

Within-bank evidence that higher AI intensity predicts lower next-year stability

Tudor-Andrei Drăgan | Wenjing Jiang | Simona Nistor | Steven Ongena

Babes-Bolyai University · Zhejiang University of Science and Technology · University of Zurich / Swiss Finance Institute · KU Leuven / NTNU / CEPR

## CORE MESSAGE

AI adoption is not prudentially neutral. Using within-bank variation, lagged design, rich controls, and 2SLS, we find that higher AI intensity predicts lower next-year bank stability.

## QUESTION AND CONTRIBUTION

### Research question:

What happens to bank risk-taking when AI becomes embedded in screening, pricing, and monitoring?


### Our contribution:

- Bank-level AI-intensity index based on annual-report disclosures.
- Within-bank panel design links AI changes to next-year stability.
- Z-score decomposition points to volatility and profitability channels.
- Endogeneity addressed explicitly through timing, controls, and IV evidence.

## HOW IS AI INTENSITY MEASURED?

Source: annual-report disclosures, CSMAR text classification framework

### AI-related dimensions:



$$\text{AI intensity} = \frac{1}{3} \times [z \cdot \ln(1+\text{Agent}) + z \cdot \ln(1+\text{Cloud}) + z \cdot \ln(1+\text{Big data})]$$

Robustness: static-count index, PCA-based index, raw-count index, and investment-based AI proxies.

35 listed Chinese commercial banks	2013–2024 sample period	227 baseline bank-year observations
---------------------------------------	----------------------------	--

Dependent variable: next-year  $\ln(z\text{-score})$ , an accounting-based distance-to-default measure.

## IDENTIFICATION STRATEGY: ADDRESSING ENDOGENEITY

### WITHIN-BANK TIMING

Lagged AI intensity predicts next-year stability after bank and year fixed effects. Identification comes from changes within the same bank over time.

### RICH CONTROLS AND ROBUSTNESS

Results survive standard bank controls, alternative Z-score windows, alternative AI measures, vast bank/province controls, and threshold designs.

