Large Language Models, Price Discovery, and Synchronized Trading

Jiawei Wang*†

October 2026

Abstract

Large language models can improve investors' understanding of disclosures while making their errors more correlated. I study this trade-off around ChatGPT's public launch using pre-launch disclosure readability and uncertainty. Within stocks, trading becomes more synchronized: order flow turns more one-sided in the hardest-to-read decile than in the easiest, by 6.4 percent of the sample mean. Across stocks, the order flow of hard-to-read firms moves more with that of other stocks, by 13 percent of the pre-launch mean per standard deviation of readability difficulty. Price discovery also improves: relative to comparable firms, a larger share of the price impact of order flow becomes permanent, and earnings news enters prices faster. Greater disclosure uncertainty predicts less one-sided trading and no improvement in price discovery. A Kyle-style model accounts for these patterns in a single framework. Investors can delegate reading to a common language model that reduces decoding error, leaves uncertainty about outcomes unresolved, and introduces shared error. If every investor had to delegate, prices could become less informative even as each reads more precisely. When adoption is a choice, crowding limits it, and greater synchronization coexists 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

†I thank 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

Large language models (LLMs) can condense long, complex corporate filings into short summaries that preserve their informative content (Kim et al., 2023). Nearly half of surveyed retail investors use generative AI, mainly to gather and interpret financial information (Blankespoor et al., 2026). Among financial institutions surveyed in Hong Kong, three quarters have implemented, are piloting, or are exploring the technology (Hong Kong Institute for Monetary and Financial Research, 2025). As investors rely on these tools to interpret disclosures, their understanding may improve, but their mistakes may also become more correlated. The European Central Bank warns that widespread use of similar AI models can make interpretations more uniform, produce systematic biases, and distort asset prices (Leitner et al., 2024). The Financial Stability Board and the Bank of England likewise identify common models as a potential source of correlated trading and amplified market stress (Financial Stability Board, 2024; Bank of England, 2025).

The same trade-off arises with analysts and shared analytical tools: better individual judgments can come with common mistakes. 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 span investor classes and stocks.

I study how this shared reliance affects trading and price discovery around ChatGPT's public launch. The first question is whether trading becomes more synchronized within a stock. Investors using LLMs may reach similar interpretations of the same disclosure, and even different models make correlated errors on the same input (Kim et al., 2025; Huang et al., 2026). The second question is whether this alignment extends across stocks. A model can carry the same reading habits and blind spots from one document to the next, so its errors may recur across firms.[1] The third question is whether price discovery improves. Investors may trade alike because they extract the same information more accurately or because they share the same mistake. Synchronization alone cannot distinguish the two.

The value of a common reader should depend on what makes a disclosure difficult to understand. 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.[2] This distinction guides the empirical design. I measure readability using the

[1] Evaluations of language models reveal tendencies that recur across inputs. For example, Liu et al. (2024) find that models use information less effectively when it appears in the middle of a long input rather than at the beginning or end, in both multi-document question answering and key-value retrieval.

[2] Language models are unreliable on facts that are rare or absent in their training data (Kandpal et al.,

Flesch–Kincaid grade level and the Gunning Fog index, which rise with sentence length and word complexity, and uncertainty using the frequency of words in the Loughran–McDonald uncertainty list (Li, 2008; Loughran and McDonald, 2011). The first measure captures difficulty in decoding the language; the second proxies for uncertainty that remains once the language is understood. Both are computed from each firm's 10-K and 10-Q filings between January 1, 2021 and November 29, 2022, the day before ChatGPT's public launch, and held fixed.[3] Using U.S. common stocks from 2021 through 2025, I compare post-launch changes across firms with different pre-launch readability and uncertainty.

For firms with harder-to-read disclosures, trading becomes more synchronized within stocks and more correlated across stocks, relative to comparable firms. The post-launch change in one-sidedness is larger in the hardest-to-read decile than in the easiest by 6.4 percent of the sample mean. Order flow in hard-to-read firms comoves more with the market, by 13 percent of the pre-launch mean per standard deviation of readability difficulty. This greater alignment accompanies better price discovery: a larger share of the price impact of order flow is permanent, and earnings news enters prices faster. Greater disclosure uncertainty predicts a different pattern. The change in one-sidedness is smaller in the most uncertain decile than in the least uncertain by 3.8 percent of the sample mean, commonality does not change, and price discovery does not improve.

I measure synchronization within a stock 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 0.28 percentage points larger than for a comparable firm, or 2.2 percent of the sample mean. The result holds for dollar and trade-count imbalance and for retail and non-retail flow, and the two groups 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' forecast dispersion falls by five percent per standard deviation of readability difficulty.

The relation is nonlinear and strengthens over time. The two easiest readability deciles show little response; a clear differential appears in the third and generally increases there-

2023; Kalai and Vempala, 2024).

[3]Using pre-launch disclosures excludes subsequent changes in firms' writing that may respond to the technology itself. Firms adapt disclosures to machine readership (Cao et al., 2023), and generative AI is increasingly used in disclosure production (Blankespoor et al., 2026).

after. This pattern is consistent with investors using reading assistance only when the text is difficult enough to make it worthwhile. Event-study coefficients are flat before December 2022 and positive in every quarter afterward. The readability differential grows every year and more than doubles from 2023 to 2025, as frontier models accumulate. Synchronization increases most in retail-heavy, small, and thinly covered firms, where investors previously had less help interpreting disclosures.

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 0.10 percentage points smaller than for a comparable firm, and analysts' forecast dispersion rises by six percent. 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.

The common reader can also link trading across stocks. 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 0.090 for non-retail flow relative to other firms, about 13 and 11 percent of the pre-launch average. The effect appears in 2023 and 2024 but fades in 2025, even as within-stock synchronization strengthens, possibly because investors spread across model families: models that share an architecture or provider make more similar mistakes (Kim et al., 2025), so fewer investors carry the same errors across firms.

To answer the third question, I examine how much of the price impact of order flow is permanent and how quickly prices adjust to earnings news. Following Hasbrouck (1991), the first test measures the information content of trading by the share of its price impact that remains a quarter later. Before the launch, about a third of the same-day price impact of one-sided order flow reversed within the following quarter for the average firm. After the launch, the transitory part roughly halves for a firm one standard deviation harder to read, so the share of the same-day impact that remains after a quarter rises from about 60 to about 80 percent. Firms with more uncertain disclosures see no such improvement. The second test, at earnings announcements, does not rely on order flow and points the same way: after the launch, hard-to-read firms see more announcement trading, a larger share of the price response in the announcement window, and less drift afterward, relative to

comparable firms, while firms with uncertain disclosures see more announcement trading without faster incorporation of the news. Greater synchronization therefore accompanies better price discovery in the firms where reading assistance should matter most.

To explain the three results in a single framework, 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, the two dimensions that the readability and uncertainty measures capture. 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.

The model drives a wedge between the precision of individual readings and the informativeness of prices. A delegated reading is more precise than an independent one whenever the model decodes better than the investor, but the error that model users share survives aggregation. If all investors must use the model, prices can become less informative even though every reading improves: a modest precision gain may not offset the shared error. Voluntary adoption avoids this outcome. As delegation becomes crowded, its profitability falls. Independent reading survives, adoption stays within what the precision gain supports, and the equilibrium combines more synchronized trading with better price discovery. In the calibration, decoding difficulty increases delegation, synchronization, cross-stock commonality, and the improvement in price discovery, while uncertainty weakens the latter three effects. The wedge reopens where adoption is pushed beyond what trading profits justify, as when a platform gives every user the same summary.

The first contribution of the paper is to show that the effects of LLMs depend on what makes a disclosure difficult to process. Information that is costly to extract from public filings is less completely reflected in prices (Bloomfield, 2002), and investors are slow to notice quiet changes in filing language (Cohen et al., 2020). Research on disclosure complexity measures these processing costs through the Fog index (Li, 2008), file size (Loughran and McDonald, 2014), and the bloat that Kim et al. (2023) use ChatGPT to remove. A large literature extracts information from disclosure text with dictionaries and machine learning (Loughran and McDonald, 2016; Gentzkow et al., 2019; Huang et al., 2023). In that literature, machines read filings for the researcher. ChatGPT put a general-purpose reader in investors' own hands. I show that synchronization and price discovery improve where language is harder to decode, but not where outcomes remain unresolved.

A second contribution is 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. By examining synchronization together with cross-stock commonality and price discovery, I provide evidence on both the gains from better reading and the shared component of investors' interpretations.

A third contribution concerns the risks of reliance on AI in financial markets (Dou et al., 2025; Cartea et al., 2026; Hirshleifer et al., 2026). Model evaluations show that language models make correlated errors across architectures and providers (Kim et al., 2025) and converge on a narrow set of analytical paths (Huang et al., 2026). My model shows that shared errors need not worsen price discovery when investors choose whether to adopt. The private return to delegation falls as adoption spreads, which limits crowding and lets more synchronized trading coexist with more informative prices. Forced adoption can reverse this result: shared errors can outweigh the precision gain and make prices less informative even though every investor's reading improves.

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

2 Related literature

One strand of related research 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 the macroeconomy and returns (Chen et al., 2025; Lopez-Lira and Tang, 2026). Studies of the Italian ban of ChatGPT examine 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 LLM providers briefly interrupt AI assistance: while access is unavailable, trading volume falls (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 (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 ChatGPT launch and find that trading volume rises, and liquidity, price efficiency, and analyst accuracy improve, more for firms with richer disclosure environments and higher 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 the convergence extends across stocks, and whether more synchronized trading leaves prices more or less informative.

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). Model evaluations provide evidence of shared errors. Huang et al. (2026) show that models concentrate on a narrow set of analytical paths and produce much 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. My model makes adoption a choice: because the private return to delegation falls as adoption spreads, voluntary adoption limits crowding, and shared errors harm price discovery only when adoption is forced.

The paper also draws on research linking disclosure attributes to trading and price discovery. Information that is costly to extract from public filings is less completely reflected in prices (Bloomfield, 2002), complex filings discourage trading by small investors (Miller, 2010), and investors are slow to notice changes in filing language (Cohen et al., 2020). My readability measure follows Li (2008), while uncertainty uses the word list of Loughran and McDonald (2011); filing length, measured using file size and word count, enters every specification as a control (Loughran and McDonald, 2014). A large literature extracts information from disclosure text using methods ranging from dictionaries to pretrained language models (Loughran and McDonald, 2016; Gentzkow et al., 2019; Huang et al., 2023). Earlier advances in disclosure technology lowered the cost of obtaining filings (Gao and Huang, 2020); this paper studies a technology that lowers the cost of interpreting them. It asks how trading and prices change when investors themselves use a common machine reader. The finding that uncertain disclosures generate more announcement trading without faster price discovery also echoes Kandel and Pearson (1995), in which different interpretations

of public information generate trading without necessarily conveying superior information.

Finally, the paper relates to work on information aggregation and on common factors in trading and prices. In Grossman and Stiglitz (1980), information acquisition and price informativeness are jointly determined. In Stein (2009), arbitrageurs who cannot tell how many peers are entering the same trade need not make markets more efficient. Shared information sources also make prices comove. Investors who buy the same information use it to price many assets (Veldkamp, 2006), and stocks that share analysts, whose forecast errors are correlated across stocks, comove more than fundamentals justify (Israelsen, 2016). Liquidity has a common component (Chordia et al., 2000), and common factors in order flows account for much of the commonality in returns (Hasbrouck and Seppi, 2001). Morck et al. (2000) read synchronous stock price movements as a sign that prices impound little firm-specific information. In this paper, the shared source is a common language model that each investor decides whether to use, so disclosure difficulty and model quality jointly determine the precision and overlap of the information investors trade on. I measure the resulting commonality with a leave-one-out order-flow beta and find that, after LLMs become widely available, it rises in firms whose pre-launch disclosures are harder to read, the same firms in which price discovery improves.

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 microcaps whose imbalance measures are dominated by microstructure noise, and at least fifteen trading days in a firm-month, so that monthly averages do not rest on a few 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 mean greater difficulty.[4] 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, from the Loughran–McDonald 10-X Summaries. It counts language that presents outcomes as contingent or

[4]Flesch–Kincaid grade $= 0.39 \times$ (words per sentence) $+ 11.8 \times$ (syllables per word) $- 15.59$; Gunning Fog $= 0.4 \times [$(words per sentence) $+ 100 \times$ (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.

imprecise, and I use it as a proxy for how far the firm's prospects remain unresolved.[5]

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 exclude later changes in disclosure practice, which may respond to the technology itself (Cao et al., 2023; 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 (Table A2).

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. Signed order imbalances carry information about individual stock returns (Chordia and Subrahmanyam, 2004). The main dependent variable is the firm-month mean of $|OIB_{id}|$. A high value means that daily trading repeatedly falls on one side of the market, whether buying or selling.

