more aggressively on more precise signals, without broader convergence among investors. One group would then trade against the other, as when institutions execute large orders against retail liquidity or non-retail investors absorb one-sided retail flow, and retail and non-retail flow would more often have opposite signs. Panel B tests this with the share of trading days on which the two have the same sign as the dependent variable. After the launch, this share rises by 0.26 percentage points per standard deviation of readability difficulty, while the uncertainty coefficient is small and insignificant. In hard-to-read firms, retail and non-retail flow more often have the same net direction on a given day, so the convergence extends across investor groups.

Figure 1 plots the event-study version of equation (1), with event time in three-month bins around the launch. Bin 0 covers December 2022 through February 2023, event time is capped at six bins on each side, with the endpoint bins absorbing the tails of the sample, and bin -6 is the omitted reference. Before the launch, the readability coefficients are flat: the five estimated pre-launch bins are close to zero, and a joint test cannot reject a flat pre-trend, with $p = 0.195$, against $p = 0.781$ for uncertainty. After the launch, all seven readability bins are positive, and four are individually significant at the five percent level. The uncertainty coefficients turn negative toward the end of the sample, but none is statistically significant.

Placebo launch dates produce no comparable differential. I assign a false launch to each of the seventeen eligible pre-launch months from April 2021 to August 2022 and estimate equation (1), with both dimensions and the ten control interactions, on pre-launch data only. None of the thirty-four placebo coefficients for dollar one-sidedness is significant at the five percent level; the placebo $|t|$ statistics have a median of 0.38 and a maximum of 1.61

for readability, and 0.36 and 1.25 for uncertainty. For trade-count imbalance, readability produces two rejections across the seventeen dates and uncertainty none, close to what chance would produce at a five percent significance level.

The actual launch also stands out among all feasible break dates in the full panel. When treatment is assigned in turn to each of the forty-eight months with at least six months of data on each side, December 2022 produces the largest $|t|$ statistic for dollar one-sidedness and the second largest for trade-count imbalance, randomization p -values of 0.021 and 0.042.¹²

4.1.2 The shape of the relation across the difficulty distribution

The linear specification imposes a constant effect per standard deviation. Table 5 relaxes it by replacing each dimension with indicators for its deciles. The first decile is omitted, so each coefficient measures the post-launch change relative to the firms with the lowest difficulty on that dimension.

The readability response is flat at the easy end: the second decile is indistinguishable from the first, at +0.12 percentage points, and the first clear increase appears in the third, at +0.54, about two-thirds of the tenth-decile estimate. Above this threshold the response rises broadly with difficulty, as if a better reader starts to matter at the threshold and matters more the harder the text. In the hardest decile, the post-launch change in one-sidedness is 0.84 percentage points larger than in the easiest, or 6.4 percent of the sample mean. The uncertainty coefficients instead decline gradually across the distribution, with no flat region or discrete jump, a further sign that the two measures capture distinct aspects of disclosure difficulty.

4.1.3 The effect over time

The differential grows over the three years after the launch. By calendar year, it rises from 0.19 percentage points in 2023 to 0.28 in 2024 and 0.42 in 2025 and is statistically significant in each year (Table 6, Panel A). At a quarterly frequency, it turns positive around the launch, builds through 2023, remains elevated with some quarter-to-quarter variation, and rises again in 2025 (Figure 2). To separate the initial break from the growth that follows, I estimate

$$y_{im} = \alpha_i + \gamma_m + \beta_1(Post_m \times C_i) + \beta_2(N(m) \times C_i) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}, \quad (2)$$

¹²Nearby treatment dates give similar estimates: dating treatment at January 2023 gives a readability coefficient of 0.28 percentage points, and dating it at the GPT-4 release gives 0.26.

where C_i is a disclosure dimension and $N(m)$ is the cumulative count of the frontier-model releases in Table 1, scaled to one in the final month. The launch term captures a discrete break in December 2022, and the stock term captures the part of the differential that grows over the following three years.

In Panel B of Table 6, the stock term is positive for both imbalance measures and accounts for most of the growth. Entered alone, it has a readability coefficient of 0.49 percentage points for dollar imbalance and 0.45 for trade-count imbalance. Entered together with the launch step, the stock coefficients remain 0.36 and 0.49, while the launch coefficients fall to 0.14 for dollar imbalance and -0.04 for trade counts. Roughly a third of the dollar effect appears at the launch and none of the trade-count effect, consistent with some large investors using the technology soon after the launch.¹³ The gradual growth fits an effect that follows improvements in model capability and their spread across investors, which can take several quarters to show up in trading.

4.1.4 Convergence in analysts' forecasts

The same convergence appears in a measure that involves no trades. Analysts' earnings forecasts outstanding before an announcement are formed after the firm's most recent 10-K or 10-Q and reflect analysts' interpretations of existing disclosures alongside other information, so their dispersion measures disagreement directly. If a common reading technology helps analysts interpret difficult text, dispersion should fall more for firms with harder-to-read disclosures. Table 7 relates the standard deviation of IBES forecasts, scaled by price, to the same post-launch interactions at the announcement level, with firm and calendar-quarter fixed effects. For the market as a whole, mean log dispersion falls from -2.18 before the launch to -2.28 afterward, a decline of about nine percent (Table A2). Relative to comparable firms, log dispersion falls by 5.1 percent per standard deviation of readability difficulty and rises by 6.0 percent per standard deviation of uncertainty, and scaling dispersion by the absolute median forecast gives -6.4 percent for readability.

4.2 Price informativeness

Greater agreement can come from better interpretation of the disclosure or from a shared error, and the two have different implications for prices. Better interpretation should reduce

¹³The results are robust to alternative release calendars. Calendars based on Epoch AI training compute include either sixteen releases in the top decile of compute between December 2022 and December 2025 or thirteen releases restricted to language and multimodal models. These calendars yield readability coefficients between 0.34 and 0.45, all significant at the 1 percent level. Weighting releases by log compute or compute rank also yields positive coefficients significant at the 1 percent level.

transitory pricing errors and speed up the incorporation of news, while a shared error can push prices away from fundamentals and be reversed later. The first test separates the price impact of order flow into a permanent part and a transitory part that is later reversed, and the second examines how quickly earnings news enters prices.

4.2.1 Price impact and its reversal

The first test asks how much of the price impact of a unit of order flow is reversed or extended over the following months. For each firm-day, I estimate

$$CAR_{i,[a,b]} = \alpha_i + \gamma_t + OIB_{it}(\kappa_1 + \kappa_2 Post_t) + \sum_d OIB_{it}(\kappa_{3d} D_{di} + \kappa_{4d} Post_t \times D_{di}) + \phi' W_{it} + e_{it}, \quad (3)$$

where $CAR_{i,[a,b]}$ is the cumulative market-adjusted log return, in percent, over trading days $[t + a, t + b]$, for five horizons: the same day $[0]$, the next day $[+1]$, the following week $[+1, +5]$, the following month $[+2, +21]$, and the following quarter $[+2, +60]$. OIB_{it} is the signed dollar imbalance defined in Section 4.1, so its coefficient measures the return associated with moving from balanced to fully one-sided flow. The vector W_{it} contains the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and their interactions with filing length. The coefficient κ_{3d} measures how the relation between signed flow and returns differed with each disclosure dimension before the launch, and the coefficient of interest, κ_{4d} , measures how its change after the launch varies with that dimension.

Table 8 reports, for each return window, the slope of returns on signed flow before the launch and its change after the launch, for a firm at the mean and per standard deviation of each disclosure dimension. For a firm one standard deviation above the mean of a dimension, the slope before the launch is $\kappa_1 + \kappa_{3d}$, and the launch adds $\kappa_2 + \kappa_{4d}$ to it. Before the launch, moving from balanced to fully one-sided flow is associated with a same-day market-adjusted return of about 2.3 percent, and after the launch this slope rises by 0.4 percentage points for the average firm. The rise is 0.18 percentage points smaller for a firm one standard deviation harder to read, about eight percent of the pre-launch impact.

The result of interest is at longer horizons. Readability does not significantly change the flow-return relation over the next day or week. Over the following quarter, the slope is negative for the average firm before the launch, at -0.71 percent per unit of imbalance, so part of the price impact of one-sided flow reverses. Reversal helps separate temporary price pressure from information: informed and uninformed orders both move prices on impact,

but temporary pressure fades while the information remains in the price. The temporary part may reflect compensation for holding inventory.¹⁴ For the average firm before the launch, a unit of imbalance moves the price by 2.3 percent on the day, of which 0.71 percentage points, about a third, reverses over the following quarter. After the launch, the average quarterly slope does not change significantly, but reversal weakens in hard-to-read firms relative to other firms: the $[+2, +60]$ coefficient κ_{4R} is +0.57 percentage points per unit of imbalance. For a firm one standard deviation harder to read, the quarterly slope before the launch is the average slope of -0.71 plus the pre-launch readability differential of -0.27 , or -0.98 . After the launch, it adds the average change of -0.11 and the readability differential of $+0.57$ and becomes -0.52 . Prices still move back against the flow, but the reversal is roughly halved while the same-day impact rises from 2.5 to 2.8 percent, so the share of the same-day impact that survives the quarter rises from about 60 to about 80 percent.

The uncertainty gradient runs the other way. Per standard deviation of uncertainty, the post-launch slope changes by -0.50 percentage points over the following month and -0.99 over the quarter, and for a firm one standard deviation more uncertain, the quarterly slope falls from -0.32 before the launch to -1.42 afterward, so the share of its same-day impact that survives the quarter falls from about 90 to about 60 percent. Relative to other firms, prices overshoot less after one-sided flow in hard-to-read firms and more in firms with uncertain disclosures.

The reversal could also shrink without more informative trading if liquidity provision became cheaper in hard-to-read firms, lowering the premium liquidity suppliers charge for absorbing one-sided flow. The spread evidence provides little support for this explanation. The five-minute realized spread, the portion of the effective spread that remains after the price movement over the next five minutes, is a proxy for liquidity suppliers' gross compensation, and in unreported results neither realized nor effective spreads change significantly in hard-to-read firms relative to other firms after the launch. Nor would a general decline in liquidity-provision costs explain why reversal weakens with readability difficulty but strengthens with uncertainty within the same sample of firms, months, and trading venues.

Figure 3 shows the pattern in raw price paths. It plots average cumulative market-adjusted returns after firm-days in the top decile of absolute dollar imbalance, when the

¹⁴Someone has to take the other side of a day of one-sided buying, and the liquidity suppliers who do so end the day with a position they did not want. They require a price above the level justified by the information in the flow to bear this exposure. As buying pressure subsides and inventories return to desired levels, the premium fades and the price falls. In [Grossman and Miller \(1988\)](#) and [Hendershott and Menkveld \(2014\)](#), this reversal is the compensation liquidity suppliers earn for accommodating the imbalance.

larger side accounts for at least 65 percent of classified dollar volume, signed in the direction of the flow, for the easiest and hardest readability terciles before and after the launch.¹⁵ Before the launch, a flow day in a hard-to-read firm moved the price by 65 basis points, and sixty days later the cumulative return was -17 basis points. After the launch, the initial move is similar, at 67 basis points, but 25 basis points remain after sixty days. The two paths track each other for the first two weeks and then diverge, and quarterly reversal falls from 82 to 42 basis points. The easy tercile shows continuation rather than reversal before the launch, with the cumulative return rising from 70 basis points on the flow day to 111 after sixty days, and afterward it is almost unchanged, at 90 and 91 basis points. Both paths are flatter after the launch, so in both groups the price on the flow day is closer to its level a quarter later.

4.2.2 Earnings announcements

The second test uses earnings announcements, a dated public signal. For abnormal volume and incorporation speed, I estimate

$$y_{ie} = \alpha_i + \gamma_{q(e)} + \sum_{d \in \{R, U\}} \beta_d (Post_e \times D_{di}) + \varphi Rank_e + \theta' X_{ie} + u_{ie}. \quad (4)$$

Here $Rank_e$ is the within-quarter percentile rank of the announcement return, scaled to $[-0.5, 0.5]$, and X_{ie} contains lagged size, volatility, and volume. Abnormal volume is mean log dollar volume over days $[-1, +1]$ minus its mean over days $[-30, -6]$. Incorporation speed is the absolute cumulative return over days $[-1, +1]$ divided by the sum of the absolute cumulative returns over $[-1, +1]$ and $[+2, +21]$. Drift is the cumulative return over days $[+2, +21]$; in the drift specification, the rank also enters interacted with $Post_e \times D_{di}$, and its coefficient measures the change in the return spread between good- and bad-news announcements.

Table 9 reports the results. After the launch, a firm one standard deviation harder to read sees a larger increase in abnormal announcement volume than a comparable firm, about 1.9 percent of dollar volume, a rise of 0.8 percentage points in the share of the price response that occurs in the announcement window, and a fall of 1.4 percentage points in the subsequent return spread between good- and bad-news announcements. Greater disclosure uncertainty also raises abnormal volume, by about 1.4 percent of dollar volume per standard deviation, but has no significant effect on incorporation speed or drift. The

¹⁵The pattern also holds using the top quintile or top five percent of absolute imbalance. These paths are unconditional averages and may also reflect differences in firm size, liquidity, retail participation, and subsequent order flow. Table 8 provides the regression evidence with controls.

additional trading in hard-to-read firms thus comes with faster incorporation of earnings news and less subsequent drift, while in firms with uncertain disclosures it comes with no evidence of better price discovery, the same contrast as in the flow-return evidence.

4.3 Comovement across hard-to-read stocks

Firm-specific attention can make trading within a stock more one-sided, but it gives trading in other stocks no reason to follow. A general-purpose model, which brings similar training, reading habits, and blind spots to every disclosure it reads, can create a component of order flow common to many stocks. If model users share part of their reading error in this way, comovement should rise most where model-assisted reading is most valuable.

For each firm-month, I first estimate the daily regression

$$OIB_{id} = a_{im} + b_{im} OIB_d^{(-i)} + u_{id}, \quad (5)$$

where $OIB_d^{(-i)}$ is the equal-weighted signed imbalance of all other sample firms on day d . The slope b_{im} is the firm's order-flow commonality. I then estimate

$$b_{im} = \alpha_i + \gamma_m + \sum_{d \in \{R, U\}} \delta_d (Post_m \times D_{di}) + \psi' (Post_m \times Z_i) + \theta' Act_{im} + \nu_{im}. \quad (6)$$

Before the launch, the average slope b_{im} is 0.83. After it, the readability differential is 0.105 for total flow and 0.091 for non-retail flow, with no significant change for retail flow (Table 10). A firm one standard deviation harder to read responds to the same 10-percentage-point market move by an additional 1.05 percentage points relative to other firms, about 13 percent of the pre-launch average. The uncertainty coefficients are negative but statistically insignificant, and placebo launch dates in October 2021, March 2022, and June 2022 produce no significant effects.¹⁶

The split between investor groups may reflect how they trade. Non-retail investors hold portfolios spanning many stocks and may adjust several positions on the same day, when a daily commonality measure can detect the alignment. If retail investors respond at different times, a similar shared component is harder to detect in their daily order flow.

The cross-stock effect is concentrated in 2023 and 2024 and fades in 2025, while within-stock synchronization keeps strengthening. One possible reason is that the two depend on

¹⁶In unreported results, allowing the post-launch change to vary also with S&P 500 membership at the launch, pre-launch institutional ownership, and the share of market value held by exchange-traded funds lowers the readability differential only from 0.105 to 0.090 for total flow and from 0.091 to 0.080 for non-retail flow, and both remain significant.

different shared errors. Investors who read the same disclosure can make correlated errors even with different models, but a shared error links trading across stocks only if it carries from one firm’s disclosure to another, which requires similar reading habits and weakens as users spread across model families. Anonymous order flow does not reveal which models investors use, so this explanation remains a conjecture.

4.4 Shared and firm-specific order flow

The price-impact regressions of Section 4.2.1 treat daily order flow as a single quantity, but part of that flow moves with trading elsewhere in the market, and Section 4.3 shows that this part rises in hard-to-read firms after the launch. I therefore split daily signed imbalance into a shared component and a firm-specific remainder. Within each firm-month, the shared component is the slope b_{im} of equation (5) times that day’s signed imbalance of all other sample firms, and the remainder is the firm’s imbalance minus this component. The two enter equation (3) in place of OIB_{it} , each with its own coefficients, and Panel B of Table 8 reports their same-day and quarterly price effects.