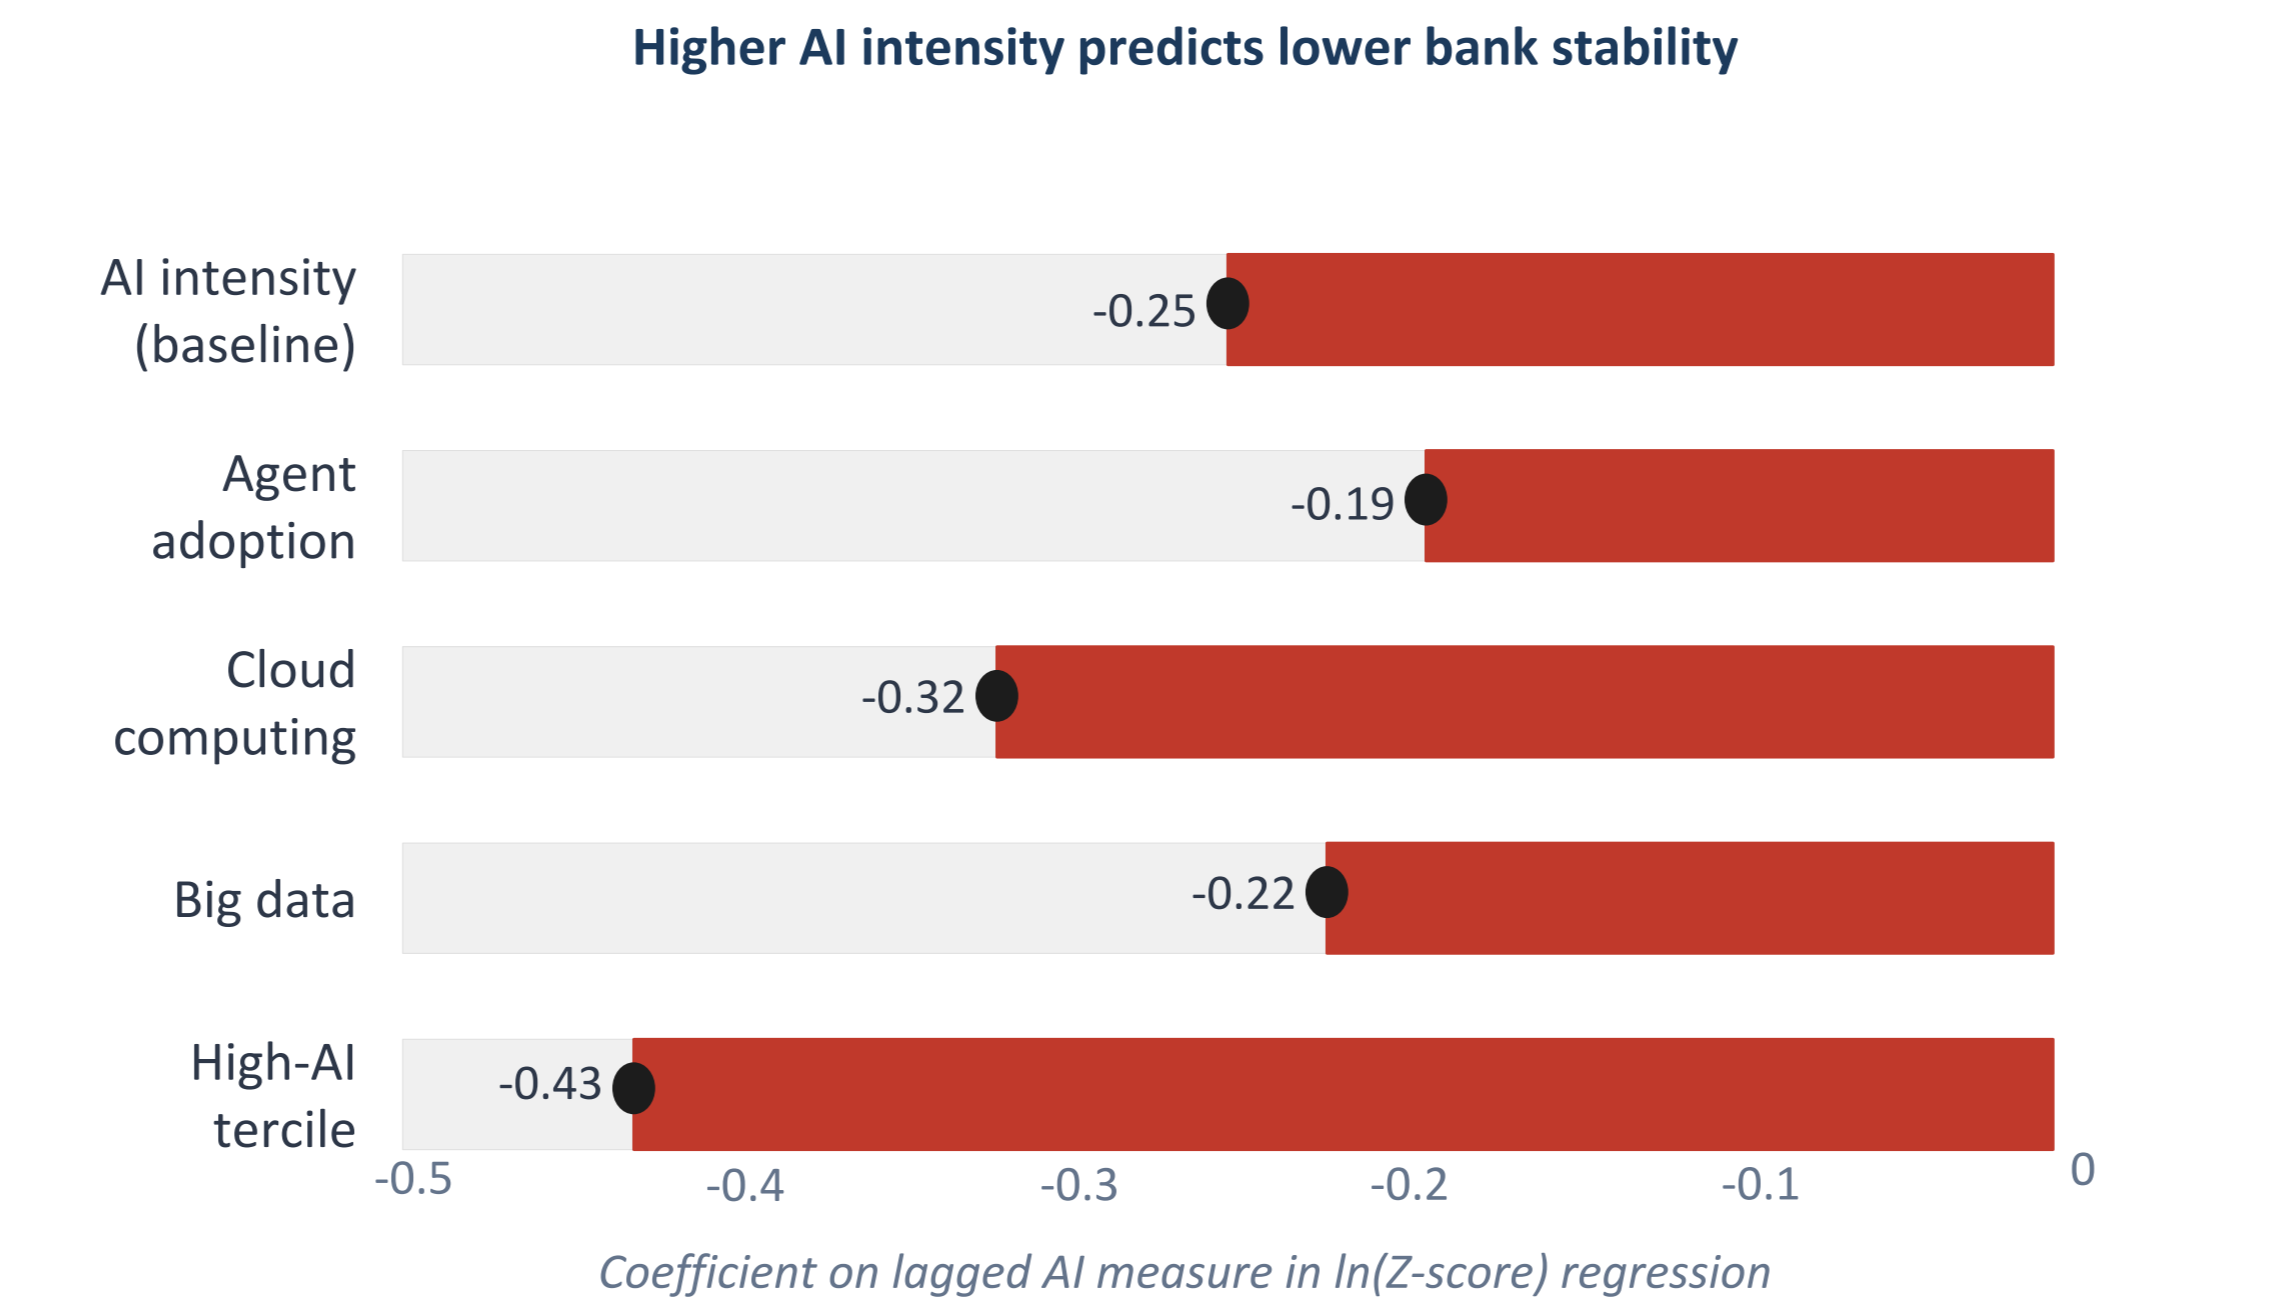
### MEASUREMENT CREDIBILITY

The baseline proxy is transparent and replicable. Findings also hold with PCA, total-count, raw-count, and investment-based AI proxies.

### 2SLS ENDOGENEITY TREATMENT

Instrument (lagged AI intensity with): i) leave-one-out peer AI adoption in same province-year; ii) growth in provincial AI firms; iii) growth in new product R&D / GDP; iv) growth in industrial R&D personnel / population. IV estimates reinforce the negative AI–stability relation.

## MAIN RESULTS




A one-standard-deviation increase in AI intensity is associated with about a 0.25 decline in next-year log Z-score (roughly a 22% lower Z-score level).

## MECHANISM



Stability weakens through higher earnings risk and lower profitability.

## WHEN IS AI MOST RISKY?



Supervisor takeaway: Risk rises when banks rely more on interbank funding and weakens when liquidity, deposit growth, and provisioning buffers are stronger.

## REGULATORY IMPORTANCE

- ✓ Treat AI scaling as a prudentially relevant organizational change.
- ✓ Supervisory attention should focus on model risk, governance, and operational resilience.
- ✓ Risks are greater when liquidity, funding stability, or provisioning buffers are weak.
- ✓ AI oversight should be integrated with liquidity and funding-risk surveillance.

## TAKE-AWAY

The conclusion does not rest on one model or one proxy; it is triangulated across fixed-effects timing, measurement robustness, richer controls, and IV evidence.

# Swiss Finance Institute Research Paper Series N°26-25

Artificial Intelligence and Risk-Taking in  
Banking



**Tudor-Andrei Drăgan**

Babeş-Bolyai University of Cluj-Napoca

**Wenjing Jiang**

Babeş-Bolyai University of Cluj-Napoca and Zhejiang University of  
Science and Technology

**Simona Nistor**

Babeş-Bolyai University of Cluj-Napoca

**Steven Ongena**

University of Zurich, KU Leuven, NTNU Business School,  
Swiss Finance Institute, and CEPR

# Artificial Intelligence and Risk-Taking in Banking

**Tudor-Andrei Drăgan**

*Babeş-Bolyai University of Cluj-Napoca, Finance Department*

E-mail: [tudor.dragan@econ.ubbcluj.ro](mailto:tudor.dragan@econ.ubbcluj.ro)

ORCID: 0009-0002-2636-3996

**Wenjing Jiang**

*Babeş-Bolyai University of Cluj-Napoca, Finance Department; Zhejiang University of Science and Technology, School of Science, Hangzhou, China*

E-mail: [jiang.wenjing@econ.ubbcluj.ro](mailto:jiang.wenjing@econ.ubbcluj.ro)

ORCID: 0009-0009-9473-1042

**Simona Nistor**

*Babeş-Bolyai University of Cluj-Napoca, Finance Department*

E-mail: [simona.nistor@econ.ubbcluj.ro](mailto:simona.nistor@econ.ubbcluj.ro)

ORCID: 0000-0001-5258-9214

**Steven Ongena**

*University of Zurich Department of Banking and Finance; Swiss Finance Institute; KU Leuven; NTNU Business School; Centre for Economic Policy Research (CEPR)*

E-mail: [steven.ongena@df.uzh.ch](mailto:steven.ongena@df.uzh.ch)

ORCID: 0000-0002-8381-0062

February, 2026

# Artificial Intelligence and Risk-Taking in Banking

## Abstract

We study how banks' adoption of artificial intelligence (AI) affects risk-taking and stability using data on major listed Chinese banks (2013–2024). We construct an AI-intensity index from agent adoption, cloud computing, and big data measures. Higher AI intensity is associated with greater risk and lower stability: a one-standard deviation increase reduces the Z-score by 0.25 units (about 22%) next year. This operates mainly via higher earnings volatility and lower profitability, not weaker capitalization. The effect is stronger for banks with low liquidity, unstable funding, and weak loss-absorption, implying AI can amplify risk and requires tighter prudential oversight. (99 words)

**JEL classification:** G01, G21, G28, O32

## 1. Introduction

As leading adopters of artificial intelligence (AI) technologies, banks are increasingly integrating these tools into their core intermediation functions.<sup>1</sup> Since banks provide essential financial infrastructure for the real economy by allocating credit, transforming maturities, and distributing risks, the diffusion of AI is not merely an incremental technological enhancement but a structural shift with significant implications for financial stability. This shift matters precisely because banks are the primary conduits through which credit and risk are priced, allocated, and transmitted, and because these functions ultimately support market confidence and economic activity. If AI impacts how risks are assessed, priced, and managed, it can reshape the distribution of risk across borrowers and institutions, with direct implications for consumer welfare and systemic resilience.

Against this background, we ask two related questions. First, what happens to bank risk-taking when AI becomes embedded in the processes through which banks allocate credit, price risk, and monitor portfolios? Second, what do these AI processes imply for prudential supervision and bank governance? Recent research and practice increasingly view AI as central to banks' operating models, enabling faster information processing, scalable decision-making, and improved adaptability in a competitive and volatile environment (Vives, 2019; FSB, 2024; ESRB, 2025). At the same time, AI integration can have a bidirectional impact on banking stability and, consequently, economic outcomes. On one side, AI can strengthen resilience through better screening, monitoring, and operational efficiency (Gambacorta et al., 2019; Berg et al., 2020; Gyau et al., 2024; Li, 2025). On the other side, AI can also have destabilizing effects by increasing model risk, opacity, bias, and cyber vulnerabilities (Danielsson et al., 2022; Jia, 2024; Elekdag et al., 2025), with potential spillovers from the financial institutions to the system. A rigorous assessment is therefore needed at both the micro and macro levels to assess how AI affects banks' risk-taking and how its implementation can guide regulatory frameworks intended to reduce the materialization of risks and minimize the adverse impacts on financial stability.