Other variables. CRSP provides returns, prices, shares outstanding, and market equity. Earnings-announcement dates come from IBES and Compustat, and factor returns come from the Fama and French (2015) five-factor series. The order-flow specification allows the post-launch change to vary with filing length and ten 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. Nine characteristics are computed from daily data over the characteristic window, January 1 through November 29, 2022. 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 cumulative return over the window. Turnover is average daily dollar volume divided by average market equity, and retail share is retail dollar volume divided by total dollar volume. The other three are price, market beta, and AI-sector beta, which captures a

[5]Disclosure length enters every specification as a control; Table A1 adds it as a third dimension.

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 on the residual.[6] The tenth characteristic is analyst coverage, the log of one plus the mean number of IBES forecasts outstanding at the firm's earnings announcements from January 2021 through November 2022, zero for firms without coverage. All ten are standardized across firms.

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 A3 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 1 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 few large trades can move dollar imbalance more than trade-count imbalance. Retail imbalance is much higher, at 24.4 percentage points. Retail trading in a single stock is thinner, so buy and sell orders offset less within a day.

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] 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 2 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 participa-

[6] The 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.

[7] The firm count differs across tests with the variables each requires. Table 1 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 ten firm characteristics and the activity controls (3,670), and the daily flow-return regressions require neither (4,009). On the 3,669 firms present in both estimation panels, the estimates are nearly identical (Table A4).

[8] A firm at the 10th percentile of readability writes sentences averaging about 32 words, of which 28.1 percent are complex, versus 41 words and 29.1 percent at the 90th percentile.

tion, and wider analyst coverage. The uncertainty measure is positively correlated with option-implied volatility and with the volatility of return on assets (Table A5). The two disclosure measures are weakly negatively correlated across firms.

4 Empirical results

The empirical analysis follows the three questions. Section 4.1 examines whether trading within a stock becomes more one-sided, Section 4.2 tests whether order flow becomes more correlated across stocks, and Section 4.3 asks whether price discovery improves. Section 4.4 shows in which firms the effects are concentrated.

4.1 Synchronized trading within stocks

4.1.1 Main specification

The baseline regression enters the two disclosure dimensions jointly:

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

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. 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.[9] The vector Z_i contains the eleven predetermined controls listed in Table 1, 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, sector exposure, and analyst coverage before 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. The identifying assumption is that, given these controls, nothing else that changed around the launch moved trading differently in firms with different disclosures.

[9] Market-wide, dollar one-sidedness fell from 13.8 to 12.5 percentage points after the launch (Table A6).

Table 3 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 0.28 percentage points larger, or 2.2 percent of the sample mean of 13.1 percentage points. A one-standard-deviation increase in uncertainty is associated with a change that is 0.10 percentage points smaller, or 0.7 percent of the sample mean.[10]

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. The negative uncertainty coefficient could therefore reflect either of two mechanisms. 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 judgment end up further apart than before.

The readability coefficient is 0.28 for dollar imbalance and 0.14 for trade-count imbalance. The larger dollar estimate means that large trades line up on one side more than small trades do. The dollar coefficients are similar for non-retail and retail flow, at 0.26 and 0.23, 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 analysts, management, and in-house research. They are also more likely to have weighed the possible scenarios before the launch, so a model that lays them out changes their view little. Retail investors have fewer such sources and 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.

[10]Both components of the readability measure, long sentences and technical vocabulary, contribute to the effect (Table A7). The readability coefficient stays positive and significant under three alternative constructions: ranking each filing within its quarter, using only filings dated before the WRDS parse changed its sentence splitter in July 2022, and replacing the composite with the complex-word share from the independent Loughran–McDonald parse (Table A8). It is also positive and significant without the firms whose writing changed after the launch (Table A9), without the retail trading episode of early 2021 (Table A10), and with negative tone as a control (Table A11); in these three checks, the uncertainty coefficient keeps its sign. In unreported regressions without the two activity controls, the readability coefficients are larger, and the uncertainty coefficients keep their signs but are significant only for retail flow.

Greater one-sidedness could also come from a few informed investors 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, and retail and non-retail flow would more often have opposite signs. Panel B tests this using the share of trading days on which the two have the same sign. 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, the endpoint bins absorb 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 does not reject a flat pre-trend. After the launch, all seven readability bins are positive, and five are individually significant at the five percent level. The uncertainty coefficients turn negative late in the sample.

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). 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 0.43 and a maximum of 1.57 for readability, and 0.42 and 1.29 for uncertainty. The actual launch also stands out among all feasible break dates. When treatment is assigned 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, a randomization p-value of 0.021.

4.1.2 The shape of the relation across the difficulty distribution

The linear specification imposes a constant effect per standard deviation. Replacing each dimension with indicators for its deciles relaxes this restriction; Figure 2 plots the coefficients, and Table A12 lists them. 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 $+0.14$ percentage points, and the first clear increase appears in the third, at $+0.55$, about two-thirds of the tenth-decile estimate. Above this threshold, the point estimates rise broadly with difficulty. The data do not reject a single step of 0.53 percentage points from the third decile, the dashed line in panel (a): the third to tenth

deciles are jointly indistinguishable ($p = 0.20$). This pattern suggests a threshold effect: model assistance begins to align trading only once disclosures are sufficiently difficult to read. 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: in the most uncertain decile, the change is 0.50 percentage points smaller than in the least uncertain, or 3.8 percent of the sample mean. The decline has 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 0.18 percentage points in 2023 to 0.28 in 2024 and 0.42 in 2025 and is significant in each year (Table 4, Panel A). At a quarterly frequency, it turns positive around the launch, builds through 2023, remains elevated with some quarterly variation, and rises again in 2025 (Figure 3). To separate the initial break from the later growth, I estimate

$$\begin{aligned} y_{im} = \alpha_i + \gamma_m + \sum_{d \in \{R,U\}} \left[\beta_{1d}(Post_m \times D_{di}) + \beta_{2d}(N(m) \times D_{di})\right] \\ + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}, \end{aligned} \tag{2}$$

where $N(m)$ is the cumulative count of frontier-model releases in Table A3, 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 next three years.

In Panel B of Table 4, the stock term is positive and accounts for most of the growth. Entered alone, it has a readability coefficient of 0.50 percentage points. Entered together with the launch step, the stock coefficient remains 0.39, while the launch coefficient falls to 0.13, so about a quarter of the effect appears at the launch. The gradual growth fits an effect that follows improvements in model capability and their spread across investors.

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 latest 10-K or 10-Q and reflect their interpretations of existing disclosures alongside other information, so their dispersion measures disagreement. If a common reader helps analysts interpret difficult text, dispersion should fall more where disclosures are harder to read. Table 5 relates the standard deviation of IBES forecasts, scaled by price, to the same post-launch interac-

tions at the announcement level, with firm and calendar-quarter fixed effects.[11] 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.

4.2 Comovement across stocks

A general-purpose model that carries the same reading habits and blind spots from one disclosure to the next 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

$$OIB_{id} = a_{im} + b_{im}\, OIB_{d}^{(-i)} + e_{id}, \tag{3}$$

where $OIB_{d}^{(-i)}$ is the equal-weighted signed imbalance of all other sample firms on day d. The slope b_{im}, the firm's order-flow commonality, is high when the firm's daily flow moves with that of the rest of the market, whether toward buying or selling. I then estimate

$$b_{im} = \alpha_i + \gamma_m + \sum_{d\in\{R,U\}} \delta_d(Post_m \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + \nu_{im}. \tag{4}$$

Table 6 summarizes the results. Before the launch, the mean slope b_{im} is 0.83. The post-launch readability differential is 0.105 for total flow and 0.090 for non-retail flow, with no significant change for retail flow. For a firm one standard deviation harder to read, a 10-percentage-point increase in market order imbalance is associated with an additional 1.05 percentage points of imbalance after the launch, relative to other firms. This is about 13 percent of the pre-launch mean slope. The difference between investor groups may reflect how they trade. Non-retail investors may adjust positions across several stocks on the same day, so their alignment shows up in daily flow. Retail investors may respond at different times, so their alignment is harder to detect in daily data. The uncertainty coefficients are negative but insignificant. The results also hold when the post-launch change is allowed to vary with S&P 500 membership, institutional ownership, and ETF ownership (Panel B).

The cross-stock effect is concentrated in 2023 and 2024 and fades in 2025 (Panel C), while within-stock synchronization continues to strengthen. A possible explanation is that two forces act on both measures. Wider adoption raises both, because more investors read

[11] 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 A6).

disclosures through a model. The spread of users across model families works against both, because models that share an architecture or provider make more similar mistakes (Kim et al., 2025), so investors who use different families share fewer errors. Frontier models came from three developers in 2023 and from six in 2025 (Table A3). The two measures differ in their sensitivity to the second force. Within a stock, investors converge on the same message as their readings improve, even when they use different models, and different models also make correlated errors on the same disclosure. Synchronization therefore depends on shared errors only in part, and wider adoption outweighs their decline. Across stocks, better reading creates no common component, because each disclosure concerns a different firm, so commonality depends only on errors that carry across disclosures. Wider adoption still raises commonality, but by 2025 it no longer offsets the decline in shared errors.

4.3 Price discovery

Sections 4.1 and 4.2 show that trading in hard-to-read firms becomes more synchronized within and across stocks after the launch. Greater synchronization 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 transitory pricing errors and speed up the incorporation of news, while a shared error can push prices away from fundamentals and be reversed later. I therefore measure how much of the price impact of order flow lasts and how fast prices absorb earnings news.

4.3.1 The price impact of order flow

For each firm-day, I estimate

$$\begin{aligned} CAR_{i,[a,b]} = \alpha_i + \gamma_t + OIB_{it}\big(\beta_1 + \beta_2 Post_t\big) \\ + \sum_{d\in\{R,U\}} OIB_{it}\big(\beta_{3d} D_{di} + \beta_{4d} Post_t \times D_{di}\big) + \phi' W_{it} + e_{it}, \end{aligned} \tag{5}$$

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 3, so each slope measures the return associated with moving from balanced to fully one-sided flow. The same-day slope is the price impact of order flow, and the slopes over later windows show what happens to that impact. A negative later slope means that part of the impact reverses; a positive slope means that the

price keeps moving in the direction of the flow. Following Hasbrouck (1991), who measures the information content of a trade by its permanent price impact, I treat the part of the same-day impact that is still in the price a quarter later as permanent and the part that reverses within the quarter as transitory.[12] The coefficient β_{3d} measures how the slope differed with each disclosure dimension before the launch, and β_{4d} measures how its change after the launch varies with that dimension. The vector W_{it} contains the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, together with the interactions of OIB_{it} and $OIB_{it} \times Post_t$ with filing length.

Table 7 reports, for each return window, the slope before the launch and its change after the launch, for a firm at the mean of both dimensions and per standard deviation of each. For a firm one standard deviation above the mean of a dimension, the slope before the launch is $\beta_1 + \beta_{3d}$, and the launch adds $\beta_2 + \beta_{4d}$ 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 for the average firm, and after the launch this impact rises by 0.4 percentage points. 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.

Readability does not significantly change the slope over the next day or week. Over the following quarter, the slope for the average firm before the launch is -0.71 percent per unit of imbalance, so about a third of the same-day impact reverses and the rest is permanent. After the launch, reversal weakens in hard-to-read firms relative to other firms: the $[+2, +60]$ coefficient β_{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 becomes -0.52. Prices still move back against the flow, but the transitory part roughly halves while the same-day impact rises from 2.5 to 2.8 percent, so the share of the impact that remains after a quarter rises from about 60 to about 80 percent. The uncertainty gradient runs the other way. For a firm one standard deviation more uncertain, the quarterly slope falls from -0.32 before the launch to -1.42 afterward. Price discovery in uncertain firms therefore does not improve relative to comparable firms.

Cheaper liquidity provision does not explain the weaker reversal. Neither the effective spread nor the five-minute realized spread, a proxy for liquidity suppliers' gross compensa-

[12] The transitory component may reflect compensation for holding inventory. Someone 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.

tion, changes significantly in hard-to-read firms relative to other firms after the launch.[13]

Figure 4 shows the pattern in raw price paths. It plots average cumulative market-adjusted returns, signed in the direction of the flow, after firm-days in the top decile of absolute dollar imbalance, on which the larger side accounts for at least 65 percent of classified dollar volume, for the easiest and hardest readability terciles before and after the launch.[14] 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 the reversal over the quarter falls from 82 to 42 basis points. In the easy tercile, the cumulative return rose from 70 basis points on the flow day to 111 after sixty days before the launch; 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.