The two components have similar same-day price impacts, and both rise less after the launch in hard-to-read firms than in comparable firms. They differ over the following quarter. Before the launch, the quarterly slope on shared flow is +0.01 for the average firm. For a firm one standard deviation harder to read, the pre-launch readability coefficient of -1.78 brings it to about -1.8 percent per unit of imbalance, so much of the firm’s same-day impact of 2.6 percent reverses. After the launch, the average change of -1.72 and the readability differential of $+2.99$ bring the slope to about -0.5 percent, and most of the reversal disappears. For a firm one standard deviation more uncertain, the slope instead moves from $+2.0$ to -3.3 percent: continuation gives way to reversal.

Relative to other firms, the quarterly slope on shared flow rises by 3.0 percentage points per standard deviation of readability difficulty and falls by 3.6 percentage points per standard deviation of uncertainty, per unit of imbalance, roughly four times the changes for firm-specific flow. Because the same-day impacts change much less, the permanent component of shared flow rises with readability difficulty and falls with uncertainty.

These results suggest that shared trading carries more information when disclosures are hard to read but creates more temporary price pressure when outcomes remain unresolved. The first pattern is consistent with investors using similar models to understand disclosures more accurately. The second pattern favors the reading of the uncertainty coefficient in Section 4.1.1 in which the model widens disagreement. A model that merely helped less where outcomes are open would not make price discovery there worse, yet after the launch a

larger share of the price impact in uncertain firms reverses, and their announcements draw more trading without faster incorporation, as expected when investors trade on different probabilities for the same scenarios.

4.5 Heterogeneity across firms

Table 11 re-estimates the main readability specifications within halves of the sample, split by pre-launch retail share, size, and analyst coverage. The increase in one-sidedness is concentrated in small and thinly covered firms: per standard deviation of readability difficulty, the differential is 0.66 percentage points in small firms and 0.02 in large firms, with a similar contrast across analyst coverage, and it is twice as large in retail-heavy stocks as in institution-heavy stocks. The weaker quarterly reversal follows the retail-share and size splits: it is significant in retail-heavy and small stocks but not in institution-heavy or large stocks. Cross-stock commonality follows a different pattern. It rises in institution-heavy stocks but not in retail-heavy stocks, consistent with the non-retail result in Section 4.3, and more in small and thinly covered firms than in large and covered ones.

The larger synchronization effect in retail-heavy stocks raises an alternative explanation: popularity lists and social feeds on trading platforms may concentrate attention and make trading more one-sided. Several findings weigh against this account. These platform features predate the launch, yet the placebo launch dates of Section 4.1.1 produce no comparable readability differential. The effect appears in non-retail flow and in institution-heavy stocks, analysts' forecasts converge in the same firms, and uncertainty predicts the opposite change in one-sidedness within the same firms and months. The price-discovery results also survive excluding the tenth of firms with the largest increases in retail share during the first half of 2021: the quarterly reversal differential is then 0.68 percentage points per unit of imbalance, compared with 0.57 in the full sample, and the post-announcement drift differential remains -0.014 (Appendix A).

The effects are largest where investors had less help in interpreting disclosures before the launch. Small and thinly covered firms receive less analyst attention, and retail investors may rely more on their own reading, so better decoding can matter more where investors have fewer other sources of interpretation.

5 A model of delegated reading

I study the trade-off between better reading and shared errors in a one-period Kyle (1985) auction with multiple informed traders and correlated signals (Lambert et al., 2018). Fol-

lowing [Kim and Verrecchia \(1994\)](#), speculators obtain private information by interpreting a public disclosure. They choose between independent reading and a common technology that reduces decoding error but introduces partly shared mistakes, while leaving unresolved uncertainty unchanged. As more investors delegate, the expected profit from delegation falls relative to independent reading, reflecting the dilution of information value studied by [Grossman and Stiglitz \(1980\)](#) and [Admati and Pfleiderer \(1986\)](#). The model links disclosure difficulty to this adoption choice and traces its implications for synchronized trading, price discovery, and commonality across stocks.

5.1 Reading technologies

There is one risky asset with liquidation value $v \sim N(0, \sigma_v^2)$. The market contains M speculators, noise traders, and a competitive market maker. Each speculator forms a view of the asset by reading the firm's disclosure. Speculator i 's interpretation is

$$s_i = v + \eta + \varepsilon_i, \quad \text{var}(\eta) = \rho\tau^2, \quad \text{var}(\varepsilon_i) = (1 - \rho)\tau^2. \quad (7)$$

The fundamental v , the shared error η , and the idiosyncratic errors ε_i are jointly normal with mean zero and mutually independent, and the ε_i are identically distributed across speculators. The total reading-error variance is τ^2 , and ρ measures the fraction of that error shared across readers. Precision and commonality are therefore separate features of a reading: total error may fall even as the errors that remain become more correlated.

A disclosure can be difficult for two reasons. The first, c , is decoding difficulty: the text contains information that readers may struggle to extract without assistance. The second, g , is irreducible uncertainty: the disclosure leaves the underlying outcome unresolved, allowing readers to reach different assessments even when they understand the text. A better reader can reduce decoding error but cannot resolve uncertainty that the disclosure itself does not settle. The two reading technologies are

$$\begin{aligned} \text{own reading:} \quad & \tau_o^2 = c + g, \quad \rho_o = 0, \\ \text{model reading:} \quad & \tau_d^2 = A + g, \quad \rho_d = \frac{C}{A + g}, \end{aligned} \quad (8)$$

where A is the model's decoding-error variance and $0 < C \leq A$.¹⁷ The parameter C is the covariance between the errors in two delegated readings of the same disclosure. I

¹⁷ A falls as model capabilities improve and a broader choice of models improves the best available model-assisted reading. The cumulative model stock $N(m)$ in Section 4.1.3 serves as an empirical proxy for these developments. Model improvement enters the results below through a lower A .

assume that each independent reader’s error is independent of all other readers’ errors. Two independent readings, or an independent and a delegated reading, therefore share only the fundamental v . Uncertainty g contributes equally to the error variance under both technologies; delegation changes decoding precision and introduces a shared error, while leaving this unresolved uncertainty unchanged.

The common error can persist even when investors use different models. To see this, decompose a model’s reading error into a component shared across model families, a family-specific component, and an idiosyncratic prompt component. Two delegated readings always share the cross-family component and also share the family component when the users choose the same model family. The resulting covariance C is therefore positive whenever either common component is present.¹⁸ This structure is consistent with evidence that model errors remain correlated across architectures and providers.¹⁹ Appendix B gives the formal decomposition.

The baseline assumes that an investor uses either her own reading or the model’s reading. Appendix B lets her combine the two: a noisier own reading puts more weight on the model and hence on its shared error, which smooths the adoption decision and leaves the trade-off between precision and commonality intact.

The market maker learns from order flow and does not observe speculators’ readings before setting prices, so information reaches prices through trading, as in Kyle (1985). Following Kim and Verrecchia (1994), private information comes from processing a public announcement. Here the model is the processing technology, and each reader turns its output into a valuation with her own prior, horizon, and prompt, so interpretations differ across users of the same model and across model families. A market maker that read through its own model would learn only the component shared across all families, while the family-specific component and the prompt residual would still reach prices through trading; Appendix B shows that the propositions are unchanged. Evidence on automated liquidity supply fits this structure: liquidity suppliers update quotes on public information and order imbalances, and informed flow adversely selects their resting orders (Hendershott et al., 2011; Brogaard et al., 2014). The empirical results also suggest that trading continues to convey investors’ interpretations: after the launch, order flow remains associated with same-day price movements, and a larger share of the price response is permanent in hard-

¹⁸The family component is likely to matter in practice: among surveyed financial institutions in Hong Kong, only seven percent build their generative AI applications solely in house, while most rely at least partly on external providers (Hong Kong Institute for Monetary and Financial Research, 2025).

¹⁹Kim et al. (2025) document substantial agreement in errors across more than three hundred models, while Huang et al. (2026) find that models concentrate on similar analytical choices in open-ended empirical tasks.

to-read firms relative to other firms.

5.2 Equilibrium

The game has two stages. In Stage 1, each of the M speculators chooses whether to read the disclosure independently or delegate the reading to the model. Let $n \in [0, M]$ denote the mass of delegating traders. The chosen technology generates the trader's signal according to (8). In Stage 2, each speculator observes her own signal and submits a market order; speculators are risk neutral and maximize expected trading profit. Noise traders submit $u \sim N(0, \sigma_u^2)$, independent of v and of every reading error. The competitive market maker observes aggregate order flow,

$$Y = \sum_{i=1}^M x_i + u,$$

and sets the Bayesian price $p = E[v \mid Y]$.

Equilibrium. Fix the primitives $\Theta = (\sigma_v^2, \sigma_u^2, M, c, g, A, C)$. A linear equilibrium determines the mass of delegating traders, the trading intensity of each type, and the price response to aggregate order flow. Let s_i denote a delegated reading and q_j an independent reading, each generated by the corresponding technology in (8). Delegators submit $x_i = \beta_d s_i$, independent readers submit $x_j = \beta_o q_j$, and the market maker sets $p = \lambda Y = E[v \mid Y]$. Each trader's order is optimal given the orders of others and the market maker's pricing rule. At an interior composition, expected trading profits are equal across the two reading technologies, $\pi_d(n; \Theta) = \pi_o(n; \Theta)$.²⁰

Before solving the delegation choice, consider a benchmark in which all speculators use the same reading technology (τ^2, ρ) . Define

$$a = \frac{\sigma_v^2}{\sigma_v^2 + \tau^2}, \quad b = \frac{\sigma_v^2 + \rho\tau^2}{\sigma_v^2 + \tau^2}, \quad D = 2 + (M - 1)b.$$

Proposition 1 (Trading equilibrium). *Given a common reading technology (τ^2, ρ) , the unique symmetric linear equilibrium is $x_i = \beta s_i$ and $p = \lambda Y$, where*

$$\lambda = \frac{\sqrt{aM\sigma_v^2}}{D\sigma_u}, \quad \beta = \sigma_u \sqrt{\frac{a}{M\sigma_v^2}}. \quad (9)$$

²⁰I treat n as continuous in the main text to obtain a closed-form solution. Appendix B provides the finite- M version.

The posterior variance of the asset's value conditional on order flow is

$$\text{var}(v | Y) = \sigma_v^2 - \frac{M\sigma_v^4}{\tilde{D}}, \quad \tilde{D} = (M+1)(\sigma_v^2 + \rho\tau^2) + 2(1-\rho)\tau^2. \quad (10)$$

Proof. See Appendix C. □

The equilibrium follows the standard Kyle logic: the market maker extracts information from aggregate order flow and adjusts the price at rate λ , speculators trade on their signals while accounting for the price impact of their own orders, and noise trading provides the camouflage that keeps informed trading profitable.

Equation (10) shows how the structure of reading errors affects price discovery. Idiosyncratic errors partly offset across investors, while shared errors remain in aggregate order flow. Each unit of common-error variance enters $\tilde{D}$ with weight $M+1$, compared with 2 for idiosyncratic-error variance, so a unit of shared error costs $(M+1)/2$ times as much price informativeness as a unit of idiosyncratic error. The reason is that the market maker observes the sum of orders. Errors that speculators make alone average out in that sum; an error they share enters it M times, as the fundamental does, and the market maker cannot tell the two apart. Variance is therefore a private cost, borne by the reader who makes the error, and covariance is a social one: the shared error is priced, stays in the price until later information reveals it, and then reverses. Model-assisted reading creates a trade-off: it lowers decoding error from c to A , while introducing the shared error component C .

5.3 Endogenous delegation

I now let traders choose between independent and model-assisted reading. For a mixed composition, define

$$S_d = \sigma_v^2 + A + g, \quad S_o = \sigma_v^2 + c + g, \quad a_o = \frac{\sigma_v^2}{S_o}.$$

Proposition 2 (Endogenous delegation). *Let $\pi_d(n)$ and $\pi_o(n)$ denote the equilibrium expected profits of a delegating trader and an independent reader when the mass of delegators is n . Delegation is a strategic substitute: $\pi_d(n)/\pi_o(n)$ is decreasing in n . Under the continuous-composition benchmark:*

- (i) if $c < A$, no trader delegates;

(ii) if $c > A$, the interior composition is

$$n^*(c, g) = \frac{(2 - a_o)\sqrt{S_o S_d} - 2S_d + \sigma_v^2 + C}{C}, \quad (11)$$

clipped at M , and $\partial n^*/\partial c > 0$;

(iii) at $c = A$, $n = 0$ and a composition with a single delegator induce the same allocation.

Hence the adoption threshold is

$$c^* = A, \quad (12)$$

and is independent of uncertainty g .

Proof. See Appendix C. □

The adoption rule has a simple intuition. The first model user shares the model's error with no other trader, so her decision depends only on the gain in precision, and delegation begins once the model's decoding error A falls below the investor's own, c . As more traders delegate, their information increasingly overlaps and the value of joining them falls. This crowding slows further adoption, while greater decoding difficulty raises the precision gain from using the model and attracts more delegators. Uncertainty enters both reading technologies equally and so does not affect the entry threshold, but a larger g always reduces the common-error share $C/(A + g)$ of a delegated reading.

Lemma 1 (The effective reader mass). *Let $t \equiv \sqrt{S_o/S_d}$, which exceeds one exactly when $c > A$, and define the effective reader mass $m \equiv M + n^*(t - 1)$. For $c > A$ with $n^* < M$, the composition and the trading intensities of Proposition 2 satisfy*

$$n^* = 1 + \frac{(t - 1)(2S_d + \sigma_v^2/t)}{C}, \quad \kappa_o = \frac{a_o}{2 + (m - 1)a_o}, \quad \kappa_d = t \kappa_o, \quad (13)$$

where $\kappa_t = \lambda \beta_t$, and the posterior variance and price impact are

$$\text{var}(v | Y) = \sigma_v^2 \frac{2 - a_o}{2 + (m - 1)a_o}, \quad \lambda = \frac{\sqrt{a_o M \sigma_v^2}}{(2 + (m - 1)a_o)\sigma_u}. \quad (14)$$

Relative to the all-own market, posterior variance and price impact scale by the same factor, $(2 + (M - 1)a_o)/(2 + (m - 1)a_o) < 1$.

Proof. See Appendix C. □

The lemma expresses the gain in price discovery as an increase in the effective number of independent readers. When $c > A$, an interior equilibrium has the same posterior variance as an all-own-reading market with reader mass $m = M + n^*(t - 1) > M$. For price informativeness, each delegator counts as t independent readers, even though the number of traders remains M . The reason is profit indifference between the two types, $\beta_d \sqrt{S_d} = \beta_o \sqrt{S_o}$: delegators and independent readers submit orders of the same expected absolute size, but a delegator's order rests on a more precise signal and so reveals more about the fundamental, as if the market contained additional independent readers.

The lemma also characterizes price impact. Delegators trade more aggressively on more precise but overlapping information, so aggregate trading profits fall even as prices become more informative, much as greater competition among informed traders lowers them in [Admati and Pfleiderer \(1988\)](#) and [Holden and Subrahmanyam \(1992\)](#). Because these profits equal the noise traders' expected loss, $\lambda \sigma_u^2$, price impact falls with them: by the same factor as the posterior variance, and by more where decoding is harder, because both the delegating share and the precision gain rise with c .

5.4 Mapping the model to the evidence

I use three model outcomes that correspond to the empirical tests. First, the model analogue of absolute order imbalance is expected absolute net order flow relative to total expected absolute submitted volume,

$$|OIB| \equiv \frac{\mathbb{E} |Y|}{\sum_{i=1}^M \mathbb{E} |x_i| + \mathbb{E} |u|}. \quad (15)$$

Second, I measure pricing error by

$$\mathcal{P} \equiv \sqrt{\text{var}(v \mid Y)}, \quad (16)$$

so a fall in $\mathcal{P}$ means better price discovery. Third, consider two otherwise identical firms $f \neq \ell$ with independent fundamentals and let $\chi^2 \leq C$ denote the covariance in delegated reading errors that carries across firms. Cross-stock order-flow commonality is

$$\varsigma \equiv \frac{\text{cov}(Y_f, Y_\ell)}{\text{var}(Y_f)}. \quad (17)$$

The full-delegation benchmark yields closed-form comparative statics for all three outcomes. Lemma 1 also characterizes pricing error under endogenous delegation. For synchronization

and cross-stock commonality, changes in c , g , and A affect both the equilibrium delegating share and trading intensities. I examine these effects numerically in Section 5.5.