<sup>1</sup> Artificial intelligence (AI) refers to computational systems that can execute tasks related to human intelligence by learning from data and generating outputs such as predictions, recommendations, decisions, or content (Yildiz & Cil, 2025).

We contribute to the debate on whether banks' investments in AI-related technologies strengthen efficiency gains and risk controls or instead amplify risk in four ways. First, we construct a bank-level measure of AI intensity that captures persistent shifts in banks' technological adoption by leveraging annual-report disclosures across three dimensions of AI-related technologies in banking: *agent adoption*, *cloud computing*, and *big data*, instead of relying on narrow or episodic digitalization proxies. Second, using within-bank variation over time, we document that increases in AI intensity are followed by declines in banks' distance-to-default. Third, we move beyond the reduced-form association by identifying where stability weakens most: Z-score decompositions show that the decline operates mainly through higher earnings volatility, and, to a lesser extent, lower profitability, with no detectable effect on capitalization. Finally, we find that the financial fragility effects of AI adoption are not uniform, being concentrated among banks with weaker liquidity buffers, greater reliance on unstable non-deposit funding, and weaker loss-absorption capacity. This finding is consistent with a risk-amplification mechanism in which AI-related technologies interact with runnable funding and limited resilience buffers to increase financial fragility.

Our analysis focuses on bank-level data capturing three dimensions of AI-related technological intensity— *agent adoption*, *cloud computing*, and *big data*—which form three foundational and complementary pillars of modern AI in banking. Banks deploy agent adoption to generate predictions and recommendations that shape credit screening and underwriting, risk-based pricing, limit setting, and continuous portfolio monitoring, applications that the empirical lending literature links to improved risk classification when richer signals are used (Berg et al., 2020). They rely on cloud computing to scale these tools (compute-on-demand, faster model deployment, integration across business lines), but this also increases operational interdependence through third-party outsourcing and potential provider concentration—an issue highlighted in policy-facing and academic work on cloud use in finance (FSB, 2019; 2023). Big-data architectures are used by banks to aggregate and process granular information, continuously feed models, and support early-warning and automated oversight. Research on big data in finance emphasizes that advances in data and processing capacity materially change information production and its economic consequences (Begenau et al., 2018; Vives, 2019). We measure banks' engagement with these technologies using disclosure-based indicators from annual reports following the China Stock Market & Accounting Research (CSMAR) Database's text-

classification framework. For each bank-year, we compute log-transformed frequencies of technology-specific keywords, standardize them, and aggregate them into a baseline composite AI-intensity index that averages the three dimensions. We complement this baseline measure with alternative index constructions, including total-intensity and principal-component indices, as well as component-specific measures, designed to capture within-bank changes in technology orientation while mitigating measurement error and dependence on any single construction. The time-series patterns show a pronounced and persistent increase in AI-related disclosures, with particularly strong growth in cloud computing and big data, consistent with a gradual build-out of AI-enabling capabilities.

We construct a panel dataset by merging bank-level financial statement data from the CSMAR China Bank Research Database and the CSMAR China Financial Technology Research Database with province-level indicators from the CSMAR China Regional Economy Research Database and the National Bureau of Statistics of China (NBSC). This setting is uniquely suited to our research question. First, CSMAR provides consistent coverage of banks' balance-sheet structure, profitability, capitalization, liquidity, and funding composition, enabling uniform measurement of distance to default (Z-score) and its components across institutions and over time. Second, and importantly, CSMAR offers structured, bank-level AI-disclosure indicators derived from official annual reports within a standardized reporting framework. While AI adoption is a global phenomenon, systematic bank-level AI data remains fragmented. Comparable bank-level AI-intensity measures are not available in official U.S. or European databases, where existing studies typically rely on researcher-constructed proxies—such as NLP-based analyses of narrative disclosures or measures derived from AI-related job vacancy postings.<sup>2</sup> The Chinese setting, therefore, enables systematic and replicable measurement of AI intensity using externally compiled disclosure data. Finally, province-level data capture time-varying local economic and digital conditions—such as innovation capacity and digital infrastructure—that shape technology adoption. Because Chinese banks' operations and technology investments are closely tied to

<sup>2</sup> For example, studies of U.S. and European banks typically construct AI measures using NLP-based analyses of narrative disclosures (Bellstam et al., 2020; Ante & Saggi, 2025; Rauf, 2025), LLM-based survey simulation (Kazinnik, 2025), or vacancy-based proxies (Babina et al., 2024). The Financial Stability Board (2024) notes that the absence of standardized AI reporting metrics in major Western economies remains a key obstacle to cross-country empirical analysis.

province-level markets, these regional conditions are economically salient for banks' AI deployment, strengthening identification and interpretation.

After imposing the rolling-window requirements needed to construct stability measures (Z-scores), the baseline sample contains 227 bank-year observations and focuses on institutions that are both of local and global relevance. Our sample includes five Financial Stability Board (FSB)–designated global systemically important banks (G-SIBs) and spans major joint-stock as well as leading city and rural commercial banks that together account for roughly 61% of Chinese banking-sector assets at end-2024. This coverage is particularly relevant because these banks dominate credit allocation in the world's largest banking system, so shifts in AI-related technological intensity reflect economically meaningful changes with direct implications for prudential policy. Finally, the sample period captures the period 2013-2024, including pronounced post-2017 acceleration and widening dispersion in AI-related activity while preserving substantial within-bank variation, allowing us to assess how a given bank's increases in AI-related technologies translate into subsequent changes in stability rather than relying on static cross-sectional differences.

We evaluate the relation between AI intensity and bank stability in a fixed-effects panel framework that absorbs unobserved heterogeneity across banks and common shocks over time. Specifically, we regress next-year distance-to-default (Z-score) on lagged AI intensity and a standard set of lagged bank controls capturing size, capitalization, asset composition, credit risk, liquidity, funding structure, and operating efficiency, with bank and year fixed effects and standard errors clustered at the bank level. This design identifies the coefficient on AI intensity based on within-bank changes over time, while the lag structure helps mitigate contemporaneous reverse causality and allows for delayed effects of technology adoption on risk. In our baseline results, the coefficient on lagged AI intensity is negative and robust across specifications and control sets. Economically, a one–standard deviation increase in the AI-intensity index is associated with a 0.25 decline in next-year log Z-score. This corresponds to approximately a 22% reduction in Z-score levels (about 45 units relative to the sample mean), or roughly 0.24 standard deviations of the Z-score distribution. The magnitude is economically meaningful and consistent with the risk-amplification view in the literature. Scaling AI-related technologies can accelerate balance-sheet adjustment and credit reallocation, increase model and operational complexity, and induce more correlated behavior through shared data, tools, and vendors—channels that raise earnings volatility

and reduce distance-to-default even if average efficiency improves (Vives, 2019; Becerra-Vicario et al., 2024; FSB, 2024; Danielsson & Uthemann, 2025).

To clarify which AI-related technologies drive this pattern, we decompose the composite AI-intensity index into its three core components— *agent adoption*, *cloud computing*, and *big data*—and re-estimate the same fixed-effects specification. Each dimension is negatively related to banks’ distance-to-default, with economically significant magnitudes, ranging from 0.11 to 0.20 units decline relative to the sample mean Z-score (approximately 14 to 28 percent of its standard deviation). This indicates that the baseline effect does not hinge on a single disclosure category. Instead, the pattern is broad-based across the complementary technologies that jointly support advanced analytics and scalable digital deployment.<sup>3</sup>

We then ask whether the stability effects we document operate through risk/volatility or through buffers. To do so, we disaggregate the Z-score into its standard components—earnings volatility, profitability, and capitalization—by examining (i) ROA volatility over the rolling window, (ii) the level of ROA, and (iii) equity-to-assets, while maintaining the same lag structure, control set, and clustered inference as in the baseline specifications to isolate the within-bank mapping from AI intensity to each channel. The evidence points clearly to volatility and profitability as the primary mechanisms. Increases in AI intensity are significantly associated with higher subsequent ROA volatility, lower ROA, and no systematic adjustment in equity-to-asset ratios. This pattern suggests that the baseline decline in Z-scores reflects an increase in earnings risk and weaker profitability rather than an erosion of capital buffers, consistent with the view that AI-related technologies first affect the risk and operational complexity of intermediation while regulatory capital adjusts more slowly (Buchak et al., 2018; Danielsson & Uthemann, 2025; Elekdag et al., 2025). From a policy perspective, the decomposition implies that supervisors should focus on AI-related risk governance and operational resilience, such as model validation and risk monitoring, because the effects of AI intensity materialize first through the volatility component.