4.3.2 Earnings announcements

The second test examines how prices incorporate earnings news. If model assistance improves investors' understanding of difficult filings, it should also help them interpret subsequent earnings announcements. More of the price response should occur at the announcement, with less drift afterward. For each announcement e of firm i, I estimate

$$y_{ie} = \alpha_i + \gamma_{q(e)} + \sum_{d \in \{R,U\}} \beta_d (Post_e \times D_{di}) + \varphi \, Rank_e + \theta' X_{ie} + \varepsilon_{ie}, \tag{6}$$

where α_i and $\gamma_{q(e)}$ are firm and calendar-quarter fixed effects, $Post_e$ equals one for announcements from December 2022 onward, $Rank_e$ is the within-quarter percentile rank of the announcement-window return, scaled to $[-0.5, 0.5]$, and X_{ie} contains lagged size, volatility, and volume, together with filing length interacted with the post indicator. Returns are market adjusted: $CAR_{[a,b]}$ sums, over trading days $[a, b]$ relative to the announcement, the firm's daily return minus the equal-weighted return of the sample firms. Abnormal volume is the log of mean dollar volume over days $[-1, +1]$ minus that over days $[-30, -6]$.

[13] The realized spread is the part of the effective spread left after the price move over the next five minutes. These spread estimates are not tabulated.

[14] The pattern also holds for the top quintile or top five percent of absolute imbalance. These unconditional averages may also reflect differences in firm size, liquidity, retail participation, and subsequent order flow. Flow persists after a flow day: over the following sixty days, daily imbalance in the direction of the day-0 flow averages 4.0 percentage points in the easiest tercile and 2.3 in the hardest before the launch. Within each cell, regressing the $[+1, +60]$ return on the day-0 return and this subsequent flow moves the easy tercile's pre-launch intercept from +0.31 to +0.02 percent and the hard tercile's from −0.75 to −1.13, so the later flow explains the continuation in the easy tercile but not the reversal in the hard tercile.

Incorporation speed is $|CAR_{[-1,+1]}|/(|CAR_{[-1,+1]}| + |CAR_{[+2,+21]}|)$, the share of the price response that occurs in the announcement window. Drift is $CAR_{[+2,+21]}$. For drift, the regression adds the interactions of the announcement-return rank with the post indicator, with each disclosure dimension and filing length, and with their products with the post indicator, so that the coefficient on $Rank_e \times Post_e \times D_{di}$ measures the change in the drift spread between good- and bad-news announcements.

Table 8 reports the estimates. After the launch, all three outcomes point to faster incorporation of earnings news in hard-to-read firms, each significant at the five percent level. Per standard deviation of readability difficulty, abnormal volume rises by 1.9 percent, the share of the price response in the announcement window by 0.79 percentage points, and the drift spread between good- and bad-news announcements narrows by 1.4 percentage points. Uncertainty raises abnormal volume by 1.4 percent per standard deviation, but changes neither incorporation speed nor drift. Investors trade more around the announcements of firms with uncertain disclosures without moving prices to the new level faster, as when they interpret the same public information differently (Kandel and Pearson, 1995).

Both tests point in the same direction. After the launch, a larger share of the price impact of order flow in hard-to-read firms is permanent, and their prices absorb earnings news faster, relative to comparable firms. The additional synchronization of Sections 4.1 and 4.2 therefore comes with an improvement in price discovery.

4.4 Heterogeneity across firms

Figure 5 shows the main specifications re-estimated within halves of the sample, split by pre-launch retail share, size, and analyst coverage, and Table A13 lists the coefficients. The increase in one-sidedness is concentrated in retail-heavy, small, and thinly covered firms. 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 4.2, and more in small and thinly covered firms than in large and widely covered ones. The weaker quarterly reversal is concentrated in retail-heavy and small stocks.[15] The uncertainty coefficients are negative in every half for one-sidedness and quarterly reversal, as in the full sample, and do not follow the readability pattern across halves.

Synchronization increases most where investors had less help in interpreting disclosures

[15] Attention from trading platforms is unlikely to explain the larger effect in retail-heavy stocks. Popularity lists and social feeds can concentrate retail attention and make trading more one-sided, but they predate the launch, and the placebo launch dates produce no comparable readability differential (Section 4.1.1). The effect also appears in non-retail flow and in institution-heavy stocks, analysts' forecasts converge in the same firms, and the results hold when the retail trading episode of early 2021 is excluded (Table A10).

before the launch. Complex filings reduce trading by small investors (Miller, 2010), and small firms receive less analyst attention, so better decoding can matter more where investors have fewer other sources of interpretation.

5 A model of delegated reading

I place the empirical results in a single framework, a one-period Kyle (1985) auction with multiple informed traders and correlated signals (Lambert et al., 2018), in which speculators obtain private information by interpreting a public disclosure (Kim and Verrecchia, 1994). Each speculator chooses between independent reading and a common technology that reduces decoding error and introduces partly shared mistakes but leaves unresolved uncertainty unchanged. The profit from delegation falls as more investors delegate, reflecting the dilution of information value studied by Grossman and Stiglitz (1980) and Admati and Pfleiderer (1986), so disclosure difficulty governs adoption and, through it, synchronization, commonality across stocks, and price discovery. The model also yields a prediction beyond these results: if every investor must delegate, prices can become less informative even though each reading is more precise, a wedge between individual precision and price informativeness that voluntary adoption avoids.

5.1 Setup

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

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

The fundamental v, the shared error η, and the idiosyncratic errors ε_i are mean-zero, jointly normal, and mutually independent, and the ε_i are identically distributed. The total reading-error variance is τ^2, and ρ is the fraction of that error shared across readers.

A disclosure can be difficult for two reasons. Decoding difficulty, r, concerns information in the text that readers struggle to extract. Irreducible uncertainty, u, remains when the disclosure leaves outcomes unresolved, even if readers understand the text. Model assistance can reduce decoding error but cannot settle those outcomes. Empirically, readability proxies for r, and the frequency of uncertainty words proxies for u.

Each investor chooses between two reading technologies:[16]

$$\begin{aligned} &\text{own reading:} && \tau_o^2 = r + u, && \rho_o = 0, \\ &\text{model reading:} && \tau_d^2 = A + u, && \rho_d = \frac{\sigma_\eta^2}{A + u}, \end{aligned} \tag{8}$$

where A is the model's decoding-error variance and σ_η^2, with $0 < \sigma_\eta^2 \leq A$, is the variance of the shared error η in a model reading. This variance equals the covariance between the errors in two delegated readings of the same disclosure.[17] I 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 u contributes equally to the error variance under both technologies; delegation changes decoding precision and introduces a shared error, but leaves this uncertainty unchanged.

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). A market maker that read the disclosure through its own model would start from a more precise prior and could at most remove the error component common to all model families. The family-specific component and each reader's idiosyncratic error would still reach prices only through trading, and Appendix B shows that the propositions are unchanged. 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 its price impact in hard-to-read firms is permanent.

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 $n \in [0, M]$ be the number 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, maximize expected trading profit, and account for the effect

[16] Appendix B allows investors to combine own and model readings. Under a variance-minimizing rule, investors place more weight on the model when their own reading is noisier. Blending smooths adoption but retains the trade-off between precision and shared error.

[17] The common error can persist even when investors use different models. Two delegated readings always share the component of reading error common to all model families, and they also share a family-specific component when the users choose the same family, so σ_η^2 is positive whenever either component is present; Appendix B gives the decomposition. The 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 (Hong Kong Institute for Monetary and Financial Research, 2025), and model errors remain correlated across architectures and providers (Kim et al., 2025).

of their own orders on the price. Noise traders submit $z \sim N(0, \sigma_z^2)$, independent of v and of every reading error. The competitive market maker observes aggregate order flow,

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

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

Equilibrium. Fix the primitives $\Theta = (\sigma_v^2, \sigma_z^2, M, r, u, A, \sigma_\eta^2)$. A linear equilibrium determines the number 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 $p = \lambda Y = \mathrm{E}[v \mid Y]$. 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, $\pi_d(n; \Theta) = \pi_o(n; \Theta)$.[18]

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

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

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*

$$\lambda = \frac{\sqrt{aM\sigma_v^2}}{D\sigma_z}, \qquad \beta = \sigma_z \sqrt{\frac{a}{M\sigma_v^2}}. \tag{9}$$

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

$$\mathrm{var}(v \mid Y) = \sigma_v^2 - \frac{M\sigma_v^4}{\widetilde{D}}, \qquad \widetilde{D} = (M+1)(\sigma_v^2 + \rho\tau^2) + 2(1-\rho)\tau^2. \tag{10}$$

Proof. See Appendix C. □

Equation (10) shows how the structure of reading errors affects price discovery. Each unit of common-error variance enters $\widetilde{D}$ with weight $M+1$, and each unit of idiosyncratic-

[18] I treat n as continuous to obtain a closed form. The indifference condition then compares the two technologies at the same composition: every trader knows n, but a trader who considers switching treats her switch as leaving n unchanged, although she accounts for her price impact. Appendix B gives the integer version, in which a switch changes n by one and the trading equilibrium is re-solved.

error variance with weight 2: errors that speculators make alone average out in aggregate order flow, while 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 r to A but adds the shared error σ_η^2.

5.3 Endogenous delegation

I now let traders choose between independent and model-assisted reading. Define

$$S_d = \sigma_v^2 + A + u, \qquad S_o = \sigma_v^2 + r + u, \qquad a_o = \frac{\sigma_v^2}{S_o}.$$

Proposition 2 (Endogenous delegation)**.** *Let $\pi_d(n)$ and $\pi_o(n)$ be the equilibrium expected profits of a delegating trader and an independent reader when n traders delegate. Delegation is a strategic substitute: $\pi_d(n)/\pi_o(n)$ decreases in n. Treating n as continuous:*

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

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

$$n^*(r,u) = \frac{(2-a_o)\sqrt{S_o S_d} - 2S_d + \sigma_v^2 + \sigma_\eta^2}{\sigma_\eta^2}, \tag{11}$$

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

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

Hence the adoption threshold is

$$r^* = A, \tag{12}$$

and is independent of uncertainty u.

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, r. 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 u always reduces the common-error share $\sigma_\eta^2/(A+u)$ of a delegated reading.

The adoption threshold also predicts the shape of the readability relation: below r^* no trader delegates and synchronization does not change, so the response is flat at the easy end of the distribution and turns positive once decoding difficulty crosses the threshold.

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

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

where $\kappa_d = \lambda\beta_d$ and $\kappa_o = \lambda\beta_o$, and the posterior variance and price impact are

$$\mathrm{var}(v \mid Y) = \sigma_v^2\,\frac{2 - a_o}{2 + (m-1)a_o}, \qquad \lambda = \frac{\sqrt{a_o M \sigma_v^2}}{\big(2 + (m-1)a_o\big)\sigma_z}. \tag{14}$$

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

Proof. See Appendix C. □

The lemma expresses the gain in price discovery as an increase in the effective number of independent readers. When $r > A$, an interior equilibrium has the same posterior variance as an all-own-reading market with $m = M + n^*(t-1) > M$ readers. 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: 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. Competition over that information lowers aggregate trading profits even as prices become more informative, as in Admati and Pfleiderer (1988) and Holden and Subrahmanyam (1992). These profits equal the noise traders' expected loss, $\lambda\sigma_z^2$, so lower profits imply lower price impact. Price impact falls by the same factor as the posterior variance. The decline is larger where decoding is harder, because both the delegating share and the precision gain increase with r.

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,

$$|OIB| \equiv \frac{\mathrm{E}\,|Y|}{\sum_{i=1}^{M} \mathrm{E}\,|x_i| + \mathrm{E}\,|z|}. \tag{15}$$

Second, for two otherwise identical firms $f \neq \ell$ with independent fundamentals, cross-stock order-flow commonality is

$$\varsigma \equiv \frac{\mathrm{cov}(Y_f, Y_\ell)}{\mathrm{var}(Y_f)}. \tag{16}$$

Third, I measure pricing error by

$$\mathcal{P} \equiv \sqrt{\mathrm{var}(v \mid Y)}, \tag{17}$$

so a fall in $\mathcal{P}$ means better price discovery.

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

$$|OIB|_o \longrightarrow \sqrt{\frac{\sigma_v^2}{\sigma_v^2 + r + u}}, \qquad |OIB|_d \longrightarrow \sqrt{\frac{\sigma_v^2 + \sigma_\eta^2}{\sigma_v^2 + A + u}}. \tag{18}$$

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

Proof. See Appendix C. □

Under independent reading, idiosyncratic errors offset in aggregate, and $|OIB|_o$ falls as $r + u$ rises. Delegation introduces a shared component σ_η^2 of reading error, which 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 r, so harder disclosures attract more delegation and expose more traders to the shared model error. Uncertainty works differently. A higher u increases the part of reading error that remains unresolved under either technology and dilutes the common-error share $\sigma_\eta^2/(A + u)$ of a delegated reading.

Proposition 4 (Cross-stock commonality)**.** *Suppose delegated readings contain a cross-firm component with covariance $\chi^2 \leq \sigma_\eta^2$. Under all-own reading, $\varsigma_o = 0$, so the change*

induced by forced full delegation is

$$\Delta\varsigma^{FD} \equiv \varsigma_d - \varsigma_o = \frac{M\chi^2}{(M+1)(\sigma_v^2+\sigma_\eta^2)+2(A+u-\sigma_\eta^2)}. \tag{19}$$

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

Proof. See Appendix C. □

Decoding difficulty affects cross-stock commonality through adoption. Below r^*, 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.