Proposition 3 (Within-stock synchronization). *Compare the all-own benchmark with the forced full-delegation benchmark. As $M \rightarrow \infty$,*

$$|OIB|_o \rightarrow \sqrt{\frac{\sigma_v^2}{\sigma_v^2 + c + g}}, \quad |OIB|_d \rightarrow \sqrt{\frac{\sigma_v^2 + C}{\sigma_v^2 + A + g}}. \quad (18)$$

For every $c > A$, the full-delegation change $\Delta|OIB|^{FD} \equiv |OIB|_d - |OIB|_o$ is positive, increasing in decoding difficulty c , and decreasing in uncertainty g .

Proof. See Appendix C. □

The benchmark makes the synchronization mechanism transparent. Under independent reading, investors make idiosyncratic errors, which tend to offset in aggregate. As $c + g$ rises, independent readings become noisier and $|OIB|_o$ falls. Delegation introduces a shared component C of reading error. This component moves model users' orders in the same direction and survives aggregation.

Decoding difficulty strengthens this effect through a second channel. Proposition 2 shows that n^* increases with c , so harder disclosures attract more delegation and expose more traders to the shared model error. Synchronization therefore rises with decoding difficulty. Uncertainty works differently. A higher g increases the part of reading error that remains unresolved under either technology and dilutes the shared component $C/(A + g)$ of a delegated reading.

The adoption threshold also provides a prediction for the shape of the readability relation. When $c < c^*$, no trader delegates and the model generates no change in synchronization. Once c crosses the threshold, delegation begins and exposure to the shared model error rises. The model therefore predicts little response at the easy end of the readability distribution, followed by a positive response once decoding difficulty becomes sufficiently high. Table 5 has this shape.

Proposition 4 (Price informativeness). *In the forced full-delegation benchmark, model use lowers pricing error $\mathcal{P}$ if and only if*

$$c > c^\dagger = A + \frac{1}{2}(M - 1)C. \quad (19)$$

Above $c^\dagger$, the reduction in pricing error becomes larger as c rises, smaller as g rises, and converges to zero as $g \rightarrow \infty$.

Proof. See Appendix C. □

The threshold shows why more precise individual readings need not produce more informative prices. When $A < c < c^\dagger$, each model-assisted reading is more precise than an independent reading, yet full delegation raises pricing error. The shared error C survives aggregation, and its cost outweighs the gain in individual precision.

Allowing investors to choose their reading technology changes the outcome. Near the adoption threshold, few traders delegate, and crowding limits adoption while the precision gain is small. Harder disclosures support more delegation, and full delegation occurs only when the gain is large enough to improve price informativeness. This gives the following result.

Corollary 1 (No deterioration at the endogenous composition). *At the equilibrium composition, $\text{var}(v \mid Y)$ is no larger than its all-own-reading value for every parameter configuration with $0 < C \leq A$, and smaller whenever $c > A$.*

Proof. See Appendix C. □

Lemma 1 makes the result transparent. The all-own-reading market has M effective readers, while an interior equilibrium with delegation has $m = M + n^*(t - 1) > M$, because delegated readings are more precise when $c > A$. Posterior variance falls with m , so price discovery improves as soon as investors begin to delegate, and the deterioration possible under forced delegation disappears.

Proposition 5 (Cross-stock commonality). *Suppose delegated readings contain a cross-firm component with covariance $\chi^2 \leq C$. Under all-own reading $\varsigma_o = 0$, so the change induced by forced full delegation is*

$$\Delta\varsigma^{FD} \equiv \varsigma_d - \varsigma_o = \frac{M\chi^2}{(M+1)(\sigma_v^2 + C) + 2(A + g - C)}. \quad (20)$$

Thus the full-delegation change in cross-stock commonality is increasing in χ^2 , decreasing in uncertainty g , and independent of c once delegation is complete.

Proof. See Appendix C. □

Decoding difficulty affects cross-stock commonality through adoption. Below c^* , no trader delegates, so the model introduces no cross-firm component into order flow. Above the threshold, harder disclosures attract more delegators, and a larger share of trading loads on the cross-firm component of model error.

The propositions describe the same comparisons as the empirical specifications in equations (1), (3), (4), and (6): how outcomes change, across firms with different disclosure characteristics, when access to model-assisted reading expands. Decoding difficulty c is a property of the reader as well as of the text. Own readings are noisiest where investors have the least help, in the retail-heavy, small, and thinly covered firms of Section 4.5, so the precision gain from delegation, and with it the predicted change in synchronization and price discovery, is largest there; where readings were already precise, delegation adds little precision and leaves the shared component.

5.5 Numerical results

Propositions 3 to 5 use full delegation to obtain closed-form comparative statics. Figure 4 evaluates the three outcomes at the endogenous composition of Lemma 1 as decoding difficulty and uncertainty vary. The calibration normalizes $\sigma_v^2 = 1$, sets $\sigma_u^2 = 4$ and $M = 20$, and chooses $A = 0.23$ and $C = 0.096$, which imply an entry threshold $c^* = 0.23$ and a full-delegation price-quality boundary $c^\dagger = 1.14$. The equilibrium delegating share rises from zero at the threshold to about a quarter at $c = 0.5$ and a half at $c = 1$, and every trader delegates above approximately $c = 2.1$.

Panel (a) shows that the increase in one-sidedness grows with decoding difficulty, matching the readability gradient in Table 4. Uncertainty weakens the effect because it adds errors that the model cannot remove and dilutes the shared component of the delegated reading. Panel (b) shows the change in pricing error. Under forced full delegation, pricing error rises when c lies between the adoption threshold and the price-quality boundary, even though each model-assisted reading is more precise. With endogenous delegation, this deterioration disappears: few investors adopt when the precision gain is small, and adoption expands as the model’s advantage grows. The improvement in price discovery is small near the adoption threshold and increases with decoding difficulty. Panel (c) shows cross-stock commonality, which is zero under independent reading. As more traders delegate, their orders load on the shared error component that carries across firms. Commonality rises with adoption and levels off once all traders delegate.

5.6 Predictions beyond the current data

The current data do not identify which investors use which models. Investor-level data linking model use to trading would allow three further predictions to be tested.

First, investors should rely more heavily on models when disclosures are harder to read.

Data linking investors, prompts, and trades could show whether the same investor uses model assistance more often for firms with greater decoding difficulty.

Second, trading comovement should depend on model family. Users of the same family should exhibit stronger within-stock comovement than users of different families, and cross-stock commonality should be greater when model use is concentrated in fewer families.

Third, model releases should produce coordinated changes in trading across hard-to-read stocks among investors who adopt the new version. Section 4.1.3 documents growth in the aggregate synchronization differential over the three years after the launch, but anonymous order flow cannot identify whose trading changes or which releases prompt those changes. Linking trades to model families and adoption dates would allow a direct test of whether users adjust their trading across several firms after adopting a new release.

6 Conclusion

Generative AI can improve investors' understanding of public information and bring their interpretations closer together. Greater agreement can reflect better decoding, shared mistakes, or both, so its implications for price discovery depend on how individual precision and shared errors change.

I distinguish two persistent features of corporate disclosures, readability and uncertainty. After the public launch of ChatGPT, analysts' forecasts converge and trading becomes more synchronized in hard-to-read firms relative to comparable easier-to-read firms, and non-retail order flow in these firms moves more with the rest of the market. Price discovery also improves: a larger share of the price response to one-sided flow is permanent, and earnings news enters prices faster. Greater disclosure uncertainty predicts less convergence and synchronization, with no corresponding improvement in price discovery. The evidence suggests that AI helps most when information is difficult to extract and that its benefits are limited when the disclosure leaves outcomes unresolved.

In the model, investors delegate when the model offers a sufficiently precise reading, but the profit from delegation falls as more traders act on overlapping information. This crowding limits adoption when the precision gain is small, so shared errors do not outweigh that gain in equilibrium, and more synchronized trading coexists with better price discovery. The empirical results are consistent with this mechanism, but the data do not observe model use at the investor-firm level. Linking investors' model choices to their trades would allow a direct test of how precision gains and shared information shape adoption and price discovery.

References

- Admati, A. R. and P. Pfleiderer (1986). A monopolistic market for information. *Journal of Economic Theory* 39(2), 400–438.
- Admati, A. R. and P. Pfleiderer (1988). A theory of intraday patterns: Volume and price variability. *The Review of Financial Studies* 1(1), 3–40.
- Beckmann, L., H. Beckmeyer, I. Filippou, S. Menze, and G. Zhou (2024). Unusual financial communication: ChatGPT, earnings calls, and financial markets. *Olin Business School Center for Finance & Accounting Research Paper* (2024/02).
- Bertomeu, J., Y. Lin, Y. Liu, and Z. Ni (2023). Capital market consequences of generative AI: Early evidence from the ban of ChatGPT in Italy. Working paper, Washington University in St. Louis.
- Bertomeu, J., Y. Lin, Y. Liu, and Z. Ni (2025). The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy. *Journal of Accounting and Economics* 80(1), 101782.
- Blankespoor, E., J. Croom, and S. M. Grant (2026). Generative AI and investor processing of financial information. *Journal of Accounting and Economics*. Forthcoming.
- Blankespoor, E., E. deHaan, and Q. Li (2026). Generative AI in financial reporting. *Journal of Accounting Research* 64(3), 1189–1232.
- Boehmer, E., C. M. Jones, X. Zhang, and X. Zhang (2021). Tracking retail investor activity. *The Journal of Finance* 76(5), 2249–2305.
- Brogaard, J., T. Hendershott, and R. Riordan (2014). High-frequency trading and price discovery. *The Review of Financial Studies* 27(8), 2267–2306.
- Bushee, B. J., J. E. Core, W. Guay, and S. J. W. Hamm (2010). The role of the business press as an information intermediary. *Journal of Accounting Research* 48(1), 1–19.
- Campbell, J. L., J. R. Stark, J. D. Warren, and Z. Wiebe (2026). Generative artificial intelligence and retail investors’ processing of earnings news. Working paper, University of Georgia.
- Cao, S., W. Jiang, B. Yang, and A. L. Zhang (2023). How to talk when a machine is listening: Corporate disclosure in the age of AI. *The Review of Financial Studies* 36(9), 3603–3642.

- Cartea, Á., P. Chang, N. Chen, and M. Zhong (2026). AI bubbles with large language models. Working paper, University of Oxford.
- Chang, A., X. Dong, X. Martin, and C. Zhou (2026). AI democratization and trading inequality. *Journal of Accounting Research* 64(3), 1287–1331.
- Chen, J., G. Tang, G. Zhou, and W. Zhu (2025). ChatGPT and DeepSeek: Can they predict the stock market and macroeconomy? *arXiv preprint arXiv:2502.10008*.
- Chen, S., L. Peng, and D. Zhou (2025). Wisdom or whims? decoding retail strategies with social media and AI. Working paper, SSRN 4867401.
- Cheng, Q., P. Lin, and Y. Zhao (2025). Does generative AI facilitate investor trading? early evidence from ChatGPT outages. *Journal of Accounting and Economics* 80, 101821.
- Chordia, T., R. Roll, and A. Subrahmanyam (2000). Commonality in liquidity. *Journal of Financial Economics* 56(1), 3–28.
- Chordia, T. and A. Subrahmanyam (2004). Order imbalance and individual stock returns: Theory and evidence. *Journal of Financial Economics* 72(3), 485–518.
- Dou, W. W., I. Goldstein, and Y. Ji (2025). AI-powered trading, algorithmic collusion, and price efficiency. Technical report, National Bureau of Economic Research.
- Epoch AI (2026). Data on AI models. Notable AI Models database, snapshot downloaded 29 July 2026.
- Even-Tov, O., B. Lourie, K. Munevar, and A. Nekrasov (2025). The effect of AI on retail investor behavior: Evidence from account-level trading data. Working paper, SSRN 5778246.
- Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of financial economics* 116(1), 1–22.
- Gao, M. and J. Huang (2020). Informing the market: The effect of modern information technologies on information production. *The Review of Financial Studies* 33(4), 1367–1411.
- Goldstein, I., Y. Xiong, and L. Yang (2025). Information sharing in financial markets. *Journal of Financial Economics* 163, 103967.

- Grossman, S. J. and M. H. Miller (1988). Liquidity and market structure. *The Journal of Finance* 43(3), 617–633.
- Grossman, S. J. and J. E. Stiglitz (1980). On the impossibility of informationally efficient markets. *The American economic review* 70(3), 393–408.
- Hansen, A. L. and S. Kazinnik (2024). Can ChatGPT decipher FedSpeak? *Available at SSRN 4399406*.
- Hasbrouck, J. and D. J. Seppi (2001). Common factors in prices, order flows, and liquidity. *Journal of Financial Economics* 59(3), 383–411.
- Healy, P. M. and K. G. Palepu (2001). Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature. *Journal of Accounting and Economics* 31(1–3), 405–440.
- Hendershott, T., C. M. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *The Journal of finance* 66(1), 1–33.
- Hendershott, T. and A. J. Menkveld (2014). Price pressures. *Journal of Financial Economics* 114(3), 405–423.
- Hirshleifer, D., L. Peng, Q. Wang, W. Zhang, and X. Zhang (2026). AI, opinion ecosystems, and finance. Working Paper 34807, National Bureau of Economic Research.
- Holden, C. W. and A. Subrahmanyam (1992). Long-lived private information and imperfect competition. *The Journal of Finance* 47(1), 247–270.
- Hong Kong Institute for Monetary and Financial Research (2025, April). Financial services in the era of generative AI: Facilitating responsible adoption. HKIMR Applied Research Report 1/2025, Hong Kong Institute for Monetary and Financial Research.
- Huang, L., M. Qiao, and D. Zhou (2024). Viral communication in financial markets: Meme images and market efficiency. Working paper, SSRN 4785481.
- Huang, W., A. J. Menkveld, and S. Yu (2026). AI “errors”. Working paper, Bank for International Settlements.
- International Monetary Fund (2024). Global financial stability report. Technical Report October 2024.

- Jha, M., J. Qian, M. Weber, and B. Yang (2024). ChatGPT and corporate policies. Technical report, National Bureau of Economic Research.
- Kalai, A. T. and S. S. Vempala (2024). Calibrated language models must hallucinate. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 160–171. ACM.
- Kandel, E. and N. D. Pearson (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy* 103(4), 831–872.
- Kandpal, N., H. Deng, A. Roberts, E. Wallace, and C. Raffel (2023). Large language models struggle to learn long-tail knowledge. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pp. 15696–15707.
- Kang, H. and X.-Y. Liu (2023). Deficiency of large language models in finance: An empirical examination of hallucination. *arXiv preprint arXiv:2311.15548*.
- Kim, A., M. Muhn, and V. Nikolaev (2023). Bloated disclosures: can ChatGPT help investors process information? *arXiv preprint arXiv:2306.10224*.
- Kim, E., A. Garg, K. Peng, and N. Garg (2025). Correlated errors in large language models. In *Proceedings of the 42nd International Conference on Machine Learning*, Volume 267 of *Proceedings of Machine Learning Research*, pp. 30038–30066.
- Kim, O. and R. E. Verrecchia (1994). Market liquidity and volume around earnings announcements. *Journal of Accounting and Economics* 17(1–2), 41–67.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315–1335.
- Lambert, N. S., M. Ostrovsky, and M. Panov (2018). Strategic trading in informationally complex environments. *Econometrica* 86(4), 1119–1157.
- Lee, C. M. C. and M. J. Ready (1991). Inferring trade direction from intraday data. *The Journal of Finance* 46(2), 733–746.
- Lehavy, R., F. Li, and K. Merkley (2011). The effect of annual report readability on analyst following and the properties of their earnings forecasts. *The Accounting Review* 86(3), 1087–1115.
- Li, F. (2008). Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics* 45(2-3), 221–247.