<sup>3</sup> While we also consider investment-based proxies—AI investment intensity (AI-related assets scaled by total assets) and the AI investment stock (the level of AI-related assets)—our primary analysis emphasizes the three disclosure-based AI-related technology dimensions (artificial intelligence, cloud computing, and big data). These measures align more closely with the operational “technology stack” through which banks deploy AI, capture adoption aspects that may not be fully capitalized on the balance sheet, and provide richer within-bank variation over time.

We subject the baseline AI–risk-taking results to a broad set of robustness exercises. First, we verify that the findings are not an artifact of how earnings risk is assessed by recomputing the Z-score using alternative rolling windows for ROA volatility (three-, five-, and eight-year horizons) and by excluding extreme observations. In all cases, higher AI intensity continues to predict lower subsequent bank stability. Second, we show that the inference does not depend on a particular way of constructing AI intensity by replacing the baseline AI-intensity index with alternative measures that emphasize disclosure breadth or depth, including total-count indices, principal-component-based indices, and component-specific measures. The estimated effect remains negative and statistically significant throughout alternative empirical models. Third, we address concerns about omitted time-varying confounders by augmenting the baseline specification with additional controls capturing asset-risk intensity, balance-sheet expansion, ownership and governance, market structure and competition, and provincial macroeconomic conditions; the AI coefficient remains tightly clustered around the baseline magnitude. Threshold designs that focus on transitions into “high-AI intensity” states also validate our main findings.

Because AI-related activities are not directly observable and disclosure-based measures may capture communication choices as well as underlying adoption, we further examine whether our inferences depend on how AI adoption is proxied. We re-estimate the baseline bank and year fixed-effects specification, including the full set of lagged control variables, but replace the primary AI-intensity index with alternative proxies that capture complementary dimensions of AI-related activity. Specifically, we use an AI investment intensity measure, defined as AI-related assets scaled by total assets, which reflects the relative balance-sheet importance of AI and is less mechanically related to bank size, and an AI investment stock measure, which captures the overall scale of AI adoption. Both proxies yield negative and statistically significant coefficients, consistent with the baseline interpretation. By contrast, a broader disclosure-based proxy for digital transformation is negative but statistically insignificant, suggesting that generic digitalization disclosures are a noisier and less targeted measure of AI-related adoption than investment-based proxies.

Another concern is that AI intensity is not randomly assigned. Banks may scale AI-related technologies in anticipation of changes in risk, profitability, or supervisory attention, and disclosure-based measures can contain non-classical noise. To address these issues, we complement the bank and year fixed-effects design with an instrumental-variables approach that

isolates variation in AI intensity driven by diffusion and local implementation capacity. Specifically, we instrument lagged bank AI intensity with (i) peer adoption in the same province-year, constructed as a leave-one-out average to proxy for local benchmarking and competitive pressure, and (ii) lagged province-level shifters that capture the depth of the local AI ecosystem and innovative capacity, including provincial growth in the number of AI firms (vendor and skilled-labor thickness), growth in new-product development expenditure scaled by GDP (innovation intensity), and growth in industrial R&D personnel scaled by population (technical capacity and absorptive capability). Because these local measures could also reflect broader development, we tighten identification by progressively conditioning on lagged provincial fundamentals and general ICT infrastructure to separate AI-specific ecosystem variation from overall digital readiness. The resulting 2SLS estimates reinforce the baseline pattern, supporting the conclusion that our main findings are not explained exclusively by reverse causality or omitted local trends.

Finally, we analyze the conditions under which investments in AI-related technologies are likely to impact bank stability by investigating moderating factors that reflect variations in banks' capacity to absorb technology-induced shocks. AI-enabled re-optimization can enhance the speed and scale of intermediation decisions and increase earnings volatility (Fuster et al., 2019; Vives, 2019; Berg et al., 2020; Berg et al., 2022). However, the extent to which this leads to fragility depends on the strength of banks' liquidity, funding, and loss-absorbing buffers. Accordingly, we build on theoretical and empirical research on bank fragility by interacting our AI-intensity index with measures of funding structure, liquidity conditions, and loss-absorbing capacity. The results demonstrate systematic heterogeneity consistent with this framework. The negative association between AI intensity and subsequent stability is more pronounced for banks that rely more on interbank and other wholesale funding sources, while it is mitigated for banks with stronger liquidity positions, higher deposit growth, and greater provisioning coverage. These findings suggest that AI-related technologies are most likely to amplify fragility when liabilities are less stable and buffers are limited, and they indicate that supervisory assessments of AI should be conducted in conjunction with evaluations of banks' liquidity, funding, and provisioning profiles, instead of treating AI adoption as a homogeneous source of risk.

Our findings have direct implications for supervision and bank governance of AI-related technologies. For regulators, our results suggest that AI adoption should be treated as a prudentially

relevant change. In practice, this implies requiring banks to inventory material AI use cases (e.g., underwriting, pricing, and monitoring), enforcing core model-governance standards, and incorporating a set of AI-linked metrics into routine supervisory frameworks and stress testing. Supervisory intensity should be higher when liquidity buffers and funding stability are weaker, including operational-resilience plans for cloud and vendor disruptions. For bank managers, our findings suggest scaling AI alongside financial resilience, i.e., strengthening model-risk frameworks, imposing portfolio “speed limits” and concentration caps on AI-driven credit expansion, and calibrating AI investments with stable funding and liquidity targets so efficiency gains do not translate into excessive volatility and lower stability.

Our analysis is also directly related to ongoing policy discussions about how prudential frameworks should adapt to the rapid diffusion of AI in the banking sector. Recent work by central banks and international regulators stresses that fast AI adoption by banks can complicate supervisory monitoring and amplify vulnerabilities through model risk and governance challenges, third-party dependencies, and correlated behavior (FSB, 2024). In Europe, recent analysis similarly highlights the potential for systemic risk amplification and calls for supervisory approaches that explicitly account for these channels (ESRB, 2025). Our evidence contributes to this debate by providing a scalable bank-level measure of AI-related technological intensity and by showing that the stability implications emerge primarily through earnings volatility and are dependent on banks’ loss-absorbency buffers. The results also complement emerging supervisory efforts to strengthen oversight of third-party and concentration risks in cloud-based ecosystems (BCBS, 2024).

The following sections outline the structure of the paper. Section 2 discusses the relevant literature and proposes the hypotheses. Section 3 introduces the data, defines the variables, and sets out the empirical framework and identification strategy. Section 4 provides the core empirical results, starting with the baseline AI–stability relation, then examining the channels through Z-score decomposition, and reporting robustness and endogeneity analyses. Section 5 provides additional evidence on AI-intensity thresholds and on state dependence via liquidity, funding structure, and loss-absorbing capacity. Section 6 summarizes the paper and concludes.

## 2. Theoretical background and hypotheses development

Our study relates to several strands of literature. First, we build on a growing body of research on banking technology literature documenting how IT expenditures and digital delivery channels reshape banks' production processes, cost structures, and performance (Beccalli, 2007; DeYoung et al., 2007; Hernando & Nieto, 2007). Second, our analysis relates to the digital-disruption perspective, which emphasizes that technological change can alter competitive dynamics and business models in banking, with direct implications for pricing, scale, and risk-taking incentives (Vives, 2019). Third, we extend the emerging academic and policy literature on AI in banking, which examines the greater incorporation of AI into core intermediation and control function, such as credit screening and underwriting, portfolio monitoring and early-warning systems, pricing and limit-setting, and compliance and fraud detection, and assesses the potential consequences for financial stability as reliance on complex and adaptive systems increases (FSB, 2024; ESRB, 2025; Yildiz & Cil, 2025). A closely related operational development is automation through Robotic Process Automation (RPA). While RPA streamlines high-volume, rule-based tasks, recent implementations increasingly incorporate AI/ML components that extend automation toward learning-based decision support (Daoud & Anaya, 2025).

These strands highlight several central mechanisms through which AI-related technologies can affect bank risk-taking and stability. On the one hand, AI may generate efficiency gains and stronger controls by improving screening and monitoring, reducing expected and unexpected credit losses, and enhancing the ability of financial institutions to handle data and distinguish between different types of borrowers (Vives, 2019). Consistent with this view, machine-learning approaches in consumer lending demonstrate that richer signals can improve default prediction and reduce default rates through better screening outcomes (Berg et al., 2020). Related survey evidence on fintech lending highlights how technology-based underwriting reshapes loan processes and risk assessment, while monitoring and early-warning tools enable timelier interventions that limit loss spikes (Berg et al., 2022). These mechanisms imply smoother provisioning and fewer unexpected losses, reinforced by AI-enabled compliance and fraud detection that can reduce operational-loss shocks (FSB, 2024) and by more granular risk-based pricing and portfolio management that better align margins with risk (Vives, 2019). In the same line, Rauf (2026) shows that U.S. banks adopting AI expand mortgage lending toward riskier

borrowers and charge higher rates, yet do not experience higher non-performing loans, suggesting that AI enhances screening precision within observably risky segments.

