Weak Recovery and the AI Disclosure-Action Wedge[1]

Dingyi Chen[1,2,*], Xiaotian Liu[1], Zijin Wang[3], Renxian Zuo[4]

1.UQ Business School, University of Queensland, Australia.

2.School of Economics and Management, China University of Geosciences (Wuhan), China

3.School of Economics, Shandong University, China

4.School of Information Management, Hubei University of Economics, China

* Corresponding author. Email addresses: uqdche10@uq.edu.au (D. Chen)

Abstract

Technology disclosure is often informative when investors can verify costly adoption inputs. We show that this link weakens when financially weak firms face retail attention that can turn salient AI language into short-horizon trading liquidity. Around the generative-AI shock, we combine a predetermined 2021 low-recovery state with leave-one-out headquarters-city peer AI talk from 2019-2021. In China, low-recovery firms in high-narrative cities increase AI MD&A disclosure and AI talk without broad AI hiring, while broad AI hiring does not rise. AI disclosure rises without observable AI hiring or broad real-input expansion; R&D is lower in the baseline and city-year specifications but is best interpreted as supporting evidence rather than the identifying mechanism. Realized AI talk without hiring is followed by larger insider share reductions among firms with high pre-shock stock-forum attention; the relation is concentrated in market-facing sales and director/executive market-facing transactions. A U.S. benchmark suggests that the insider share-reduction margin is not a generic feature of AI disclosure but is sharper in China's retail-attention setting. Informativeness depends on the sender's financing state and the market's investor structure.

Keywords: AI disclosure; technology narratives; weak recovery; retail attention; insider share reductions; China.

JEL Classification: G14, G30, G32, G34, M41, O33.

---

1 Acknowledgments and Disclosures. We thank Qi Deng, Zhong-guo Zhou, and Jundong Hou for helpful comments and suggestions. This research received no project-specific funding. We have no financial relationships or other potential conflicts of interest to disclose. We have read the NBER research dissemination disclosure policy and attest that this acknowledgment discloses all sources of research support and all material and relevant financial relationships.

1. Introduction

In April 2026, Allbirds offered an out-of-sample illustration of the mechanism studied in this paper. The firm had agreed to sell its footwear brand and assets to American Exchange Group for an estimated $39 million, then announced a $50 million convertible financing facility and a planned pivot into AI compute infrastructure under the NewBird AI name (Allbirds 2026a; Allbirds 2026b). The legacy business was far from the AI frontier, but the technology narrative was immediately tradable: Bloomberg reported that the stock rose 582 percent around the rebranding announcement (Bloomberg News 2026). The episode illustrates the empirical tension studied here. A firm with weak recovery prospects can be far from real technology implementation and still obtain a short-horizon market response by attaching itself to a salient technology category. This episode is only a motivating out-of-sample illustration; the empirical design below fixes firm states before the generative-AI shock and uses Chinese listed-firm data through 2024.

We document the same pattern at scale in a setting where real AI activity should leave traces. Prior work measures AI adoption through specialized labor demand, online vacancies, and firm outcomes (Alekseeva et al. 2021; Acemoglu et al. 2022; Babina, Fedyk, He, and Hodson 2024). Cao, Jiang, Yang, and Zhang (2023) establish AI disclosure as a natural object because firms change disclosure production when machine readers become important. The question here is where the link between words and investment confirmation breaks. AI adoption requires scarce technical workers, data infrastructure, system integration, and organizational redesign; AI language is cheap to add to a filing and hard to verify quickly.

The separation is most likely when recovery prospects are weak. Financing pressure can deter investment when residual claimants capture too little upside (Myers 1977), and innovative activity is especially sensitive to financing capacity and long-horizon experimentation (Brown, Fazzari, and Petersen 2009; Manso 2011). Under weak recovery, costly AI experimentation has little time to change survival probabilities. Disclosure positioning carries scrutiny risk, but limited liability truncates downside value and makes right-tail attention or liquidity more valuable. Local technology narratives amplify this margin because nearby firms have already made AI vocabulary familiar and tradable.

China links disclosure, hiring, real inputs, trading, and insider shareholding changes in a retail-attention market. Retail investors buy attention-grabbing stocks (Barber and Odean 2008), and attention is observable in search and forum activity (Da, Engelberg, and Gao 2011). Technology narratives can generate trading volume even without a positive price response, creating the liquidity conditions under which technology talk without observable investment confirmation can have private value to insiders.

We identify the wedge using financial state and local narrative conditions fixed before the generative-AI shock. LowRecovery_i is measured in 2021 from distance to an industry-year valuation frontier, historical rates of reaching that frontier from similar distance bins, survival pressure, and top-manager horizon. Local narrative availability is proxied by leave-one-out headquarters-city peer AI discussion in MD&A during 2019-2021. The baseline design interacts LowRecovery_i, predetermined local narrative exposure, and the post-shock period, with firm, industry-year, and province-year fixed effects and city-level inference.

Weak-recovery firms in high-narrative cities increase AI language after the shock. The triple interaction raises AI talk without broad AI hiring by 4.46 percentage points and AI MD&A disclosure by 4.90 percentage points. Broad AI hiring does not increase. A four-state decomposition shows probability mass moving from no-talk/no-hire into talk/no-hire; talk-with-hire does not move. R&D intensity falls in the main and city-year designs, while capex, intangible assets, long-term deferred expenses, selling and administrative expenses, patents, employment, and graduate employee share do not increase.

The market evidence connects the realized wedge to insider liquidity use. Reduced-form DDD estimates for insider share reductions are close to zero, as expected when disclosure at most opens a conditional post-filing trading window. In post-shock filings, realized AI talk without broad AI hiring is followed by more insider share reductions, larger reduction amounts, and higher positive net selling; purchases do not increase. The relation is stronger among firms with high pre-shock Eastmoney/Guba stock-forum attention and is concentrated in market-facing sales and director/executive market-facing transactions. Equity-incentive transactions do not move, and pre-filing placebo windows are flat.

The institutional contrast is sharpest in insider liquidity use. Before the generative-AI shock, U.S. 10-K AI language and AI Key Developments move with AI workforce, while AI talk without AI workforce falls with AI workforce. After the shock, a U.S. retail-attention benchmark does not produce a stable insider-sale relation. The evidence identifies the state in which emerging-technology disclosure becomes less informative about investment and more financially relevant in a retail-attention market.

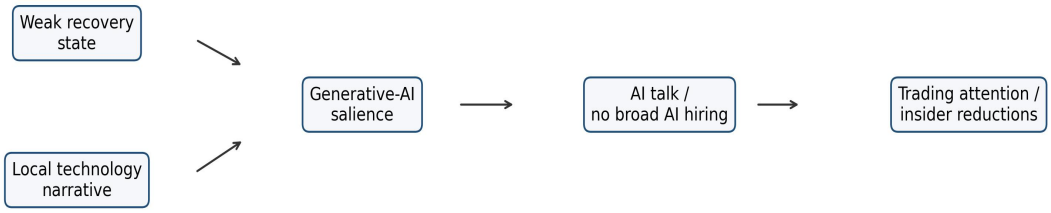


Figure 1. Economic Mechanism

Notes: Figure 1 summarizes the conceptual mechanism. The empirical tests separately examine the disclosure-action wedge, observable real inputs, trading attention, and insider share reductions.

2. Related Literature and Contribution

We begin with the standard interpretation of the AI-adoption literature: AI activity can be real, measurable, and economically important. Specialized labor demand, online vacancies, and firm-level innovation outcomes make adoption observable outside the text of corporate filings (Alekseeva et al. 2021; Acemoglu et al. 2022; Babina, Fedyk, He, and Hodson 2024). Cao, Jiang, Yang, and Zhang (2023) provide the published disclosure benchmark by showing that firms adjust disclosure when AI-based machine readers become important. Concurrent work studies the determinants and informativeness of voluntary AI disclosure (Barrios, Campbell, Johnson, and Liu 2025; Cao, Liu, Qiu, and Zhao 2026). We keep that benchmark but change the conditioning variable: weak recovery changes how informative AI language is about investment.

This prediction follows naturally from the voluntary-disclosure literature because disclosure is shaped by incentives, credibility, and audience beliefs, not only by fundamentals. Strategic communication begins with sender-receiver incentive conflict (Crawford and Sobel 1982), while the disclosure choice itself depends on proprietary costs, information incentives,

and the broader reporting environment (Verrecchia 1983; Dye 1985; Beyer et al. 2010). Corporate text can reveal economically meaningful information only when measurement fits context (Li 2010; Loughran and McDonald 2011). Emerging technologies make the context especially fragile because investors can observe words before they can verify systems, workers, products, or productivity gains. Cao, Jiang, Yang, and Zhang (2023) show that firms change disclosure when machines become important readers of filings. We study a different audience problem: when a salient technology category becomes tradable, a weak firm can use technology language even when costly adoption is unlikely to be confirmed by investment.

The closest market precedent is the dot-com and blockchain name-change literature. Cooper, Dimitrov, and Rau (2001) document large abnormal returns when firms added dot-com names during the internet boom, and Cooper, Khorana, Osobov, Patel, and Rau (2005) show that name changes during the dot-com decline also carried valuation effects, while Akyildirim, Corbet, Sensoy, and Yarovaya (2020) report related effects for blockchain-related name changes. Our setting documents a different transmission. The disclosure-action wedge is not rewarded by positive annual-report returns; instead, filing-window price-impact measures fall while insider exit rises. The economic margin therefore operates through tradable liquidity, not through a clean valuation revision. It is observed inside annual-report language, not in formal name changes.

The insider share-reduction evidence links this disclosure wedge to corporate governance and market attention. Voluntary disclosure and insider transactions are connected in time (Noe 1999), and insider-trading incentives depend on the surrounding disclosure regime (Brochet 2010; Lenkey 2014). Retail attention links this governance setting to observable trading conditions. Category-based attention, limited information processing, salience-driven retail trading, and search-based attention all predict that visible narratives can draw trading activity even when

fundamentals are hard to verify (Peng and Xiong 2006; Hirshleifer, Lim, and Teoh 2009; Barber and Odean 2008; Da, Engelberg, and Gao 2011). We identify a narrower governance margin: realized AI talk without broad hiring can coincide with post-filing windows in which insiders reduce exposure, especially when filings attract abnormal trading volume.

The financing-state mechanism draws on debt overhang, innovation finance, and limited-liability logic. Myers's (1977) debt-overhang logic explains why financing pressure can deter investment when residual claimants capture too little upside. Innovation magnifies this friction because R&D depends on external finance and because experimentation requires tolerance for agency frictions and delayed payoffs (Brown, Fazzari, and Petersen 2009; Holmstrom 1989; Manso 2011). Limited-liability models connect residual claims to downside boundaries (Merton 1974; Leland 1994). We use this logic to separate two margins often pooled in emerging-technology disclosure: costly innovation input and low-cost technology positioning.

3. Conceptual Framework and Hypothesis Development

3.1 Technology Shocks as Dual Opportunities

A salient technology shock changes both what firms can do and what firms can say. The real margin is costly: AI adoption requires specialized workers, data infrastructure, software integration, organizational redesign, and time before productivity gains become visible. The disclosure margin is cheaper and faster: the firm can describe existing activities through AI vocabulary, announce exploratory plans, or position itself near a salient technology category. In many firms these two margins move together, so AI language is informative about AI activity. The wedge studied here arises when the narrative margin remains open while the real investment margin is constrained.

The financial state determines where this wedge is most likely to appear. A firm under weak recovery conditions faces a short horizon and limited slack. Fixed-cost experimentation has less time to affect survival, and the benefits of a successful technology project may arrive too late to change the firm's recovery path. This logic follows from debt overhang and innovation finance: Myers (1977) shows that financing pressure can deter positive-NPV investment, Brown, Fazzari, and Petersen (2009) link R&D to external finance, and Manso (2011) emphasizes that innovation requires tolerance for long-horizon experimentation. In this state, the fact that AI is valuable in general does not imply that AI experimentation is privately attractive for the constrained firm.

Technology language has a different cost structure. It has low direct production cost, but it carries expected scrutiny and reputation costs when observable activity fails to support the claim. Weak firms need not face lower scrutiny. The relevant force is limited liability on the downside: when residual value is already compressed, a small probability of attracting attention, liquidity, or stakeholder relevance can be valuable relative to the remaining downside. Low-cost positioning becomes more attractive precisely when costly real action is unlikely to restore the firm.

3.2 Local Narrative Availability

Local technology narratives shift the disclosure margin. A firm headquartered in a city where peer firms already discuss AI inherits vocabulary, examples, and investor familiarity. The empirical exposure measures pre-shock narrative availability, not local productivity: how easy it is for a focal firm to make AI language sound plausible in its local information environment.

This distinction determines the falsification tests. If pre-shock peer AI talk captured only real AI opportunity, exposed firms should respond after the shock by hiring AI workers or expanding other real inputs. If it captures narrative availability, the response can appear in AI language without broad AI hiring. The adjacent-digital scope test is useful for the same reason. It shows that the active local exposure is a broader pre-existing digital narrative environment activated by the generative-AI shock, not an isolated AI-specific productivity opportunity.

3.3 Predictions

The first prediction is a disclosure-action wedge. After the generative-AI shock, weak-recovery firms in high-narrative cities should increase AI disclosure, but the increase should be concentrated in AI talk without broad AI hiring. AI talk accompanied by AI hiring should not be the main response.

The second prediction concerns real-input confirmation. The framework predicts lack of confirmation, not mechanical substitution. Broad AI hiring and conventional input measures should not expand with the disclosure response. R&D may fall when financing pressure and short horizons make long-run experimentation less attractive, but the central empirical object is the joint pattern of more AI language and no observable investment expansion.

The third prediction concerns insider liquidity use. The predetermined DDD state should predict the disclosure-action wedge, not insider selling by itself. Insider share reductions should be concentrated in realized talk/no-hire filings when the firm's investor base is more prone to trade on salient narratives. The main tests use pre-shock stock-forum attention as this retail-attention proxy; filing-window trading volume is treated as descriptive evidence on realized liquidity, not as a causal treatment.

The framework has several direct falsification margins. If AI talk without broad hiring were simply an early announcement of implementation, broad AI hiring or other real inputs should rise in the same cells after the shock. If insider share reductions reflected mechanical compensation activity, equity-incentive transactions should move with the talk/no-hire state. If the insider evidence were driven by a general selling trend, analogous pre-filing reductions should appear before the annual-report release. The empirical sections test these margins directly.