- Li, K., F. Mai, R. Shen, C. Yang, and T. Zhang (2026). Dissecting corporate culture using generative AI. *The Review of financial studies* 39(1), 253–296.
- Lopez-Lira, A. and Y. Tang (2026). Can ChatGPT forecast stock price movements? return predictability and large language models. *Journal of Financial Economics*. Forthcoming.
- Loughran, T. and B. McDonald (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-Ks. *The Journal of Finance* 66(1), 35–65.
- Loughran, T. and B. McDonald (2014). Measuring readability in financial disclosures. *The Journal of Finance* 69(4), 1643–1671.
- Lu, F., L. Huang, and S. Li (2023). ChatGPT, generative AI, and investment advisory. *Available at SSRN 4519182*.
- Lu, R., L. Vulicevic, and Y. Xu (2026). Do LLMs make markets more efficient? Working paper, University of Toronto.
- Morck, R., B. Yeung, and W. Yu (2000). The information content of stock markets: Why do emerging markets have synchronous stock price movements? *Journal of Financial Economics* 58(1–2), 215–260.
- Shaffer, M. and C. C. Y. Wang (2024). Scaling core earnings measurement with large language models. Working paper, SSRN 4979501.
- Sheng, J., Z. Sun, B. Yang, and A. L. Zhang (2026). Generative AI and asset management. *Review of Financial Studies*. Forthcoming.
- Stanco, J. and K. H. Chung (2025). Generative AI and the transmission of public information in the stock market: Evidence from the release of ChatGPT. Working paper, SSRN 5171287.
- Stein, J. C. (2009). Presidential address: Sophisticated investors and market efficiency. *The Journal of Finance* 64(4), 1517–1548.
- Xu, Y., Y. Xuan, and G. Zheng (2025). How AI agents mislead investors: Early evidence from automated stock recommendations. *Available at SSRN 5651130*.

A Additional empirical results

Table A1 re-estimates both disclosure coefficients of each test on the 3,669 firms common to the two estimation panels. The order-flow coefficients are identical, and the flow-return coefficients are close to their full-sample values.

Table A2 reports the market-wide means of one-sidedness and of analyst forecast dispersion before and after the launch, the changes that the month and calendar-quarter fixed effects absorb. Every imbalance measure and log dispersion is lower after the launch, while mean dispersion scaled by price is unchanged.

Table A3 enters the two components of readability, words per sentence and the complex-word share, separately and together. Both carry the effect, and the complex-word share carries the larger coefficient per standard deviation whether or not length and uncertainty are included.

Table A4 compares the baseline readability exposure with three alternatives. The WRDS SEC Analytics Suite, which supplies the sentence-level parse behind the baseline, revised its sentence splitter in the third quarter of 2022, inside the exposure window, and again in the second quarter of 2023. The alternatives rank each filing among the filings of its quarter, use only the filings dated before July 2022, or replace the composite with the complex-word share from the Loughran–McDonald 10-X Summaries, an independent parse. Each gives a positive and significant coefficient for dollar imbalance, smaller than the baseline’s. The Loughran–McDonald share captures vocabulary alone, while sentence length accounts for most of the variation in the composite, so it agrees least with the baseline and gives the smallest coefficient; the measure built from the fewer filings dated before July 2022 is not significant for trade-count imbalance.

Table A5 measures the persistence of the text measures on the Loughran–McDonald 10-X Summaries, whose word counts, unlike the WRDS parse, are constructed the same way throughout the sample. Adjacent-year rank correlations range from 0.76 to 0.89, and pre- and post-launch firm means correlate at 0.79 to 0.83.

Table A6 relates the two disclosure measures to pre-launch measures of uncertainty about fundamentals: option-implied volatility, the volatility of return on assets, idiosyncratic volatility, analyst forecast dispersion, and the share of litigious words in the same filings. The uncertainty measure is positively correlated with implied volatility, the volatility of return on assets, and litigious language, also after removing size, while readability is negatively correlated with every volatility measure.

Table A7 re-estimates the main specification separately for the firms whose quintile of the Loughran–McDonald complex-word share moved by two or more between 2021–22 and

2023–25 and for the rest. The readability coefficient for dollar imbalance is 0.26 among the firms that did not move, and 0.26 to 0.29 when the firms that became easier or harder to read are excluded.

Table A8 checks whether the results are driven by the retail trading episode in the first half of 2021. I use three approaches: start the pre-launch period in July 2021, exclude the tenth of firms with the largest increases in retail share during the episode, or control for that increase interacted with the post-launch indicator. Excluding those months leaves the readability differentials in one-sidedness, commonality, and quarterly reversal at least as large as in the full sample. Average pre-launch reversal is weaker in the shorter sample, with a quarterly slope of -0.51 percent per unit of imbalance compared with -0.71 in the full sample. Thus, the episode contributes to average reversal but does not drive the readability differential. The announcement coefficients retain their signs but lose significance when those months are excluded. Both firm-level checks preserve their significance, suggesting that the loss is due to the shorter pre-launch period rather than exposure to the episode. The uncertainty differential in quarterly reversal remains negative and significant in all three specifications.

Table A9 re-estimates the synchronization regressions with filing length, a control in the main specifications, entered as a third dimension alongside readability and uncertainty. The readability and uncertainty coefficients are unchanged, and the length coefficients are small and insignificant at the five percent level in every row. In unreported results, length is also insignificant in the decile and commonality specifications, and its coefficient for announcement incorporation speed disappears once the post-launch change is allowed to vary with the firm characteristics. A long filing takes time to process, and a model can reduce that time. But faster processing synchronizes trading only if it also aligns investors’ interpretations. When the additional text is repetitive or only weakly related to valuation, a shared summary does little to reduce disagreement. After the launch of public LLMs, investors reach the same conclusions they would have reached without AI, just faster.

B Model extensions

Mixed-composition outcomes used in the quantitative illustration. The analytical propositions above use full delegation. For the quantitative analysis, let n traders delegate and define $\kappa_t = \lambda\beta_t$ for $t \in \{d, o\}$. Proposition 2 determines the equilibrium values of κ_d and κ_o . Define

$$K \equiv [n\kappa_d + (M - n)\kappa_o]\sigma_v^2. \quad (21)$$

This term summarizes how much information about the fundamental is revealed through aggregate order flow. Since

$$\text{cov}(v, Y) = [n\beta_d + (M - n)\beta_o]\sigma_v^2 = \frac{K}{\lambda}$$

and Bayesian pricing implies $\lambda = \text{cov}(v, Y) / \text{var}(Y)$, I obtain

$$\text{var}(Y) = \frac{K}{\lambda^2}$$

and therefore

$$\text{var}(v \mid Y) = \sigma_v^2 - K. \quad (22)$$

Using centered normality in equation (15), the common factor $\sqrt{2/\pi}$ cancels and the mixed-composition one-sidedness measure is

$$|OIB| = \frac{\sqrt{K}}{n\kappa_d\sqrt{S_d} + (M - n)\kappa_o\sqrt{S_o} + \lambda\sigma_u}. \quad (23)$$

If delegated readings contain the cross-firm component χ^2 , only delegators load on it, so

$$\text{cov}(Y_f, Y_\ell) = (n\beta_d)^2\chi^2.$$

Combining this with $\text{var}(Y_f) = K/\lambda^2$ gives

$$\varsigma = \frac{(n\kappa_d)^2\chi^2}{K}. \quad (24)$$

Evaluating equations (22), (23), and (24) at $n = n^*$ gives the endogenous-equilibrium outcomes reported in Figure 4.

Finite- M interpretation and observability. Equation (11) treats the mass of delegators as continuous, so a trader compares the two reading technologies holding the aggregate composition n fixed. With a finite number of traders, switching technologies changes the composition and therefore changes the continuation trading equilibrium. If there are n delegators, a delegator who switches to independent reading faces composition $n - 1$, while an independent reader who switches to delegation faces composition $n + 1$. An integer equilibrium therefore satisfies

$$\pi_d(n) \geq \pi_o(n - 1), \quad \pi_o(n) \geq \pi_d(n + 1), \quad (25)$$

with the trading equilibrium re-solved after each deviation. This formulation requires traders to observe the aggregate composition n before trading, but not the identities of individual delegators. If reading choices were private, the continuation game would instead involve uncertainty about trader types. Under the baseline calibration, the integer equilibrium is close to the continuous solution. For example, $n = 4$ compared with $n^* = 4.7$ at $c = 0.5$, $n = 10$ compared with $n^* = 10.3$ at $c = 1$, and $n = 19$ compared with $n^* = 19.2$ at $c = 2$.

Combining own and model readings. The baseline treats own reading and model reading as separate choices. Suppose instead that an investor combines the two. If she places weight w on the model reading, the variance of the decodable component of her error is

$$w^2 A + (1 - w)^2 c. \quad (26)$$

Differentiating,

$$2wA - 2(1 - w)c = 0,$$

so the variance-minimizing weight is

$$w^*(c) = \frac{c}{c + A}. \quad (27)$$

At that weight,

$$w^{*2} A + (1 - w^*)^2 c = \frac{Ac}{A + c} = Aw^*.$$

Including the irreducible uncertainty component g , total error variance and the common-error share are

$$\tau^2(w^*) = \frac{Ac}{A + c} + g = Aw^* + g, \quad \rho(w^*) = \frac{w^{*2} C}{Aw^* + g}. \quad (28)$$

Thus $Aw^* + g$ is the minimized variance evaluated at w^* ; an arbitrary blend has a higher variance.

The weight $w^*(c) = c/(c + A)$ increases with decoding difficulty. Investors therefore rely more heavily on the model when their own reading is noisier. Blending improves precision, but greater model weight also increases exposure to the shared error, whose covariance across users is $w^{*2} C$. The same precision and commonality trade-off therefore remains.

The weight w^* minimizes reading-error variance. In the trading equilibrium, investors also account for how their information overlaps with that of other traders, so the profit-maximizing weight can differ. Allowing blending also smooths the sharp adoption threshold

because $w^* > 0$ whenever $c > 0$. The key comparative static survives: model reliance increases with decoding difficulty, and greater reliance brings greater exposure to the common model error.

Model families as a reduced-form interpretation. The common-error parameter C can be interpreted through a simple model-family structure. Suppose a delegated reading contains three error components: a component ω shared across all model families, a family-specific component ψ_k , and an idiosyncratic prompt residue with variance s_ε^2 . If two users choose the same family with probability H , then

$$A = \omega^2 + \psi^2 + s_\varepsilon^2, \quad C = \omega^2 + H\psi^2. \quad (29)$$

The parameter A therefore captures the total error of an individual delegated reading, while C captures the error shared across users.

This decomposition separates model quality from model fragmentation. A fall in A improves the precision of a delegated reading and lowers the adoption threshold. A fall in H reduces C while leaving A unchanged, so greater dispersion across model families weakens the common component without changing the precision of an individual reading. For cross-stock commonality, only part of this shared error needs to carry across firms. Denoting its covariance by $\chi^2 \leq C$ gives the cross-stock component used in Proposition 5.

A market maker with its own model reading. Suppose the market maker also obtains a delegated reading, $s_0 = v + \omega + \psi_0 + \xi_0$, where ψ_0 is the family component of the model it uses and ξ_0 its own prompt residue, and sets $p = E[v \mid s_0, Y]$. Conditioning on s_0 before orders arrive replaces the fundamental by its residual $\tilde{v} = v - E[v \mid s_0]$, with variance $\text{var}(v \mid s_0) < \sigma_v^2$, and replaces each speculator's reading by its residual given s_0 . In the limiting case in which s_0 reveals the cross-family component exactly, ω is in the quote before trading and drops out of the covariance between two delegated readings, while the family component that two users share with probability H and their prompt residues do not: the covariance falls from $\omega^2 + H\psi^2$ to $H\psi^2$. The trading game is then the one analyzed above with σ_v^2 replaced by $\text{var}(v \mid s_0)$ and C by $H\psi^2$. Because c and g enter only through the speculators' own reading error, every comparative static in decoding difficulty and uncertainty is unchanged; a market maker that reads through a model shrinks the prior and the shared error without altering the structure of the equilibrium.

C Proofs

Proof of Proposition 1. Under (7),

$$s_i = v + \eta + \varepsilon_i, \quad \text{var}(s_i) = \sigma_v^2 + \tau^2.$$

For $j \neq i$,

$$\text{cov}(v, s_i) = \sigma_v^2, \quad \text{cov}(s_j, s_i) = \sigma_v^2 + \rho\tau^2.$$

Because the variables are jointly Gaussian,

$$\mathbb{E}[z \mid s_i] = \frac{\text{cov}(z, s_i)}{\text{var}(s_i)} s_i.$$

Therefore

$$\mathbb{E}[v \mid s_i] = a s_i, \quad \mathbb{E}[s_j \mid s_i] = b s_i, \quad (30)$$

where

$$a = \frac{\sigma_v^2}{\sigma_v^2 + \tau^2}, \quad b = \frac{\sigma_v^2 + \rho\tau^2}{\sigma_v^2 + \tau^2}.$$

Suppose every trader other than i follows $x_j = \beta s_j$, and conjecture that the market maker uses $p = \lambda Y$. Trader i chooses x_i after observing s_i . Her conditional expected profit is

$$\begin{aligned} \mathbb{E}[(v - p)x_i \mid s_i] &= x_i \mathbb{E} \left[v - \lambda \left(x_i + \sum_{j \neq i} x_j + u \right) \mid s_i \right] \\ &= x_i \mathbb{E}[v \mid s_i] - \lambda x_i^2 - \lambda x_i \sum_{j \neq i} \mathbb{E}[x_j \mid s_i], \end{aligned} \quad (31)$$

where $\mathbb{E}[u \mid s_i] = 0$ because noise trading is independent of all signals. Using (30),

$$\mathbb{E}[x_j \mid s_i] = \beta \mathbb{E}[s_j \mid s_i] = \beta b s_i.$$

There are $M - 1$ other speculators, so

$$\mathbb{E}[(v - p)x_i \mid s_i] = x_i \{a - \lambda(M - 1)\beta b\} s_i - \lambda x_i^2. \quad (32)$$

The second derivative with respect to x_i is $-2\lambda < 0$. The first-order condition is

$$\{a - \lambda(M - 1)\beta b\} s_i - 2\lambda x_i = 0,$$

or

$$x_i = \frac{a - \lambda(M-1)\beta b}{2\lambda} s_i.$$

In a symmetric equilibrium $x_i = \beta s_i$. Hence

$$2\lambda\beta = a - \lambda(M-1)\beta b,$$

so

$$\lambda\beta D = a, \quad D \equiv 2 + (M-1)b, \quad \beta = \frac{a}{\lambda D}. \quad (33)$$

It remains to impose competitive Bayesian pricing. Let

$$V \equiv \text{var}\left(\sum_{i=1}^M s_i\right).$$

There are M signal variances and $M(M-1)$ ordered covariance terms, hence

$$V = M(\sigma_v^2 + \tau^2) + M(M-1)(\sigma_v^2 + \rho\tau^2). \quad (34)$$

Since

$$Y = \beta \sum_{i=1}^M s_i + u,$$

I have

$$\text{cov}(v, Y) = \beta M \sigma_v^2, \quad \text{var}(Y) = \beta^2 V + \sigma_u^2.$$

The competitive market maker therefore sets

$$\lambda = \frac{\text{cov}(v, Y)}{\text{var}(Y)} = \frac{\beta M \sigma_v^2}{\beta^2 V + \sigma_u^2}. \quad (35)$$

Substitute $\beta = a/(\lambda D)$ from (33):

$$\lambda = \frac{a M \sigma_v^2 / (\lambda D)}{a^2 V / (\lambda^2 D^2) + \sigma_u^2}.$$

Multiplying both sides by λD^2 gives

$$a^2 V + \lambda^2 D^2 \sigma_u^2 = a M \sigma_v^2 D. \quad (36)$$

To simplify the right-hand side, use

$$a = \frac{\sigma_v^2}{\sigma_v^2 + \tau^2}, \quad b(\sigma_v^2 + \tau^2) = \sigma_v^2 + \rho\tau^2.$$

Then

$$\begin{aligned} MD(\sigma_v^2 + \tau^2) &= M\{2 + (M - 1)b\}(\sigma_v^2 + \tau^2) \\ &= 2M(\sigma_v^2 + \tau^2) + M(M - 1)(\sigma_v^2 + \rho\tau^2), \end{aligned}$$

whereas (34) gives

$$V = M(\sigma_v^2 + \tau^2) + M(M - 1)(\sigma_v^2 + \rho\tau^2).$$

Thus

$$MD(\sigma_v^2 + \tau^2) - V = M(\sigma_v^2 + \tau^2),$$

and therefore

$$M\sigma_v^2 D - aV = a\{MD(\sigma_v^2 + \tau^2) - V\} = M\sigma_v^2.$$

Equation (36) consequently reduces to

$$\lambda^2 D^2 \sigma_u^2 = aM\sigma_v^2.$$

Taking the unique positive root,

$$\lambda = \frac{\sqrt{aM\sigma_v^2}}{D\sigma_u}.$$

Using (33),

$$\beta = \frac{a}{\lambda D} = \sigma_u \sqrt{\frac{a}{M\sigma_v^2}},$$

which proves (9).