Proposition 5 (Price informativeness)**.** *In the forced full-delegation benchmark, model use lowers pricing error $\mathcal{P}$ if and only if*

$$r > \bar{r} = A + \tfrac{1}{2}(M-1)\sigma_\eta^2. \tag{20}$$

Above $\bar{r}$, the reduction in pricing error becomes larger as r rises, smaller as u rises, and converges to zero as $u \to \infty$.

Proof. See Appendix C. □

More precise individual readings need not produce more informative prices. When $A < r < \bar{r}$, each model-assisted reading is more precise than an independent reading, yet full delegation raises pricing error. The shared error η survives aggregation, and its cost outweighs the precision gain. Letting investors 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.

Corollary 1 (No deterioration at the endogenous composition)**.** *At the equilibrium composition,* $\mathrm{var}(v \mid Y)$ *is no larger than its all-own-reading value for every parameter configuration with $0 < \sigma_\eta^2 \le A$, and smaller whenever $r > A$.*

Proof. See Appendix C. □

Lemma 1 shows that the all-own-reading market has M effective readers, while an interior equilibrium with delegation has $m = M + n^*(t-1) > M$. Posterior variance falls with m, so price discovery improves as soon as investors begin to delegate, and the deterioration possible under forced delegation disappears.

5.5 Numerical results

Propositions 3 to 5 use full delegation to obtain closed-form comparative statics. Figure 6 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 $\sigma_z^2 = 4$ and $M = 20$, and chooses $A = 0.23$ and $\sigma_\eta^2 = 0.096$, which imply an entry threshold $r^* = 0.23$ and a full-delegation price-quality boundary $\bar{r} = 1.14$. The equilibrium delegating share rises from zero at the threshold to about a quarter at $r = 0.5$ and a half at $r = 1$, and every trader delegates above about $r = 2.1$.

Panel (a) shows that the increase in one-sidedness grows with decoding difficulty and shrinks with uncertainty. Panel (b) shows that cross-stock commonality, zero under independent reading, rises as more traders delegate and their orders load on the error component that carries across firms. Panel (c) shows that forced full delegation raises pricing error between the adoption threshold and the price-quality boundary, while at the endogenous composition it falls, by more as decoding difficulty rises. These patterns match the evidence: harder-to-read disclosures are associated with greater synchronization, stronger cross-stock commonality, and better price discovery, while uncertainty limits the gains.

5.6 Predictions beyond the current data

Three further predictions require investor-level data linking model use to trading. First, the same investor should use model assistance more often for firms whose disclosures are harder to read. Second, users of the same model family should comove more within a stock than users of different families, and cross-stock commonality should be greater when model use is concentrated in fewer families. Third, investors who adopt a new model release should change their trading in several hard-to-read stocks at the same time. Section 4.1.3 documents the aggregate growth of the synchronization differential, but anonymous order flow cannot show whose trading changes or which release prompts the change.

6 Conclusion

After ChatGPT's public launch, trading in firms with hard-to-read disclosures becomes more synchronized within stocks and more correlated across stocks, relative to comparable firms, and analysts' forecasts converge. Price discovery also improves: a larger share of the price impact of their order flow becomes permanent, and their prices incorporate earnings news faster. Greater disclosure uncertainty predicts less convergence and synchronization,

with no improvement in price discovery. The distinction suggests that LLMs help investors extract information from difficult prose but help less when the disclosure leaves the outcome unresolved. The model shows how greater synchronization can coexist with better price discovery: the profit from delegation falls as other traders adopt, so crowding limits adoption while the precision gain is small, and voluntary adoption improves price discovery.

Policy reports warn that widespread use of common models could correlate investors' mistakes and trading decisions (Leitner et al., 2024; International Monetary Fund, 2024; Financial Stability Board, 2024; Bank of England, 2025; Hong Kong Institute for Monetary and Financial Research, 2025). In the model, whether this correlation harms price discovery depends on why investors adopt. When adoption is driven by trading profits, crowding limits it, but other incentives can lead investors to rely more on a common model. Career concerns, for example, can make an investor reluctant to depart from her peers, even when they are wrong (Scharfstein and Stein, 1990). A platform that shows every user the same AI summary can likewise give its users a common reading they did not choose. Both can push adoption beyond what trading profits justify. If every investor delegates, shared errors can outweigh a modest precision gain and make prices less informative. A common model can help each investor read better without helping the market price better.

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.

Bank of England (2025). Artificial intelligence in the financial system. Financial Stability in Focus April 2025, Bank of England, Financial Policy Committee.

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.

Bloomfield, R. J. (2002). The "incomplete revelation hypothesis" and financial reporting. *Accounting Horizons 16*(3), 233–243.

Boehmer, E., C. M. Jones, X. Zhang, and X. Zhang (2021). Tracking retail investor activity. *The Journal of Finance 76*(5), 2249–2305.

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.

Cohen, L., C. Malloy, and Q. Nguyen (2020). Lazy prices. *The Journal of Finance 75*(3), 1371–1415.

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.

Financial Stability Board (2024). The financial stability implications of artificial intelligence. Report November 2024, Financial Stability Board.

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.

Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature 57*(3), 535–574.

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. (1991). Measuring the information content of stock trades. *The Journal of Finance 46*(1), 179–207.

Hasbrouck, J. and D. J. Seppi (2001). Common factors in prices, order flows, and liquidity. *Journal of Financial Economics 59*(3), 383–411.

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, A. H., H. Wang, and Y. Yang (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research 40*(2), 806–841.

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.

Israelsen, R. D. (2016). Does common analyst coverage explain excess comovement? *Journal of Financial and Quantitative Analysis 51*(4), 1193–1229.

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.

Leitner, G., J. Singh, A. van der Kraaij, and B. Zsámboki (2024). The rise of artificial intelligence: Benefits and risks for financial stability. Financial Stability Review special feature May 2024, European Central Bank.

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.

Liu, N. F., K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, and P. Liang (2024). Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics 12*, 157–173.

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.

Loughran, T. and B. McDonald (2016). Textual analysis in accounting and finance: A survey. *Journal of Accounting Research 54*(4), 1187–1230.

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.

Miller, B. P. (2010). The effects of reporting complexity on small and large investor trading. *The Accounting Review 85*(6), 2107–2143.

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.

Scharfstein, D. S. and J. C. Stein (1990). Herd behavior and investment. *American Economic Review 80*(3), 465–479.

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.

Veldkamp, L. L. (2006). Information markets and the comovement of asset prices. *Review of Economic Studies 73*(3), 823–845.

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 adds filing length as a third disclosure dimension. The readability and uncertainty coefficients are unchanged, while the length coefficients are small and insignificant at the five percent level. Length is also insignificant at the five percent level in the commonality tests of Table 6 and the price-discovery tests of Table 7. Models may save time on long filings, but faster reading need not align interpretations, especially when the additional text is repetitive or has little bearing on valuation.

Table A2 measures the persistence of the text measures on the Loughran–McDonald 10-X Summaries. 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 A3 lists the twenty-three frontier model releases in the cumulative model stock $N(m)$, with their dates and developers.

Table A4 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 A5 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 net of size, while readability is negatively correlated with every volatility measure.

Table A6 reports the market-wide means of one-sidedness and of analyst forecast dispersion before and after the launch. Every imbalance measure and log dispersion is lower after the launch, while mean dispersion scaled by price is unchanged.

Table A7 enters the two components of readability, words per sentence and the complex-word share, separately and together. Both contribute, and the complex-word share has the larger coefficient per standard deviation.

Table A8 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.

Table A9 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.

Table A10 checks whether the results reflect the retail trading episode in early 2021 by excluding those months, dropping the tenth of firms with the largest increases in retail share, or controlling for that increase interacted with the post-launch indicator. Starting the pre-launch period in July 2021 leaves the readability differentials in one-sidedness, commonality, and quarterly reversal at least as large; the quarterly differential rises to 0.65 percentage points per unit of imbalance, significant at the five percent level. It is 0.68 without the top decile of meme intensity, significant at the one percent level, and 0.46 with the control, significant at the ten percent level. The readability coefficients in the announcement tests keep their signs in every variant and remain significant in the two firm-level checks, but not in the shorter sample, where incorporation speed rises by 0.69 and drift narrows by 0.95 percentage points per standard deviation. The uncertainty differential in quarterly reversal remains negative and significant in every variant.

Table A11 adds pre-launch negative tone as a control. Tone is positively correlated with uncertainty but unrelated to readability. The readability estimates change little; the quarterly differential, for example, is 0.57 without tone and 0.67 with it. The uncertainty effects on retail one-sidedness and forecast dispersion remain significant, while those on total one-sidedness and quarterly reversal keep their signs but lose significance.

Table A12 lists the decile coefficients plotted in Figure 2. It also reports the two one-parameter shapes drawn there, a step from the third decile and a linear decile score, and tests that the second decile equals zero and that the third to tenth deciles are equal.

Table A13 lists the coefficients within halves of the sample plotted in Figure 5.

B Model extensions

Mixed-composition outcomes in the numerical results. Let n traders delegate and define $\kappa_d = \lambda\beta_d$ and $\kappa_o = \lambda\beta_o$, with equilibrium values given by Proposition 2. Let

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

Since $\operatorname{cov}(v, Y) = K/\lambda$ and Bayesian pricing implies $\lambda = \operatorname{cov}(v, Y)/\operatorname{var}(Y)$, order-flow variance is K/λ^2 and posterior variance is

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

Centered normality allows the factor $\sqrt{2/\pi}$ to cancel in equation (15), giving

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

Only delegators load on the cross-firm error component, whose covariance is χ^2. Hence $\mathrm{cov}(Y_f, Y_\ell) = (n\beta_d)^2\chi^2$, and cross-stock commonality is

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

Evaluating these expressions at $n = n^*$ gives the outcomes in Figure 6.

Integer composition. Equation (11) treats n as continuous and holds composition fixed when comparing technologies. With integer traders, switching changes n by one. An equilibrium therefore requires

$$\pi_d(n) \geq \pi_o(n-1), \qquad \pi_o(n) \geq \pi_d(n+1), \tag{25}$$

where each condition applies when the corresponding type is present and the trading equilibrium is re-solved after each deviation. At every calibration point in Figure 6, the integer equilibrium is unique and lies within one trader of the continuous solution. At $r = 0.5$, for example, $n = 4$ while $n^* = 4.7$; at $r = 1$, $n = 10$ while $n^* = 10.3$.

Combining own and model readings. An investor who places weight w on the model reading has decoding-error variance

$$w^2A + (1-w)^2r. \tag{26}$$

The variance-minimizing weight is

$$w^*(r) = \frac{r}{r+A}. \tag{27}$$

If irreducible uncertainty remains u, total error variance and the common-error share are

$$\tau^2(w^*) = \frac{Ar}{A+r} + u = Aw^* + u, \qquad \rho(w^*) = \frac{w^{*2}\sigma_\eta^2}{Aw^* + u}. \tag{28}$$

As decoding becomes harder, investors put more weight on the model and gain precision but take on more shared error. Under this variance-minimizing rule, blending smooths the

adoption threshold because $w^* > 0$ whenever $r > 0$. The profit-maximizing weight may differ because traders also care about overlap with others' information.

Model families. The shared-error variance σ_η^2 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 component with variance s_ε^2. If two users choose the same family with probability H, then

$$A = \omega^2 + \psi^2 + s_\varepsilon^2, \qquad \sigma_\eta^2 = \omega^2 + H\psi^2. \tag{29}$$

The parameter A therefore captures the total error of an individual delegated reading, while σ_η^2 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 σ_η^2 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 $\chi^2 \leq \sigma_\eta^2$ gives the cross-stock component used in Proposition 4.

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 idiosyncratic error, and sets $p = \mathrm{E}[v \mid s_0, Y]$. Conditioning on s_0 before orders arrive replaces the fundamental by its residual $\tilde{v} = v - \mathrm{E}[v \mid s_0]$, with variance $\mathrm{var}(v \mid s_0) < \sigma_v^2$, 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 idiosyncratic errors 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 $\mathrm{var}(v \mid s_0)$ and σ_η^2 by $H\psi^2$. Because r and u 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),

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

For $j \neq i$,

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

Therefore

$$\mathrm{E}[v \mid s_i] = a s_i, \qquad \mathrm{E}[s_j \mid s_i] = b s_i, \tag{30}$$

where

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

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

$$\begin{aligned}
\mathrm{E}[(v-p)x_i \mid s_i] &= x_i \,\mathrm{E}\left[v - \lambda\left(x_i + \sum_{j \neq i} x_j + z \right) \,\middle|\, s_i \right] \\
&= x_i \,\mathrm{E}[v \mid s_i] - \lambda x_i^2 - \lambda x_i \sum_{j \neq i} \mathrm{E}[x_j \mid s_i] \\
&= x_i \{ a - \lambda(M-1)\beta b \} s_i - \lambda x_i^2.
\end{aligned} \tag{31}$$

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

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

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

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

so

$$\lambda\beta D = a, \qquad D \equiv 2 + (M-1)b, \qquad \beta = \frac{a}{\lambda D}. \tag{32}$$

It remains to impose competitive Bayesian pricing. Let

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

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

$$V = M(\sigma_v^2 + \tau^2) + M(M-1)(\sigma_v^2 + \rho\tau^2). \tag{33}$$

Since

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

I have

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

The competitive market maker therefore sets

$$\lambda = \frac{\text{cov}(v, Y)}{\text{var}(Y)} = \frac{\beta M \sigma_v^2}{\beta^2 V + \sigma_z^2}. \tag{34}$$

Substitute $\beta = a/(\lambda D)$ from (32):

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

Multiplying both sides by λD^2 gives

$$a^2 V + \lambda^2 D^2 \sigma_z^2 = aM\sigma_v^2 D. \tag{35}$$

To simplify the right-hand side, use

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

Then

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

Comparing this with (33) gives

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

and therefore

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

Equation (35) consequently reduces to

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

Taking the unique positive root,

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

Using (32),

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

which proves (9).