4. Data and Measurement

4.1 Sample

The sample consists of Chinese listed firm-years from 2014 to 2024. The panel combines annual-report AI text measures, job-posting measures of AI labor demand, financial statements, over-debt data, innovation input and output data, productivity measures, employee composition, and executive-horizon data. The empirical design fixes the weak-recovery state and local exposure before the shock and studies outcomes around the post-2023 generative-AI period. R&D and patent measures are available through 2023 in the current innovation data extract, while disclosure, hiring, financial, and employment outcomes extend through 2024.

Table 1 reports the sample construction and summary statistics. The initial AI-disclosure and job-posting panel contains 44,660 firm-years for 5,660 listed firms. Requiring a nonmissing 2021 weak-recovery state leaves 42,275 firm-years for 4,772 firms. Requiring nonmissing headquarters-city peer exposure leaves 40,955 firm-years for 4,631 firms. The preferred peer-support rule leaves 34,266 firm-years for 3,896 firms, and the R&D sample contains 28,757 firm-years for 3,790 firms. In the preferred sample, 52.5 percent of firm-years have AI MD&A

disclosure, 22.4 percent have broad AI hiring, and 34.5 percent have AI talk without broad AI hiring.

Table 1. Sample and Summary Statistics

Panel A. Sample construction.

Sample	Firm-years	Firms
AI-disclosure/job-posting panel, 2014-2024	44660	5660
Nonmissing 2021 weak-recovery state	42275	4772
Nonmissing HQ-city peer AI narrative exposure	40955	4631
Preferred sample: HQ-city peer support >= 20	34266	3896
Main disclosure-regression sample	33875	3895
R&D-regression sample	28757	3790

Panel B. Summary statistics.

Variable	N	Mean	SD	P25	Median	P75
AI talk without broad AI hiring	34266	0.345	0.475	0.000	0.000	1.000
AI MD&A disclosure	33875	0.525	0.499	0.000	1.000	1.000
Broad AI hiring	34266	0.224	0.417	0.000	0.000	0.000
Log(1 + broad AI postings)	34266	0.396	0.895	0.000	0.000	0.000
R&D intensity (RD1)	28757	0.024	0.023	0.006	0.020	0.033
R&D intensity	28757	0.051	0.060	0.012	0.037	0.062

Variable	N	Mean	SD	P25	Median	P75
(RD2)						
Patent output (Patent1)	28757	2.798	1.762	1.609	2.996	4.025
Patent output (Patent2)	28757	3.459	2.011	2.303	3.761	4.844
TFP (OP)	30293	6.788	0.916	6.162	6.670	7.302
TFP (LP)	30293	9.012	1.141	8.234	8.894	9.673
Log employment	34187	7.577	1.352	6.683	7.478	8.377
Graduate employee share	33990	0.096	0.142	0.016	0.045	0.117
Log assets	34190	22.342	1.480	21.316	22.084	23.092
Leverage	34190	0.433	0.247	0.257	0.415	0.586
ROA	34190	0.028	0.145	0.010	0.034	0.065
Over-debt degree	26734	-0.002	0.161	-0.110	-0.006	0.103
Weak-recovery proxy, 2021	30229	0.071	0.141	0.007	0.018	0.041
Low weak-recovery state, 2021	34266	0.234	0.424	0.000	0.000	0.000
HQ-city peer AI disclosure, 2019-2021	34266	0.541	0.126	0.478	0.570	0.638
HQ-city peer support	34266	432.822	422.745	86.000	198.000	995.000

Panel C. Predetermined weak-recovery state.

Variable	Low recovery mean	Other firms mean	Difference
AI talk without broad AI hiring	0.350	0.344	0.007
AI MD&A disclosure	0.520	0.527	-0.006
Broad AI hiring	0.218	0.226	-0.008
R&D intensity (RD2)	0.037	0.056	-0.018
Log employment	7.855	7.492	0.363
Leverage	0.594	0.384	0.210
ROA	-0.000	0.036	-0.037
HQ-city peer AI disclosure, 2019-2021	0.546	0.539	0.007

Notes: Panel A reports sample construction. Panel B reports summary statistics in the preferred sample. Panel C compares firms in the lowest 2021 weak-recovery quartile with other firms. The weak-recovery state is predetermined relative to the generative-AI shock.

4.2 AI Disclosure and AI Hiring

The main disclosure outcome is an indicator for whether the firm's MD&A contains AI-related discussion. This measure captures the extensive margin of AI narrative use in the annual report. It follows the textual-finance logic in Li (2010), who measures forward-looking language in filings, and Loughran and McDonald (2011), who establish the value of domain-specific dictionaries for financial text. The main real-action outcome is an indicator for whether the firm posts broad AI-related jobs in a firm-year. We also use the log of one plus broad AI postings. The labor-demand measure follows Alekseeva et al. (2021) and Acemoglu et al. (2022), who treat AI skills and online vacancies as observable adoption inputs. Babina, Fedyk, He, and Hodson (2024) link AI investment to firm-level outcomes.

Our primary wedge outcome is

$$\text{TalkNoHire}_{it} = 1[\text{AI MD\&A}_{it} = 1, \text{Broad AI Hiring}_{it} = 0].$$

This measure isolates firm-years in which AI is discussed in MD&A but no broad AI hiring is observed. It is a disclosure-action state.

4.3 Predetermined Weak-Recovery State

We construct the continuous RecoveryProspects_i score in 2021. Higher values indicate stronger recovery prospects. The score has three components. The first measures distance from an industry-year valuation frontier. We winsorize an approximate Tobin's q, compute the 90th percentile within industry-year cells, and calculate each firm's distance to that frontier. Using pre-2020 data, we estimate historical two-year probabilities of reaching the frontier from distance bins. The second component measures survival pressure using leverage, over-debt, profitability, and liquidity. The third component measures top-manager horizon using remaining years to age 65.

The main state variable, LowRecovery_i, equals one for firms in the bottom quartile of RecoveryProspects_i. LowRecovery_i therefore identifies weak-recovery firms: firms with the weakest 2021 recovery prospects. The indicator is predetermined relative to the generative-AI shock and fixed across all post-shock outcomes. It is a reduced-form financial-state measure. Component tests below show that the survival-pressure component carries most of the economic content.

Table 1, Panel C validates the state variable. Weak-recovery firms are more levered than other firms in the preferred sample, with mean leverage of 0.594 compared with 0.384. They are less profitable, with mean ROA close to zero compared with 0.036 for other firms. They also have lower RD2 in levels, 0.037 compared with 0.056. These level differences are not the source

of identification, because the regressions include firm fixed effects and event-study tests, but they show that the state variable captures the intended financial condition.

An ex post validation sorts firms into 2021 RecoveryProspects_i deciles and follows realized outcomes over 2022-2023. Firms in the top decile reach the industry-year valuation frontier in 57.6 percent of cases, compared with much lower rates through the middle of the distribution. The lowest decile has a 17.5 percent ST-or-suspended rate and a 64.3 percent probability of negative ROA over the next two years, compared with 2.5 percent and 21.3 percent in the top decile. The continuous RecoveryProspects_i score has a 0.385 correlation with realized frontier-reaching, while LowRecovery_i is positively correlated with future ST or suspension. These validations support the interpretation of the score as a financial-state measure that separates high-feasibility firms from firms under financial pressure.

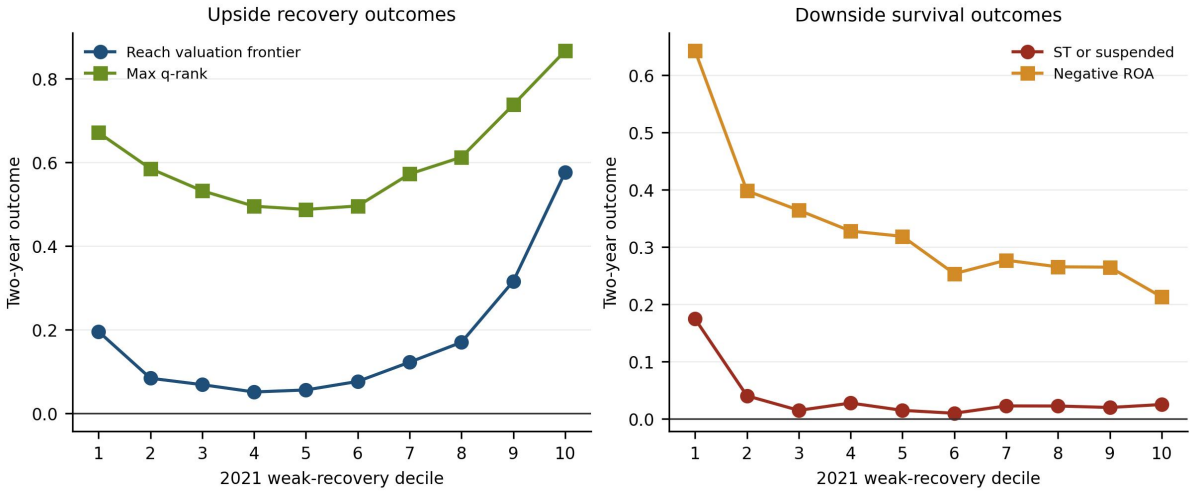


Figure 2. Validation of the Weak-Recovery State

Notes: Figure 2 plots two-year outcomes by 2021 recovery-prospects decile. Higher deciles indicate stronger recovery prospects. LowRecovery_i denotes the bottom quartile of RecoveryProspects_i. The figure validates the state variable against subsequent upside recovery and downside survival outcomes; it is not used to estimate the main treatment effect.

The component audit clarifies the economic content of the state variable. The composite result is driven mainly by the survival-pressure component, not by top-manager horizon. Within that component, over-debt and liquidity pressure are more closely aligned with disclosure, while leverage is most closely aligned with R&D retrenchment. We therefore interpret LowRecovery_i as a weak-recovery state whose empirical content is primarily financial pressure, not a pure leverage or managerial-horizon measure.

4.4 Local Technology Narrative Availability

For each firm, we fix headquarters city in 2021 and compute leave-one-out local narrative exposure among other listed firms in the same city over 2019-2021. The empirical proxy is the average of peer firms' AI MD&A indicators, excluding the focal firm, and is standardized in the analysis sample. The preferred sample requires at least 20 peer firm-year observations in the city-peer exposure calculation. The proxy uses AI language because the generative-AI shock is the focal event, but the theoretical object is broader local narrative availability.

We fix headquarters city in 2021 because annual city fields can change or be reported inconsistently. The peer-count screen addresses noise in small leave-one-out peer cells. Variable audits show that the main results are strongest when local exposure is measured with adequate peer support; thin cells attenuate the disclosure estimates.

A split-half reliability check supports the peer-count screen. We split 2019-2021 city peers into two deterministic halves and compute the correlation between the two pre-shock city AI-narrative measures. The split-half correlation is 0.088 in city cells with fewer than ten peer observations and 0.055 in cells with 10-19 observations. It rises to 0.424 in cells with 20-49

observations and 0.732 in cells with at least 100 observations. The preferred peer-count threshold therefore functions as a measurement-quality screen, not as an outcome-based sample choice.

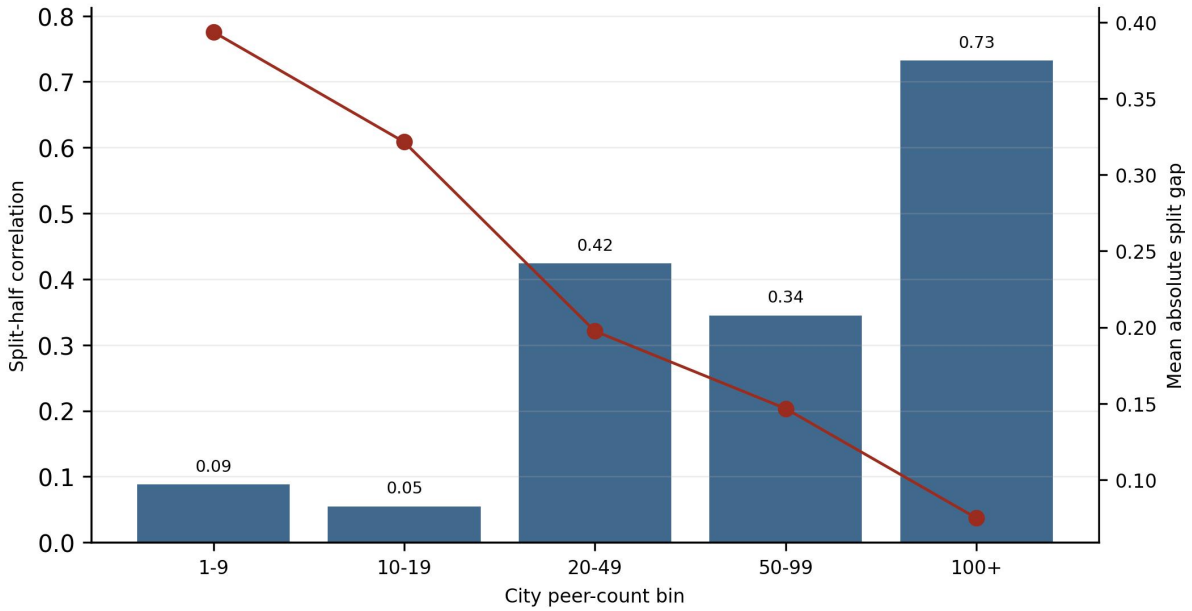


Figure 3. Reliability of Leave-One-Out City-Peer Narrative Exposure

Notes: Figure 3 reports split-half reliability for the pre-shock leave-one-out city-peer AI-talk exposure. The preferred peer-count screen is a measurement-quality rule motivated by reliability, not an outcome-based sample choice.

4.5 Pre-Shock Retail Attention

For the insider-liquidity tests, we construct a predetermined firm-level retail-attention proxy from Eastmoney/Guba stock-forum activity. The platform provides stock-code-specific discussion boards used heavily by individual investors, which makes it a natural pre-shock measure of retail attention. We aggregate daily stock-forum reads, posts, and comments during calendar year 2021 and merge the firm-level totals to the annual-report event panel. The 2021 file covers 4,798 listed stocks. The median covered stock has about 2.84 million reads, 7,346 posts, 10,523 comments, and forum activity on 350 days during the year; the interquartile range for reads is about 1.53 million to 5.48 million. The main heterogeneity measure is the z-score of

log one plus 2021 forum reads. A binary version equals one for firms above the preferred-sample median. Because these measures are fixed before the generative-AI shock, they do not condition on filing-window trading outcomes.