Finally, because (v, Y) are jointly Gaussian,

$$\begin{aligned} \text{var}(v \mid Y) &= \sigma_v^2 - \frac{\text{cov}(v, Y)^2}{\text{var}(Y)} \\ &= \sigma_v^2 - \frac{\beta^2 M^2 \sigma_v^4}{\beta^2 V + \sigma_u^2}. \end{aligned} \tag{37}$$

From the expression for β ,

$$\beta^2 = \frac{\sigma_u^2}{M(\sigma_v^2 + \tau^2)}, \quad \frac{\sigma_u^2}{\beta^2} = M(\sigma_v^2 + \tau^2).$$

Dividing numerator and denominator of the second term in (37) by β^2 ,

$$\text{var}(v \mid Y) = \sigma_v^2 - \frac{M^2 \sigma_v^4}{V + M(\sigma_v^2 + \tau^2)}.$$

Now

$$\begin{aligned} \frac{V}{M} &= (\sigma_v^2 + \tau^2) + (M - 1)(\sigma_v^2 + \rho\tau^2) \\ &= M(\sigma_v^2 + \rho\tau^2) + (1 - \rho)\tau^2, \end{aligned}$$

and

$$\sigma_v^2 + \tau^2 = (\sigma_v^2 + \rho\tau^2) + (1 - \rho)\tau^2.$$

Hence, after dividing the denominator by M ,

$$\tilde{D} = (M + 1)(\sigma_v^2 + \rho\tau^2) + 2(1 - \rho)\tau^2,$$

which gives (10). □

Proof of Proposition 2. Let n traders delegate and $M - n$ do not. Write a delegated signal as

$$s_i = v + \eta + \varepsilon_i^d, \quad \text{var}(\eta) = C, \quad \text{var}(\varepsilon_i^d) = A + g - C,$$

and an independent reading as

$$q_j = v + \varepsilon_j^o, \quad \text{var}(\varepsilon_j^o) = c + g.$$

The idiosyncratic errors are mutually independent and independent of (v, η) . Define

$$S_d = \text{var}(s_i) = \sigma_v^2 + A + g, \quad S_o = \text{var}(q_j) = \sigma_v^2 + c + g,$$

and

$$a_d = \frac{\sigma_v^2}{S_d}, \quad a_o = \frac{\sigma_v^2}{S_o}, \quad b = \frac{\sigma_v^2 + C}{S_d}.$$

For a delegator i ,

$$\text{cov}(v, s_i) = \sigma_v^2, \quad \text{cov}(s_k, s_i) = \sigma_v^2 + C \quad (k \neq i \text{ also delegating}),$$

while an own reader shares only the fundamental with the delegator,

$$\text{cov}(q_j, s_i) = \sigma_v^2.$$

Joint normality therefore gives

$$\mathbb{E}[v \mid s_i] = a_d s_i, \quad \mathbb{E}[s_k \mid s_i] = b s_i, \quad \mathbb{E}[q_j \mid s_i] = a_d s_i. \quad (38)$$

For an own reader j , her reading error is independent of every other reading error. Hence

$$\text{cov}(v, q_j) = \text{cov}(s_i, q_j) = \text{cov}(q_k, q_j) = \sigma_v^2,$$

and

$$\mathbb{E}[v \mid q_j] = \mathbb{E}[s_i \mid q_j] = \mathbb{E}[q_k \mid q_j] = a_o q_j. \quad (39)$$

Suppose delegators use $x_i = \beta_d s_i$, own readers use $x_j = \beta_o q_j$, and the market maker uses $p = \lambda Y$. Consider first a delegator i . There are $n - 1$ other delegators and $M - n$ own readers. Conditional on s_i ,

$$\begin{aligned} \mathbb{E}[(v - p)x_i \mid s_i] &= x_i \mathbb{E} \left[v - \lambda \left(x_i + \sum_{\substack{k \neq i \\ k \in d}} x_k + \sum_{j \in o} x_j + u \right) \mid s_i \right] \\ &= x_i a_d s_i - \lambda x_i^2 - \lambda x_i (n - 1) \beta_d b s_i - \lambda x_i (M - n) \beta_o a_d s_i \\ &= x_i \left\{ a_d - \lambda [(n - 1) \beta_d b + (M - n) \beta_o a_d] \right\} s_i - \lambda x_i^2. \end{aligned} \quad (40)$$

Differentiating with respect to x_i and imposing the type-symmetric strategy $x_i = \beta_d s_i$ gives

$$2\lambda \beta_d = a_d - \lambda [(n - 1) \beta_d b + (M - n) \beta_o a_d]. \quad (41)$$

Now consider an own reader j . She faces n delegators and $M - n - 1$ other own readers. Using (39),

$$\begin{aligned} \mathbb{E}[(v - p)x_j \mid q_j] &= x_j a_o q_j - \lambda x_j^2 - \lambda x_j n \beta_d a_o q_j - \lambda x_j (M - n - 1) \beta_o a_o q_j \\ &= x_j \left\{ a_o - \lambda [n \beta_d a_o + (M - n - 1) \beta_o a_o] \right\} q_j - \lambda x_j^2. \end{aligned} \quad (42)$$

Differentiating with respect to x_j and imposing $x_j = \beta_o q_j$ yields

$$2\lambda \beta_o = a_o - \lambda [n \beta_d a_o + (M - n - 1) \beta_o a_o]. \quad (43)$$

Define

$$\kappa_d = \lambda\beta_d, \quad \kappa_o = \lambda\beta_o.$$

Then (41)–(43) become the linear system

$$[2 + (n - 1)b]\kappa_d + (M - n)a_d\kappa_o = a_d, \quad (44)$$

$$na_o\kappa_d + [2 + (M - n - 1)a_o]\kappa_o = a_o. \quad (45)$$

Let $r \equiv \kappa_d/\kappa_o$. Divide both equations by κ_o . The first equation gives

$$[2 + (n - 1)b]r + (M - n)a_d = \frac{a_d}{\kappa_o},$$

and the second gives

$$na_or + 2 + (M - n - 1)a_o = \frac{a_o}{\kappa_o}.$$

Therefore

$$\begin{aligned} r &= \frac{(2 - a_o)a_d}{a_o[2 + (n - 1)b - na_d]} \\ &= \frac{(2 - a_o)S_o/S_d}{2 - b + nC/S_d}, \end{aligned} \quad (46)$$

where the second equality uses

$$\frac{a_d}{a_o} = \frac{S_o}{S_d}, \quad 2 + (n - 1)b - na_d = 2 - b + n(b - a_d), \quad b - a_d = \frac{C}{S_d}.$$

I next derive the profit expression before imposing technology indifference. The objective of either type has the generic form

$$x_th_tz_t - \lambda x_t^2,$$

where z_t is that type's signal. At the optimum the first-order condition gives $h_tz_t = 2\lambda x_t$. Substituting this back into the objective,

$$x_t(2\lambda x_t) - \lambda x_t^2 = \lambda x_t^2.$$

Taking expectations and using $x_t = \beta_t z_t$,

$$\pi_t = \lambda \mathbb{E}[x_t^2] = \lambda \beta_t^2 S_t = \frac{\kappa_t^2 S_t}{\lambda}, \quad t \in \{d, o\}. \quad (47)$$

Hence an interior continuous-composition equilibrium satisfies

$$\pi_d(n) = \pi_o(n) \iff \kappa_d^2 S_d = \kappa_o^2 S_o.$$

All trading intensities are positive, so

$$\frac{\kappa_d}{\kappa_o} = \sqrt{\frac{S_o}{S_d}}. \quad (48)$$

Combining (46) and (48),

$$\frac{(2 - a_o)S_o/S_d}{2 - b + nC/S_d} = \sqrt{\frac{S_o}{S_d}}.$$

Multiply by S_d and use $bS_d = \sigma_v^2 + C$:

$$\begin{aligned} (2 - a_o)\sqrt{S_o S_d} &= 2S_d - bS_d + nC \\ &= 2S_d - \sigma_v^2 - C + nC. \end{aligned}$$

Solving for n gives

$$n^* = \frac{(2 - a_o)\sqrt{S_o S_d} - 2S_d + \sigma_v^2 + C}{C},$$

which is (11).

To establish strategic substitution, combine (46) with (47):

$$\frac{\pi_d(n)}{\pi_o(n)} = \left(\frac{\kappa_d}{\kappa_o}\right)^2 \frac{S_d}{S_o} = \frac{(2 - a_o)^2 S_o/S_d}{(2 - b + nC/S_d)^2}. \quad (49)$$

Because $C > 0$, the denominator is increasing in n , so the profit ratio is decreasing in n . Delegation is therefore a strategic substitute and there can be at most one interior crossing.

At the precision boundary $c = A$, I have $S_o = S_d$ and $a_o = a_d$. Substituting into (11) gives

$$n^* = \frac{(2 - a_o)S_d - 2S_d + \sigma_v^2 + C}{C} = \frac{-a_o S_d + \sigma_v^2 + C}{C} = 1,$$

because $a_o S_d = \sigma_v^2$. Economically, a single delegator does not share the model error with any other trader. When $c = A$, her delegated signal therefore has the same precision as an independent reading, making the $n = 0$ and $n = 1$ allocations equivalent. The main text selects $n = 0$ at this knife edge. If $c < A$, delegation is less precise even for the first user, so no trader delegates. If $c > A$, delegation improves precision at entry and adoption begins. Hence the threshold is $c^* = A$.

Differentiating (11) with respect to c gives

$$C \frac{\partial n^*}{\partial c} = \frac{\sigma_v^2}{S_o^2} \sqrt{S_o S_d} + \frac{2 - a_o}{2} \sqrt{\frac{S_d}{S_o}} > 0. \quad (50)$$

Thus harder disclosures attract more delegation above the threshold. $\square$

Proof of Lemma 1. Start from the interior conditions in the proof of Proposition 2. Profit indifference (48) gives $\kappa_d/\kappa_o = \sqrt{S_o/S_d} = t$. Substituting $\kappa_d = t\kappa_o$ into the own-reader condition (45),

$$na_ot\kappa_o + [2 + (M - n - 1)a_o]\kappa_o = a_o,$$

and collecting terms,

$$\kappa_o[2 + a_o\{M + n(t - 1) - 1\}] = a_o, \quad \text{so} \quad \kappa_o = \frac{a_o}{2 + (m - 1)a_o}, \quad m = M + n(t - 1).$$

For the composition, use $\sqrt{S_o S_d} = tS_d$ and $a_otS_d = (\sigma_v^2/S_o)(S_o/t) = \sigma_v^2/t$ in (11):

$$n^* = \frac{(2 - a_o)tS_d - 2S_d + \sigma_v^2 + C}{C} = \frac{C + 2(t - 1)S_d + \sigma_v^2 - \sigma_v^2/t}{C} = 1 + \frac{(t - 1)(2S_d + \sigma_v^2/t)}{C}.$$

For the posterior variance, (21) gives

$$K = [n\kappa_d + (M - n)\kappa_o]\sigma_v^2 = \kappa_o m \sigma_v^2,$$

hence, by (22),

$$\text{var}(v \mid Y) = \sigma_v^2 - K = \sigma_v^2 \left[1 - \frac{ma_o}{2 + (m - 1)a_o} \right] = \sigma_v^2 \frac{2 - a_o}{2 + (m - 1)a_o}.$$

Setting $m = M$ reproduces the all-own market: from Proposition 1 with $\rho = 0$ and $\tau^2 = c + g$,

$$\text{var}(v \mid Y) = \sigma_v^2 - \frac{M\sigma_v^4}{(M - 1)\sigma_v^2 + 2S_o} = \sigma_v^2 \frac{2S_o - \sigma_v^2}{2S_o + (M - 1)\sigma_v^2} = \sigma_v^2 \frac{2 - a_o}{2 + (M - 1)a_o},$$

which proves the m -representation. Profit indifference $\lambda\beta_d^2 S_d = \lambda\beta_o^2 S_o$ implies $\beta_d \sqrt{S_d} = \beta_o \sqrt{S_o}$: delegators and independent readers submit orders with the same unconditional size. The effect of the equilibrium composition on posterior variance is therefore summarized by the effective reader mass m .

For price impact, write $Y = \beta_d \sum_{i \in d} s_i + \beta_o \sum_{j \in o} q_j + u$ and $\beta_t = \kappa_t/\lambda$. Then $\lambda^2 \text{var}(Y) =$

$Q + \lambda^2 \sigma_u^2$, where

$$Q = \kappa_d^2 [nS_d + n(n-1)(\sigma_v^2 + C)] + \kappa_o^2 [(M-n)S_o + (M-n)(M-n-1)\sigma_v^2] + 2\kappa_d\kappa_o n(M-n)\sigma_v^2.$$

Bayesian pricing gives $\lambda^2 \text{var}(Y) = \lambda \text{cov}(v, Y) = K$, so $\lambda^2 \sigma_u^2 = K - Q$. Substituting $\kappa_d = t\kappa_o$ and $t^2 S_d = S_o$,

$$Q = \kappa_o^2 W, \quad W = MS_o + \sigma_v^2 (m^2 - t^2 n - (M-n)) + t^2 n(n-1)C.$$

Write $D_m \equiv 2 + (m-1)a_o$, so $\kappa_o = a_o/D_m$ and $K = a_o \sigma_v^2 m/D_m$. Using $a_o S_o = \sigma_v^2$, $m - M = n(t-1)$, and $S_o/t^2 = S_d$,

$$a_o W = M\sigma_v^2 + a_o \sigma_v^2 (m^2 - m - t(t-1)n) + a_o t^2 n(n-1)C.$$

The interior composition satisfies $(n-1)C = (t-1)(2S_d + \sigma_v^2/t)$, hence $a_o t^2 n(n-1)C = n(t-1)(2\sigma_v^2 + a_o \sigma_v^2 t)$, and therefore

$$a_o W = M\sigma_v^2 + a_o \sigma_v^2 (m^2 - m) + 2\sigma_v^2 (m - M) = \sigma_v^2 (mD_m - M).$$

Then

$$\lambda^2 \sigma_u^2 = \frac{a_o \sigma_v^2 m}{D_m} - \frac{a_o^2 W}{D_m^2} = \frac{a_o}{D_m^2} (D_m \sigma_v^2 m - a_o W) = \frac{a_o M \sigma_v^2}{D_m^2},$$

which gives the expression for λ in (14). Since the all-own market has $\lambda_o = \sqrt{a_o M \sigma_v^2} / ((2 + (M-1)a_o)\sigma_u)$ and $\text{var}(v | Y)_o = \sigma_v^2(2 - a_o)/(2 + (M-1)a_o)$, both objects scale by $(2 + (M-1)a_o)/D_m$ at the interior equilibrium. $\square$

Proof of Proposition 3. Under a homogeneous reading technology (τ^2, ρ) , Proposition 1 implies

$$x_i = \beta s_i, \quad \beta = \frac{\sigma_u}{\sqrt{M(\sigma_v^2 + \tau^2)}}.$$

From (15), centered normality gives $E|Z| = \sqrt{2/\pi} \sqrt{\text{var}(Z)}$,

$$|OIB| = \frac{\sqrt{\beta^2 \text{var}(\sum_i s_i) + \sigma_u^2}}{M\beta \sqrt{\sigma_v^2 + \tau^2 + \sigma_u}}. \quad (51)$$

The signal sum has variance

$$\text{var}\left(\sum_{i=1}^M s_i\right) = M(\sigma_v^2 + \tau^2) + M(M-1)(\sigma_v^2 + \rho\tau^2).$$

Divide the numerator and denominator of (51) by $M\beta$. Since $M\beta \rightarrow \infty$ and

$$\frac{1}{M^2} \text{var} \left(\sum_{i=1}^M s_i \right) \rightarrow \sigma_v^2 + \rho\tau^2,$$

I obtain

$$|OIB| \rightarrow \sqrt{\frac{\sigma_v^2 + \rho\tau^2}{\sigma_v^2 + \tau^2}}. \quad (52)$$

For all-own reading, $\rho = 0$ and $\tau^2 = c + g$, which gives

$$|OIB|_o \rightarrow \sqrt{\frac{\sigma_v^2}{\sigma_v^2 + c + g}}.$$

For forced full delegation, $\rho\tau^2 = C$ and $\tau^2 = A + g$, which gives

$$|OIB|_d \rightarrow \sqrt{\frac{\sigma_v^2 + C}{\sigma_v^2 + A + g}}.$$

These are the two expressions in (18).