Finally,

$$\begin{aligned} \operatorname{var}(v \mid Y) &= \sigma_v^2 - \frac{\operatorname{cov}(v, Y)^2}{\operatorname{var}(Y)} \\ &= \sigma_v^2 - \frac{\beta^2 M^2 \sigma_v^4}{\beta^2 V + \sigma_z^2}. \end{aligned} \tag{36}$$

From the expression for β,

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

Dividing numerator and denominator of the second term in (36) by β^2,

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

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,

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

which gives (10). □

Proof of Proposition 2. Let n traders delegate and $M-n$ do not. A delegated signal is

$$s_i = v + \eta + \varepsilon_i^d, \qquad \operatorname{var}(\eta) = \sigma_\eta^2, \qquad \operatorname{var}(\varepsilon_i^d) = A + u - \sigma_\eta^2,$$

and an independent reading is

$$q_j = v + \varepsilon_j^o, \qquad \operatorname{var}(\varepsilon_j^o) = r + u.$$

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

$$S_d = \operatorname{var}(s_i) = \sigma_v^2 + A + u, \qquad S_o = \operatorname{var}(q_j) = \sigma_v^2 + r + u,$$

and

$$a_d = \frac{\sigma_v^2}{S_d}, \qquad a_o = \frac{\sigma_v^2}{S_o}, \qquad b = \frac{\sigma_v^2 + \sigma_\eta^2}{S_d}.$$

For a delegator i,

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

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

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

Joint normality therefore gives

$$\mathrm{E}[v \mid s_i] = a_d s_i, \qquad \mathrm{E}[s_k \mid s_i] = b s_i, \qquad \mathrm{E}[q_j \mid s_i] = a_d s_i. \tag{37}$$

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

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

and

$$\mathrm{E}[v \mid q_j] = \mathrm{E}[s_i \mid q_j] = \mathrm{E}[q_k \mid q_j] = a_o q_j. \tag{38}$$

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,

$$\begin{aligned}\mathrm{E}[(v-p)x_i \mid s_i] &= x_i\,\mathrm{E}\left[v-\lambda\left(x_i+\sum_{\substack{k\neq i\\ k\in d}}x_k+\sum_{j\in o}x_j+z\right)\middle| s_i\right]\\ &= x_i a_d s_i-\lambda x_i^2-\lambda x_i(n-1)\beta_d b s_i-\lambda x_i(M-n)\beta_o a_d s_i\\ &= x_i\Big\{a_d-\lambda[(n-1)\beta_d b+(M-n)\beta_o a_d]\Big\}s_i-\lambda x_i^2. \end{aligned}\tag{39}$$

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

$$2\lambda\beta_d = a_d-\lambda\big[(n-1)\beta_d b+(M-n)\beta_o a_d\big].\tag{40}$$

Now consider an own reader j. She faces n delegators and $M-n-1$ other own readers. Using (38),

$$\begin{aligned}\mathrm{E}[(v-p)x_j \mid q_j] &= x_j a_o q_j-\lambda x_j^2-\lambda x_j n\beta_d a_o q_j-\lambda x_j(M-n-1)\beta_o a_o q_j\\ &= x_j\Big\{a_o-\lambda[n\beta_d a_o+(M-n-1)\beta_o a_o]\Big\}q_j-\lambda x_j^2. \end{aligned}\tag{41}$$

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

$$2\lambda\beta_o = a_o-\lambda\big[n\beta_d a_o+(M-n-1)\beta_o a_o\big].\tag{42}$$

Define

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

Then (40)–(42) become the linear system

$$[2+(n-1)b]\kappa_d+(M-n)a_d\kappa_o = a_d,\tag{43}$$

$$na_o\kappa_d+[2+(M-n-1)a_o]\kappa_o = a_o.\tag{44}$$

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

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

and the second gives

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

Therefore

$$\begin{aligned}\mu &= \frac{(2-a_o)a_d}{a_o[2+(n-1)b-na_d]} \\ &= \frac{(2-a_o)S_o/S_d}{2-b+n\sigma_\eta^2/S_d},\end{aligned} \tag{45}$$

where the second equality uses

$$\frac{a_d}{a_o}=\frac{S_o}{S_d}, \qquad 2+(n-1)b-na_d = 2-b+n(b-a_d), \qquad b-a_d=\frac{\sigma_\eta^2}{S_d}.$$

The objective of either type has the form

$$x\,h\,s-\lambda x^2,$$

where x is the trader's order, s her signal, and h the term in braces in (39) or (41). At the optimum, the first-order condition gives $hs = 2\lambda x$. Substituting this back into the objective,

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

Taking expectations and using $x_i=\beta_d s_i$ for a delegator and $x_j=\beta_o q_j$ for an own reader,

$$\pi_d=\lambda\beta_d^2 S_d=\frac{\kappa_d^2 S_d}{\lambda}, \qquad \pi_o=\lambda\beta_o^2 S_o=\frac{\kappa_o^2 S_o}{\lambda}. \tag{46}$$

Hence an interior continuous-composition equilibrium satisfies

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

All trading intensities are positive, so

$$\frac{\kappa_d}{\kappa_o}=\sqrt{\frac{S_o}{S_d}}. \tag{47}$$

Combining (45) and (47),

$$\frac{(2-a_o)S_o/S_d}{2-b+n\sigma_\eta^2/S_d}=\sqrt{\frac{S_o}{S_d}}.$$

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

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

Solving for n gives

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

which is (11).

To show that delegation is a strategic substitute, combine (45) with (46):

$$\frac{\pi_d(n)}{\pi_o(n)} = \left(\frac{\kappa_d}{\kappa_o}\right)^2 \frac{S_d}{S_o} = \frac{(2 - a_o)^2 S_o / S_d}{(2 - b + n\sigma_\eta^2 / S_d)^2}. \tag{48}$$

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 $r = A$, $S_o = S_d$ and $a_o = a_d$. Substituting into (11) gives

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

because $a_o S_d = \sigma_v^2$. Economically, a single delegator does not share the model error with any other trader. When $r = A$, 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 $r < A$, delegation is less precise even for the first user, so no trader delegates. If $r > A$, delegation improves precision at entry and adoption begins. Hence the threshold is $r^* = A$.

Differentiating (11) with respect to r gives

$$\sigma_\eta^2 \frac{\partial n^*}{\partial r} = \frac{\sigma_v^2}{S_o^2}\sqrt{S_o S_d} + \frac{2 - a_o}{2}\sqrt{\frac{S_d}{S_o}} > 0. \tag{49}$$

Thus harder disclosures attract more delegation above the threshold. □

Proof of Lemma 1. Start from the interior conditions in the proof of Proposition 2. Profit indifference (47) gives $\kappa_d / \kappa_o = \sqrt{S_o / S_d} = t$. Substituting $\kappa_d = t\kappa_o$ into the own-reader condition (44),

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

and collecting terms,

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

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

$$\begin{aligned} n^* &= \frac{(2-a_o)tS_d-2S_d+\sigma_v^2+\sigma_\eta^2}{\sigma_\eta^2}=\frac{\sigma_\eta^2+2(t-1)S_d+\sigma_v^2-\sigma_v^2/t}{\sigma_\eta^2} \\ &= 1+\frac{(t-1)(2S_d+\sigma_v^2/t)}{\sigma_\eta^2}. \end{aligned}$$

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),

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

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

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

which proves the m-representation. Profit indifference $\lambda\beta_d^2S_d=\lambda\beta_o^2S_o$ 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 number of readers m.

For price impact, write $Y=\beta_d\sum_{i\in d}s_i+\beta_o\sum_{j\in o}q_j+z$ with $\beta_d=\kappa_d/\lambda$ and $\beta_o=\kappa_o/\lambda$. Then $\lambda^2\operatorname{var}(Y)=Q+\lambda^2\sigma_z^2$, where

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

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

$$Q=\kappa_o^2W, \qquad W=MS_o+\sigma_v^2\big(m^2-t^2n-(M-n)\big)+t^2n(n-1)\sigma_\eta^2.$$

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 $S_o/t^2 = S_d$,

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

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

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

Then

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

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

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

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

For a mean-zero normal variable Z, $\mathrm{E}\,|Z| = \sqrt{2/\pi}\sqrt{\mathrm{var}(Z)}$, so (15) becomes

$$|OIB| = \frac{\sqrt{\beta^2\,\mathrm{var}(\sum_i s_i) + \sigma_z^2}}{M\beta\sqrt{\sigma_v^2 + \tau^2} + \sigma_z}. \tag{50}$$

The signal sum has variance

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

Divide the numerator and denominator of (50) by $M\beta$. Since $M\beta \to \infty$ and

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

I obtain

$$|OIB| \longrightarrow \sqrt{\frac{\sigma_v^2 + \rho\tau^2}{\sigma_v^2 + \tau^2}}. \tag{51}$$

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

$$|OIB|_o \longrightarrow \sqrt{\frac{\sigma_v^2}{\sigma_v^2 + r + u}}.$$

For forced full delegation, $\rho\tau^2 = \sigma_\eta^2$ and $\tau^2 = A + u$, which gives

$$|OIB|_d \longrightarrow \sqrt{\frac{\sigma_v^2 + \sigma_\eta^2}{\sigma_v^2 + A + u}}.$$

These are the two expressions in (18).

Full delegation is more one-sided if

$$r(\sigma_v^2 + \sigma_\eta^2) + u\sigma_\eta^2 > \sigma_v^2(A - \sigma_\eta^2), \tag{52}$$

which holds whenever $r > A$. The full-delegation term is independent of r, while the all-own term decreases with r, so $\partial\Delta|OIB|^{FD}/\partial r > 0$.

For uncertainty,

$$\frac{\partial\Delta|OIB|^{FD}}{\partial u} = -\frac{\sqrt{\sigma_v^2 + \sigma_\eta^2}}{2(\sigma_v^2 + A + u)^{3/2}} + \frac{\sigma_v}{2(\sigma_v^2 + r + u)^{3/2}} < 0. \tag{53}$$

The first term is larger in absolute value because it has a larger numerator and, when $r > A$, a smaller denominator. □

Proof of Proposition 4. 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

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

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

$$\text{cov}(Y_f, Y_\ell) = (M\beta)^2\chi^2. \tag{54}$$

For delegated signals,

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

so

$$\text{var}(Y_f) = \beta^2 \left[M(\sigma_v^2 + A + u) + M(M-1)(\sigma_v^2 + \sigma_\eta^2)\right] + \sigma_z^2.$$

Proposition 1 gives

$$\beta^2 = \frac{\sigma_z^2}{M(\sigma_v^2 + A + u)}, \qquad \frac{\sigma_z^2}{\beta^2 M} = \sigma_v^2 + A + u.$$

Therefore, using (16),

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

which, together with $\varsigma_o = 0$, is (19). The expression is linear and increasing in χ^2, independent of r, and decreasing in u. □

Proof of Proposition 5. Under all-own reading and full delegation, respectively,

$$\widetilde{D}_o = (M+1)\sigma_v^2 + 2(r+u), \qquad \widetilde{D}_d = (M+1)\sigma_v^2 + (M-1)\sigma_\eta^2 + 2(A+u).$$

Hence

$$\widetilde{D}_d - \widetilde{D}_o = 2A + (M-1)\sigma_\eta^2 - 2r. \tag{55}$$

Posterior variance $V = \sigma_v^2 - M\sigma_v^4/\widetilde{D}$ increases with $\widetilde{D}$, so full delegation improves price discovery if and only if

$$r > A + \frac{M-1}{2}\sigma_\eta^2 \equiv \bar{r}.$$

For $r > \bar{r}$, the reduction in pricing error is $\sqrt{V_o} - \sqrt{V_d} > 0$. It increases with r because V_o increases while V_d is unchanged. Moreover,

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

As u rises, the numerator is unchanged and both denominator terms increase, so $V_o - V_d$ shrinks. Both variances also increase, and thus