5. Empirical Design

The baseline specification is

$$Y_{it} = \beta(LowRecovery_i \times LocalNarrative_i \times Post_t) + \gamma(LowRecovery_i \times Post_t) + \delta(LocalNarrative_i \times Post_t) + \alpha_i + \lambda_{j(i),t} + \mu_{p(i),t} + \varepsilon_{it}$$

where Post_t equals one for 2023 and 2024, j(i) denotes industry, and p(i) denotes province. The coefficient of interest is beta. It compares low-recovery and other firms before and after the shock across locations with different predetermined AI narrative exposure.

Industry-year and province-year fixed effects absorb broad sectoral and regional shocks, and firm fixed effects absorb persistent firm differences. The identifying variation is within firm, relative to other firms in the same industry-year and province-year, as the generative-AI shock arrives. Because $LowRecovery_i$ and $LocalNarrative_i$ are predetermined and time-invariant, their main effects are absorbed by firm fixed effects; the specification includes $LowRecovery_i$ x $Post_t$ and $LocalNarrative_i$ x $Post_t$ to absorb differential post-shock changes by weak-recovery state and by local narrative exposure.

The main threat is that local narrative exposure may proxy for local AI booms, local industrial policy, or unobserved city-level demand shocks. We address this threat directly with headquarters-city-by-year fixed effects. Because local exposure is fixed at the city level, this specification does not use within-city variation in exposure levels. It asks whether the low-versus-high recovery differential within the same city-year is larger in cities with stronger pre-

shock technology narratives. Any post-shock force common to all firms headquartered in the same 2021 city is absorbed.

A second concern is that local peer exposure may capture real AI adoption opportunity rather than narrative availability. The outcome pattern helps distinguish these interpretations. A real-opportunity channel predicts higher AI hiring, and plausibly higher innovation input, among exposed firms after the shock. The disclosure-action-wedge channel predicts more AI language without broad AI hiring and without real-input expansion. We therefore interpret the disclosure, hiring, and R&D estimates jointly rather than as separate outcomes.

A third concern is inference. Local narrative exposure is measured at the headquarters-city peer level, so the preferred inference clusters by 2021 headquarters city. Bertrand, Duflo, and Mullainathan (2004) show that serial correlation can make standard difference-in-differences inference too liberal. Cameron, Gelbach, and Miller (2011) motivate multiway clustering when shocks operate along more than one dimension, and MacKinnon and Webb (2017) provide the wild-bootstrap logic for uneven cluster structures. We therefore report city clustering, two-way firm-city clustering, and wild city bootstrap p-values. These choices are more conservative than firm-only clustering for the local-exposure design.

A fourth concern is fiscal-year timing. ChatGPT arrived at the end of 2022, while Chinese annual reports for fiscal-year 2022 were generally written and released in 2023. We therefore use actual CSMAR annual-report announcement dates to distinguish fiscal-year outcomes from filing-date exposure. The baseline keeps the conventional 2023-2024 post period; robustness checks reclassify reports by actual filing date and treat fiscal-year 2022 as a transition year.

The event-study specification replaces $Post_t$ with year indicators interacted with $LowRecovery_i \times LocalNarrative_i$, omitting 2022 in the conventional fiscal-year design. Filing-date robustness checks show that the main AI disclosure result survives when annual-report announcement dates are used directly, and that excluding fiscal-year 2022 strengthens the talk-without-hiring result.

6. Main Results

The main table asks a simple question: after the generative-AI shock, do financially weak firms in high-narrative cities move on the language margin or on the investment margin? The dependent variables are ordered to make that comparison explicit. The first two columns measure AI language, the next two columns measure broad AI labor demand, and the final columns measure R&D intensity. The coefficient of interest is the interaction of $LowRecovery_i$, pre-shock local narrative exposure, and the post-shock period.

Table 2 reports the main estimates. Because local narrative exposure is standardized, the coefficient measures the differential post-shock response of low-recovery firms associated with a one-standard-deviation increase in predetermined local narrative exposure. The empirical proxy is pre-shock city-peer AI talk. The triple interaction increases AI talk without broad AI hiring by 0.0446, with a city-clustered standard error of 0.0175. It increases AI MD&A disclosure by 0.0490, with a city-clustered standard error of 0.0129.

Table 2. Main Disclosure-Action Wedge

Outcome	Coef.	SE	p	N	Mean	Wild city p
AI talk, no broad AI	0.0446**	(0.0175)	0.011	34266	0.345	0.023

Outcome	Coef.	SE	p	N	Mean	Wild city p
hiring						
AI MD&A disclosure	0.0490***	(0.0129)	0.000	33875	0.525	0.001
Broad AI hiring indicator	-0.0005	(0.0135)	0.970	34266	0.224	0.974
Log broad AI postings	0.0099	(0.0240)	0.680	34266	0.396	—–
R&D intensity (RD2)	-0.0041***	(0.0016)	0.008	28757	0.051	0.007

Notes: The coefficient is on LowRecovery_i x LocalNarrative_i x Post_t. LowRecovery_i is the 2021 bottom-quartile indicator of $\text{RecoveryProspects}_i$; weak-recovery firms are therefore those with the weakest recovery prospects. LocalNarrative_i is standardized pre-shock leave-one-out headquarters-city peer AI talk. Post_t equals one in fiscal years 2023 and 2024. Regressions include firm, industry-year, province-year, LowRecovery_i x Post_t , and LocalNarrative_i x Post_t fixed effects/controls as specified in the text. Standard errors are clustered by 2021 headquarters city. Wild city p-values are from the corresponding wild cluster bootstrap; the wild city p-value for Log broad AI postings is not reported.

The magnitudes are economically meaningful. The effect on AI talk without broad AI hiring is 12.9 percent of the preferred-sample mean of 0.345. The effect on AI MD&A disclosure is 9.3 percent of the mean of 0.525. The estimates capture a material shift in the margin where weak recovery interacts with local narrative availability.

The same specification shows no increase in broad AI hiring. The coefficient on the broad AI hiring indicator is -0.0005, and the coefficient on log broad AI postings is 0.0099; both are statistically insignificant. The point estimates do not suggest that Weak-recovery firms in high-exposure cities are building AI labor capacity after the shock.

Real innovation input moves in the opposite direction in the main design. The coefficient is -0.0014 for RD1 and -0.0041 for RD2, statistically significant at conventional levels under

city-clustered inference. The RD2 estimate is 8.1 percent of the dependent-variable mean. Wild city bootstrap p-values are 0.023 for talk without broad hiring, 0.001 for AI MD&A disclosure, 0.974 for broad AI hiring, and 0.007 for RD2. AI disclosure therefore rises without observable AI hiring or broad real-input expansion. R&D is lower in the baseline and city-year specifications but is best interpreted as supporting evidence rather than the identifying mechanism.

7. Timing, Robustness, and Identification

Figure 4 reports the event-study estimates, and Internet Appendix Table IA.7 reports the joint pretrend tests. With 2022 as the omitted fiscal year, joint pretrend tests have p-values of 0.467 for AI talk without broad hiring, 0.894 for AI MD&A disclosure, 0.424 for broad AI hiring, and 0.797 for RD2. The 2024 coefficient for AI MD&A disclosure is positive and statistically significant, while broad AI hiring remains indistinguishable from zero. The R&D effect appears in 2023, the last year available for the R&D series in the current extract.

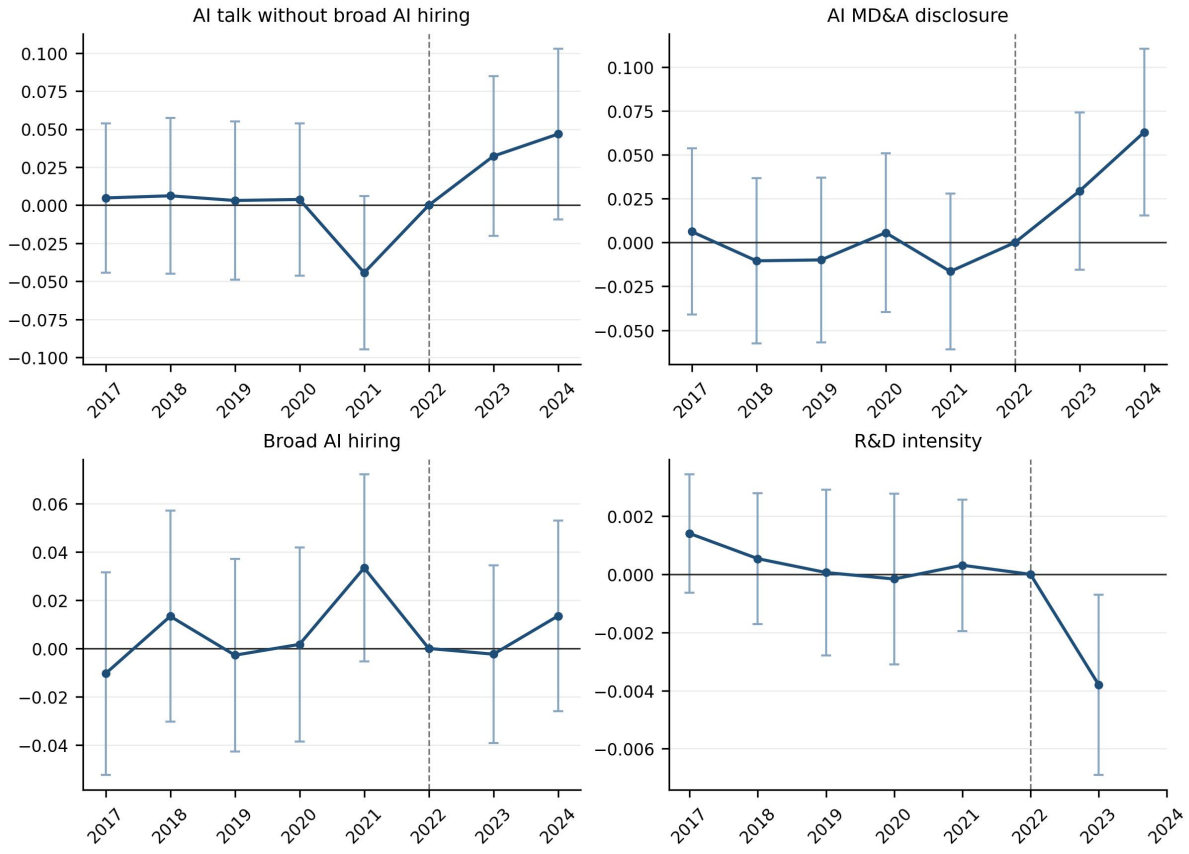


Figure 4. Event-Study Estimates around the Generative-AI Shock

Notes: Figure 4 plots event-study coefficients on the interaction of weak-recovery status and local narrative availability. Fiscal-year 2022 is the omitted year in the conventional design. Vertical bars are 95 percent confidence intervals using the reported standard errors.

Table 3 reports the headquarters-city-year specification, which addresses the concern that local narrative exposure reflects city-level post-shock forces. The table keeps the preferred peer-count sample and replaces province-year fixed effects with headquarters-city-by-year fixed effects. The coefficient on AI talk without broad hiring remains positive and significant at 0.0471. AI MD&A disclosure remains positive and significant at 0.0470. Broad AI hiring remains close to zero, and RD2 remains negative and significant at -0.0036. These estimates show that the main pattern is not explained by city-level post-shock shocks common to firms headquartered in the same city.

Table 3. Timing and Identification

Specification	Outcome	Coef.	SE	p	N
City-year FE, city cluster	AI talk, no broad AI hiring	0.0471***	(0.0174)	0.007	34266
City-year FE, city cluster	AI MD&A disclosure	0.0470***	(0.0129)	0.000	33875
City-year FE, city cluster	Broad AI hiring indicator	-0.0038	(0.0133)	0.776	34266
City-year FE, city cluster	R&D intensity (RD2)	-0.0036**	(0.0016)	0.028	28757
Filing-date post	AI talk, no broad AI hiring	0.0303*	(0.0182)	0.096	34190
Filing-date post	AI MD&A disclosure	0.0401***	(0.0154)	0.009	33875
Filing-date post	Broad AI hiring indicator	0.0092	(0.0122)	0.452	34190
Filing-date post, exclude FY2022	AI talk, no broad AI hiring	0.0476**	(0.0194)	0.014	30327
Filing-date post, exclude FY2022	AI MD&A disclosure	0.0529***	(0.0150)	0.000	30033
Filing-date post, exclude FY2022	Broad AI hiring indicator	0.0019	(0.0135)	0.886	30327
City-industry-year FE, AI exposure	AI talk, no broad AI hiring	0.0568*	(0.0300)	0.058	34266
City-industry-year FE, AI exposure	AI MD&A disclosure	0.0511**	(0.0238)	0.032	33875

Specification	Outcome	Coef.	SE	p	N
City-industry-year FE, AI exposure	Broad AI hiring indicator	-0.0117	(0.0209)	0.578	34266
City-industry-year FE, AI exposure	R&D intensity (RD2)	-0.0033	(0.0032)	0.295	28757

Notes: The reported coefficient is the triple interaction unless otherwise stated. City-year specifications absorb headquarters-city-by-year shocks. Filing-date specifications define post-shock exposure using actual annual-report announcement dates. The city-industry-year specification absorbs headquarters-city-by-industry-by-year fixed effects. Standard errors are clustered by 2021 headquarters city.

A stricter local-shock check absorbs headquarters-city-by-industry-by-year fixed effects. This specification is demanding because the treatment state varies only across firms within narrow city-industry-year cells. AI MD&A remains positive and significant at 0.0511, and AI talk without broad hiring remains positive with a p-value of 0.058. Broad AI hiring remains insignificant. The RD2 decline is not robust in this narrow-cell specification, so the city-industry-year evidence is strongest for the disclosure-action wedge rather than for R&D.

Table 3 also reports filing-date timing. Actual annual-report announcement dates show that fiscal-year 2022 annual reports were disclosed after the ChatGPT shock and should be treated as a transition year. When $Post_t$ is defined from announcement dates, AI MD&A disclosure remains significant and talk without broad hiring is marginally significant. When fiscal-year 2022 is excluded, AI talk without broad hiring is 0.0476 and AI MD&A disclosure is 0.0529, both statistically significant under city clustering. These tests address the concern that the omitted fiscal-year 2022 benchmark is contaminated by post-shock annual-report writing.

Appendix Figure IA.1 repeats the event-study with fiscal-year 2021 as the omitted year, so fiscal-year 2022 appears as an estimated transition coefficient rather than as the benchmark.

The pre-2021 coefficients remain jointly insignificant for the disclosure, hiring, and R&D outcomes. The disclosure coefficients grow in 2023 and 2024, while broad AI hiring remains close to zero.

Figure 3 and Internet Appendix Table IA.5 report peer-support reliability. The reliability of leave-one-out local narrative exposure improves sharply once peer support reaches the preferred threshold, and thin peer cells have large split-half gaps. This supports treating the peer-count rule as a measurement-quality screen rather than an outcome-based sample choice.