Full delegation is more one-sided when

$$\frac{\sigma_v^2 + C}{\sigma_v^2 + A + g} > \frac{\sigma_v^2}{\sigma_v^2 + c + g}.$$

Cross-multiplying gives

$$c(\sigma_v^2 + C) + gC > \sigma_v^2(A - C). \quad (53)$$

If $c > A$, then

$$c(\sigma_v^2 + C) > A(\sigma_v^2 + C) > \sigma_v^2(A - C),$$

so (53) holds. The full-delegation term is independent of c while the all-own term is decreasing in c , hence $\partial \Delta|OIB|^{FD} / \partial c > 0$.

For uncertainty, write

$$P = \sigma_v^2 + C, \quad Z = A + g - C, \quad X = c + g.$$

Then

$$\Delta|OIB|^{FD} = \frac{\sqrt{P}}{\sqrt{P+Z}} - \frac{\sigma_v}{\sqrt{\sigma_v^2 + X}},$$

and

$$\frac{\partial \Delta|OIB|^{FD}}{\partial g} = -\frac{\sqrt{P}}{2(P+Z)^{3/2}} + \frac{\sigma_v}{2(\sigma_v^2 + X)^{3/2}} < 0. \quad (54)$$

The inequality follows because $P > \sigma_v^2$ and $c > A$ implies $\sigma_v^2 + X > P + Z = \sigma_v^2 + A + g$. $\square$

Proof of Proposition 4. Under all-own reading, $\rho = 0$ and $\tau^2 = c + g$. Equation (10) therefore uses

$$\tilde{D}_o = (M+1)\sigma_v^2 + 2(c+g).$$

Under forced full delegation, $\rho\tau^2 = C$ and $(1-\rho)\tau^2 = A+g-C$, so

$$\tilde{D}_d = (M+1)(\sigma_v^2 + C) + 2(A+g-C).$$

Subtracting,

$$\tilde{D}_d - \tilde{D}_o = 2A + (M-1)C - 2c. \quad (55)$$

From Proposition 1,

$$\text{var}(v | Y) = \sigma_v^2 - \frac{M\sigma_v^4}{\tilde{D}},$$

which is increasing in $\tilde{D}$. Hence forced full delegation improves price discovery if and only if the right-hand side of (55) is negative, or

$$c > A + \frac{M-1}{2}C \equiv c^\dagger.$$

The uncertainty term g cancels from the threshold.

To study magnitudes, let

$$V_o \equiv \text{var}(v | Y)_o, \quad V_d \equiv \text{var}(v | Y)_d.$$

Then

$$V_d - V_o = \frac{M\sigma_v^4(\tilde{D}_d - \tilde{D}_o)}{\tilde{D}_o\tilde{D}_d}. \quad (56)$$

Above $c^\dagger$ this difference is negative. As c rises, the numerator becomes more negative while $\tilde{D}_d$ is unchanged; equivalently, V_d is independent of c and V_o increases in c . Therefore the pricing-error reduction $\mathcal{P}_o - \mathcal{P}_d = \sqrt{V_o} - \sqrt{V_d}$ increases with c . As g rises, the numerator in (56) is unchanged while both denominator terms increase, so $V_d - V_o$ moves toward zero.

Since both V_o and V_d also increase with g , the identity

$$\sqrt{V_o} - \sqrt{V_d} = \frac{V_o - V_d}{\sqrt{V_o} + \sqrt{V_d}}$$

shows that the reduction in pricing error shrinks with g and converges to zero as $g \rightarrow \infty$. $\square$

Proof of Corollary 1. The parameter $a_o = \sigma_v^2/S_o$ depends only on the own-reading technology and is therefore the same in the equilibrium and the all-own benchmark. By Lemma 1, both posterior variances take the form $\sigma_v^2(2 - a_o)/(2 + (m - 1)a_o)$, which is decreasing in the effective mass m . The benchmark sits at $m = M$; the interior equilibrium sits at $m = M + n^*(t - 1) \geq M$, with equality only when no trader delegates, that is, when $c \leq A$. This proves the claim whenever $n^* < M$. If instead $n^* \geq M$, the equilibrium is full delegation, and the closed form of n^* bounds the parameters:

$$n^* \geq M \implies (M - 1)C \leq (t - 1) \left(2S_d + \frac{\sigma_v^2}{t} \right) < 2S_d(t^2 - 1) = 2(c - A),$$

where the strict inequality follows from $\sigma_v^2/t < 2S_d t$, equivalent to $\sigma_v^2 < 2S_o$, which always holds. Hence the corner requires $c > A + \frac{1}{2}(M - 1)C = c^\dagger$, and Proposition 4 shows that full delegation lowers pricing error in that region. $\square$

Proof of Proposition 5. Consider two otherwise identical firms $f \neq \ell$ with independent fundamentals. Let the component of delegated reading error carried across the two firms have covariance χ^2 . Under all-own reading, own-reading errors and fundamentals are independent across firms, so

$$\text{cov}(Y_f, Y_\ell) = 0 \implies \varsigma_o = 0.$$

Under full delegation, every one of the M speculators in a firm loads on the cross-firm component with trading coefficient β . Hence

$$\text{cov}(Y_f, Y_\ell) = (M\beta)^2 \chi^2. \tag{57}$$

For delegated signals,

$$\text{var} \left(\sum_{i=1}^M s_i \right) = M(\sigma_v^2 + A + g) + M(M - 1)(\sigma_v^2 + C),$$

so

$$\text{var}(Y_f) = \beta^2 [M(\sigma_v^2 + A + g) + M(M - 1)(\sigma_v^2 + C)] + \sigma_u^2.$$

Proposition 1 gives

$$\beta^2 = \frac{\sigma_u^2}{M(\sigma_v^2 + A + g)}, \quad \frac{\sigma_u^2}{\beta^2 M} = \sigma_v^2 + A + g.$$


Therefore, using (17),

$$\begin{aligned} \varsigma_d &= \frac{M^2 \beta^2 \chi^2}{\text{var}(Y_f)} \\ &= \frac{M \chi^2}{(\sigma_v^2 + A + g) + (M - 1)(\sigma_v^2 + C) + \sigma_u^2 / (\beta^2 M)} \\ &= \frac{M \chi^2}{(M + 1)(\sigma_v^2 + C) + 2(A + g - C)}, \end{aligned}$$

which, together with $\varsigma_o = 0$, is (20). The expression is linear and increasing in χ^2 , independent of c , and decreasing in g . $\square$


Figures and tables

Figure 1: Event study for the synchronization result




Notes: Coefficients β_{dk} from $y_{im} = \alpha_i + \gamma_m + \sum_d \sum_{k \neq -6} \beta_{dk}(\mathbf{1}\{K_m = k\} \times D_{di}) + \sum_{k \neq -6} \psi'_k(\mathbf{1}\{K_m = k\} \times Z_i) + \theta' Act_{im} + e_{im}$, estimated in one regression. The dependent variable is dollar $|OIB|$ in raw units of the imbalance share. D_{di} is the firm's standardized pre-launch readability, uncertainty, or filing length, all three entered jointly, and the figure plots the readability and uncertainty coefficients. Event time K_m is measured in three-month bins from the launch; bin 0 runs from December 2022 through February 2023. Event time is clipped at six bins on each side, and bin -6, the earliest bin, is the omitted category. Z_i contains the nine firm characteristics, Act_{im} contains log dollar volume and the log number of intraday intervals with trading activity, and α_i and γ_m are firm and month fixed effects. Bars are 95 percent confidence intervals from standard errors clustered by firm and month. Joint tests that the five estimated pre-launch coefficients are zero give $p = 0.195$ for readability and $p = 0.781$ for uncertainty. The sample contains 148,984 firm-months for 3,670 firms from January 2021 through December 2025.

Figure 2: The readability differential by quarter




Notes: Coefficients β_{Rk} on readability $\times$ three-month-interval indicators from $y_{im} = \alpha_i + \gamma_m + \sum_d \sum_{k \neq -1} \beta_{dk}(\mathbf{1}\{Q_m = k\} \times D_{di}) + \sum_z \sum_{k \neq -1} \psi_{zk}(\mathbf{1}\{Q_m = k\} \times Z_{zi}) + \theta' Act_{im} + e_{im}$, estimated for dollar $|OIB|$ with both disclosure dimensions entered jointly. Event time Q_m is measured in three-month intervals from the ChatGPT launch and is clipped at six intervals before and twelve after, so the leftmost point pools the earlier pre-launch months; interval -1 is the omitted category. Z_{zi} are the ten predetermined controls, Act_{im} contains the two contemporaneous activity controls, and α_i and γ_m are firm and month fixed effects. The shaded area shows 95 percent confidence intervals from standard errors clustered by firm and month. Vertical lines mark the ChatGPT launch and six major model releases. The sample contains 148,984 firm-months for 3,670 firms from January 2021 through December 2025.

Figure 3: The price path after one-sided flow days



Notes: Each line plots the mean cumulative market-adjusted return after the flow days in its cell, at horizons h from 0 to 60 trading days. A flow day is a firm-day whose absolute dollar imbalance is in the top decile of the panel, $|OIB|$ of at least 0.295, so that the larger side accounts for at least 65 percent of classified dollar volume. The return at horizon h is the sum over days $[0, +h]$ of the firm's daily log return minus the equal-weighted log return of the sample firms, in percent, winsorized at the 1st and 99th percentiles across the panel and multiplied by -1 on net-selling days. The easy and hard terciles are the lowest and highest thirds of the panel's firm-days by the firm's pre-launch readability, and the period splits the flow days at the launch of ChatGPT on November 30, 2022. The daily panel contains the firm-days of U.S. common stocks with a lagged price of at least five dollars from January 2021 through December 2025, 3,199,135 firm-days for 4,034 firms. The cells contain 90,764 flow days (easy tercile, before the launch), 86,753 (easy, after), 26,875 (hard, before), and 31,016 (hard, after); the sixty-day horizon loses the last three months of the sample.

Figure 4: Numerical evaluation of the model



Notes: Each panel evaluates the model at the endogenous composition n^* , the mass of delegating traders at which a delegating trader and an independent reader earn the same expected profit. The solid curve uses the bottom and left axes and varies decoding difficulty c with uncertainty fixed at $g = 0$. The dashed curve uses the top and right axes and varies uncertainty g , holding $c = 1$ in panels (a) and (c) and $c = 3$ in panel (b). Panel (a) plots the change in within-stock one-sidedness, panel (b) the change in pricing error, $\mathcal{P} = \sqrt{\text{var}(v | Y)}$, and panel (c) the change in cross-stock order-flow commonality. The gray dashed curve in panel (b) is the forced full-delegation benchmark, in which every trader delegates. Dotted vertical lines mark the entry threshold $c^* = 0.23$ and, in panel (b), the price-quality boundary $c^\dagger = 1.14$. The calibration sets $\sigma_v^2 = 1$, $\sigma_u^2 = 4$, $M = 20$, $A = 0.23$, $C = 0.096$, and $\chi^2 = 0.0235$, under which every trader delegates above approximately $c = 2.1$.

Table 1: Frontier model releases used in $N(m)$

Date	Model	Developer
2022-11-30	ChatGPT (GPT-3.5), the launch	OpenAI
2023-03-14	GPT-4	OpenAI
2023-07-11	Claude 2	Anthropic
2023-11-06	GPT-4 Turbo	OpenAI
2023-12-06	Gemini 1.0 Ultra	Google DeepMind
2024-02-15	Gemini 1.5 Pro	Google DeepMind
2024-03-04	Claude 3 Opus	Anthropic
2024-04-18	Llama 3	Meta
2024-05-13	GPT-4o	OpenAI
2024-06-20	Claude 3.5 Sonnet	Anthropic
2024-08-13	Grok-2	xAI
2024-09-12	o1-preview	OpenAI
2024-12-11	Gemini 2.0	Google DeepMind
2025-01-20	DeepSeek-R1	DeepSeek
2025-02-24	Claude 3.7 Sonnet	Anthropic
2025-02-27	GPT-4.5	OpenAI
2025-03-25	Gemini 2.5 Pro	Google DeepMind
2025-04-05	Llama 4	Meta
2025-05-22	Claude 4	Anthropic
2025-07-09	Grok 4	xAI
2025-08-07	GPT-5	OpenAI
2025-11-18	Gemini 3 Pro	Google DeepMind
2025-11-24	Claude Opus 4.5	Anthropic

Notes: A release enters the calendar if it meets four criteria. First, the model is publicly available at launch through an interface that investors can use to analyze a filing. This excludes research announcements, benchmark-only systems, and models without public access. Second, the developer is at the capability frontier at the time of release. Third, the release introduces the developer’s leading general-purpose model, excluding size variants, distilled or lower-cost versions, and maintenance updates. Fourth, the release date is verified against Epoch AI’s Notable AI Models database.

Table 2: Summary statistics

	Mean	SD	P25	Median	P75
Panel A: absolute order imbalance, firm-month averages					
Absolute imbalance, dollar	13.1	10.7	7.2	9.6	13.7
non-retail flow	13.9	11.5	7.6	10.1	14.4
retail flow	24.4	18.7	11.3	17.8	31.0
Absolute imbalance, trade count	10.1	8.1	5.5	7.5	10.9
Panel B: disclosure dimensions, one observation per firm					
Readability	0.05	0.96	−0.58	0.10	0.73
Uncertainty	0.02	1.01	−0.66	−0.24	0.33
uncertainty words (percent)	1.38	0.41	1.10	1.27	1.51
Panel C: ten predetermined controls, one observation per firm					
Filing length	0.04	1.00	−0.68	−0.03	0.67
words per filing (thousands)	34.9	18.6	22.0	30.2	42.4
Market equity (\$mn)	7,118	22,151	172	727	3,383
Price (\$)	43.4	69.8	6.0	17.5	47.1
Turnover (percent per day)	1.37	2.57	0.40	0.70	1.25
Idiosyncratic volatility (percent)	0.33	0.41	0.06	0.13	0.45
Effective spread (bps)	48.0	62.4	8.2	18.7	63.1
Momentum (return over the window)	−0.22	0.41	−0.53	−0.21	0.02
Retail share of volume (percent)	8.9	8.5	3.2	5.0	11.8
Market beta	1.02	0.63	0.55	0.95	1.42
AI-sector beta	0.02	0.36	−0.15	0.02	0.21
Panel D: activity controls, firm-month					
Log dollar volume	16.15	2.31	14.61	16.31	17.82
Log intraday intervals	4.26	0.28	4.34	4.36	4.36

Notes: The sample contains 151,022 firm-months for 3,876 firms from January 2021 through December 2025. It requires a lagged price of at least five dollars, at least fifteen trading days in the month, and nonmissing values for readability, uncertainty, and length. Panel A reports the firm-month average of absolute daily order imbalance, in percentage points, for total dollar flow, for its non-retail and retail components, and for total trade-count flow. Panel B reports the two disclosure dimensions, constructed as firm-level averages across 10-K and 10-Q filings from January 2021 through November 29, 2022 and standardized across firms; the memo row gives the uncertainty measure in raw units. Panel C reports the ten predetermined controls whose interactions with the post indicator enter every specification: filing length, constructed from the same filings as the disclosure dimensions and standardized, with the raw word count as a memo row, and nine firm characteristics computed from daily TAQ and CRSP data from January 1, 2022 through November 29, 2022 and shown in raw units. Market equity and price are averages of daily values and enter the regressions in logs. Turnover is average daily dollar volume divided by average market equity. Idiosyncratic volatility is the average daily volatility of residual returns from the factor model. Effective spread is the average daily percent effective spread. Momentum is the cumulative return over the window. Retail share is average daily retail dollar volume, classified as in [Boehmer et al. \(2021\)](#), divided by average daily total dollar volume. Market beta is the slope of the firm’s daily return on the equal-weighted return of sample firms, and AI-sector beta is the slope on the value-weighted return of sample firms in SIC 7372 and 3674 after removing its market component. When these characteristics enter the regressions, each is standardized across firms. Panel D reports the two contemporaneous firm-month activity controls, log dollar volume and the log number of intraday intervals with trading activity. All variables are winsorized at the 1st and 99th percentiles.