On the other hand, AI may amplify risk through speed, cyclical, and complexity. More generally, prior research shows that changes in banks' activities and business models can alter the risk-return profile and increase profit volatility even when average performance improves (Stiroh, 2004). Evidence from U.S. financial holding companies similarly documents that diversification can have adverse effects by raising volatility and lowering risk-adjusted performance (Stiroh & Rumble, 2006). Within this framework, faster AI-enabled decision cycles can facilitate rapid credit scaling and repricing of exposures, increasing procyclicality and sensitivity to shocks (Becerra-Vicario et al., 2024; FSB, 2024). Technology-driven competition can compress spreads, making profitability more sensitive to credit-quality and funding shocks (Vives, 2019). From a systemic perspective, reliance on common models and shared data sources can synchronize banks' responses and amplify correlated behavior in stress episodes (FSB, 2024), while fintech-driven credit intermediation can reallocate risk across institutions and market segments and may facilitate regulatory arbitrage (Buchak et al., 2018). Recent systemic-risk frameworks further argue that AI can raise tail risk through speed advantages, common information sources, and strategic complementarities (Danielsson & Uthemann, 2025). Supervisory assessments emphasize additional vulnerabilities operating through opacity, governance challenges, third-party dependencies (including shared cloud and AI providers), cyber and operational risks, and provider concentration, all of which can generate clustered losses in periods of stress (ESRB, 2025; FSB, 2025). Governance frictions may intensify these effects when decision-makers over-rely on automated recommendations and delay human intervention as models deteriorate (FSB, 2025). Finally, AI adoption may require restructuring and "catch-up" investment expenses, which increase short-term disruption to profitability, particularly during the deployment and learning phases (Babina et al., 2024), while governance and risk frameworks highlight how incentives and controls shape banks' risk profiles and buffer choices (Laeven & Levine, 2009).

Overall, the literature yields competing predictions for the net effect of AI intensity on bank risk-taking. If AI-related technologies primarily strengthen information production, screening, and monitoring, they should reduce unexpected losses and smooth performance, thereby increasing distance-to-insolvency (Berg et al., 2020; FSB, 2024). If instead AI mainly increases the speed and scale of balance-sheet adjustment and heightens operational and model complexity,

potentially synchronizing behavior through common data, models, or vendors, then it may amplify cyclical and tail risk and reduce distance-to-insolvency (FSB, 2024; Danielsson & Uthemann, 2025; ESRB, 2025). We summarize these competing views in two testable hypotheses:

*H1a (Stability-enhancing): Higher AI intensity increases bank stability (higher Z-score) by improving screening, monitoring, and operational controls, which lowers unexpected losses and stabilizes profitability.*

*H1b (Risk-amplifying): Higher AI intensity decreases bank stability (lower Z-score) by accelerating risk-taking, increasing operational and model complexity, and fostering correlated behavior, which raises earnings volatility and tail risk.*

### **3. Data and empirical methodology**

#### **3.1. Sample and data sources**

We construct an unbalanced panel of 35 publicly listed Chinese commercial banks over 2013–2024 by combining bank financial statement data from CSMAR with annual-report disclosures used to measure AI-related technological intensity.<sup>4</sup> Due to the rapid rise of AI technologies in recent years, comprehensive data for unlisted banks is limited. This setting is a great laboratory to investigate the relation between AI and bank risk-taking because it features economically and systemically important intermediaries operating in a large, rapidly digitizing financial system. The sample includes five Financial Stability Board–designated global systemically important banks (G-SIBs), together with major joint-stock banks and leading city and rural commercial banks spanning multiple provinces and business models. Collectively, these institutions account for roughly 61% of Chinese banking-sector assets at end-2024—approximately CNY 271.2 trillion (about USD 37.2 trillion).<sup>5</sup> For comparison, EU-headquartered credit institutions held €33.12 trillion in total assets at end-September 2024 (≈USD 37.1 trillion using the EUR/USD reference rate of 1.1196), while FDIC-insured U.S. banks reported USD 24.1

<sup>4</sup> Bank-level data are drawn from the CSMAR China Bank Research Database and CSMAR China Financial Technology Research Database. Province-level controls are obtained from the CSMAR China Regional Economy Research Database and the National Bureau of Statistics (China).

<sup>5</sup> National Financial Regulatory Administration (China), sector statistics of the banking: <https://www.nfra.gov.cn/>

trillion in total assets in Q4 2024.<sup>6, 7</sup> Therefore, variation in AI-related technology intensity reflects changes among banks that are central to credit allocation and systemic risk. Moreover, China is one of the global leaders in AI patents, making its banking market a high-salience setting for studying how AI-related technologies map into bank risk (WIPO, 2019; Maslej et al., 2025).

Our focus on Chinese banks is also driven by data availability. The CSMAR database uniquely reports bank-level data on AI-related technologies in a standardized format that is not available for banks in most other jurisdictions. In Europe and the United States, comparable AI-intensity measures are not reported in regulatory datasets and are typically inferred from researcher-constructed proxies, such as textual analyses of disclosures (Bellstam et al., 2020; Ante & Saggi, 2025; Rauf, 2025) or AI-related hiring activity (Babina et al., 2024). The availability of standardized indicators in CSMAR enables consistent and replicable measurement of AI intensity across banks and over time. To ensure data integrity, we restrict our sample to 35 listed banks with sufficiently complete annual reports and financial-statement information over the 2013–2024 period. Imposing the time-series coverage required to construct rolling distance-to-default measures for each bank yields a baseline sample of 241 bank-year observations.

This composition of our sample is relevant for both prudential supervision and bank management. From a supervisory perspective, these banks play a crucial role in transmitting credit and macro-financial conditions, which makes them an ideal choice for assessing whether increased use of AI-related technologies is linked to changes in financial resilience. From a managerial perspective, they are the institutions facing the most significant trade-offs in scaling AI, as they are among the most active in global technology investment and disclosure. They also provide rich narrative disclosures in their annual reports, allowing for the construction of consistent, comparable text-based measures of AI-related adoption over time.

<sup>6</sup> European Central Bank (ECB) Press Release: <https://www.ecb.europa.eu/press/pr/date/2025/html/ecb.pr250207~cb9caa3836.ro.html>

<sup>7</sup> Federal Deposit Insurance Corporation (FDIC): <https://www.fdic.gov/quarterly-banking-profile/quarterly-banking-profile-fourth-quarter-2024>

### 3.2. Measuring Bank Stability

We measure bank stability using the accounting-based Z-score, a standard proxy for banks' distance-to-default that combines profitability, capitalization, and earnings volatility into a single solvency indicator. The Z-score is a commonly used measure in the banking literature and is closely aligned with supervisory and policy definitions of financial stability (Boyd & Runkle, 1993; Demirgüç-Kunt & Huizinga, 2010). Following Laeven and Levine (2009), higher Z-scores indicate a diminished risk of insolvency and stronger bank stability.

Formally, the Z-score for bank  $i$  in year  $t$  is constructed as:

$$Z - score_{it} = \frac{ROA_{it} + Equity/Assets_{it}}{\sigma(ROA_i)},$$

where ROA represents return on assets, Equity/Assets captures the capital buffer, and  $\sigma(ROA_i)$  is the standard deviation of ROA. This measure corresponds closely to the World Bank's definition of bank Z-score as the value of standard deviations by which returns would need to decrease in order to deplete equity, and is widely used in cross-country bank stability assessments (World Bank Global Financial Development Database).

Consistent with prior studies, we employ the natural logarithm of the Z-score to reduce skewness and minimize the influence of extreme observations (Laeven & Levine, 2009; Houston et al., 2010; Beck et al., 2013). Our main dependent variable is therefore the log Z-score.

*Baseline Z-score construction.* In the main specification, earnings volatility is computed using a four-year rolling window covering years  $t-4$  to  $t-1$ . Importantly, the current year is excluded from the volatility calculation. This choice is economically and econometrically motivated. Using a window that includes year  $t$  (e.g.,  $t-3$  to  $t$ ) can mechanically amplify the impact of contemporaneous shocks by affecting both the numerator (current ROA) and the denominator (earnings volatility). By contrast, the  $t-4$  to  $t-1$  window yields an ex-ante measure of risk, capturing how risky the bank's current profitability and capital position are given past earnings volatility. This approach addresses simultaneity and look-ahead concerns while aligning with best practices in the empirical banking literature (e.g., Houston et al., 2010; Hafeez et al., 2022). Summary statistics for the baseline log Z-score are reported in **Table 1**. The mean value of the log Z-score is 4.96, with a standard deviation of 0.77, and ranges from 3.58 to 7.33 across our sample banks. Comprehensive definitions of each variable are included in **Appendix A**.