Additional checks use winsorized local exposure and alternative peer-support thresholds in the headquarters-city exposure design. Broader appendix tests from the full audit consider alternative state definitions, balanced panels, and exclusions of finance firms, top cities, high-AI industries, and nonmanufacturing firms. A local adjacent-technology scope test based on pre-shock city peer mentions of cloud computing, IoT, VR/AR, big data, and related digital terms also predicts the wedge. The scope test clarifies that the local exposure should be read as a broader pre-existing digital narrative environment that makes AI positioning credible after the generative-AI shock, not as an AI-specific productivity opportunity. Consistent with that interpretation, the AI exposure residualized against the adjacent-digital exposure no longer carries the talk-without-hiring wedge.

8. Mechanisms and Measurement

We organize the mechanism around the main alternatives to the disclosure-action interpretation. The four-state decomposition addresses a mechanical-outcome concern: the estimates show movement into talk/no-hire rather than a generic shift in either component. Total and non-AI postings address a missing-hiring-data concern, since they rise in some specifications,

so the null AI-hiring result is not a general disappearance of job-posting coverage. Internet Appendix Table IA.18 further shows that these non-AI posting increases are concentrated in customer/operations and administrative jobs and not in AI postings or a statistically significant non-AI technical expansion. Capex, intangible assets, long-term deferred expenses, selling and administrative expenses, amortization, patents, productivity, employment, and graduate employee share address an outsourced-adoption alternative. The component audit in Internet Appendix Table IA.1 addresses a managerial-horizon alternative: the horizon component does not drive the effect. The insider share-reduction section then tests two governance alternatives: equity-incentive transactions do not move with talk/no-hire, and the retail-attention interaction is measured with 2021 stock-forum activity rather than filing-window trading volume.

Table 4 decomposes the disclosure-hiring wedge into four states: no talk/no hire, talk/no hire, talk/hire, and no talk/hire. The triple interaction lowers no talk/no hire by 0.0464 and raises talk/no hire by 0.0464. The talk/hire and no talk/hire states do not move. This decomposition shows that the wedge is not a mechanical artifact of combining two outcomes; the probability mass moves into the talk-without-hiring state.

Table 4. Four-State Disclosure-Hiring Decomposition

Outcome/state	Coef.	SE	p	N	Mean
No talk, no hire	-0.0464***	(0.0147)	0.002	33875	0.427
Talk, no hire	0.0464***	(0.0177)	0.009	33875	0.349
Talk and hire	0.0025	(0.0116)	0.826	33875	0.176

Outcome/state	Coef.	SE	p	N	Mean
No talk, hire	-0.0026	(0.0072)	0.720	33875	0.048

Notes: The table decomposes firm-years into mutually exclusive disclosure-hiring states. The reported coefficient is the preferred weak-recovery × local-narrative × post coefficient. Standard errors are clustered by 2021 headquarters city.

Internet Appendix Table IA.17 reports text-specificity checks based on broad AI dictionaries, following the textual-finance emphasis on domain-specific language in Loughran and McDonald (2011). The effect appears in core AI mentions and in core AI talk without broad hiring, but it is weaker for literal generative-AI wording. The main text therefore relies on observable investment confirmation not fine-grained sentence labels: AI language rises without matching AI labor demand or conventional input expansion.

The adjacent-digital scope test in Internet Appendix Table IA.2 reinforces this interpretation and also disciplines the R&D evidence. When pre-shock AI exposure is residualized against adjacent digital narrative exposure, the talk/no-hire estimate attenuates and the R&D decline disappears: the RD2 coefficient is -0.0004 with a p-value of 0.831. The active local exposure therefore does not isolate AI-specific productivity opportunities. It is better interpreted as a broader pre-existing digital narrative environment that the generative-AI shock made newly salient. Under this reading, the R&D estimates are supporting evidence of weak real-input confirmation, not the core mechanism.

Table 5 reports additional real outcomes and cheap-adoption proxies. The mechanism is not a broad expansion of AI-specific activity. Weak-recovery firms exposed to local narrative availability do not increase broad AI hiring. They also do not show significant increases in capex, selling and administrative expenses, patent output, TFP, total employment, or graduate employee share. Internet Appendix Table IA.16 extends this test to buy-vs-make financial-statement

proxies: net intangible assets, long-term deferred expenses, capex cash paid for fixed/intangible/long-term assets, selling and administrative expenses, intangible amortization, and long-term deferred amortization do not increase. The data-resource intangible item is too sparse to support inference. Total and non-AI job postings rise in some specifications, which indicates that the absence of AI hiring is not a mechanical consequence of disappearing job-posting coverage or a general hiring collapse. Internet Appendix Table IA.18 decomposes the raw hiring file: the cleanest non-AI increases appear in customer/operations and administrative postings, while AI postings remain negative and insignificant and non-AI technical postings are not statistically significant. This composition is more consistent with ordinary operating labor than with AI implementation. R&D intensity falls in the main and city-year specifications, but the central real-input evidence is the absence of observable investment confirmation.

Table 5. Observable Real-Input Confirmation

Outcome	Coef.	SE	p	N	Mean
Cash capex/assets	0.0016	(0.0016)	0.323	34141	0.044
Selling expense/assets	0.0017	(0.0016)	0.287	32929	0.042
Admin expense/assets	0.0012	(0.0011)	0.287	33454	0.040
SG&A/assets	0.0029	(0.0021)	0.178	32929	0.082
Log total postings	0.1335***	(0.0468)	0.004	34266	3.394
Non-AI hiring indicator	-0.0037	(0.0142)	0.797	34266	0.809

Outcome	Coef.	SE	p	N	Mean
Log non-AI postings	0.1335***	(0.0472)	0.005	34266	3.384
Patent output (Patent1)	0.0320	(0.0414)	0.440	31672	2.793
Patent output (Patent2)	0.0436	(0.0485)	0.368	31672	3.454
TFP (OP)	0.0056	(0.0160)	0.727	33508	6.774
TFP (LP)	0.0019	(0.0189)	0.918	33508	9.000
Log employment	-0.0192	(0.0186)	0.301	37649	7.587
Graduate employee share	-0.0016	(0.0037)	0.660	37435	0.093

Notes: The coefficient is on the preferred triple interaction. Outcomes include capital expenditures, selling and administrative expenses, total and non-AI postings, patent output, productivity, employment, and graduate employee share. These tests examine whether AI language is accompanied by observable real-input expansion. Standard errors are clustered by 2021 headquarters city where available.

The real-input tests ask whether observable investment confirms the AI language. The state that predicts AI disclosure does not predict broad AI labor demand, conventional input expansion, or balance-sheet evidence of a large outsourced-AI response. The same null rules out several alternatives. A buy-rather-than-make response would leave traces in capex, intangible assets, deferred expenses, selling and administrative expenses, amortization, or other operating inputs; the available proxies do not move. A pure managerial-horizon story would predict stronger effects for short-horizon managers, but the horizon component does not drive the main decomposition. A customer or partner signaling story would predict subsequent real follow-through; the hiring, R&D, patent, employment, and productivity margins do not show it. The

equity-incentive falsification in Table 6 further separates market-facing insider share reductions from mechanical compensation transactions.

9. Market Attention and Insider Liquidity Use

The market tests ask whether the realized wedge has financial relevance. They are not part of the DDD identification of the disclosure response. Table 6 first shows that the predetermined DDD state has no positive short-window CAR around annual-report filings; the market-adjusted CAR over days 0 to 1 is about -0.46 percentage points. Annual reports contain many pieces of information, so this reduced-form return is not a clean price reaction to AI language. The attention response appears in trading volume and liquidity, not in price appreciation.

Table 6. Market Attention and Insider Share Reductions

Panel	Specification	Outcome	Variable	Coef.	SE	p	N
A. Filing returns	DDD, CAR	CAR [0,+1]	DDD state	-0.0046**	(0.0022)	0.041	29506
A. Filing returns	DDD, CAR, city-year FE	CAR [0,+1]	DDD state	-0.0046**	(0.0022)	0.039	29506
B. Trading attention	Post-filing talk/no-hire	Abnormal log trading value [0,+1]	Talk/no-hire	0.0541***	(0.0201)	0.007	7605
B. Trading attention	Post-filing talk/no-hire, firm FE	Abnormal log trading value [0,+1]	Talk/no-hire	0.0189	(0.0389)	0.627	7605
C. Talk/no-hire state	Observed talk/no-hire, post filings	Insider share-reduction	Talk/no-hire	0.0148**	(0.0059)	0.012	7635

Panel	Specification	Outcome	Variable	Coef.	SE	p	N
		indicator [0,+60]					
C. Talk/no-hire state	Observed talk/no-hire, post filings	Log insider share-reduction amount [0,+60]	Talk/no-hire	0.2071**	(0.0910)	0.023	7635
C. Talk/no-hire state	Observed talk/no-hire, post filings	Insider purchase indicator [0,+60]	Talk/no-hire	0.0062	(0.0073)	0.396	7635
D. Forum attention	2021 forum reads (z-score)	Insider share-reduction indicator [0,+60]	Talk/no-hire × 2021 forum reads	0.0083**	(0.0038)	0.029	7340
D. Forum attention	2021 forum reads (z-score)	Log insider share-reduction amount [0,+60]	Talk/no-hire × 2021 forum reads	0.1367**	(0.0593)	0.021	7340
D. Forum attention	2021 forum reads (z-score)	Log positive net insider share reductions [0,+60]	Talk/no-hire × 2021 forum reads	0.1689***	(0.0572)	0.003	7340
D. Forum attention	2021 forum reads (z-score)	Insider purchase indicator [0,+60]	Talk/no-hire × 2021 forum reads	-0.0119**	(0.0059)	0.045	7340
E. High-forum attention	High 2021 forum reads	Insider share-reduction indicator [0,+60]	Talk/no-hire × high 2021 forum	0.0170*	(0.0100)	0.090	7340

Panel	Specification	Outcome	Variable	Coef.	SE	p	N
			reads				
E. High-forum attention	High 2021 forum reads	Log insider share-reduction amount [0,+60]	Talk/no-hire × high 2021 forum reads	0.2690*	(0.1570)	0.087	7340
E. High-forum attention	High 2021 forum reads	Log positive net insider share reductions [0,+60]	Talk/no-hire × high 2021 forum reads	0.3179**	(0.1575)	0.044	7340
E. High-forum attention	High 2021 forum reads	Insider purchase indicator [0,+60]	Talk/no-hire × high 2021 forum reads	-0.0285**	(0.0144)	0.047	7340
F. Pre-filing placebo	Pre-filing placebo	Insider share-reduction indicator [-90,-1]	Talk/no-hire × high 2021 forum reads	-0.0090	(0.0098)	0.361	7340
F. Pre-filing placebo	Pre-filing placebo	Log insider share-reduction amount [-90,-1]	Talk/no-hire × high 2021 forum reads	-0.0758	(0.1494)	0.612	7340

Panel	Specification	Outcome	Variable	Coef.	SE	p	N
F. Pre-filing placebo	Pre-filing placebo	Log positive net insider share reductions [-90,-1]	Talk/no-hire × high 2021 forum reads	-0.1372	(0.1498)	0.360	7340
F. Pre-filing placebo	Pre-filing placebo	Insider purchase indicator [-90,-1]	Talk/no-hire × high 2021 forum reads	-0.0035	(0.0128)	0.782	7340
G. Market-facing sales	Market-facing sales	Market-sale share-reduction indicator [0,+60]	Talk/no-hire	0.0113**	(0.0044)	0.011	7472
G. Market-facing sales	Market-facing sales	Log market-sale share-reduction amount [0,+60]	Talk/no-hire	0.1621**	(0.0671)	0.016	7472
G. Market-facing sales	Market-facing sales	Log positive net market-sale selling [0,+60]	Talk/no-hire	0.1608**	(0.0653)	0.014	7472
G. Market-facing sales	Market-facing sales	Market-sale purchase indicator [0,+60]	Talk/no-hire	-0.0007	(0.0045)	0.872	7472
H. Director/executive sales	Director/executive market-facing sales	Director/executive market-sale indicator [0,+60]	Talk/no-hire	0.0082**	(0.0040)	0.039	7472
H. Director/executive sales	Director/executive market-facing sales	Log director/executive market-sale amount [0,+60]	Talk/no-hire	0.1242**	(0.0607)	0.041	7472

Panel	Specification	Outcome	Variable	Coef.	SE	p	N
H. Director/executive sales	Director/executive market-facing sales	Log positive net director/executive selling [0,+60]	Talk/no-hire	0.1217**	(0.0600)	0.043	7472
H. Director/executive sales	Director/executive market-facing sales	Director/executive market-purchase indicator [0,+60]	Talk/no-hire	-0.0021	(0.0044)	0.639	7472
I. Equity-incentive check	Equity-incentive transactions	Equity-incentive share-reduction indicator [0,+60]	Talk/no-hire	0.0010	(0.0014)	0.486	7472
I. Equity-incentive check	Equity-incentive transactions	Log equity-incentive share-reduction amount [0,+60]	Talk/no-hire	0.0169	(0.0174)	0.333	7472

Notes: Panel A reports annual-report return outcomes around filing dates. Panel B reports abnormal trading-value outcomes for realized talk/no-hire filings. Panel C reports insider shareholding changes after filing dates using CSMAR event-level data. Panels D-E replace post-treatment filing-window attention with predetermined 2021 Eastmoney/Guba stock-forum attention. The continuous measure is the z-score of log 2021 forum reads; the binary measure equals one for firms above the preferred-sample median. Panel F reports pre-filing placebo outcomes. Panels G-H decompose insider share reductions by transaction method and insider position using CSMAR transaction reasons and positions; market-facing sales include bidding transactions, secondary-market trading, block trades, and negotiated transfers. Panel I reports the equity-incentive falsification. Insider outcomes are measured over the indicated filing windows. The evidence is interpreted as consistent with insider liquidity use, not as a reduced-form causal effect of the predetermined DDD state on insider trading.

The baseline insider evidence uses CSMAR event-level changes in shareholdings of directors, supervisors, executives, and related persons. Noe (1999) documents the timing link between voluntary disclosure and insider transactions, and Lenkey (2014) shows how disclosure rules can change insider trading incentives. In post-shock filings, talk/no-hire raises the probability of an insider share reduction over the next 60 days by 1.48 percentage points and the log amount of insider share reductions over 60 days by 0.207. Purchases do not increase. In post-2023 talk/no-hire filings at firms with high 2021 forum reads, observed insider share reductions over the next 60 days sum to about RMB 385 million; a simple comparison with high-forum no-talk/no-hire filings implies about RMB 208 million more share reductions in the treated state.