Table 3: Firm characteristics by readability and uncertainty tercile

	Readability tercile			Uncertainty tercile		
	Easy	Middle	Hard	Low	Middle	High
Panel A: disclosure text						
Words per sentence	33.6	36.7	39.7	37.7	36.9	36.3
Complex-word share (percent)	28.4	28.8	29.1	28.8	28.9	28.7
Uncertainty words (percent)	1.32	1.31	1.20	1.04	1.27	1.75
Words per filing (thousands)	28.4	31.3	30.6	27.1	27.5	38.6
Panel B: pre-launch firm characteristics						
Market equity (\$mn)	213	814	2,756	837	721	667
Price (\$)	7.7	16.6	39.7	22.2	19.3	12.8
Turnover (percent per day)	0.68	0.76	0.68	0.62	0.68	0.85
Idiosyncratic volatility (percent)	0.36	0.12	0.07	0.11	0.13	0.15
Effective spread (bps)	51.9	17.6	8.8	16.5	18.9	21.7
Momentum (return over the window)	−0.35	−0.27	−0.11	−0.13	−0.18	−0.38
Retail share of volume (percent)	9.6	5.3	3.7	4.8	4.7	5.8
Market beta	1.03	1.11	0.80	0.78	0.88	1.26
AI-sector beta	−0.05	0.05	0.08	0.03	0.01	0.02
Panel C: analyst coverage						
Analysts per announcement	1.8	3.2	4.3	2.7	2.8	3.1
Share of firms with coverage (percent)	72	88	91	77	83	91

Notes: Each cell reports the median across the firms in the tercile, except the last row, which reports a share. The sample contains one observation per firm for the 3,876 firms with nonmissing readability, uncertainty, and length in the January 2021 through December 2025 sample. Readability is the standardized sum of the Flesch–Kincaid grade level and the Gunning Fog index, and uncertainty is the frequency of words in the Loughran–McDonald uncertainty list, each averaged over the firm’s 10-K and 10-Q filings from January 1, 2021 through November 29, 2022 and standardized across firms. Words per sentence and the complex-word share, the share of words with three or more syllables, come from the WRDS SEC Analytics Suite parse; uncertainty words and words per filing come from the Loughran–McDonald 10-X Summaries; all four are firm averages over the same filings. Panel B reports the nine pre-launch characteristics computed from daily TAQ and CRSP data from January 1, 2022 through November 29, 2022, in raw units and winsorized at the 1st and 99th percentiles: market equity and price are averages of daily values; turnover is average daily dollar volume divided by average market equity; idiosyncratic volatility is the average daily volatility of residual returns from the factor model; effective spread is the average daily percent effective spread; momentum is the cumulative return over the window; retail share is average daily retail dollar volume divided by average daily total dollar volume; market beta is the slope of the firm’s daily return on the equal-weighted return of sample firms, and AI-sector beta the slope on the value-weighted return of sample firms in SIC 7372 and 3674 after removing its market component. Analysts per announcement is the mean number of IBES estimates outstanding at the firm’s earnings announcements from January 2021 through November 29, 2022; the last row is the share of firms with at least one. Across firms, the correlation between the readability and uncertainty measures is -0.18 .

Table 4: Disclosure difficulty and synchronized trading

	Mean	Readability	Uncertainty
Panel A: order imbalance, by measure and investor type			
Dollar $ OIB $, total flow	13.1	+0.29***	−0.09**
each dimension entered alone		+0.30***	−0.14***
non-retail flow	13.9	+0.26***	−0.07
retail flow	24.4	+0.24**	−0.37***
Trade-count $ OIB $, total flow	10.1	+0.15***	−0.06
Signed dollar OIB , total flow		+0.02	+0.03
Panel B: agreement between retail and non-retail flow			
Share of days with the same sign	49.9	+0.26***	+0.11

Notes: Each cell reports the coefficient β_d from $y_{im} = \alpha_i + \gamma_m + \sum_{d \in \{R,U\}} \beta_d(Post_m \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, estimated with the row's variable as y_{im} and expressed in percentage points: the post-launch differential per standard deviation of the dimension named in the column. $Post_m$ equals one from December 2022 onward, D_{di} is the standardized pre-launch text measure, Z_i contains the ten predetermined controls (filing length and nine firm characteristics), Act_{im} contains log dollar volume and the log number of intraday intervals with trading activity, and α_i and γ_m are firm and month fixed effects. Standard errors are clustered by firm and month. Both disclosure dimensions enter jointly except in the memo row of Panel A. The Mean column reports the firm-month mean of the dependent variable, in percentage points. Panel A varies the imbalance measure and the investor group. Panel B uses the share of trading days in the month on which retail and non-retail signed flow have the same sign. The sample contains 148,984 firm-months for 3,670 firms. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 5: The response across the difficulty distribution

Decile	Readability	Uncertainty
2	+0.12	+0.18
3	+0.54***	−0.11
4	+0.44**	−0.12
5	+0.50**	−0.20
6	+0.70***	−0.31*
7	+0.67***	+0.03
8	+0.58***	−0.06
9	+0.71***	−0.29*
10	+0.84***	−0.49**

Notes: Coefficients β_d from $y_{im} = \alpha_i + \gamma_m + \sum_{d=2}^{10} \beta_d(Post_m \times \mathbf{1}\{decile_i = d\}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, where y_{im} is the firm-month mean of dollar $|OIB|$ for total flow, in percentage points, and each column uses the deciles of one disclosure dimension across the firm-months of the sample. The first decile is omitted, so each coefficient is the post-launch change for that decile relative to the easiest decile of the corresponding dimension. All specifications include firm and month fixed effects, post interactions with the ten predetermined controls, and controls for log dollar volume and log intraday intervals. Standard errors are clustered by firm and month. The sample contains 148,984 firm-months for 3,670 firms from January 2021 through December 2025. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 6: The timing of the differential

Panel A: the year path of the differential			
	2023	2024	2025
Dollar $ OIB $			
Readability	+0.19***	+0.28***	+0.42***
Uncertainty	−0.02	−0.13**	−0.10*
Trade-count $ OIB $			
Readability	+0.01	+0.16**	+0.36***
Uncertainty	−0.06	−0.09*	−0.02
Panel B: the launch step against the cumulative model stock			
	Stock alone	Entered together	
	$N(m) \times D$	$Post \times D$	$N(m) \times D$
Dollar $ OIB $			
Readability	+0.49***	+0.14**	+0.36***
Uncertainty	−0.14**	−0.05	−0.10
Trade-count $ OIB $			
Readability	+0.45***	−0.04	+0.49***
Uncertainty	−0.01	−0.09**	+0.07

Notes: Panel A reports coefficients β_{yd} from $y_{im} = \alpha_i + \gamma_m + \sum_y \sum_d \beta_{yd}(\mathbf{1}\{year_m = y\} \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$ for $y \in \{2023, 2024, 2025\}$, with the pre-launch months of 2021 and 2022 as the base period. Panel B reports coefficients β_1 and β_2 from $y_{im} = \alpha_i + \gamma_m + \beta_1(Post_m \times C_i) + \beta_2(N(m) \times C_i) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$; its first column includes only the model-stock interactions, and its last two columns include both the post-launch and model-stock interactions. In both panels, D_{di} and C_i are the firm's standardized pre-launch readability, uncertainty, and filing length, all three entered jointly, and the table reports the readability and uncertainty coefficients. y_{im} is the firm-month mean of absolute dollar or trade-count imbalance in percentage points, $Post_m$ equals one from December 2022 onward, $N(m)$ is the cumulative count of frontier language-model releases from the ChatGPT launch through month m , twenty-three in all, scaled to equal one in the final month of the sample, Z_i contains the nine firm characteristics, Act_{im} contains log dollar volume and the log number of intraday intervals with trading activity, and α_i and γ_m are firm and month fixed effects. Standard errors are clustered by firm and month. The sample contains 148,984 firm-months for 3,670 firms from January 2021 through December 2025. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 7: Analyst forecast dispersion

Dependent variable	Readability	Uncertainty
Log dispersion, scaled by price	−0.051**	+0.060**
Dispersion, scaled by price (percent)	−0.025**	+0.003
Log dispersion, scaled by median forecast	−0.064**	+0.020

Notes: Coefficients β_d from $y_{ie} = \alpha_i + \gamma_{q(e)} + \sum_{d \in \{R,U\}} \beta_d (Post_e \times D_{di}) + \psi'(Post_e \times Z_i) + \theta \log N_{ie} + u_{ie}$, estimated at the announcement level with both disclosure dimensions entered jointly. Here y_{ie} is the dispersion measure named in each row for firm i 's announcement e , α_i and $\gamma_{q(e)}$ are firm and calendar-quarter fixed effects, D_{di} is the standardized pre-launch text measure named in the column, Z_i contains the ten predetermined controls, and N_{ie} is the number of forecasts outstanding. Dispersion is the standard deviation of the IBES forecasts outstanding before the announcement, scaled by the share price or by the absolute median forecast and winsorized at the 1st and 99th percentiles. Announcements with fewer than two forecasts are excluded. $Post_e$ equals one for announcements from December 2022 onward. Standard errors are clustered by firm and calendar quarter. The sample contains 34,255 announcements from 2,870 firms. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 8: Price impact and its reversal

Return window	Before the launch			Change after the launch		
	Average firm κ_1	Readability κ_{3R}	Uncertainty κ_{3U}	Average firm κ_2	Readability κ_{4R}	Uncertainty κ_{4U}
Panel A: signed dollar imbalance						
[0], impact	+2.288***	+0.245***	+0.564***	+0.416***	−0.177***	+0.101
[+1]	+0.000	−0.017	+0.037*	−0.055**	+0.003	−0.042
[+1, +5]	−0.055	−0.029	+0.031	−0.001	+0.053	−0.078
[+2, +21]	−0.410***	−0.024	+0.163	+0.152	+0.135	−0.498***
[+2, +60]	−0.707***	−0.268*	+0.385	−0.112	+0.568**	−0.986***
Panel B: the shared and firm-specific parts of signed imbalance						
[0], shared	+2.609***	+0.033	+0.607***	+0.425***	−0.176**	+0.034
[0], firm-specific	+2.580***	+0.066	+0.538***	+0.349***	−0.180***	+0.065
[+2, +60], shared	+0.012	−1.780**	+2.019*	−1.719	+2.988***	−3.638**
[+2, +60], firm-specific	−0.796***	−0.237	+0.272	−0.178	+0.671**	−0.938**

Notes: Coefficients from $CAR_{i,[a,b]} = \alpha_i + \gamma_t + OIB_{it}(\kappa_1 + \kappa_2 Post_t) + \sum_d OIB_{it}(\kappa_{3d} D_{di} + \kappa_{4d} Post_t \times D_{di}) + \phi' W_{it} + e_{it}$, estimated separately for each return window with both disclosure dimensions entered jointly. The dependent variable $CAR_{i,[a,b]}$ is the cumulative market-adjusted log return over trading days $[t+a, t+b]$, the stated window relative to the day t of the flow: the sum over those days of the firm's daily log return minus the equal-weighted log return of the sample firms, in percent; [0] is the same-day return and measures price impact. $Post_t$ equals one from December 2022 onward, and D_{di} is the standardized pre-launch text measure. OIB_{it} is the signed Lee–Ready dollar imbalance $(B - S)/(B + S)$ for total flow, so a coefficient is the change in the percentage return associated with moving from balanced to fully one-sided flow. The first three columns describe the slope before the launch: κ_1 on OIB_{it} is the slope for a firm with both dimensions at zero, their mean across firms, and κ_{3d} on $OIB_{it} \times D_{di}$ is its difference per standard deviation of the dimension named in the column. The last three describe the change after the launch: κ_2 on $OIB_{it} \times Post_t$ for a firm with both dimensions at zero, and κ_{4d} on $OIB_{it} \times Post_t \times D_{di}$ per standard deviation of the dimension named in the column. The slope for a firm one standard deviation above the mean of dimension d , with the other dimension at zero, is $\kappa_1 + \kappa_{3d}$ before the launch and $\kappa_1 + \kappa_2 + \kappa_{3d} + \kappa_{4d}$ after it. W_{it} contains the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and the corresponding interactions with filing length; α_i and γ_t are firm and date fixed effects. Standard errors are clustered by firm and date. The panel contains the firm-days of U.S. common stocks with a lagged price of at least five dollars from January 2021 through December 2025, 3,189,685 firm-days for 4,034 firms at the impact window; the 60-day window loses the last three months of the sample. Panel B splits OIB_{it} into a shared part and a firm-specific part. Within each firm-month with at least fifteen trading days, daily signed imbalance is regressed on the equal-weighted signed imbalance of all other sample firms that day; the shared part is the fitted slope times that day's imbalance of the other firms, and the firm-specific part is the remainder, each winsorized at the 1st and 99th percentiles. The two parts enter the regression in place of OIB_{it} , each with its own set of coefficients, on 3,118,077 firm-days at the impact window. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 9: Earnings-announcement tests for the two text dimensions

Dependent variable	Readability	Uncertainty
Abnormal volume	+0.0192**	+0.0140**
Incorporation speed	+0.0079**	+0.0028
Drift	−0.0141**	−0.0051

Notes: Coefficients β_d from $y_{ie} = \alpha_i + \gamma_q + \sum_d \beta_d (Post_e \times D_{di}) + \theta' X_{ie} + u_{ie}$, estimated at the announcement level with both dimensions entered jointly; for drift, the coefficient on the announcement-return rank interacted with $Post_e \times D_{di}$. Abnormal volume is mean log dollar volume over days $[-1, +1]$ minus its mean over days $[-30, -6]$. Incorporation speed is $|CAR_0|/(|CAR_0| + |DRIFT|)$, where CAR_0 is measured over $[-1, +1]$ and $DRIFT$ over $[+2, +21]$. The announcement-return rank is the within-quarter percentile rank of CAR_0 , scaled to $[-0.5, 0.5]$. $Post_e$ equals one for announcements from December 2022 onward, and D_{di} is the standardized pre-launch text measure named in the column. The vector X includes the announcement-return rank, filing length interacted with the post indicator, and lagged size, volatility, and volume; the drift specification additionally interacts the rank with each disclosure dimension and the post indicator. All specifications include firm and calendar-quarter fixed effects. Standard errors are clustered by firm and announcement day. The sample contains 39,098 announcements from 3,100 firms. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 10: Cross-stock flow commonality

	Readability	Uncertainty
Panel A: the post-launch differential		
Total flow	+0.105***	−0.009
Non-retail	+0.091***	−0.010
Retail	+0.033	−0.043
Panel B: the year path of the differential		
Total flow		
2023	+0.148***	−0.010
2024	+0.147***	−0.013
2025	−0.036	+0.035
Non-retail		
2023	+0.136***	−0.007
2024	+0.139***	−0.021
2025	−0.035	+0.028
Retail		
2023	+0.028	−0.076
2024	+0.089	−0.019
2025	−0.044	−0.033

Notes: Panel A reports coefficients δ_d from $b_{im} = \alpha_i + \gamma_m + \sum_{d \in \{R,U\}} \delta_d(Post_m \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + \nu_{im}$, and Panel B coefficients δ_{yd} from $b_{im} = \alpha_i + \gamma_m + \sum_y \sum_d \delta_{yd}(\mathbf{1}\{year_m = y\} \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + \nu_{im}$ with the pre-launch months of 2021 and 2022 as the base period, each for the flow type named in the row with readability and uncertainty entered jointly. The dependent variable b_{im} is firm-month commonality in signed order flow, the slope from a within-month regression of firm i 's daily signed OIB on the equal-weighted signed OIB of all other sample firms that day, in months with at least fifteen trading days. $Post_m$ equals one from December 2022 onward, and D_{di} is the standardized pre-launch text measure. Both panels include firm and month fixed effects, post interactions with the ten predetermined controls, and controls for log dollar volume and log intraday intervals. Standard errors are clustered by firm and month. The samples contain 149,413 total-flow, 149,400 non-retail, and 146,978 retail firm-months from January 2021 through December 2025. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 11: The readability effects within halves of the sample