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

decreases with u and converges to zero as $u \to \infty$. □

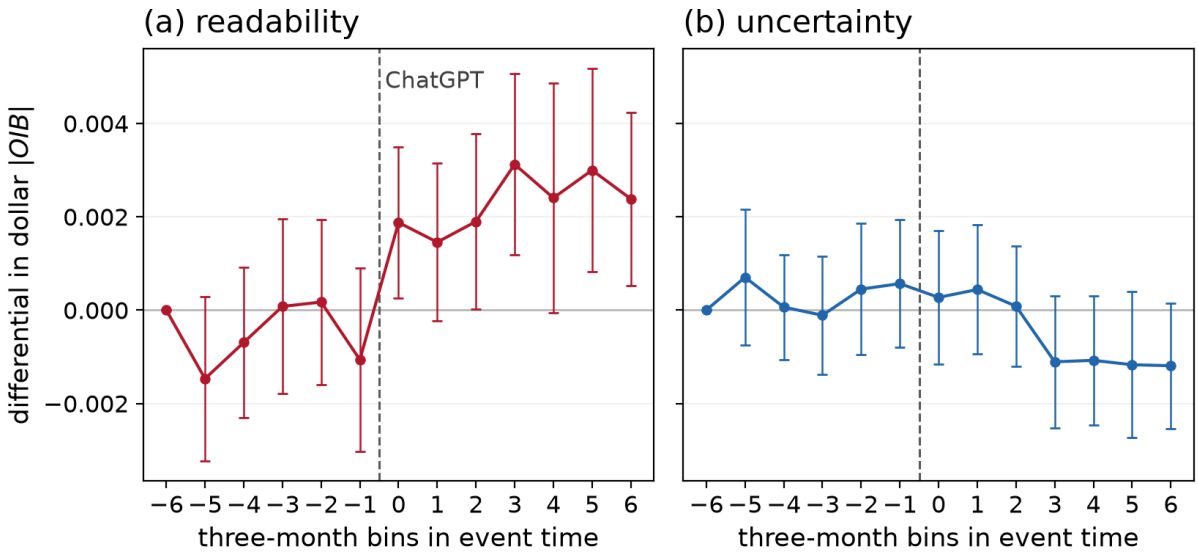
Proof of Corollary 1. The parameter $a_o = \sigma_v^2/S_o$ depends only on the own-reading technology and is therefore the same in the equilibrium and the all-own benchmark. By Lemma 1, both posterior variances take the form $\sigma_v^2(2-a_o)/(2+(m-1)a_o)$, which is decreasing in the effective number of readers 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 $r \leq A$. 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:

$$n^* \geq M \quad \Longrightarrow \quad (M-1)\sigma_\eta^2 \leq (t-1)\left(2S_d + \frac{\sigma_v^2}{t}\right) < 2S_d(t^2-1) = 2(r-A),$$

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 $r > A + \frac{1}{2}(M-1)\sigma_\eta^2 = \bar{r}$, and Proposition 5 shows that full delegation lowers pricing error in that region. □

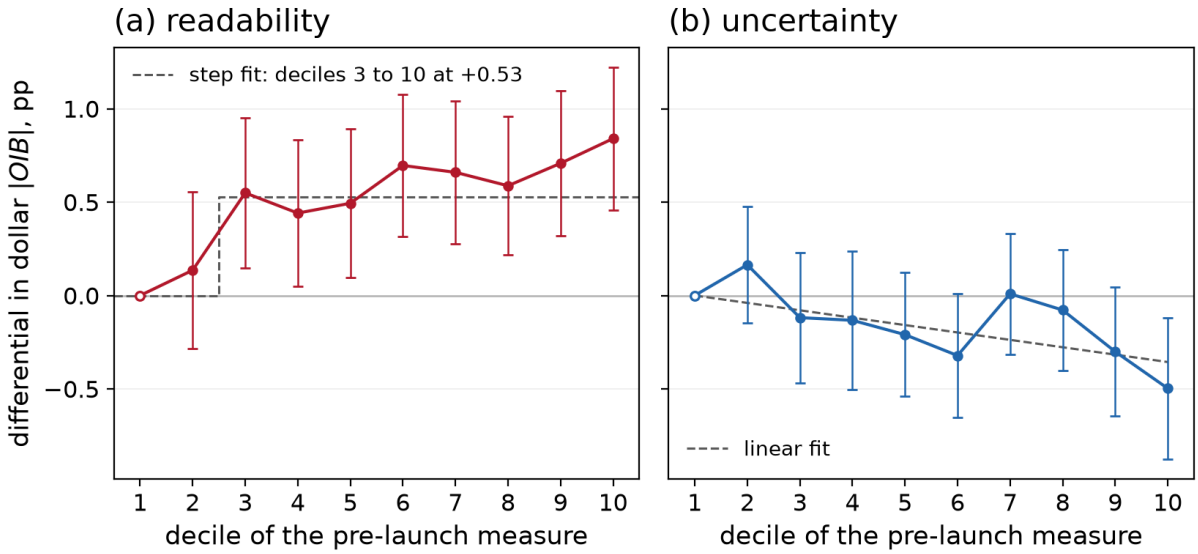
Figures and tables

Figure 1: Event study for the synchronization result



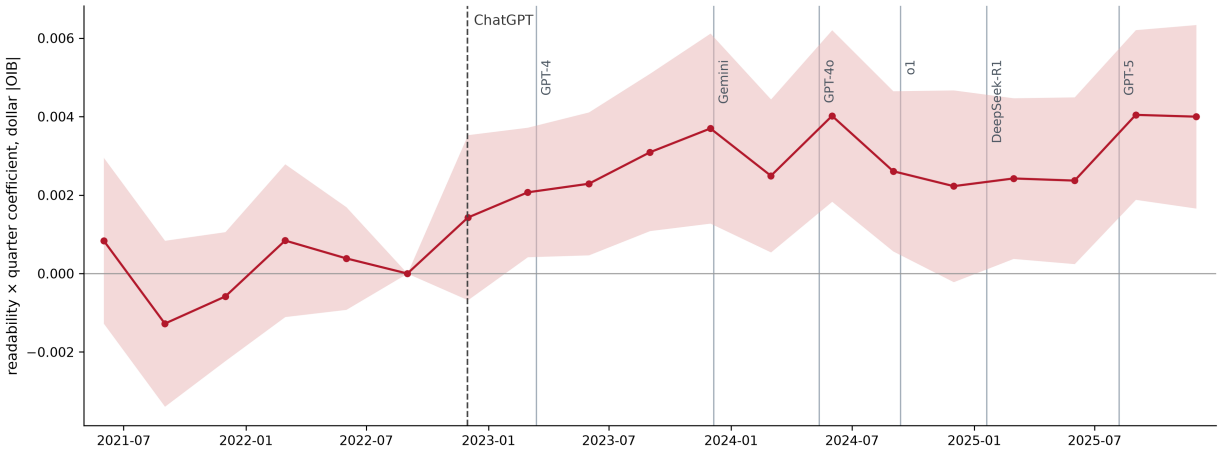
Notes: Coefficients β_{dk} from $y_{im} = \alpha_i + \gamma_m + \sum_{d\in\{R,U\}} \sum_{k\neq -6} \beta_{dk}(\mathbf{1}\{K_m = k\} \times D_{di}) + \sum_{k\neq -6} \psi'_k(\mathbf{1}\{K_m = k\} \times Z_i) + \theta' Act_{im} + e_{im}$. The dependent variable is dollar $|OIB|$ in raw units of the imbalance share. D_{di} is the firm's standardized pre-launch readability (R) or uncertainty (U), both entered jointly. Event time K_m is in three-month bins from the launch, capped at six bins on each side, so the endpoint bins absorb the tails of the sample; bin 0 runs from December 2022 through February 2023, and bin -6 is omitted. Z_i contains the eleven predetermined controls, 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. The sample contains 148,984 firm-months for 3,670 firms from January 2021 through December 2025.

Figure 2: The response across the difficulty distribution



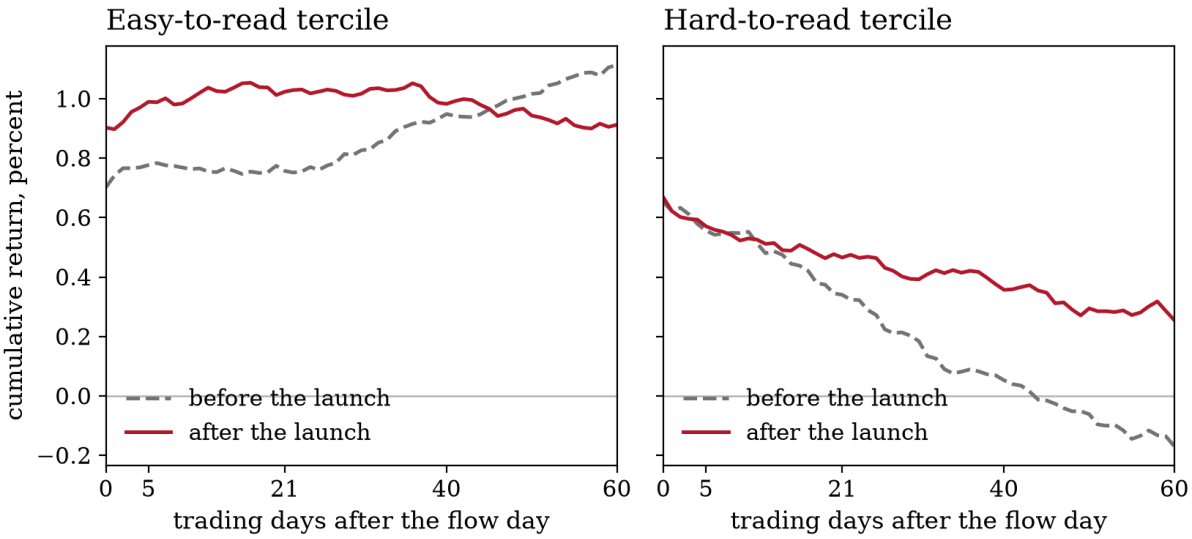
Notes: Coefficients β_j from $y_{im} = \alpha_i + \gamma_m + \sum_{j=2}^{10} \beta_j \left(Post_m \times \mathbf{1}\{decile_i = j\}\right) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, where y_{im} is the firm-month mean of dollar $|OIB|$ for total flow, in percentage points, and each panel uses the deciles of one disclosure dimension across firm-months. The first decile (open circle) is omitted, so each coefficient is relative to the lowest decile of the same dimension. Bars are 95 percent confidence intervals from standard errors clustered by firm and month. The dashed line in panel (a) is a single step, the coefficient on $Post_m \times \mathbf{1}\{decile_i \geq 3\}$ when it replaces the decile indicators; the dashed line in panel (b) is the fit of a linear decile score, $Post_m \times (decile_i - 5.5)/2.87$, drawn relative to the first decile. Z_i contains the eleven predetermined controls, 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. The sample contains 148,984 firm-months for 3,670 firms from January 2021 through December 2025.

Figure 3: The readability differential by quarter



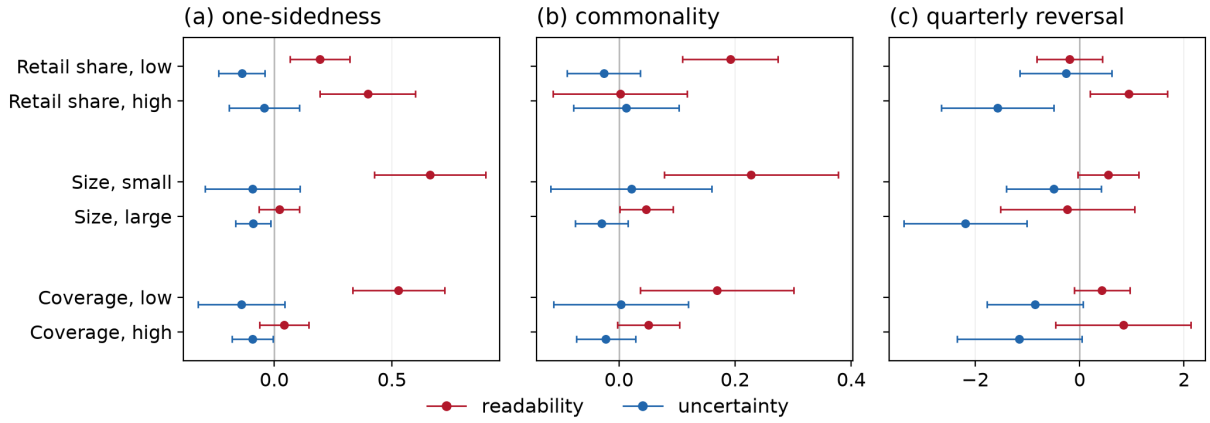
Notes: Coefficients β_{Rk} on readability × three-month-interval indicators from $y_{im} = \alpha_i + \gamma_m + \sum_{d\in\{R,U\}} \sum_{k\neq -1} \beta_{dk}(\mathbf{1}\{Q_m = k\} \times D_{di}) + \sum_{k\neq -1} \psi'_k(\mathbf{1}\{Q_m = k\} \times Z_i) + \theta' Act_{im} + e_{im}$, estimated for dollar $|OIB|$ with both disclosure dimensions entered jointly. Event time Q_m is in three-month intervals from the ChatGPT launch, clipped at six intervals before and twelve after, so the leftmost point pools the earlier pre-launch months; interval -1 is omitted. Z_i contains the eleven predetermined controls, Act_{im} contains log dollar volume and log intraday intervals, 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 4: The price path after one-sided flow days