The retail-attention evidence uses firm-level stock-forum activity measured before the shock. A one-standard-deviation increase in 2021 Eastmoney/Guba forum reads strengthens the talk/no-hire association with 60-day insider share reductions by 0.83 percentage points, with log insider share-reduction amount by 0.137, and with log positive net insider selling by 0.169. The interaction for insider purchases is negative. A binary high-forum-read specification gives the same sign pattern, and pre-filing placebo outcomes do not show the same interaction.

Direct counterparty and market-quality diagnostics use newly acquired CSMAR shareholder-count, Level-2 small-trade, Dragon-Tiger investor-type, daily high-low, limit-status, and pledge-detail data. The 20-day small-trade share is positive among realized talk/no-hire filings with high pre-shock forum attention, and Amihud price-impact measures fall after talk/no-hire filings, while abnormal absolute returns and return volatility do not increase. Short-window breadth measures are noisier, Dragon-Tiger evidence is weak, and pledge-detail tests do not show a relief channel. We therefore treat direct counterparty evidence as supporting diagnostics. The main liquidity-use interpretation rests on predetermined retail-attention heterogeneity and insider transaction composition.

The transaction-level CSMAR data clarify who sells and through what channel. The talk/no-hire association is concentrated in market-facing sales: over the 60 days after filing, talk/no-hire is associated with a 1.13 percentage point higher probability of market-facing reductions and a 0.162 increase in log market-sale reduction amount. Within director and executive market-facing transactions, the corresponding estimates are 0.82 percentage points and 0.124. Equity-incentive transactions do not move. Controls for pre-filing insider sales, lagged sales, returns, size, baseline trading value, and pledge pressure do not eliminate the relation.

The cross-country comparison clarifies the interpretation. The U.S. evidence should not be read as showing that the wedge is absent in magnitude: the post-shock U.S. talk/no-KeyDev coefficient is the same order as the Chinese talk/no-hire coefficient, but it is estimated less precisely. The sharp contrast is insider liquidity use. In China, the pre-shock retail-attention interaction predicts positive net market-facing insider share reductions; in the U.S. retail-attention benchmark, the insider-sale relation is noisy and unstable across specifications. The insider-liquidity margin is therefore not a generic feature of AI disclosure but is sharper in China's retail-attention setting.

Table 7. Cross-Country Institutional Contrast

Market	Margin	Outcome	Estimate	p	N	Interpretation
China	Institutional contrast: insider share reductions	Log positive net insider selling [0,+60]	0.1689***	0.003	7340	Chinese talk/no-hire filings are followed by larger insider share reductions when pre-shock forum attention is higher.
United States	U.S. retail-attention benchmark	Log insider sale value [0,+60]	0.6556	0.267	1410	The U.S. retail-attention interaction is noisy and does not deliver a stable insider-

Market	Margin	Outcome	Estimate	p	N	Interpretation
						sale margin.
China	Non-exit benchmark	Equity-incentive share-reduction indicator [0,+60]	0.0010	0.486	7472	Equity-incentive transactions do not move in the same way.
United States	Investment-confirmation benchmark	AI disclosure	0.0945***	0.000	12073	U.S. AI language maps positively to observable AI workforce before the generative-AI shock.
United States	Investment-confirmation benchmark	AI KeyDev	0.0226***	0.000	12073	Capital IQ AI Key Developments also map positively to AI workforce.
United States	Investment-confirmation benchmark	Talk/no-workforce	-0.0359***	0.000	12073	U.S. talk without workforce falls as observable AI workforce increases.
China	Disclosure-action wedge	AI talk without broad AI hiring	0.0446**	0.011	34266	China shows the talk/no-hire response

Market	Margin	Outcome	Estimate	p	N	Interpretation
						in the main predetermined DDD design.
United States	Post-ChatGPT wedge benchmark	AI talk without AI KeyDev	0.0352	0.280	3193	The U.S. wedge estimate is similar in scale but much less precise; it is not the main institutional contrast.
China	Real-input confirmation	Broad AI hiring	-0.0005	0.970	34266	Broad AI hiring does not increase.
China	Supporting real-input evidence	R&D intensity	-0.0041***	0.008	28757	R&D falls in the main design; this is supporting evidence, not the sole mechanism.

Notes: This table summarizes the cross-country institutional contrast. The main contrast is insider liquidity use, not the raw wedge magnitude. The U.S. investment-confirmation rows benchmark whether AI language maps to observable AI workforce before the generative-AI shock. The U.S. post-shock wedge row is reported for transparency; its point estimate is similar in scale to the Chinese wedge but statistically weak. The table is not a pooled cross-country treatment design.

CSMAR's share pledge data add weaker but relevant context. The accessible WRDS table reports the stock of shares pledged, frozen, or under custody, not contract-level pledge dates, liquidation lines, or release dates, so we use it only as a pre-existing pledge-pressure measure. The pledge interaction is positive but marginal; the stronger finance result is the insider share-

reduction pattern around talk/no-hire filings among firms with greater pre-shock retail attention and in market-facing exit transactions.

10. Conclusion

The evidence identifies a breakdown in investment confirmation: AI language rises in precisely the state where observable AI hiring and other real-input proxies do not expand. The empirical wedge is defined by AI disclosure and observable broad AI hiring. Vendor use, employee redeployment, small-scale trials, and customer exposure can all generate AI language without broad AI job postings. The result is therefore not a statement-level falsity classification. It is a systematic divergence between technology language and the most visible investment margin that should accompany broad implementation.

This distinction matters for interpreting emerging-technology disclosure. AI disclosure can be informative on average and still become less informative for firms whose financial state makes costly real innovation unattractive. In our setting, disclosure rises, broad AI hiring does not increase, conventional real inputs do not expand, R&D intensity is lower in the main and city-year specifications, and realized talk/no-hire filings are followed by larger insider share reductions where pre-shock retail attention is high. The share-reduction relation is concentrated in market-facing sales and director/executive market-facing sales. Event-study estimates, filing-date checks, city-clustered inference, wild city bootstrap tests, and headquarters-city-by-year fixed effects support the disclosure-action wedge.

The broader implication is that technology shocks can widen the gap between innovation language and innovation input in retail-attention capital markets. The informativeness of technology disclosure is not a property of the technology alone. It depends on the sender's

financing state, local narrative availability, and investor structure, which together determine whether attention can become tradable liquidity.

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "Artificial Intelligence and Jobs: Evidence from Online Vacancies." Journal of Labor Economics 40 (S1): S293-S340.

Akyildirim, Erdinc, Shaen Corbet, Ahmet Sensoy, and Larisa Yarovaya. 2020. "The Impact of Blockchain-Related Name Changes on Corporate Performance." Journal of Corporate Finance 65:101759.

Alekseeva, Liudmila, Jose Azar, Mireia Gine, Sampsa Samila, and Bledi Taska. 2021. "The Demand for AI Skills in the Labor Market." Labour Economics 71:102002.

Allbirds, Inc. 2026a. "Allbirds, Inc. Executes $50M Convertible Financing Facility Agreement; Announces Expansion into AI Compute Infrastructure." Exhibit 99.1 to Form 8-K, filed April 20, 2026.

Allbirds, Inc. 2026b. "Additional Definitive Proxy Soliciting Materials." Exhibit 99.2 to Form 8-K, filed April 20, 2026.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." Journal of Financial Economics 151:103745.

Barber, Brad M., and Terrance Odean. 2008. "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors." Review of Financial Studies 21 (2): 785-818.

Barrios, John Manuel, John L. Campbell, Ryan G. Johnson, and Christine Liu. 2025. "Signals or Smoke? The Determinants and Informativeness of Corporate Artificial Intelligence (AI) Disclosures." Working paper.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" Quarterly Journal of Economics 119 (1): 249-275.

Beyer, Anne, Daniel A. Cohen, Thomas Z. Lys, and Beverly R. Walther. 2010. "The Financial Reporting Environment: Review of the Recent Literature." Journal of Accounting and Economics 50 (2-3): 296-343.

Bloomberg News. 2026. "Allbirds Soars 582% after Sneaker Firm Rebrands as AI Stock." April 15.

Brown, James R., Steven M. Fazzari, and Bruce C. Petersen. 2009. "Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom." Journal of Finance 64 (1): 151-185.

Brochet, Francois. 2010. "Information Content of Insider Trades before and after the Sarbanes-Oxley Act." The Accounting Review 85 (2): 419-446.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2011. "Robust Inference with Multiway Clustering." Journal of Business and Economic Statistics 29 (2): 238-249.

Cao, Yang, Miao Liu, Jiaping Qiu, and Ran Zhao. 2026. "Determinants and Information Content of Voluntary Artificial Intelligence Disclosure." Working paper.

Cao, Sean S., Wei Jiang, Baozhong Yang, and Alan L. Zhang. 2023. "How to Talk When a Machine Is Listening: Corporate Disclosure in the Age of AI." Review of Financial Studies 36 (9): 3603-3642.

Cooper, Michael J., Orlin Dimitrov, and P. Raghavendra Rau. 2001. "A Rose.com by Any Other Name." Journal of Finance 56 (6): 2371-2388.

Cooper, Michael J., Ajay Khorana, Igor Osobov, Ajay Patel, and P. Raghavendra Rau. 2005. "Managerial Actions in Response to a Market Downturn: Valuation Effects of Name Changes in the Dot.com Decline." Journal of Corporate Finance 11 (1-2): 319-335.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." Econometrica 50 (6): 1431-1451.

Da, Zhi, Joseph Engelberg, and Pengjie Gao. 2011. "In Search of Attention." Journal of Finance 66 (5): 1461-1499.

Dye, Ronald A. 1985. "Disclosure of Nonproprietary Information." Journal of Accounting Research 23 (1): 123-145.

Holmstrom, Bengt. 1989. "Agency Costs and Innovation." Journal of Economic Behavior and Organization 12 (3): 305-327.

Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh. 2009. "Driven to Distraction: Extraneous Events and Underreaction to Earnings News." Journal of Finance 64 (5): 2289-2325.

Leland, Hayne E. 1994. "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure." Journal of Finance 49 (4): 1213-1252.

Lenkey, Stephen L. 2014. "Advance Disclosure of Insider Trading." Review of Financial Studies 27 (8): 2504-2537.

Li, Feng. 2010. "The Information Content of Forward-Looking Statements in Corporate Filings--A Naive Bayesian Machine Learning Approach." Journal of Accounting Research 48 (5): 1049-1102.

Loughran, Tim, and Bill McDonald. 2011. "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks." Journal of Finance 66 (1): 35-65.

MacKinnon, James G., and Matthew D. Webb. 2017. "Wild Bootstrap Inference for Wildly Different Cluster Sizes." Journal of Applied Econometrics 32 (2): 233-254.

Manso, Gustavo. 2011. "Motivating Innovation." Journal of Finance 66 (5): 1823-1860.

Merton, Robert C. 1974. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." Journal of Finance 29 (2): 449-470.

Myers, Stewart C. 1977. "Determinants of Corporate Borrowing." Journal of Financial Economics 5 (2): 147-175.

Noe, Christopher F. 1999. "Voluntary Disclosures and Insider Transactions." Journal of Accounting and Economics 27 (3): 305-326.

Peng, Lin, and Wei Xiong. 2006. "Investor Attention, Overconfidence and Category Learning." Journal of Financial Economics 80 (3): 563-602.

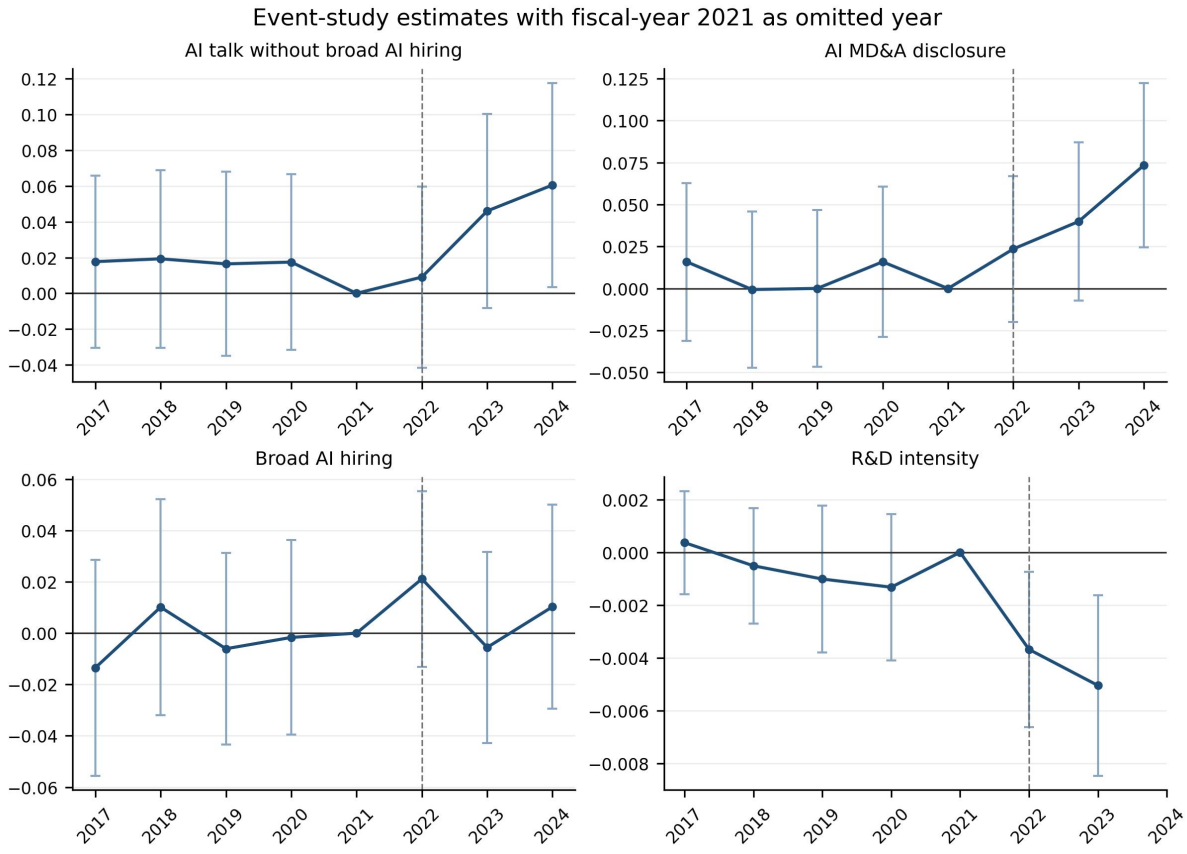
Verrecchia, Robert E. 1983. "Discretionary Disclosure." Journal of Accounting and Economics 5:179-194.