	One-sidedness	Quarterly reversal	Commonality
Retail share, low	+0.20***	−0.19	+0.195***
Retail share, high	+0.40***	+0.94**	+0.001
Size, small	+0.66***	+0.55*	+0.234***
Size, large	+0.02	−0.24	+0.047**
Coverage, low	+0.53***	+0.43	+0.170**
Coverage, high	+0.04	+0.84	+0.051*

Notes: Each cell reports the post-launch differential per standard deviation of readability, re-estimated within the stated half of the sample, with both disclosure dimensions entered jointly and *Post* equal to one from December 2022 onward. Firms are split at the median of the pre-launch retail share of dollar volume, of pre-launch market equity, and of the average number of IBES forecasts outstanding before their pre-launch announcements, with firms without coverage in the low half. One-sidedness is the firm-month mean of absolute dollar imbalance, in percentage points, regressed on $Post \times D$ with firm and month fixed effects, post interactions with the ten predetermined controls, and the two activity controls. Quarterly reversal is the change in the slope of the cumulative market-adjusted log return over trading days $[+2, +60]$ on signed dollar imbalance, in percent per unit of imbalance, with firm and date fixed effects and the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread as controls; a positive value means less reversal of the price impact. Commonality is the firm-month slope of daily signed imbalance on the equal-weighted signed imbalance of the other sample firms, regressed on $Post \times D$ with the same fixed effects and controls as one-sidedness. Standard errors are clustered by firm and by month or date. The order-flow panel contains 148,984 firm-months and the daily panel 2,958,342 firm-days; the retail-share and size halves omit the few firms without pre-launch characteristics. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Appendix tables

Table A1: Estimates on the firms common to both panels

Test and coefficient	Full sample		Common firms	
	Readability	Uncertainty	Readability	Uncertainty
Dollar $ OIB $ on $Post \times D$	+0.285***	−0.093**	+0.285***	−0.093**
Trade-count $ OIB $ on $Post \times D$	+0.146***	−0.062	+0.146***	−0.062
$CAR_{[0]}$ on $OIB \times Post \times D$	−0.177***	+0.101	−0.183***	+0.105*
$CAR_{[+2,+60]}$ on $OIB \times Post \times D$	+0.568**	−0.986***	+0.546**	−1.043***

Notes: Coefficients on $Post \times D$ in the first two rows and on $OIB \times Post \times D$ in the last two, estimated on the full panels and on the 3,669 firms that appear in both the order-flow panel (3,670 firms) and the flow-return panel (4,009 firms), with 148,980 firm-months and the corresponding firm-days. The first two rows regress the firm-month mean of absolute dollar or trade-count imbalance, in percentage points, on $Post \times D$, with firm and month fixed effects, post interactions with the ten predetermined controls, and log dollar volume and log intraday intervals. The next two rows regress the cumulative market-adjusted log return over the stated window of trading days, in percent, on signed dollar imbalance and its interactions with $Post$ and D , with firm and date fixed effects, the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and their interactions with filing length. Both disclosure dimensions enter jointly, $Post$ equals one from December 2022 onward, and D is the standardized pre-launch text measure named in the column. Standard errors are clustered by firm and by month or date. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A2: Market-wide means before and after the launch

	Before	After
Panel A: absolute order imbalance, firm-month means in percentage points		
Dollar $ OIB $	13.79	12.50
Trade-count $ OIB $	10.68	9.68
Dollar $ OIB $, non-retail	14.71	13.35
Dollar $ OIB $, retail	26.99	22.42
Panel B: analyst forecast dispersion, announcement-level means		
Log dispersion, scaled by price	−2.18	−2.28
Dispersion, scaled by price (percent)	0.31	0.31

Notes: Means before and from December 2022. Panel A uses the firm-months with nonmissing imbalance measures, the three text measures, the nine firm characteristics, and the two activity controls, 148,984 firm-months for 3,670 firms (148,973 for retail flow). Panel B uses the announcements with at least two outstanding IBES forecasts and a share price, 34,606 announcements; dispersion is the standard deviation of the forecasts scaled by the share price, winsorized at the 1st and 99th percentiles.

Table A3: The two components of readability

$Post \times$ component	Dollar $ OIB $	Trade-count $ OIB $
Composite readability, alone	+0.30***	+0.17***
Words per sentence, alone	+0.15***	+0.07*
Complex-word share, alone	+0.27***	+0.19***
Both components:		
words per sentence	+0.15***	+0.07*
complex-word share	+0.27***	+0.18***
Both, with length and uncertainty:		
words per sentence	+0.11**	+0.05
complex-word share	+0.34***	+0.19***
uncertainty	-0.16***	-0.10**

Notes: Coefficients β_k on $Post_m \times$ the standardized firm-level component, one or both components entered, from $y_{im} = \alpha_i + \gamma_m + \sum_k \beta_k (Post_m \times Comp_{ki}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, where y_{im} is the firm-month mean of absolute dollar or trade-count imbalance in percentage points and $Post_m$ equals one from December 2022 onward. Words per sentence and the complex-word share are firm means over the exposure window from the WRDS parse, standardized across firms. All specifications include firm and month fixed effects, post interactions with the nine firm characteristics, and the activity controls; the last three rows add the post interactions with length and uncertainty, so that Z_i contains all ten predetermined controls, and the uncertainty row reports the coefficient on $Post_m \times$ uncertainty from that regression. Standard errors are clustered by firm and month. The sample contains 148,984 firm-months for 3,670 firms. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A4: Alternative constructions of readability

	In use (level)	Within-quarter rank	Before July 2022	Loughran– McDonald
Panel A: rank agreement with the exposure in use (Spearman)				
Correlation	1.00	0.97	0.95	0.40
Panel B: $Post \times C_i$				
Dollar $ OIB $	+0.30***	+0.25***	+0.23***	+0.14***
Trade-count $ OIB $	+0.16***	+0.11**	+0.08	+0.11***

Notes: Panel A reports the Spearman rank correlation of each exposure with the exposure in use, and Panel B the coefficient on $Post \times C_i$, in percentage points of $|OIB|$, with each exposure entered separately. Exposures are measured over the filings from January 2021 to November 29, 2022, except as stated, and standardized across firms. The baseline readability measure is the sum of the standardized firm means of Flesch–Kincaid grade level and Gunning Fog from the WRDS SEC Analytics Suite parse. The within-quarter rank measure assigns each filing its percentile rank among filings from the same quarter and then averages these ranks across the firm’s filings. The before-July-2022 measure uses the baseline construction on the filings dated before July 1, 2022, when the suite revised the sentence splitter behind these scores. The Loughran–McDonald measure is the firm’s mean complex-word share from the Loughran–McDonald 10-X Summaries, a parse independent of the WRDS suite. $Post$ equals one from December 2022 onward. All specifications include firm and month fixed effects, post interactions with the ten predetermined controls, and controls for log dollar volume and log intraday intervals. Standard errors are clustered by firm and month. The estimation sample contains 148,697 firm-months with nonmissing values for all four exposures, filing length, the nine firm characteristics, and both activity controls. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A5: Persistence of disclosure difficulty

Panel A: year-over-year Spearman rank correlation of firm-year means						
	19→20	20→21	21→22	22→23	23→24	24→25
Complex-word share	0.84	0.87	0.86	0.87	0.89	0.89
Log net file size	0.82	0.77	0.78	0.79	0.79	0.81
Log word count	0.81	0.76	0.77	0.79	0.78	0.80
Uncertainty share	0.82	0.79	0.79	0.83	0.83	0.85
Panel B: pre-launch firm mean against post-launch firm mean						
	Spearman		same or adjacent quintile (%)			
Complex-word share	0.81		89.9			
Log net file size	0.80		88.9			
Log word count	0.79		88.2			
Uncertainty share	0.83		90.0			
Panel C: Spearman correlation with the exposure-window ranking						
	2020		2023		2025	
Complex-word share	0.83		0.84		0.73	
Log net file size	0.76		0.79		0.73	
Log word count	0.75		0.79		0.73	
Uncertainty share	0.78		0.83		0.76	

Notes: The data come from the Loughran–McDonald 10-X Summaries and cover 10-K and 10-Q filings from 2019 through 2025, matched to the sample by accession number and CIK. Panels A and B use 110,435 filings from 4,778 firms. The pre-launch period ends on November 29, 2022, and the post-launch period begins on November 30, 2022. Panel C correlates each firm’s annual mean with its mean over the exposure window from January 2021 through November 2022 and uses 111,082 filings from 4,237 firms.

Table A6: Uncertainty language and uncertainty about fundamentals

Pre-launch measure	Firms	Uncertainty		Readability	
		Spearman	Partial	Spearman	Partial
Option-implied volatility	3,056	+0.27***	+0.30***	−0.58***	−0.36***
Volatility of return on assets	2,910	+0.19***	+0.19***	−0.44***	−0.29***
Idiosyncratic volatility	3,949	+0.05***	+0.07***	−0.53***	−0.09***
Analyst forecast dispersion	2,699	+0.08***	−0.01	−0.29***	−0.07***
Litigious-word share	3,949	+0.16***	+0.16***	+0.08***	−0.04**

Notes: Each cell reports a correlation across firms between the standardized pre-launch text measure named in the column and the pre-launch measure named in the row. Spearman is the rank correlation, and Partial is the correlation between the two ranks after each is regressed on the rank of log market equity. Uncertainty is the frequency of words in the Loughran–McDonald uncertainty list and readability the standardized sum of the Flesch–Kincaid grade level and the Gunning Fog index, each averaged over the firm’s 10-K and 10-Q filings from January 1, 2021 through November 29, 2022. Option-implied volatility is the mean 30-day at-the-money implied volatility of calls and puts on the OptionMetrics volatility surface over the same period, for firms with at least 60 daily observations. The volatility of return on assets is the standard deviation of quarterly income before extraordinary items divided by lagged total assets, over fiscal quarters ending from January 2016 through November 2022, for firms with at least twelve quarters. Idiosyncratic volatility is the average daily volatility of residual returns from the factor model, and log market equity the log of average daily market equity, both from January 1 through November 29, 2022. Analyst forecast dispersion is the standard deviation of the IBES forecasts outstanding before each earnings announcement from January 2021 through November 29, 2022, scaled by the share price and averaged over the firm’s announcements, using announcements with at least two forecasts and a price of at least five dollars. The litigious-word share is the frequency of words in the Loughran–McDonald litigious list in the same filings as the text measures. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A7: Estimates by change in vocabulary complexity

	Dollar $ OIB $		Trade-count $ OIB $	
	Readability	Uncertainty	Readability	Uncertainty
All firms with Loughran–McDonald ranks	+0.29***	−0.09**	+0.16***	−0.06
Firms that did not move	+0.26***	−0.09**	+0.14**	−0.05
Firms that moved	+0.48**	−0.11	+0.36*	−0.21
Excluding firms that became easier	+0.29***	−0.09**	+0.16***	−0.05
Excluding firms that became harder	+0.26***	−0.09**	+0.13**	−0.06

Notes: Coefficients β_d from $y_{im} = \alpha_i + \gamma_m + \sum_{d \in \{R,U\}} \beta_d (Post_m \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, estimated on the firms named in each row, where y_{im} is the firm-month mean of absolute dollar or trade-count imbalance in percentage points, $Post_m$ equals one from December 2022 onward, and D_{di} is the standardized pre-launch text measure. A firm moved if its within-year quintile of the Loughran–McDonald complex-word share, averaged over 2023–25, differs by two or more from its average over 2021–22; of the 3,844 firms with Loughran–McDonald ranks, 392 moved, 190 toward easier text and 202 toward harder. All specifications include firm and month fixed effects, post interactions with the ten predetermined controls, and the activity controls. Standard errors are clustered by firm and month. The rows contain 145,372, 134,710, 10,662, 139,778, and 140,304 firm-months in the order listed. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A8: Estimates without the retail trading episode of early 2021

	One-sidedness	Commonality	Quarterly reversal	Incorporation speed	Drift
Panel A: readability					
Full sample	+0.285***	+0.109***	+0.568**	+0.0079**	−0.0141**
Pre-launch period from July 2021	+0.292***	+0.128***	+0.650**	+0.0069	−0.0094
Without the top meme-intensity decile	+0.251***	+0.134***	+0.676***	+0.0075*	−0.0144**
With $Post \times$ meme intensity	+0.298***	+0.105***	+0.463*	+0.0079**	−0.0140**
Panel B: uncertainty					
Full sample	−0.093**	−0.011	−0.986***	+0.0028	−0.0051
Pre-launch period from July 2021	−0.104**	+0.001	−1.328***	+0.0012	−0.0006
Without the top meme-intensity decile	−0.111**	−0.030	−1.241***	+0.0032	−0.0050
With $Post \times$ meme intensity	−0.072	−0.017	−0.933**	+0.0028	−0.0052

Notes: Each cell reports the post-launch differential per standard deviation of the dimension named in the panel, with both disclosure dimensions entered jointly and $Post$ equal to one from December 2022 onward. The first row uses the full sample, January 2021 through December 2025. The second begins the sample in July 2021 in every test. The third drops the 320 firms in the top decile of meme intensity, the firm's retail share of dollar volume from January through June 2021 minus the same share from January 1 through November 29, 2022. The fourth keeps all firms and adds the interaction of meme intensity with $Post$ to the controls. One-sidedness is the firm-month mean of absolute dollar imbalance, in percentage points, regressed on $Post \times D$ with firm and month fixed effects, post interactions with the ten predetermined controls, and log dollar volume and log intraday intervals. Commonality is the firm-month slope of daily signed imbalance on the equal-weighted signed imbalance of the other sample firms, regressed on $Post \times D$ with the same fixed effects and controls. Quarterly reversal is the coefficient on $OIB \times Post \times D$ in a regression of the cumulative market-adjusted log return over trading days $[+2, +60]$, in percent, on signed dollar imbalance and its interactions with $Post$ and D , with firm and date fixed effects and the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and their interactions with filing length; a positive value means less reversal of the price impact. Incorporation speed is the absolute cumulative return over days $[-1, +1]$ around an earnings announcement divided by the sum of the absolute cumulative returns over $[-1, +1]$ and $[+2, +21]$, regressed on $Post \times D$, and drift is the coefficient on the within-quarter rank of the announcement return interacted with $Post \times D$ in a regression of the cumulative return over $[+2, +21]$; both are estimated at the announcement level with firm and calendar-quarter fixed effects, the announcement-return rank, filing length interacted with the post indicator, and lagged size, volatility, and volume. Standard errors are clustered by firm and by month, date, or announcement day. The full sample contains 148,984 firm-months, 2,958,342 firm-days at the quarterly window, and 39,098 announcements; the sample from July 2021 contains 132,222, 2,607,137, and 34,724; the sample without the top decile of meme intensity contains 137,950, 2,741,386, and 37,032; the fourth row uses 2,958,045 firm-days at the quarterly window. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A9: Length entered as a third dimension

	Readability	Length	Uncertainty
Panel A: specification, dependent variable total dollar $ OIB $			
Each dimension alone	+0.30***	−0.01	−0.13***
Three dimensions jointly	+0.29***	+0.05	−0.09**
Panel B: absolute imbalance, by measure and investor type			
Dollar $ OIB $, total flow	+0.29***	+0.05	−0.09**
non-retail	+0.26***	+0.03	−0.07
retail	+0.24**	−0.11	−0.37***
Trade-count $ OIB $, total flow	+0.15***	−0.05	−0.06
non-retail	+0.13**	−0.07	−0.06
retail	+0.06	−0.14*	−0.21***

Notes: Coefficients β_d from $y_{im} = \alpha_i + \gamma_m + \sum_{d \in \{R, L, U\}} \beta_d (Post_m \times D_{di}) + \psi' (Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, with length entered as a third dimension L and Z_i containing the nine firm characteristics, in percentage points of order imbalance. $Post_m$ equals one from December 2022 onward, and D_{di} is the standardized pre-launch text measure. All specifications include firm and month fixed effects, post interactions with the nine firm characteristics, and controls for log dollar volume and log intraday intervals. Standard errors are clustered by firm and month. The three dimensions enter jointly except in the first row of Panel A. The sample contains 148,984 firm-months for 3,670 firms. ***, **, and * denote significance at the 1, 5, and 10 percent levels.