*Alternative Z-score measures.* To evaluate robustness, we develop alternative Z-scores using three-year, five-year, and eight-year rolling windows for ROA volatility, again excluding the current year. These alternative windows allow for different assumptions regarding the persistence of earnings risk. The resulting log Z-score measures exhibit similar distributions and substantial overlap in their ranges, as reported in **Table 1**. Mean values range from approximately 4.63 to 5.16, with comparable dispersion across specifications.

Across all window choices, the empirical results remain qualitatively consistent, indicating that our outcomes are not reliant on a particular volatility horizon. In further analyses, we also investigate the individual components of the Z-score—earnings volatility, profitability, and capitalization—to identify the channels through which AI-related technologies affect overall bank stability. Definitions and construction details for all stability measures are summarized in **Appendix A**.

### **3.3. Measuring AI Intensity**

AI intensity is measured at the bank level, using text-based indicators derived from annual-report disclosures, following the CSMAR text classification framework. We construct four AI-intensity indices: a standardized mean index, a standardized total index, a principal-component-based index, and a raw word-frequency count index. Using word-frequency measures from annual-report disclosures as proxies for AI adoption follows established practice in the text-based finance literature (e.g., Loughran & McDonald, 2014; Loughran & McDonald, 2015; Hoberg & Phillips, 2016; Li, 2025; Rauf, 2026). This approach is designed to capture time-varying shifts in banks' adoption of AI-related technologies, while mitigating concerns that any single disclosure item or aggregation method drives the results. Text-based measures are particularly suitable in this setting because they reflect banks' strategic orientation and organizational emphasis on AI adoption, as opposed to isolated investment events.

For each bank-year, we compute the frequency of AI-related keywords disclosed in annual reports. Keywords are classified into three categories—agent adoption, cloud computing, and big data—based on the CSMAR fintech dictionary. Definitions and construction details are provided in **Appendix 1**, and **Appendix 2** reports the full keyword dictionary. Given the strong right-skewness of raw word counts, we log-transform all word-frequency measures and standardize

them to have a mean of zero and unit variance.<sup>8</sup> Let  $z_{\ln\_wordfreq\_k}$  denote the standardized, log-transformed frequency for technology  $k \in \{\text{Agent, Cloud, Data}\}$ . Our baseline AI-intensity index is then defined as the unweighted average of the three standardized components:

$$AI\text{-intensity index}_{i,t} = \frac{1}{3} (z_{\ln\_wordfreq\_Agent,i,t} + z_{\ln\_wordfreq\_Cloud,i,t} + z_{\ln\_wordfreq\_Data,i,t}).$$

This standardized mean-based index, centered near zero, is used as our primary measure in the empirical analysis.

To ensure that our findings are not sensitive to a particular aggregation choice, we complement the baseline AI-intensity index with three alternative specifications: (i) a standardized total-count index that sums the standardized components, (ii) a principal component-based index extracted from the three standardized measures, and (iii) a raw-count index defined as the mean of the untransformed disclosure frequencies for agent adoption, cloud computing, and big data.<sup>9</sup> These alternative measures capture different aspects of AI-related technologies—depth, co-movement, and scale—and are employed in robustness analyses.<sup>10</sup> As shown in **Table 1**, all indices exhibit substantial cross-sectional and time-series variation.

In **Graph 1**, we plot the time-series evolution of the standardized components together with the composite AI-intensity index. We observe a pronounced and persistent upward trend in the overall AI-intensity index, consistent with a steady intensification of banks’ adoption of AI-related technologies. The increase is driven primarily by earlier and stronger growth in cloud computing and big data relative to agent adoption, suggesting that scalable infrastructure and data capabilities expanded ahead of more explicit AI applications. **Graph 2** decomposes AI-related measures by technology category and shows that cloud computing and big data account for an

<sup>8</sup> We apply a log transformation,  $\ln(x + 1)$ , to reduce skewness and limit the influence of extreme observations. This approach follows standard practice in empirical finance and econometrics for handling skewed distributions and improving statistical properties of the estimators (Wooldridge, 2010).

<sup>9</sup> **Table A2** provides a detailed description of the keyword dictionary and formulas used to construct the word-frequency measures.

<sup>10</sup> We interpret cloud computing and big data as AI-enabling technologies that are necessary inputs to AI adoption at scale.

increasing share of total AI-related discussion over time, consistent with their role as foundational enablers.<sup>11</sup>

**Figure 3** compares the distribution of the AI-intensity index before and after 2017. The post-2017 distribution shifts markedly to the right and displays a thicker upper tail, indicating both higher average AI intensity and greater dispersion across banks. While AI-related activity was relatively limited and concentrated before 2017, the later period is characterized by substantial heterogeneity, with a subset of banks emerging as clear technology leaders. These dynamics are consistent with a diffusion process in which adoption accelerates over time but remains uneven across institutions.

### 3.4. Empirical Methodology

#### 3.4.1. Identification

To assess whether our empirical strategy exploits meaningful variation in the data, in **Table 2** we report a within–between decomposition of the main variables. The dependent variable, bank stability measured by the logarithm of the Z-score, exhibits substantial time-series variation, with approximately 45% of its total variation occurring within banks over time. Importantly, AI intensity displays even stronger within-bank dynamics, as nearly 61% of the variation in the composite AI-intensity index is driven by within-bank changes. A similar pattern emerges for the individual AI-related word-frequency measures, for which between 50% and 70% of total variation is within banks, indicating that banks actively adjust their AI-related disclosures over time. These pronounced time-series patterns motivate the inclusion of year fixed effects to absorb common macroeconomic conditions, regulatory changes, and economy-wide technological trends that simultaneously affect banks’ stability and digitalization decisions. Overall, these patterns justify the use of bank fixed effects, ensuring that identification relies on within-bank changes in AI intensity and year fixed effects, which capture common macroeconomic, regulatory, and technological shocks affecting banks’ stability and digitalization decisions.

<sup>11</sup> *The patterns reported in are robust to using raw (non-logged) word-frequency counts instead of log-transformed measures.*

In **Table 3**, we report Pearson pairwise correlations for the main variables used in the baseline regressions. All variables are measured at the bank–year level and are defined in **Appendix A**. The correlations are generally moderate, indicating that multicollinearity is unlikely to materially affect the multivariate analysis. Bank stability, proxied by the logarithm of the Z-score, is positively related to capital adequacy and cost efficiency and negatively associated with nonperforming loans, consistent with standard banking theory. AI intensity is weakly correlated with stability in the univariate correlations, reinforcing the need for specifications that condition on bank fundamentals and exploit within-bank variation over time. Among the control variables, bank size, liquidity, and deposit reliance are positively related, consistent with scale and funding-structure patterns commonly observed in the banking sector.<sup>12</sup>

### 3.4.2. Baseline Specification

To investigate the relationship between AI intensity and bank stability, we estimate bank-year panel regressions of the following form:

$$\ln(Z - score_{i,t}) = \beta \times AI-intensity\ index_{i,t-1} + \mathbf{X}'_{i,t-1} \times \gamma + \mu_i + \lambda_t + \varepsilon_{i,t},$$

where  $\ln(Z - score_{i,t})$  denotes the logarithm of the bank Z-score for bank  $i$  in year  $t$ ,  $AI-intensity\ index_{i,t-1}$  is the lagged measure of AI intensity,  $\mathbf{X}'_{i,t-1}$  is a vector of lagged bank-level control variables,  $\mu_i$  and  $\lambda_t$  denote bank and year fixed effects. We cluster standard errors at the bank level.

The coefficient  $\beta$  is identified from within-bank changes in the AI-intensity index over time. All time-invariant institutional heterogeneity is absorbed by bank fixed effects (e.g., persistent differences in business models, ownership, governance, and regional characteristics), while year fixed effects capture common macroeconomic conditions, regulatory changes, and economy-wide technology trends. Lagging AI intensity mitigates concerns related to reverse causality, whereby contemporaneous changes in bank stability could influence disclosure practices or technology-related investments. The lag structure is also consistent with the view that AI-related technology adoption influences banks' risk profiles with a delay. Importantly, as documented in

<sup>12</sup> To further address multicollinearity, we compute variance inflation factors (VIFs) from a pooled OLS similar to the baseline model estimated on the same sample and lagged covariates. The mean VIF is 1.47 and the largest VIF is 2.13 (bank size); all remaining VIFs lie between 1.21 and 1.63. These magnitudes are below conventional cutoffs (e.g., 5 or 10), suggesting that multicollinearity is unlikely to meaningfully affect the inference of our empirical framework.

the within-between variance decomposition, AI intensity exhibits substantial within-bank variation over time, ensuring that the inclusion of both bank and year fixed effects does not mechanically absorb the AI signal.