Appendix A. Illustrative Firm-Year

An illustrative weak-recovery service-sector firm filed its 2022 annual report after the generative-AI shock, described AI and related digital technologies as upgrading its core services, had no broad AI job postings in the matched hiring data, experienced abnormal filing-window trading value, and recorded RMB 24.9 million of insider share reductions over the next 60 days.

Appendix B. Timing Event Study with Fiscal-Year 2021 Omitted

The baseline event-study omits fiscal-year 2022 to match the conventional shock window in the main specification. Because fiscal-year 2022 annual reports were generally written and filed after the ChatGPT shock, Appendix Figure IA.1 instead omits fiscal-year 2021 and estimates the fiscal-year 2022 coefficient directly. This timing variant leaves the pretrend inference intact: joint tests over 2017-2020 have p-values of 0.928 for AI talk without broad hiring, 0.864 for AI MD&A disclosure, 0.835 for broad AI hiring, and 0.851 for RD2. The 2022 disclosure coefficients are small and statistically insignificant, while the AI disclosure coefficients become larger in 2023 and 2024. Broad AI hiring remains statistically indistinguishable from zero throughout the post-shock period.



Appendix Figure IA.1. Event-Study Estimates with Fiscal-Year 2021 Omitted

Internet Appendix Table IA.1. Component Decomposition of LowRecovery_i

State component	Outcome	Coef.	SE	p
LowRecovery_i composite	AI talk, no broad AI hiring	0.0446**	(0.0175)	0.011
LowRecovery_i composite	AI MD&A disclosure	0.0490***	(0.0129)	0.000
LowRecovery_i composite	Broad AI hiring indicator	-0.0005	(0.0135)	0.970
LowRecovery_i composite	R&D intensity (RD2)	-0.0041***	(0.0016)	0.008
High survival pressure	AI talk, no broad AI hiring	0.0382*	(0.0197)	0.052
High survival pressure	AI MD&A disclosure	0.0468***	(0.0177)	0.008
High survival pressure	Broad AI hiring indicator	-0.0005	(0.0153)	0.975
High survival pressure	R&D intensity (RD2)	-0.0037**	(0.0015)	0.011
High frontier distance	AI talk, no broad AI hiring	0.0388*	(0.0214)	0.069
High frontier distance	AI MD&A disclosure	0.0270	(0.0207)	0.191
High frontier distance	Broad AI hiring indicator	-0.0076	(0.0151)	0.615
High frontier distance	R&D intensity (RD2)	0.0018	(0.0018)	0.326
Short horizon	AI talk, no broad AI hiring	0.0411	(0.0492)	0.403
Short horizon	AI MD&A disclosure	0.0656	(0.0540)	0.225
Short horizon	Broad AI hiring indicator	0.0040	(0.0315)	0.900
Short horizon	R&D intensity (RD2)	0.0018	(0.0051)	0.730

Notes: This table reports the baseline LowRecovery_i x LocalNarrative_i x Post_t coefficient using the preferred composite and the separate 2021 components of LowRecovery_i. Regressions follow the baseline fixed-effect structure with firm, industry-year, and province-year fixed effects; standard errors are clustered by 2021 headquarters city.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.2. Local Narrative Scope

Specification	Outcome	Coef.	SE	p	N
Adjacent digital narrative scope test	AI talk, no broad AI hiring	0.0418***	(0.0149)	0.005	34266
Adjacent digital narrative scope test	AI MD&A disclosure	0.0360***	(0.0119)	0.002	33875
Adjacent digital narrative scope test	Broad AI hiring indicator	-0.0111	(0.0110)	0.314	34266
Adjacent digital narrative scope test	R&D intensity (RD2)	-0.0041***	(0.0015)	0.008	28757
Adjacent digital scope test, city-year FE	AI talk, no broad AI hiring	0.0411***	(0.0152)	0.007	34266
Adjacent digital scope test, city-year FE	AI MD&A disclosure	0.0314***	(0.0116)	0.007	33875
Adjacent digital scope test, city-year FE	Broad AI hiring indicator	-0.0148	(0.0109)	0.173	34266
Adjacent digital scope test, city-year FE	R&D intensity (RD2)	-0.0038**	(0.0017)	0.028	28757
AI exposure residualized vs. digital	AI talk, no broad AI hiring	0.0162	(0.0223)	0.467	34266
AI exposure residualized vs. digital	AI MD&A disclosure	0.0480***	(0.0169)	0.004	33875
AI exposure residualized vs. digital	Broad AI hiring indicator	0.0341**	(0.0167)	0.041	34266
AI exposure residualized vs. digital	R&D intensity (RD2)	-0.0004	(0.0018)	0.831	28757
AI residualized vs. digital, city-year FE	AI talk, no broad AI hiring	0.0092	(0.0227)	0.684	34266
AI residualized vs. digital,	AI MD&A disclosure	0.0386**	(0.0158)	0.015	33875

Specification	Outcome	Coef.	SE	p	N
city-year FE					
AI residualized vs. digital, city-year FE	Broad AI hiring indicator	0.0351**	(0.0177)	0.047	34266
AI residualized vs. digital, city-year FE	R&D intensity (RD2)	0.0002	(0.0020)	0.934	28757

Notes: This table tests the scope of the local narrative exposure. Adjacent digital narrative rows use pre-shock local peer digital-transformation language, and residualized rows remove adjacent digital narrative exposure from the AI exposure measure. City-year rows absorb headquarters-city-by-year shocks; standard errors are clustered by 2021 headquarters city.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.3. Retail-Participation Context

Specification	Outcome	Coef.	SE	p	N
High-retail provinces	AI talk, no broad AI hiring	0.0625***	(0.0217)	0.004	28738
High-retail provinces	AI MD&A disclosure	0.0552***	(0.0157)	0.000	28489
High-retail provinces	Broad AI hiring indicator	-0.0125	(0.0160)	0.437	28738
High-retail provinces	R&D intensity (RD2)	-0.0041**	(0.0019)	0.032	24219
Low-retail provinces	AI talk, no broad AI hiring	0.0424	(0.0429)	0.322	5452
Low-retail provinces	AI MD&A disclosure	0.0714*	(0.0398)	0.073	5386
Low-retail provinces	Broad AI hiring indicator	0.0358	(0.0262)	0.172	5452
Low-retail provinces	R&D intensity (RD2)	-0.0025	(0.0022)	0.240	4538
High-retail provinces, city-year FE	AI talk, no broad AI hiring	0.0628***	(0.0207)	0.002	28738
High-retail provinces, city-year FE	AI MD&A disclosure	0.0530***	(0.0153)	0.001	28489
High-retail provinces, city-year FE	Broad AI hiring indicator	-0.0144	(0.0159)	0.364	28738
High-retail provinces, city-year FE	R&D intensity (RD2)	-0.0035*	(0.0020)	0.078	24219
Low-retail provinces, city-year FE	AI talk, no broad AI hiring	0.0513	(0.0418)	0.220	5452
Low-retail provinces, city-year FE	AI MD&A disclosure	0.0685*	(0.0366)	0.061	5386
Low-retail provinces, city-year FE	Broad AI hiring indicator	0.0235	(0.0245)	0.339	5452
Low-retail provinces, city-year FE	R&D intensity (RD2)	-0.0029	(0.0019)	0.132	4538

Notes: This table splits the baseline design by province-level retail-participation context. Coefficients are the LowRecovery_i x LocalNarrative_i x Post_t estimates for the indicated outcomes. Specifications follow the baseline fixed-effect structure, with city-year rows adding headquarters-city-by-year fixed effects; standard errors are clustered by 2021 headquarters city.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.4. RecoveryProspects_i Validation

Decile	Firms	Mean RecoveryProspects_i	Reach frontier	Max q rank	ST/suspended	Negative ROA	Mean ROA	Leverage
1	397	0.001	0.195	0.671	0.175	0.643	-0.059	0.724
2	397	0.004	0.084	0.585	0.041	0.398	-0.002	0.617
3	396	0.007	0.068	0.532	0.015	0.365	0.007	0.528
4	397	0.012	0.051	0.495	0.028	0.328	0.014	0.466
5	396	0.016	0.056	0.488	0.015	0.319	0.018	0.375
6	397	0.021	0.076	0.496	0.010	0.254	0.030	0.305

Decile	Firms	Mean RecoveryProspects_i	Reach frontier	Max q rank	ST/suspended	Negative ROA	Mean ROA	Leverage
7	396	0.029	0.122	0.572	0.023	0.277	0.030	0.299
8	397	0.043	0.170	0.613	0.023	0.266	0.034	0.306
9	396	0.108	0.316	0.738	0.020	0.265	0.034	0.395
10	397	0.461	0.576	0.867	0.025	0.213	0.059	0.278

Notes: This table validates RecoveryProspects_i by decile using post-2021 recovery outcomes and baseline financial characteristics. Lower deciles indicate weaker predicted recovery prospects. Reach frontier, maximum q rank, ST/suspended status, negative ROA, and mean ROA are measured over the validation window; leverage is measured at the 2021 baseline.

Internet Appendix Table IA.5. Peer-Count Split-Half Reliability

Peer-count bin	Cities	Mean peer count	Split-half corr.	Mean abs. gap
1-9	65	6.8	0.088	0.394
10-19	59	13.9	0.055	0.322
20-49	53	26.9	0.424	0.197
50-99	24	74.8	0.345	0.147
100+	25	295.6	0.732	0.075

Notes: This table reports split-half reliability for the pre-shock headquarters-city peer exposure. Cities are grouped by peer-count bins. Split-half correlation compares exposure measures built from split local peer sets, and mean absolute gap is the average absolute difference between the two split-half exposure measures.

Internet Appendix Table IA.6. Event-Study Pretrend Tests with Fiscal-Year 2021 Omitted

Outcome	Omitted year	Pre years	Wald stat.	p	N	Firms
AI talk, no broad AI hiring	2021	2017,2018,2019,2020	0.878	0.928	34266	3896
AI MD&A disclosure	2021	2017,2018,2019,2020	1.287	0.864	33875	3895
R&D intensity (RD2)	2021	2017,2018,2019,2020	1.360	0.851	28757	3790
Broad AI hiring indicator	2021	2017,2018,2019,2020	1.451	0.835	34266	3896

Notes: This table reports joint Wald tests for pre-shock event-study coefficients when fiscal-year 2021 is the omitted year. The pretrend years are 2017-2020. Outcomes follow the timing event-study specification, and standard errors are clustered by 2021 headquarters city.

Internet Appendix Table IA.7. Event-Study Pretrend Tests with Fiscal-Year 2022 Omitted

Outcome	Omitted year	Pre years	Wald stat.	p	N	Firms
AI talk, no broad AI hiring	2022	2017,2018,2019,2020,2021	4.597	0.467	34266	3896
AI MD&A disclosure	2022	2017,2018,2019,2020,2021	1.660	0.894	33875	3895
R&D intensity (RD2)	2022	2017,2018,2019,2020,2021	2.363	0.797	28757	3790
Broad AI hiring indicator	2022	2017,2018,2019,2020,2021	4.933	0.424	34266	3896

Notes: This table reports joint Wald tests for pre-shock event-study coefficients when fiscal-year 2022 is the omitted year. The pretrend years are 2017-2021. This timing variant estimates the 2022 coefficient directly because fiscal-year 2022 annual reports were generally filed after the ChatGPT shock; standard errors are clustered by 2021 headquarters city.

Internet Appendix Table IA.8. Pre-Filing Insider Share-Reduction Placebo

Specification	Outcome	Variable	Coef.	SE	p	N
Pre-filing placebo	Insider share-reduction indicator [-90,-1]	Talk/no-hire x high forum attention	-0.0103	(0.0166)	0.536	7472
Pre-filing placebo	Log insider share-reduction amount [-90,-1]	Talk/no-hire x high forum attention	-0.0769	(0.2491)	0.758	7472
Pre-filing placebo	Log positive net insider selling [-90,-1]	Talk/no-hire x high forum attention	-0.0752	(0.2505)	0.764	7472
Pre-filing placebo	Insider purchase indicator [-90,-1]	Talk/no-hire x high forum attention	0.0145	(0.0127)	0.255	7472
Pre-filing placebo	Log insider purchase amount [-90,-1]	Talk/no-hire x high forum attention	0.2339	(0.1806)	0.195	7472

Notes: This table repeats the realized talk/no-hire x high forum attention specification using insider share-reduction and purchase outcomes measured over the [-90,-1] window before annual-report filing. The placebo tests whether comparable insider trading appears before the filing-window disclosure event; standard errors are clustered by 2021 headquarters city.