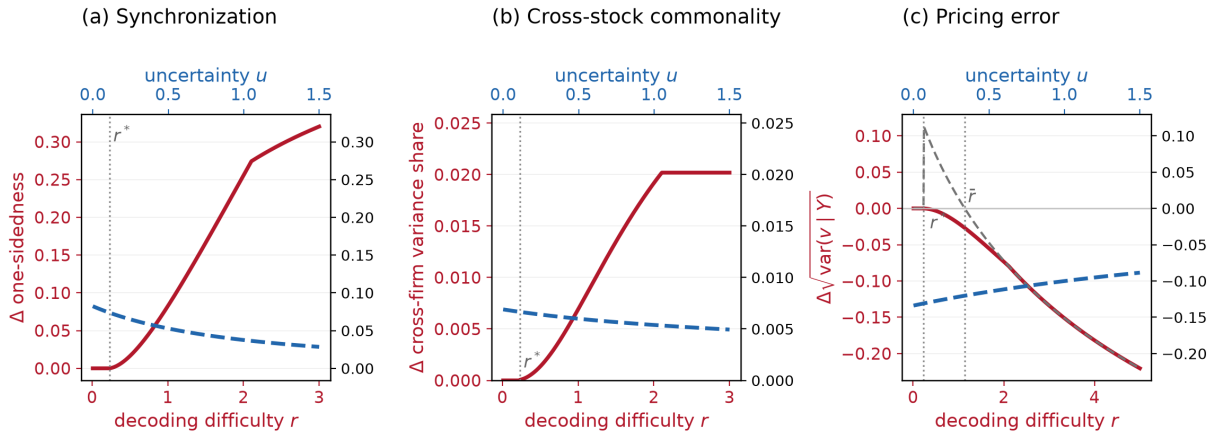
Notes: Each line plots the mean cumulative market-adjusted return after the flow days in its cell, at horizons h from 0 to 60 trading days. A flow day is a firm-day whose absolute dollar imbalance is in the top decile of the panel, $|OIB|$ of at least 0.295 (the larger side accounts for at least 65 percent 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 ChatGPT launch 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 drops the last three months.

Figure 5: The effects within halves of the sample



Notes: Each point is the post-launch differential per standard deviation of readability (red) or uncertainty (blue), re-estimated within the stated half of the sample with both disclosure dimensions entered jointly; *Post* equals one from December 2022 onward, and bars are 95 percent confidence intervals. 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, the firm-month mean of absolute dollar imbalance in percentage points, and commonality, the firm-month slope of daily signed imbalance on the equal-weighted signed imbalance of the other sample firms, are regressed on $Post \times D$ with firm and month fixed effects, post interactions with the eleven predetermined controls, and log dollar volume and log intraday intervals. Quarterly reversal is the post-launch 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, the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and the interactions of signed imbalance and of its product with *Post* with filing length; a positive value means less reversal of the price impact. 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.

Figure 6: Numerical evaluation of the model



Notes: Each panel evaluates the model at the endogenous composition n^*, the number of delegating traders. The solid curve uses the bottom and left axes and varies decoding difficulty r with uncertainty fixed at $u = 0$. The dashed curve uses the top and right axes and varies uncertainty u, holding $r = 1$ in panels (a) and (b) and $r = 3$ in panel (c). Each panel plots the value of an outcome at n^* minus its value when every trader reads independently. Panel (a) plots the change in within-stock one-sidedness, $|OIB| = \mathrm{E}\,|Y|/(\sum_{i=1}^{M} \mathrm{E}\,|x_i| + \mathrm{E}\,|z|)$, where x_i is the order of speculator i, z is noise trading, and $Y = \sum_{i=1}^{M} x_i + z$ is aggregate order flow; panel (b) the change in cross-stock order-flow commonality, $\varsigma = \mathrm{cov}(Y_f, Y_\ell)/\,\mathrm{var}(Y_f)$, for two otherwise identical firms f and ℓ; and panel (c) the change in pricing error, $\mathcal{P} = \sqrt{\mathrm{var}(v \mid Y)}$, the standard deviation of the asset's value v given aggregate order flow. The gray dashed curve in panel (c) is the forced full-delegation benchmark, in which every trader delegates. Dotted vertical lines mark the entry threshold $r^* = 0.23$ and, in panel (c), the price-quality boundary $\bar{r} = 1.14$. The calibration sets $\sigma_v^2 = 1$, $\sigma_z^2 = 4$, $M = 20$, $A = 0.23$, $\sigma_\eta^2 = 0.096$, and $\chi^2 = 0.0235$.

Table 1: 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: eleven 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
Analysts per announcement	4.0	3.9	1.3	3.0	5.4

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 eleven predetermined controls whose interactions with the post indicator enter the order-flow and dispersion specifications: filing length, constructed from the same filings and standardized, with the raw word count as a memo row, and ten firm characteristics in raw units, nine of them computed from daily TAQ and CRSP data from January 1, 2022 through November 29, 2022. 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 five-factor model; effective spread is the average daily percent effective spread; momentum is the cumulative return over the window; and 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. Analysts per announcement is the mean number of IBES forecasts outstanding at the firm's earnings announcements from January 2021 through November 2022, zero for firms without coverage; it enters the regressions as the log of one plus this number. In the regressions, each characteristic is standardized across firms. All variables are winsorized at the 1st and 99th percentiles.

Table 2: 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
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 the 3,876 firms with nonmissing readability, uncertainty, and length, one observation per firm. 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 ten pre-launch characteristics in raw units. The first nine are computed from daily TAQ and CRSP data from January 1, 2022 through November 29, 2022 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 five-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 forecasts outstanding at the firm's earnings announcements from January 2021 through November 2022, zero for firms without coverage, and 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 3: Disclosure difficulty and synchronized trading

	Mean	Readability	Uncertainty
Panel A: order imbalance, by measure and investor type			
Dollar $\|OIB\|$, total flow	13.1	$+0.28^{***}$	-0.10^{**}
each dimension entered alone		$+0.30^{***}$	-0.15^{***}
non-retail flow	13.9	$+0.26^{***}$	-0.07
retail flow	24.4	$+0.23^{**}$	-0.38^{***}
Trade-count $\|OIB\|$, total flow	10.1	$+0.14^{***}$	-0.06
Signed dollar OIB, total flow		$+0.03$	$+0.04$
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 $y_{im} = \alpha_i + \gamma_m + \sum_{d\in\{R,U\}} \beta_d(Post_m \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, with the row's variable as y_{im}, 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 eleven predetermined controls, 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 4: The timing of the differential

Panel A: the year path of the differential			
	2023	2024	2025
Readability	$+0.18^{***}$	$+0.28^{***}$	$+0.42^{***}$
Uncertainty	-0.05	-0.12^{**}	-0.05

Panel B: the launch step against the cumulative model stock			
	Stock alone	Entered together	
	$N(m) \times D$	$Post \times D$	$N(m) \times D$
Readability	$+0.50^{***}$	$+0.13^{*}$	$+0.39^{***}$
Uncertainty	-0.05	-0.11^{**}	$+0.03$

Notes: Panel A reports coefficients β_{yd} from $y_{im} = \alpha_i + \gamma_m + \sum_y \sum_{d\in\{R,U\}} \beta_{yd}(\mathbf{1}\{year_m = y\} \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$ for $y \in \{2023, 2024, 2025\}$, with the pre-launch months of 2021 and 2022 as the base period. Panel B reports coefficients β_{1d} and β_{2d} from $y_{im} = \alpha_i + \gamma_m + \sum_{d\in\{R,U\}}[\beta_{1d}(Post_m \times D_{di}) + \beta_{2d}(N(m) \times D_{di})] + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$; 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} is the firm's standardized pre-launch readability (R) or uncertainty (U), both entered jointly. y_{im} is the firm-month mean of absolute dollar 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, scaled to equal one in the final month of the sample, Z_i contains the eleven predetermined controls, 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 5: 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 $y_{ie} = \alpha_i + \gamma_{q(e)} + \sum_{d\in\{R,U\}} \beta_d(Post_e \times D_{di}) + \psi'(Post_e \times Z_i) + \theta \log N_{ie} + \varepsilon_{ie}$, 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 eleven 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 6: Cross-stock flow commonality

	Readability	Uncertainty
Panel A: the post-launch differential		
Total flow	$+0.105^{***}$	-0.010
Non-retail	$+0.090^{***}$	-0.011
Retail	$+0.033$	-0.044
Panel B: with index and ownership controls		
Total flow	$+0.090^{***}$	$+0.008$
Non-retail	$+0.080^{**}$	$+0.008$
Panel C: the year path of the differential, total flow		
2023	$+0.150^{***}$	$+0.005$
2024	$+0.150^{***}$	$+0.006$
2025	-0.040	$+0.015$

Notes: Panel A reports coefficients δ_d from $b_{im} = \alpha_i + \gamma_m + \sum_{d\in\{R,U\}} \delta_d(Post_m \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + \nu_{im}$ for the flow type named in the row. Panel B adds three predetermined controls to Z_i: an indicator for S&P 500 membership on November 29, 2022, the firm's mean 13F institutional ownership share from January 1 through November 29, 2022, and the market value of exchange-traded fund holdings on June 30, 2022 divided by the firm's market equity in June 2022, the last two standardized across firms. Panel C reports coefficients δ_{yd} from $b_{im} = \alpha_i + \gamma_m + \sum_y \sum_{d\in\{R,U\}} \delta_{yd}(\mathbf{1}\{year_m = y\} \times D_{di}) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + \nu_{im}$ for total flow, with the pre-launch months of 2021 and 2022 as the base period. Readability and uncertainty enter 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. All panels include firm and month fixed effects, post interactions with the eleven 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; Panel B uses the 135,397 total-flow and 135,394 non-retail firm-months with all three additional controls. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 7: Price impact and its persistence

	Before the launch			Change after the launch		
Return window	Average firm β_1	Readability β_{3R}	Uncertainty β_{3U}	Average firm β_2	Readability β_{4R}	Uncertainty β_{4U}
[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^{***}

Notes: Coefficients from $CAR_{i,[a,b]} = \alpha_i + \gamma_t + OIB_{it}(\beta_1 + \beta_2 Post_t) + \sum_{d\in\{R,U\}} OIB_{it}(\beta_{3d} D_{di} + \beta_{4d} Post_t \times D_{di}) + \phi' W_{it} + e_{it}$, estimated separately for each return window with both disclosure dimensions entered jointly. The dependent variable $CAR_{i,[a,b]}$ is the sum over trading days $[t+a, t+b]$ of the firm's daily log return minus the equal-weighted log return of the sample firms, in percent, where t is the day of the flow; [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. The first three columns describe the slope before the launch: β_1 on OIB_{it} is the slope for a firm with both dimensions and filing length at zero, their means across firms, and β_{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: β_2 on $OIB_{it} \times Post_t$ for the same firm, and β_{4d} on $OIB_{it} \times Post_t \times D_{di}$ per standard deviation of the dimension named in the column. W_{it} contains the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and the interactions of OIB_{it} and $OIB_{it} \times Post_t$ with filing length, the standardized sum of log net file size and log word count of the same filings; 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 4,009 firms at the impact window; the 60-day window drops the last three months. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 8: Earnings-announcement tests

Dependent variable	Readability	Uncertainty
Abnormal volume	$+1.92^{**}$	$+1.40^{**}$
Incorporation speed	$+0.79^{**}$	$+0.28$
Drift	-1.41^{**}	-0.53

Notes: Coefficients β_d from $y_{ie} = \alpha_i + \gamma_{q(e)} + \sum_{d\in\{R,U\}} \beta_d (Post_e \times D_{di}) + \varphi\, Rank_e + \theta' X_{ie} + \varepsilon_{ie}$, estimated at the announcement level with both dimensions entered jointly; for drift, the coefficient on the announcement-return rank interacted with $Post_e \times D_{di}$, the change in the return spread over days $[+2, +21]$ between good- and bad-news announcements. Abnormal volume is the log of mean dollar volume over days $[-1, +1]$ minus that over days $[-30, -6]$. Incorporation speed is $|CAR_{[-1,+1]}|/(|CAR_{[-1,+1]}| + |CAR_{[+2,+21]}|)$, where $CAR_{[a,b]}$ is the sum over trading days $[a, b]$ around the announcement of the firm's daily return minus the equal-weighted return of the sample firms, and drift is $CAR_{[+2,+21]}$. All three outcomes are multiplied by 100. The announcement-return rank $Rank_e$ is the within-quarter percentile rank of $CAR_{[-1,+1]}$, 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. X_{ie} contains lagged size, volatility, and volume, and filing length interacted with the post indicator; the drift specification additionally interacts the rank with each disclosure dimension, filing length, 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.