We also include a standard set of lagged bank-level controls capturing size, capitalization, asset composition, credit risk, liquidity, funding structure, and operational efficiency, which are core determinants of bank stability in banking theory. Bank size reflects diversification gains and potential risk-taking incentives associated with market power and implicit guarantees (Demirgüç-Kunt & Huizinga, 2013; DeYoung, 2014). Capital adequacy strengthens loss-absorbing capacity and directly enhances resilience to shocks (Beltratti & Stulz, 2012; Berger et al., 2016). Asset composition, proxied by the loan share, captures exposure to credit risk inherent in traditional intermediation (Acharya, 2009; Buchak et al., 2018). Credit quality, measured by nonperforming loans, reflects realized borrower risk and is a key driver of bank fragility (Jiang et al., 2016). Liquidity buffers mitigate funding stress and rollover risk (Krishnamurthy et al., 2014; Jondeau et al., 2020), while deposit-based funding is generally more stable than wholesale financing (Egan et al., 2022). Finally, the cost-to-income ratio captures operational efficiency, which is closely linked to profitability and downside risk (Fiordelisi et al., 2011; Demirgüç-Kunt et al., 2020).<sup>13</sup>

Although we do not claim a fully causal interpretation, this framework helps limit reverse causality and omitted-variable bias, and the estimated coefficients can be interpreted as systematic within-bank associations between changes in AI intensity and subsequent stability. We further assess robustness by considering alternative AI measures and stability metrics, expanding the control set, and implementing additional designs, including endogeneity treatment.

<sup>13</sup> The role of bank size, capitalization, asset quality, liquidity, and efficiency in shaping risk and stability is well documented for Chinese banks (Allen et al., 2005; Berger et al., 2009). Capital adequacy, credit risk, and loan composition are central determinants of fragility and performance in China's bank-dominated system (Bailey et al., 2011). Funding structure, liquidity regulation, and operational efficiency further condition risk-taking and resilience in Chinese commercial banks (Ma & Li, 2020).

## 4. Main Results

### 4.1. Baseline Relationship Between AI Intensity and Bank Stability

In this section, we present our main results. **Table 4** documents a robust negative association between AI intensity and bank stability, consistent with H1b. In column (1), where we relate banks' distance to default (Z-score) to lagged AI intensity and year effects, the coefficient on lagged AI intensity is negative and statistically significant ( $\beta = -0.264$ ,  $t = -3.70$ ), indicating that banks increasing AI-related disclosure intensity subsequently exhibit lower Z-scores. This result is not sensitive to conditioning information. After adding the full set of lagged bank controls in column (2), our baseline specification, the estimate remains economically similar ( $\beta = -0.253$ ,  $t = -4.08$ ) and the within- $R^2$  rises from 0.30 to 0.40, consistent with controls capturing meaningful time-varying determinants of stability without attenuating the AI signal. Columns (3)–(4) further confirm that the negative relation is not an artifact of the fixed-effects transformation. Under random effects, lagged AI intensity remains negative and significant ( $\beta = -0.189$ ,  $t = -3.13$ ), and adding province fixed effects leaves the inference unchanged ( $\beta = -0.221$ ,  $t = -3.65$ ), suggesting that the pattern is not driven by time-invariant regional differences.<sup>14</sup>

The magnitude is economically meaningful. Because the AI-intensity index is standardized, the baseline estimate in column (2) implies that a one-standard-deviation increase in AI-intensity index is followed by a 0.25 decline in next-year log Z-score, corresponding to a 22% lower Z-score ( $e^{-0.25} - 1$ ) and approximately 0.24 standard deviations change in Z-score. This pattern is consistent with a risk-amplification mechanism in which rapid scaling of AI-related technologies coincides with higher risk-taking and greater earnings volatility. One interpretation is that AI-enabled decision speed facilitates faster balance-sheet adjustment and expansion, while increased model and operational complexity raises earnings risk. In parallel, technology-driven competition may compress margins, amplifying the sensitivity of profitability to shocks—channels that map naturally into the Z-score's dependence on profitability and earnings volatility (Boyd & Runkle, 1993; Acharya, 2009; Laeven & Levine, 2009; Buchak et al., 2018).

<sup>14</sup> As a robustness check, we winsorize all variables at the 1<sup>st</sup> and 99<sup>th</sup> percentiles. Results reported in **Appendix 3** are qualitatively unchanged. We therefore report specifications using the raw (non-winsorized) variables in the main tables for transparency.

The control variables generally enter with signs consistent with prior evidence. Tier 1 capital is positively across specifications, consistent with capital’s loss-absorbing role in strengthening resilience and lowering default risk (Beltratti & Stulz, 2012; Berger et al., 2016). A higher loan share in total assets predicts lower subsequent stability, in line with greater exposure to credit risk and cyclical intermediation (Buchak et al., 2018). The deposit funding share enters negatively. In our setting, shifts in deposit reliance likely reflect business-model adjustments and competition for retail funding, which can accompany balance-sheet expansion that raises asset risk (Egan et al., 2022). The cost-to-income ratio is positive in the random-effects specifications but not under bank fixed effects, suggesting that cross-sectional differences in efficiency correlate with stability, whereas within-bank changes in efficiency are less informative once time-invariant bank characteristics are absorbed (De Silva et al., 2023).

**Figure 4** provides an illustration of our baseline results. We plot a binned scatter of residualized AI intensity against residualized log Z-scores after removing both bank and year fixed effects. By construction, this visualization isolates within-bank variation over time and abstracts from time-invariant heterogeneity across institutions as well as common macroeconomic shocks. The plot displays a clear downward slope, indicating that increases in AI intensity within a given bank are associated with lower stability. The relation appears approximately linear over the support of the data, suggesting that the estimates are not driven by outliers or strong nonlinearities.

## **4.2. AI Subcomponents and Enabling Technologies**

From a supervisory standpoint, it is relevant to identify which dimensions of AI-related technologies are most closely linked to bank stability. We therefore disaggregate our composite AI intensity measure into its underlying technology components and related enabling technologies. This decomposition assesses whether the baseline effect reflects broad engagement with agent adoption technologies or is concentrated in specific infrastructures—such as cloud computing and data capabilities—that can facilitate faster scaling, increase model and operational complexity, and introduce new third-party dependencies (including outsourcing and vendor concentration). Our approach follows the broader text-based finance literature that uses disclosure signals to proxy for firms’ technological orientation and transformation (Li, 2010; Hoberg & Phillips, 2016) and is

consistent with evidence that innovation and technology-driven competition can reshape intermediaries' risk-taking incentives (Buchak et al., 2018).

In **Table 5**, we report bank and year fixed-effects estimates in which we replace the composite AI-intensity index with lagged, standardized disclosure measures for each component. Columns (1)–(3) show that the negative association with subsequent stability is broad-based across the three building blocks of our baseline AI-intensity index. Agent adoption, cloud computing, and big data each predict lower Z-scores, with statistically significant coefficients. A one-standard-deviation increase in the agent adoption index is associated with a 0.22-unit reduction in the one-year-ahead Z-score, representing an approximate 28% decline relative to its standard deviation. The estimated effects for cloud computing and big data are smaller in magnitude but remain economically meaningful, yielding Z-score reductions of 0.11 and 0.13 units (which corresponds to 14% and 15% standard deviation change), respectively. This pattern implies that the baseline effect is not explained by a single keyword category but reflects a common association across complementary technologies that jointly support advanced analytics and rapid digital scaling.

Finally, column (4) examines an additional AI-related enabling technology—blockchain—that falls outside our baseline AI-intensity index. Because blockchain disclosures are available for a smaller subset of bank-years, including this component materially reduces our estimation sample. For this reason, we do not incorporate blockchain into the main AI intensity measure. Accordingly, the estimates should be interpreted with caution. Even so, the coefficient on blockchain intensity is negative, indicating that greater blockchain-related engagement is also associated with lower bank stability.