Internet Appendix Table IA.9. Full Timing and Local-Shock Robustness

Specification	Outcome	Coef.	SE	p	N
City-year FE, city cluster	AI talk, no broad AI hiring	0.0471***	(0.0174)	0.007	34266
City-year FE, city cluster	AI MD&A disclosure	0.0470***	(0.0129)	0.000	33875
City-year FE, city cluster	Broad AI hiring indicator	-0.0038	(0.0133)	0.776	34266
City-year FE, city cluster	R&D intensity (RD2)	-0.0036**	(0.0016)	0.028	28757
Exclude FY2022	AI talk, no broad AI hiring	0.0455**	(0.0192)	0.018	30403
Exclude FY2022	AI MD&A disclosure	0.0529***	(0.0150)	0.000	30033
Exclude FY2022	Broad AI hiring indicator	0.0023	(0.0138)	0.868	30403
Exclude FY2022	R&D intensity (RD2)	-0.0050***	(0.0017)	0.004	25105
Filing-date post	AI talk, no broad AI hiring	0.0303*	(0.0182)	0.096	34190
Filing-date post	AI MD&A disclosure	0.0401***	(0.0154)	0.009	33875
Filing-date post	Broad AI hiring indicator	0.0092	(0.0122)	0.452	34190
Filing-date post	Log broad AI postings	0.0321	(0.0227)	0.158	34190
Filing-date post, city-year FE	AI talk, no broad AI hiring	0.0349*	(0.0179)	0.051	34190
Filing-date post, city-year FE	AI MD&A disclosure	0.0387**	(0.0155)	0.013	33875
Filing-date post, city-year FE	Broad AI hiring	0.0050	(0.0118)	0.673	34190

Specification	Outcome	Coef.	SE	p	N
	indicator				
Filing-date post, city-year FE	Log broad AI postings	0.0247	(0.0227)	0.276	34190
Filing-date post, exclude FY2022	AI talk, no broad AI hiring	0.0476**	(0.0194)	0.014	30327
Filing-date post, exclude FY2022	AI MD&A disclosure	0.0529***	(0.0150)	0.000	30033
Filing-date post, exclude FY2022	Broad AI hiring indicator	0.0019	(0.0135)	0.886	30327
Filing-date post, exclude FY2022	Log broad AI postings	0.0174	(0.0249)	0.484	30327
City-industry-year FE, AI exposure	AI talk, no broad AI hiring	0.0568*	(0.0300)	0.058	34266
City-industry-year FE, AI exposure	AI MD&A disclosure	0.0511**	(0.0238)	0.032	33875
City-industry-year FE, AI exposure	Broad AI hiring indicator	-0.0117	(0.0209)	0.578	34266
City-industry-year FE, AI exposure	R&D intensity (RD2)	-0.0033	(0.0032)	0.295	28757
City-broad-industry-year FE, AI exposure	AI talk, no broad AI hiring	0.0476**	(0.0198)	0.016	34266
City-broad-industry-year FE, AI exposure	AI MD&A disclosure	0.0398**	(0.0165)	0.016	33875
City-broad-industry-year FE, AI exposure	Broad AI hiring indicator	-0.0083	(0.0147)	0.575	34266
City-broad-industry-year FE, AI exposure	R&D intensity (RD2)	-0.0036**	(0.0018)	0.047	28757

Notes: This table reports robustness to city-year fixed effects, excluding fiscal-year 2022, filing-date post definitions, and city-industry-year local-shock controls. Coefficients are on LowRecovery_i x LocalNarrative_i x Post_t unless otherwise stated; standard errors are clustered by 2021 headquarters city.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.10. Additional Insider Share-Reduction Robustness

Specification	Outcome	Variable	Coef.	SE	p	N
High-attention main effect	Insider share-reduction indicator [0,+60]	High attention	0.0072	(0.0108)	0.505	7472
High-attention main effect	Log insider share-reduction amount [0,+60]	High attention	0.0657	(0.1561)	0.674	7472
Talk/hire x high attention	Insider share-reduction indicator [0,+60]	Talk/hire x high attention	-0.0163	(0.0194)	0.403	7472
Talk/hire x high attention	Log insider share-reduction amount [0,+60]	Talk/hire x high attention	-0.1359	(0.2913)	0.641	7472
No-talk/hire x high attention	Insider share-reduction indicator [0,+60]	No-talk/hire x high attention	-0.0120	(0.0296)	0.684	7472
No-talk/hire x high attention	Log insider share-reduction amount [0,+60]	No-talk/hire x high attention	-0.3830	(0.2739)	0.162	7472
Reduced-form DDD state	Insider share-reduction indicator [0,+60]	DDD state	-0.0079	(0.0085)	0.355	30414
Reduced-form DDD state	Log insider share-reduction amount	DDD state	-0.0805	(0.1411)	0.568	30414

Specification	Outcome [0,+60]	Variable	Coef.	SE	p	N
Reduced-form DDD state	Insider purchase indicator [0,+60]	DDD state	-0.0152*	(0.0084)	0.069	30414

Notes: This table reports additional insider share-reduction specifications around annual-report filing. The first rows separate high-attention main effects and alternative disclosure-hiring states; the reduced-form rows use the predetermined DDD state. Insider outcomes are measured over [0,+60] after filing unless the row label states otherwise; standard errors are clustered by 2021 headquarters city.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.11. U.S. External Validation

Panel	Outcome	Variable	Coef.	SE	p	N	Firms
U.S. workforce benchmark	10-K AI talk on AI workforce	Log(1 + AI employment)	0.0945***	(0.0096)	0.000	12073	1727
U.S. workforce benchmark	10-K AI focus on AI workforce	Log(1 + AI employment)	0.0040***	(0.0011)	0.000	8456	1619
U.S. workforce benchmark	Market-facing AI talk on AI workforce	Log(1 + AI employment)	0.0563***	(0.0073)	0.000	12073	1727
U.S. workforce benchmark	AI KeyDev on AI workforce	Log(1 + AI employment)	0.0226***	(0.0059)	0.000	12073	1727
U.S. workforce benchmark	AI KeyDev count on AI workforce	Log(1 + AI employment)	0.0304***	(0.0095)	0.001	12073	1727
U.S. workforce benchmark	Talk/no-workforce on AI workforce	Log(1 + AI employment)	-0.0359***	(0.0055)	0.000	12073	1727
U.S. workforce benchmark: financial state	Talk/no-workforce on distress	Distress score	-0.0033	(0.0024)	0.180	11215	1690
U.S. workforce benchmark: financial state	10-K AI talk on distress	Distress score	0.0037	(0.0038)	0.325	11215	1690
U.S. workforce benchmark: financial state	AI workforce on distress	Distress score	0.0016	(0.0043)	0.707	11215	1690
U.S. workforce benchmark: financial state	AI talk-workforce mapping x distress	Log(1 + AI employment) x distress	0.0035	(0.0031)	0.263	11215	1690
U.S. workforce benchmark: financial state	Market AI talk-workforce mapping x distress	Log(1 + AI employment) x distress	-0.0013	(0.0028)	0.652	11215	1690

Notes: This table is an external-validation exercise using U.S. data. The workforce benchmark uses 2012-2021 firm-years from the AI2021 workforce file merged to 10-K AI text and Capital IQ Key Developments. Regressions include firm and fiscal-year fixed effects and cluster standard errors by firm. The table is not a parallel post-2023 identification design because the U.S. AI workforce data end in 2021. Instead, it shows that, in a pre-generative-AI U.S. benchmark, AI language and AI KeyDev activity are strongly associated with AI workforce, while talk/no-workforce falls with AI workforce.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.12. U.S. 13F Retail-Attention Benchmark

Specification	Outcome	Variable	Coef.	SE	p	N	Firms	Clusters
U.S. 13F retail, industry-year FE	Insider sale indicator [0,+60]	Talk/no-KeyDev x retail-investor	0.0187	(0.0256)	0.466	1410	522	49

Specification	Outcome	Variable	Coef.	SE	p	N	Firms	Clusters
U.S. 13F retail, industry-year FE	Log insider sale value [0,+60]	Talk/no-KeyDev x retail-investor exposure	0.3341	(0.3724)	0.370	1410	522	49
U.S. 13F retail, industry-year FE	Insider sale indicator [-60,-1]	Talk/no-KeyDev x retail-investor exposure	-0.0084	(0.0084)	0.322	1410	522	49
U.S. 13F retail, industry-year FE	Log insider sale value [-60,-1]	Talk/no-KeyDev x retail-investor exposure	0.1632	(0.1109)	0.141	1410	522	49
U.S. 13F retail, firm FE	Insider sale indicator [0,+60]	Talk/no-KeyDev x retail-investor exposure	0.0334	(0.0374)	0.373	1410	522	49
U.S. 13F retail, firm FE	Log insider sale value [0,+60]	Talk/no-KeyDev x retail-investor exposure	0.6556	(0.5912)	0.267	1410	522	49
U.S. 13F retail, firm FE	Insider sale indicator [-60,-1]	Talk/no-KeyDev x retail-investor exposure	-0.0129	(0.0106)	0.220	1410	522	49
U.S. 13F retail, firm FE	Log insider sale value [-60,-1]	Talk/no-KeyDev x retail-investor exposure	0.2679	(0.1632)	0.101	1410	522	49

Notes: This table is a U.S. institutional benchmark, not a parallel causal design. The sample covers post-2023 U.S. firm-years in which 10-K AI language, Capital IQ Key Developments, insider Form 4 transactions, and a corrected pre-2022 13F institutional-ownership snapshot can be linked. The retail-investor exposure proxy equals one minus pre-shock institutional ownership. Talk/no-KeyDev denotes 10-K AI talk without a Capital IQ AI Key Development. The interaction between talk/no-KeyDev and pre-shock retail-investor exposure is not stable across U.S. specifications; post-filing sale estimates weaken under firm fixed effects and pre-filing placebo sale values move in some industry-year specifications. The contrast disciplines the interpretation: the Chinese insider liquidity-use result is a retail-attention market margin rather than a universal implication of AI talk/no-hire.

Internet Appendix Table IA.13. Cross-Country Institutional Contrast

Block	Market	Outcome	Estimate	p	N	Interpretation
A. U.S. investment-confirmation benchmark	United States	AI disclosure	0.0945***	0.000	12073	U.S. AI language maps positively to observable AI workforce before the generative-AI shock.
A. U.S. investment-confirmation benchmark	United States	AI KeyDev	0.0226***	0.000	12073	Capital IQ AI Key Developments also map positively to AI workforce.
A. U.S. investment-confirmation benchmark	United States	Talk/no-workforce	-0.0359***	0.000	12073	U.S. talk without workforce falls as observable AI workforce increases.
B. U.S. post-ChatGPT wedge	United	AI talk without AI	-0.0250	0.345	3193	The U.S. post-

Block	Market	Outcome	Estimate	p	N	Interpretation
benchmark	States	KeyDev				ChatGPT distress x industry-narrative design does not reproduce the Chinese wedge.
B. U.S. post-ChatGPT wedge benchmark	United States	AI talk without AI KeyDev	0.0352	0.280	3193	The stricter U.S. industry-year specification is also statistically weak.
C. China disclosure-action wedge	China	AI talk without broad AI hiring	0.0446**	0.011	34266	China shows the disclosure-action wedge in the main predetermined DDD design.
C. China disclosure-action wedge	China	AI MD&A disclosure	0.0490***	0.000	33875	The same state-narrative interaction increases AI disclosure.
C. China disclosure-action wedge	China	Broad AI hiring	-0.0005	0.970	34266	Broad AI hiring does not increase.
C. China disclosure-action wedge	China	R&D intensity	-0.0041***	0.008	28757	R&D falls in the main design; this is supporting evidence, not the sole mechanism.
D. China four-state decomposition	China	No talk / no hire	-0.0464***	0.002	33875	Probability mass moves out of the no-talk/no-hire state.
D. China four-state decomposition	China	Talk / no hire	0.0464***	0.009	33875	Probability mass moves into the talk/no-hire state.
D. China four-state decomposition	China	Talk / hire	0.0025	0.826	33875	Talk accompanied by hiring does not move.
E. China insider liquidity use	China	Insider share-reduction indicator [0,+60]	0.0146**	0.015	7503	Realized talk/no-hire filings are followed by more insider share reductions.
E. China insider liquidity use	China	Log insider share-reduction amount [0,+60]	0.2034**	0.027	7503	The reduction amount also rises after realized talk/no-hire filings.
E. China insider liquidity use	China	Log positive net insider selling [0,+60]	0.1689***	0.003	7340	The insider-liquidity relation is stronger where pre-shock forum attention is higher.
F. China insider transaction type	China	Market-sale share-reduction indicator [0,+60]	0.0113**	0.011	7472	The China insider response is concentrated in market-facing sales.
F. China insider transaction type	China	Director/executive market-sale amount [0,+60]	0.1242**	0.041	7472	Director/executive market-facing sale amounts rise after realized talk/no-hire filings.
F. China insider transaction type	China	Equity-incentive share-reduction	0.0010	0.486	7472	Equity-incentive transactions do not

Block	Market	Outcome	Estimate	p	N	Interpretation
		indicator [0,+60]				move in the same way.
G. U.S. 13F retail-attention benchmark	United States	Insider sale indicator [0,+60]	0.0187	0.466	1410	The U.S. 13F retail interaction is statistically weak.
G. U.S. 13F retail-attention benchmark	United States	Log insider sale value [0,+60]	0.3341	0.370	1410	U.S. sale values do not deliver a stable insider-sale relation.
G. U.S. 13F retail-attention benchmark	United States	Log insider sale value [-60,-1]	0.1632	0.141	1410	Pre-filing placebo values move in the industry-year specification, reinforcing descriptive use only.
G. U.S. 13F retail-attention benchmark	United States	Insider sale indicator [0,+60]	0.0334	0.373	1410	The U.S. 13F retail interaction is statistically weak.
G. U.S. 13F retail-attention benchmark	United States	Log insider sale value [0,+60]	0.6556	0.267	1410	U.S. sale values do not deliver a stable insider-sale relation.
G. U.S. 13F retail-attention benchmark	United States	Log insider sale value [-60,-1]	0.2679	0.101	1410	Pre-filing placebo values move in the industry-year specification, reinforcing descriptive use only.

Notes: This table consolidates estimates from existing U.S. and Chinese exercises. It is an institutional contrast, not a new pooled cross-country causal design. The U.S. workforce benchmark uses pre-generative-AI 2012-2021 firm-years and shows that AI disclosure and AI Key Developments map to AI workforce. The Chinese rows report the main predetermined DDD and realized insider liquidity-use results. The U.S. 13F rows show that the retail-attention interaction does not deliver a stable insider-sale relation and has placebo movement in some specifications. The table supports the paper's interpretation: emerging-technology disclosure is investment-confirming in the U.S. benchmark, while weak financing state and retail-attention market structure in China make the disclosure-action wedge financially relevant around insider share reductions.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.14. Full-Text Regulatory-Scrutiny Diagnostic

Panel A. One-sample event-window diagnostics

Sample	Outcome	N	Mean	p
Clean regulatory events	CAR [0,+1]	187	0.0157	0.001
Clean regulatory events	Regulatory-event return [0,+5]	194	0.0137	0.031
Clean regulatory events	Regulatory-event return [0,+20]	194	0.0073	0.442
Clean regulatory events	Insider share reduction before regulatory event [-90,-1]	194	0.0619	0.000
Clean regulatory events	Insider share reduction after regulatory event [0,+60]	194	0.0361	0.008
Clean regulatory events	Crash indicator after regulatory event [0,+20]	194	0.0412	0.004
Talk/no-hire regulatory-event subset	CAR [0,+1]	90	0.0372	0.000
Talk/no-hire regulatory-event subset	Regulatory-event return [0,+5]	94	0.0142	0.164
Talk/no-hire regulatory-event subset	Regulatory-event return [0,+20]	94	0.0202	0.129

Sample	Outcome	N	Mean	p
Talk/no-hire regulatory-event subset	Insider share reduction before regulatory event [-90,-1]	94	0.0213	0.158
Talk/no-hire regulatory-event subset	Insider share reduction after regulatory event [0,+60]	94	0.0213	0.158
Talk/no-hire regulatory-event subset	Crash indicator after regulatory event [0,+20]	94	0.0319	0.083

Panel B. Matched-control event study

Sample	Outcome	Events	Treated minus controls	p
Clean regulatory events	Matched event return [0,+5]	191	0.0216	0.002
Clean regulatory events	Matched event return [+1,+20]	191	0.0047	0.675
Clean regulatory events	Matched crash indicator [+1,+20]	191	-0.0974	0.000
Clean regulatory events	Matched insider share reduction before event [-90,-1]	191	0.0136	0.452
Clean regulatory events	Matched insider share reduction after event [0,+60]	191	0.0042	0.781
Clean regulatory events	Matched log insider share-reduction amount [0,+60]	191	0.0818	0.705
Talk/no-hire regulatory-event subset	Matched event return [0,+5]	94	0.0231	0.052
Talk/no-hire regulatory-event subset	Matched event return [+1,+20]	94	0.0256	0.173
Talk/no-hire regulatory-event subset	Matched crash indicator [+1,+20]	94	-0.1255	0.000
Talk/no-hire regulatory-event subset	Matched insider share reduction before event [-90,-1]	94	-0.0106	0.538
Talk/no-hire regulatory-event subset	Matched insider share reduction after event [0,+60]	94	-0.0043	0.801
Talk/no-hire regulatory-event subset	Matched log insider share-reduction amount [0,+60]	94	-0.0460	0.854

Notes: Panel A uses full PDF text extracted from CNINFO regulatory and exchange documents downloaded from the local CNINFO regulatory-document archive. Clean events require inquiry, review, abnormal-trading, or scrutiny language together with AI, digital, concept, or business-substance language in the document text. Panel B matches each treated regulatory event to up to five non-event firm-years in the same fiscal year and industry-year, or the same broad industry-year when exact cells are thin; controls are evaluated on the treated firm's regulatory-event date. The diagnostic does not provide a clean main-text quasi-experiment: matched event firms have positive short-window returns, no significant 20-day correction, fewer 20-day crash realizations than controls, and no significant matched insider share-reduction differential. We therefore keep the regulatory material as an institutional diagnostic rather than as the paper's main identification design.