Appendix tables

Table A1: Length entered as a third dimension

	Readability	Length	Uncertainty
Panel A: specification, dependent variable total dollar $\lvert OIB \rvert$			
Each dimension alone	$+0.30^{***}$	-0.02	-0.13^{***}
Three dimensions jointly	$+0.28^{***}$	$+0.04$	-0.10^{**}
Panel B: absolute imbalance, by measure and investor type			
Dollar $\lvert OIB \rvert$, total flow	$+0.28^{***}$	$+0.04$	-0.10^{**}
non-retail	$+0.26^{***}$	$+0.01$	-0.07
retail	$+0.23^{**}$	-0.12	-0.38^{***}
Trade-count $\lvert OIB \rvert$, total flow	$+0.14^{***}$	-0.06	-0.06
non-retail	$+0.13^{**}$	-0.07^{*}	-0.06
retail	$+0.05$	-0.16^{*}	-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 ten 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 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.

Table A2: 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 A3: 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 A4: Estimates on the firms common to both panels

	Full sample		Common firms	
Test and coefficient	Readability	Uncertainty	Readability	Uncertainty
Dollar $\lvert OIB \rvert$ on $Post \times D$	$+0.282^{***}$	-0.095^{**}	$+0.282^{***}$	-0.095^{**}
Trade-count $\lvert OIB \rvert$ on $Post \times D$	$+0.145^{***}$	-0.064	$+0.145^{***}$	-0.064
$CAR_{[0]}$ 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 eleven predetermined controls, and log dollar volume and log intraday intervals. The last two rows regress the cumulative market-adjusted log return over the stated window, 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 the interactions of signed imbalance and of its product with *Post* 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 A5: Uncertainty language and uncertainty about fundamentals

		Uncertainty		Readability	
Pre-launch measure	Firms	Spearman	Net of size	Spearman	Net of size
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 text measure named in the column and the pre-launch measure named in the row. Spearman is the rank correlation, and Net of size 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 five-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 A6: Market-wide means before and after the launch

	Before	After
Panel A: absolute order imbalance, firm-month means in percentage points		
Dollar $\lvert OIB\rvert$	13.79	12.50
Trade-count $\lvert OIB\rvert$	10.68	9.68
Dollar $\lvert OIB\rvert$, non-retail	14.71	13.35
Dollar $\lvert OIB\rvert$, 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 ten 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 A7: The two components of readability

$Post\times$ component	Dollar $\lvert OIB\rvert$	Trade-count $\lvert OIB\rvert$
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	$+0.33^{***}$	$+0.19^{***}$
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 ten firm characteristics, and log dollar volume and log intraday intervals; the last three rows add the post interactions with length and uncertainty, so that Z_i contains all eleven 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 A8: 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: $Post \times D_{Ri}$				
Dollar $\lvert OIB \rvert$	$+0.30^{***}$	$+0.25^{***}$	$+0.23^{***}$	$+0.14^{***}$
Trade-count $\lvert OIB \rvert$	$+0.16^{***}$	$+0.11^{**}$	$+0.07$	$+0.11^{***}$

Notes: Panel A reports the Spearman rank correlation of each exposure with the exposure in use, and Panel B the coefficient on $Post \times D_{Ri}$, where D_{Ri} is the readability exposure named in the column, in percentage points of $|OIB|$, with each exposure entered separately. Exposures use the filings from January 2021 to November 29, 2022, except as stated, and are 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 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 eleven predetermined controls, and controls for log dollar volume and log intraday intervals. Standard errors are clustered by firm and month. The sample contains 148,697 firm-months with nonmissing values for all four exposures, filing length, the ten firm characteristics, and both activity controls. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A9: Estimates by change in vocabulary complexity

	Dollar $\lvert OIB \rvert$		Trade-count $\lvert OIB \rvert$	
	Readability	Uncertainty	Readability	Uncertainty
All firms with Loughran–McDonald ranks	$+0.29^{***}$	-0.09^{**}	$+0.15^{***}$	-0.06
Firms that did not move	$+0.26^{***}$	-0.09^{**}	$+0.14^{**}$	-0.05
Firms that moved	$+0.51^{**}$	-0.09	$+0.38^{*}$	-0.20
Excluding firms that became easier	$+0.29^{***}$	-0.09^{**}	$+0.16^{***}$	-0.05
Excluding firms that became harder	$+0.26^{***}$	-0.10^{**}	$+0.13^{**}$	-0.06
All firms, with the current complex-word share	$+0.24^{***}$	-0.08^{*}	$+0.12^{**}$	-0.05

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 eleven predetermined controls, and log dollar volume and log intraday intervals. 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 order. The last row adds to the baseline the firm's latest Loughran–McDonald complex-word share as of each month, standardized on the exposure window, and its interaction with $Post_m$; it uses the 148,113 firm-months with a filing in the preceding 400 days. The current share enters with a coefficient of -0.10 for dollar imbalance and -0.11 for trade-count imbalance, and its post-launch interaction with $+0.10$ and $+0.08$, all significant at the 1 percent level. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A10: Estimates without the retail trading episode of early 2021

	One-sidedness	Commonality	Quarterly reversal	Incorporation speed	Drift
Panel A: readability					
Full sample	$+0.282^{***}$	$+0.108^{***}$	$+0.568^{**}$	$+0.79^{**}$	-1.41^{**}
Pre-launch period from July 2021	$+0.290^{***}$	$+0.127^{***}$	$+0.650^{**}$	$+0.69$	-0.95
Without the top meme-intensity decile	$+0.249^{***}$	$+0.134^{***}$	$+0.676^{***}$	$+0.75^{*}$	-1.44^{**}
With *Post*× meme intensity	$+0.294^{***}$	$+0.105^{***}$	$+0.463^{*}$	$+0.79^{**}$	-1.41^{**}
Panel B: uncertainty					
Full sample	-0.095^{**}	-0.011	-0.986^{***}	$+0.28$	-0.53
Pre-launch period from July 2021	-0.106^{**}	$+0.000$	-1.328^{***}	$+0.12$	-0.10
Without the top meme-intensity decile	-0.113^{**}	-0.030	-1.241^{***}	$+0.32$	-0.52
With *Post*× meme intensity	-0.074^{*}	-0.017	-0.933^{**}	$+0.28$	-0.55

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, the firm-month mean of absolute dollar imbalance in percentage points, and commonality, the firm-month slope of daily signed imbalance on the equal-weighted signed imbalance of the other sample firms, are regressed on $Post \times D$ with firm and month fixed effects, post interactions with the eleven predetermined controls, and log dollar volume and log intraday intervals. 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, the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and the interactions of signed imbalance and of its product with *Post* with filing length; in the fourth row, signed imbalance and its product with *Post* are also interacted with meme intensity. A positive value means less reversal of the price impact. Incorporation speed is the absolute cumulative market-adjusted 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 $Post \times D$; 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 multiplied by 100 and 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; the drift regression also interacts the rank with each disclosure dimension, filing length, and the post indicator. In the fourth row, the rank and its product with *Post* are further interacted with meme intensity. 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 and 39,098 announcements. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A11: Estimates with negative tone as a control

	Readability		Uncertainty		Tone
	Baseline	With tone	Baseline	With tone	
Dollar $\lvert OIB \rvert$, total flow	$+0.282^{***}$	$+0.283^{***}$	-0.095^{**}	-0.090	-0.008
Dollar $\lvert OIB \rvert$, retail flow	$+0.232^{**}$	$+0.241^{**}$	-0.377^{***}	-0.328^{***}	-0.077
Commonality, total flow	$+0.105^{***}$	$+0.097^{***}$	-0.010	-0.050	$+0.064$
Log forecast dispersion	-0.051^{**}	-0.052^{**}	$+0.060^{**}$	$+0.056^{*}$	$+0.006$
Price impact, [0]	-0.177^{***}	-0.147^{***}	$+0.101$	$+0.240^{***}$	-0.288^{***}
Quarterly reversal, $[+2, +60]$	$+0.568^{**}$	$+0.669^{***}$	-0.986^{***}	-0.430	-1.159^{***}

Notes: Each cell reports a post-launch differential per standard deviation of the dimension named in the column, with both dimensions entered jointly and *Post* equal to one from December 2022 onward. The columns with tone add negative tone, the frequency of Loughran–McDonald negative words in the firm's 10-K and 10-Q filings from January 1, 2021 through November 29, 2022, standardized across firms, wherever filing length enters, and the last column reports its coefficient. The $\lvert OIB \rvert$ rows use the firm-month mean of absolute daily dollar imbalance, in percentage points, and the commonality row the firm-month slope of daily signed imbalance on the equal-weighted imbalance of the other sample firms, each with firm and month fixed effects, post interactions with the eleven predetermined controls, and log dollar volume and log intraday intervals. The dispersion row uses the log standard deviation of the IBES forecasts outstanding before an earnings announcement, scaled by the share price, with firm and calendar-quarter fixed effects, post interactions with the eleven predetermined controls, and the log number of forecasts. The last two rows report the coefficients on $OIB \times Post \times D$ in regressions of the cumulative market-adjusted log return, in percent, over the same day and over trading days $[+2, +60]$ on signed dollar imbalance and its interactions with *Post*, D, and filing length, with firm and date fixed effects and the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread. The samples contain 148,984 firm-months (148,973 for retail flow), 149,413 for commonality, 34,255 announcements for dispersion, and 3,189,685 and 2,958,342 firm-days for the two return windows. Standard errors are clustered by firm and by month, date, or calendar quarter. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A12: Post-launch differential by decile of disclosure difficulty

	Readability	Uncertainty
Panel A: decile coefficients		
Decile 2	+0.14	+0.16
Decile 3	+0.55***	−0.12
Decile 4	+0.44**	−0.13
Decile 5	+0.49**	−0.21
Decile 6	+0.70***	−0.32*
Decile 7	+0.66***	+0.01
Decile 8	+0.59***	−0.08
Decile 9	+0.71***	−0.30*
Decile 10	+0.84***	−0.50**
Panel B: one-parameter shapes and tests		
Step, deciles 3 to 10 against 1 and 2	+0.53***	−0.27***
Linear decile score, per standard deviation	+0.22***	−0.11***
p-value, decile 2 equal to zero	0.53	0.30
p-value, deciles 3 to 10 equal	0.20	0.10

Notes: Panel A reports the coefficients β_j from $y_{im} = \alpha_i + \gamma_m + \sum_{j=2}^{10} \beta_j \left(Post_m \times \mathbf{1}\{decile_i = j\}\right) + \psi'(Post_m \times Z_i) + \theta' Act_{im} + e_{im}$, 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 relative to the lowest decile of the same dimension. Panel B replaces the decile indicators with a single step, $Post_m \times \mathbf{1}\{decile_i \geq 3\}$, or with a linear decile score, $Post_m \times (decile_i - 5.5)/2.87$, which has unit standard deviation across deciles, and reports Wald tests on the Panel A coefficients that the second decile equals zero and that the third to tenth deciles are equal. All specifications include firm and month fixed effects, post interactions with the eleven 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 A13: Post-launch differentials within halves of the sample

	One-sidedness	Commonality	Quarterly reversal
Panel A: readability			
Retail share, low	$+0.19^{***}$	$+0.192^{***}$	-0.19
Retail share, high	$+0.40^{***}$	$+0.002$	$+0.94^{**}$
Size, small	$+0.66^{***}$	$+0.228^{***}$	$+0.55^{*}$
Size, large	$+0.02$	$+0.047^{**}$	-0.24
Coverage, low	$+0.53^{***}$	$+0.169^{**}$	$+0.43$
Coverage, high	$+0.04$	$+0.051^{*}$	$+0.84$
Panel B: uncertainty			
Retail share, low	-0.14^{***}	-0.026	-0.26
Retail share, high	-0.04	$+0.012$	-1.57^{***}
Size, small	-0.09	$+0.021$	-0.50
Size, large	-0.09^{**}	-0.030	-2.19^{***}
Coverage, low	-0.14	$+0.004$	-0.86^{*}
Coverage, high	-0.09^{**}	-0.022	-1.16^{*}

Notes: Each cell reports the post-launch differential per standard deviation of the dimension named in the panel, 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, the firm-month mean of absolute dollar imbalance in percentage points, and commonality, the firm-month slope of daily signed imbalance on the equal-weighted signed imbalance of the other sample firms, are regressed on $Post \times D$ with firm and month fixed effects, post interactions with the eleven predetermined controls, and log dollar volume and log intraday intervals. Quarterly reversal is the post-launch 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, the lagged return, log dollar volume, log market equity, idiosyncratic volatility, and effective spread, and the interactions of signed imbalance and of its product with *Post* with filing length; a positive value means less reversal of the price impact. 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.