Internet Appendix Table IA.15. Insider Share-Reduction Magnitudes and Pre-Shock Forum-Attention Gradient

Panel A. RMB magnitude

Amount type	Treated filings	Control filings	Observed treated sum (RMB mn)	Treated mean (RMB mn)	Control mean (RMB mn)	Simple incremental (RMB mn)
Director/executive sale amount [0,+60]	935	814	384.7	0.411	0.189	208.4

Panel B. Pre-shock forum-read gradient

Forum quartile	Talk/no-hire N	No-talk/no-hire N	Talk/no-hire sale rate	No-talk/no-hire sale rate	Rate diff.	Log amount diff.
1	542	429	0.024	0.019	0.005	0.062
2	505	451	0.014	0.020	-0.006	-0.070
3	496	403	0.010	0.005	0.005	0.087
4	408	380	0.022	0.008	0.014	0.185

Notes: Panel A summarizes observed 60-day insider share-reduction amounts for post-2023 talk/no-hire filings at firms with high 2021 Eastmoney/Guba forum reads. The simple incremental amount compares treated high-forum talk/no-hire filings with high-forum no-talk/no-hire filings and is descriptive, not causal. Panel B bins firms by pre-shock forum reads and reports raw market-facing sale rates and log share-reduction amounts by realized disclosure-hiring state. The quartile gradient is used as a stress test for the pre-shock forum-attention proxy; the paper does not require a monotone decile pattern.

Internet Appendix Table IA.16. Buy-vs-Make Financial-Statement Proxies

Outcome	Coef.	SE	p	N	Mean	Note
Net intangible assets/assets	-0.0012	(0.0012)	0.334	37464	0.042	—
Data-resource intangibles/assets	—	—	—	71	—	insufficient sample
Long-term deferred expenses/assets	0.0001	(0.0003)	0.829	33337	0.006	—
Cash capex/assets	0.0009	(0.0013)	0.490	37626	0.044	—
Selling expense/assets	0.0002	(0.0009)	0.861	36375	0.040	—
Admin expense/assets	0.0003	(0.0006)	0.620	36938	0.039	—
Intangible amortization/assets	-0.0001	(0.0001)	0.278	37289	0.003	—
Long-term deferred amortization/assets	-0.0001	(0.0001)	0.562	34160	0.002	—
R&D expenses/assets (CSMAR)	-0.0006**	(0.0003)	0.028	24498	0.026	—

Notes: The table reports the baseline LowRecovery_i x LocalNarrative_i x Post_t coefficient for financial-statement proxies that could capture a buy-rather-than-make response to AI adoption. The specification uses firm, industry-year, and province-year fixed effects and city-level clustering. The data-resource intangible item is reported for transparency but has too few nonmissing observations to support inference; em dashes indicate that the coefficient is not reported.Stars denote * p<0.10, ** p<0.05, and *** p<0.011.

Internet Appendix Table IA.17. Text-Specificity Diagnostics

Specification	Outcome	Coef.	SE	p	N	Mean
Text specificity, composite state	Core AI mention	0.0687***	(0.0143)	0.000	34039	0.337
Text specificity, composite state	Broad AI mention	0.0339*	(0.0174)	0.052	34039	0.506
Text specificity, composite state	GenAI mention	0.0028	(0.0129)	0.825	34039	0.072
Text specificity, composite state	Core AI talk, no hire	0.0593***	(0.0141)	0.000	34039	0.200
Text specificity, composite state	Broad AI talk, no hire	0.0372**	(0.0179)	0.038	34039	0.335
Text specificity, composite state	GenAI talk, no hire	0.0157	(0.0099)	0.111	34039	0.038

Notes: This diagnostic reports dictionary-based text-specificity checks. It is not used as fine-grained sentence-level mechanism evidence. The main-text mechanism relies on observable investment confirmation: AI language rises without matching AI labor demand or conventional real-input expansion.Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.18. Non-AI Posting Composition Diagnostic

Outcome	Coef.	SE	p	N	Mean
Log total postings (raw hiring file)	0.0818	(0.0503)	0.104	37756	1.435
Log AI postings (raw hiring file)	-0.0158	(0.0167)	0.343	37756	0.156
Log non-AI postings (raw hiring file)	0.0828	(0.0506)	0.102	37756	1.428

Outcome	Coef.	SE	p	N	Mean
Log non-AI sales postings	0.0772*	(0.0467)	0.098	37756	0.974
Log non-AI customer/operations postings	0.0986**	(0.0459)	0.032	37756	1.077
Log non-AI administrative postings	0.0869**	(0.0415)	0.036	37756	0.880
Log non-AI technical postings	0.0747	(0.0489)	0.127	37756	1.271
Log other non-AI postings	0.0586*	(0.0311)	0.060	37756	0.519
Any total postings (raw hiring file)	-0.0042	(0.0092)	0.653	37756	0.390
Any AI postings (raw hiring file)	-0.0080	(0.0099)	0.417	37756	0.100
Any non-AI postings (raw hiring file)	-0.0059	(0.0094)	0.534	37756	0.390
Any non-AI sales postings	0.0152	(0.0125)	0.223	37756	0.347
Any non-AI customer/operations postings	0.0187*	(0.0109)	0.086	37756	0.363
Any non-AI administrative postings	0.0124	(0.0114)	0.279	37756	0.340
Any non-AI technical postings	-0.0018	(0.0094)	0.847	37756	0.378
Any other non-AI postings	0.0211	(0.0131)	0.108	37756	0.276

Notes: The table decomposes raw CSMAR job postings into AI and non-AI occupation categories. The coefficient is the preferred LowRecovery_i x LocalNarrative_i x Post_t coefficient with firm, industry-year, and province-year fixed effects and headquarters-city clustering. The diagnostic addresses the positive total/non-AI posting coefficients in Table 5: the strongest non-AI posting responses are in customer/operations and administrative jobs, while AI postings remain negative and insignificant and non-AI technical postings are not statistically significant. Coefficients are on the relevant triple interaction unless otherwise stated. Standard errors in parentheses are clustered by 2021 headquarters city except where a two-way or firm-fixed-effect specification is explicitly named. Stars denote * p<0.10, ** p<0.05, and *** p<0.01.

Internet Appendix Table IA.19. Retail Counterparty Diagnostics

Diagnostic	Target	Coeff.	SE	p-value	N	Mean Y	Role
Level-2 small-order OIB, [0,+4]	Talk/no-hire	0.0012	0.0007	0.086	9149	-0.0046	Weak supporting evidence
Level-2 small-order OIB x pre-shock forum, [0,+4]	Talk/no-hire x high forum	-0.0012	0.0013	0.377	9000	-0.0046	Null or not emphasized
A-share holder-count growth	Talk/no-hire	0.0095	0.0047	0.043	8598	0.0005	Supporting diagnostic
A-share holder-count growth x pre-shock forum	Talk/no-hire x high forum	-0.0030	0.0057	0.601	8598	0.0005	Null or not emphasized
Dragon-Tiger broker-department net buying x pre-shock forum, [0,+4]	Talk/no-hire x high forum	0.0030	0.0018	0.089	9158	-0.0012	Weak supporting evidence
Dragon-Tiger inclusion, [0,+4]	Talk/no-hire	0.0040	0.0058	0.483	9323	0.0597	Null or not emphasized

Notes: This table reports direct retail-counterparty diagnostics. Level-2 tests use CSMAR buy-sell imbalance data; small orders are A-share transactions below RMB 50,000 and active buy/sell direction follows Lee and Ready (1991). Shareholder-count tests use changes in A-share holder counts from the closest pre-filing report date to the next post-filing report date. Dragon-Tiger tests use CSMAR abnormal-trading list records and broker-department net buying scaled by abnormal-trading amount. The diagnostics support only a narrow auxiliary statement: talk/no-hire filings are followed by modest holder-count growth and weak small-order imbalance, while the stronger pre-shock forum interaction is not robust. Accordingly, these tests are reported as transparency diagnostics rather than as main evidence identifying the buyers behind insider share reductions.

Internet Appendix Table IA.20. Market Quality and Liquidity Diagnostics

Outcome	Coeff.	SE	p-value	N	Mean Y

Outcome	Coeff.	SE	p-value	N	Mean Y
Abnormal log trading value, [0,+1]	0.0147	0.0113	0.196	9099	0.2770
Abnormal log trading value, [0,+4]	0.0109	0.0113	0.335	9103	0.1698
Abnormal value turnover, [0,+4]	0.0004	0.0003	0.237	9103	0.0029
Abnormal Amihud price impact, [0,+1]	-0.0037	0.0012	0.003	9099	0.0041
Abnormal Amihud price impact, [0,+4]	-0.0016	0.0015	0.285	9103	0.0020
Abnormal Amihud price impact, [0,+19]	-0.0020	0.0008	0.019	9111	-0.0009
Abnormal absolute return, [0,+4]	-0.0002	0.0003	0.607	9103	0.0030
Abnormal return volatility, [0,+4]	-0.0003	0.0004	0.431	9103	-0.0013
Cumulative return, [+2,+20]	0.0069	0.0021	0.001	9113	0.0009
Cumulative return, [+2,+60]	0.0102	0.0027	0.000	9132	-0.0465

Notes: This table reports market-quality diagnostics around annual-report filings. The sample is post-generative-AI annual-report filings with daily trading data. The target coefficient is the realized talk/no-hire state, controlling for the other disclosure-hiring states, firm size, pre-filing trading value, pre-filing returns, industry-year fixed effects, and province-year fixed effects; standard errors are clustered by 2021 headquarters city. Abnormal measures subtract the [-20,-1] trading-day pre-filing average. Amihud is absolute daily return divided by RMB 100 million of trading value; negative coefficients indicate lower price impact. The estimates support a narrow liquidity-window interpretation: price impact declines after talk/no-hire filings, while abnormal absolute-return volatility does not rise.

Internet Appendix Table IA.21. Additional Microstructure Diagnostics

This table summarizes additional diagnostics using the newly purchased CSMAR Level-2 order-breadth files, daily high-low and limit-status data, Dragon-Tiger investor-type records, and pledge transaction details. The tests audit whether the realized talk/no-hire state is accompanied by stronger retail breadth, intraday price pressure, investor-type imbalance, or pledge relief. The results are used as diagnostics rather than main evidence because the strongest patterns are weak, window-specific, or economically ambiguous.

Outcome	Sample	Target	b	s.e.	p	N	Interpretation
Abnormal log small-order count [0,+4]	Baseline	Talk/no-hire	0.0021	0.0043	0.6259	8877	Short-window small-trade count breadth is null.
Abnormal small-order count share [0,+19]	Talk/no-hire x high forum	Talk/no-hire x high forum	0.0049	0.0027	0.0674	8795	Long-window small-trade share is weakly positive under high pre-shock forum attention.
Abnormal small-order amount share [0,+19]	Talk/no-hire x high forum	Talk/no-hire x high forum	0.0046	0.0022	0.0413	8795	Long-window small-trade amount share is positive under high

Outcome	Sample	Target	b	s.e.	p	N	Interpretation
							pre-shock forum attention.
Abnormal intraday high-low range [0,+4]	Baseline	Talk/no-hire	-0.0005	0.0003	0.0987	8943	Talk/no-hire filings do not increase intraday range; point estimate is slightly negative.
Limit-up indicator [0,+4]	Baseline	Talk/no-hire	0.0071	0.0067	0.2874	9128	No increase in limit-up probability.
Limit-down indicator [0,+4]	Baseline	Talk/no-hire	-0.0097	0.0046	0.0349	9128	Down-limit probability falls; no evidence of immediate crash risk.
Broker-department net buying [0,+4]	Talk/no-hire x high forum	Talk/no-hire x high forum	0.0030	0.0018	0.0893	9158	Ordinary broker-department net buying is weakly positive only in the short window.
Institution-special-seat net buying [0,+4]	Talk/no-hire x high forum	Talk/no-hire x high forum	-0.0009	0.0007	0.1991	9158	Institution-special-seat net buying is not positive.
Pledge-release indicator [0,+60]	Baseline	Talk/no-hire	0.0127	0.0078	0.1033	9323	No robust post-filing pledge-release relief.
Net pledged-share change/total shares [0,+60]	Talk/no-hire x high forum	Talk/no-hire x high forum	0.0008	0.0003	0.0069	8974	High-forum interaction predicts more net pledged shares, which is ambiguous and not a relief channel.

Notes: Level-2 breadth outcomes are abnormal changes relative to the [-20,-1] trading-day pre-filing window. Intraday range and limit-hit outcomes use Hiprc, Loprc, and LimitStatus. Dragon-Tiger outcomes separate ordinary broker-department and institution-special-seat net buying. Pledge-detail outcomes use DeclareDate when available and ChangeDate otherwise. The evidence supports only a cautious liquidity-window interpretation and does not identify a direct retail counterparty or a pledge-relief mechanism.