### **4.3. Channels: Z-Score Decomposition**

In **Table 6**, we disaggregate the Z-score into its underlying components: ROA volatility (sd ROA over the rolling window), ROA level, and equity-to-assets. This decomposition is standard in the banking-stability literature and allows separate “risk/volatility” effects from “buffer” effects (Boyd & Runkle, 1993; Laeven & Levine, 2009; Beltratti & Stulz, 2012). In all three regressions, we keep the same lag structure, controls, and clustered standard errors as in Table 3 to isolate how within-bank changes in AI intensity translate into the underlying drivers of stability.

The output points to volatility and profitability as the primary channels. Lagged AI intensity is positively related to subsequent ROA volatility (col. (1)): a one-standard-deviation increase in the AI-intensity index is followed by a statistically significant increase in earnings volatility ( $\beta=0.01$ ,  $t=3.07$ ). Consistent with this, we also find a decline in the ROA level (col. (2),  $\beta=-0.03$ ,  $t=-2.21$ ). In contrast, the association with equity-to-assets is small and statistically indistinguishable from zero (col. (3)).

These estimates suggest that the negative link between AI intensity and the Z-score in the baseline results is driven mainly by higher earnings risk and weaker profitability, rather than by an erosion of capital buffers. This pattern is consistent with mechanisms in which rapid scaling of advanced analytics increases model and operational complexity and accelerates the repricing and reallocation of risk, effects that appear first in profit volatility, while capital adjusts more gradually (Acharya, 2009; Danielsson et al., 2022; Danielsson & Uthemann, 2025).

From a policy perspective, the decomposition suggests that supervision should place particular weight on AI-related risk governance and earnings-risk dynamics (e.g., model risk management, validation, and operational resilience), since the stability effects emerge primarily through earnings volatility.

#### 4.4. Robustness Analyses

We assess robustness along several dimensions. First, we vary the construction of the stability measure by recomputing the log Z-score using three-, five-, and eight-year rolling windows for ROA volatility (excluding the current year), and re-estimate the baseline bank fixed-effects specification with the full control set, year fixed effects, and bank-clustered standard errors. Because the Z-score is mechanically sensitive to the horizon over which earnings risk is measured, this exercise addresses the concern that our results might reflect an arbitrary volatility-window choice. **Table 7** shows that the coefficient on lagged AI intensity remains negative and statistically significant across all window definitions.<sup>15</sup>

Second, we replace the baseline mean-based standardized AI-intensity with alternative AI measures that capture intensity and breadth in different ways: (i) a standardized total-count index,

<sup>15</sup> The small changes in sample size are driven by the rolling-window requirements.

(ii) a principal component–based index (first-component scores), and (iii) a raw word-count index. Because textual AI measures can differ in whether they capture the depth of discussion (counts), the breadth/co-movement across technologies (PCA), or simple disclosure volume (raw counts), this exercise verifies that our inferences are not sensitive to a particular scaling or aggregation choice. Across all variants, we maintain the baseline control set, bank and year fixed effects, and bank-clustered standard errors.<sup>16</sup> The AI coefficient in **Table 8** remains negative and highly significant in all three specifications.

As a third robustness check, we investigate whether our baseline outcomes depend on the measurement of AI-related activity. We re-estimate our main bank fixed-effects specification (with year fixed effects and the full set of lagged controls), retaining the baseline set of lagged bank controls, but replacing the baseline disclosure-based AI-intensity index with alternative proxies capturing complementary dimensions of technology adoption. In Column (1) of **Table 9**, we proxy AI adoption by AI investment intensity, defined as a bank’s investments in AI technologies (i.e., AI investment stock) scaled by total assets. This measure reflects the relative importance of AI investment on the bank’s balance sheet and is less likely to be mechanically driven by cross-sectional differences in bank size. In Column (2), we use the natural logarithm of the AI investment stock to proxy for the overall scale of AI adoption. Both investment-based proxies yield negative and statistically significant coefficients, consistent with the baseline interpretation. In Column (3), we consider a broader disclosure-based proxy for digital transformation based on digital-related term frequency in annual reports. This textual measure reflects banks’ communication and strategic emphasis on digitalization, which may be correlated with broader technology adoption and organizational change. The corresponding coefficient is negative but statistically insignificant, suggesting that digitalization-related disclosures may be a noisier and less targeted measure of AI adoption than investment-based proxies.

Finally, in **Table 10**, we assess whether the baseline AI–stability relation is driven by omitted, time-varying factors correlated with both AI intensity and bank risk. Columns (1)–(7) extend the baseline specification by adding additional controls one at a time to capture distinct sets of potential confounders: (i) asset risk and “risk density” (the ratio of risk-weighted assets to

<sup>16</sup> The PCA specification is estimated on a smaller sample, as constructing the principal component requires a balanced set of the underlying technology measures.

earning assets), (ii) balance-sheet expansion incentives (loan growth), (iii) ownership and governance (top-5 shareholding concentration and the share of independent directors), (iv) market structure and competitive pressure (the top five-bank Herfindahl index), and (v) local macroeconomic conditions (provincial income per capita and inflation). These factors are standard determinants of risk-taking in the banking literature (Boyd & Runkle, 1993; Iannotta et al., 2007; Laeven & Levine, 2009; Beck et al., 2013; Berger & Bouwman, 2013). In Column (8), we include all additional controls jointly, and in Column (9), we substitute bank fixed effects with bank-type fixed effects (e.g., state-owned/joint-stock/city/rural) to absorb time-invariant differences across bank types while retaining additional cross-sectional variation. Across all specifications, the coefficient on lagged AI intensity remains negative and highly significant, with magnitudes tightly clustered around the baseline model ( $\beta$  is approximately  $-0.26$  to  $-0.28$  in columns (1)–(8)), and it remains economically meaningful under the bank-type fixed-effects specification ( $\beta = -0.21$  in column (9)). These results reflect that the negative association between AI intensity and subsequent Z-scores is unlikely to be influenced by concurrent changes in asset risk, growth, governance, competitive conditions, or local macro fundamentals.

#### **4.5. Endogeneity Analysis**

A potential concern of our framework is that bank AI intensity is unlikely to be exogenous. Banks may expand AI in anticipation of changes in risk, performance, or supervisory scrutiny; unobserved province-level shocks may jointly influence adoption and outcomes; and our disclosure-based proxy may contain non-classical noise. These forces can bias fixed-effects OLS and attenuate coefficients when AI intensity is measured with error. We therefore estimate 2SLS models that instrument lagged bank-level AI intensity with lagged variables designed to move banks' incentives and costs to adopt AI while remaining plausibly orthogonal to bank outcomes once we condition on fixed effects and controls.

Our instrument choice follows two mechanisms emphasized in the technology-adoption and finance literatures. First, adoption diffuses through peer learning and competitive pressure, creating interdependence in firms' policies (Manski, 1993; Bramoullé et al., 2009; Leary & Roberts, 2014; Al Mamun et al., 2023). Second, implementing advanced IT and, by extension, AI, requires complementary inputs (skilled labor, vendors, organizational capital), so local ecosystems

can shift adoption costs and feasibility (Brynjolfsson & Hitt, 2000; Bresnahan et al., 2002). In credit markets, recent evidence shows that data-driven and machine-learning tools affect screening and underwriting, reinforcing that adoption hinges on both incentives and implementation capacity (Fuster et al., 2019; Berg et al., 2020; Fuster et al., 2022). Finally, because our AI-intensity index is text-based, IV is also useful to mitigate attenuation from noisy disclosure proxies, consistent with the broader “text-as-data” finance literature (Loughran & McDonald, 2011).

We use four lagged instruments capturing peer diffusion and local AI-relevant inputs. The first, the *AI-intensity peers index*, is the leave-one-out mean AI intensity of other banks in the same province-year. It captures benchmarking, imitation, and competitive responses, while excluding bank  $i$ 's own AI measure limits mechanical correlation (Manski, 1993). Peer channels are a standard mechanism through which firms adjust policies in response to their local competitors (Leary & Roberts, 2014). The remaining instruments proxy the local availability of complementary inputs that lower implementation frictions. *AI firms growth ratio* measures growth in the number of AI firms in the province, capturing the thickness of the local vendor and specialized labor market. A deeper ecosystem can reduce search and contracting costs and accelerate rollout, consistent with evidence that IT adoption and its returns depend on complements in skills and organization (Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002). *New product development intensity growth ratio* reflects changes in new-product development expenditures scaled by GDP, capturing an innovation environment that facilitates experimentation and speeds diffusion (Griliches, 1957). Finally, the *R&D personnel intensity growth ratio* is the growth ratio in industrial R&D personnel scaled by population, capturing local technical capacity and absorptive capability (Cohen & Levinthal, 1990; Babina et al., 2024; Wu et al., 2025) and reflecting localized knowledge bases that support implementation (Jaffe et al., 1993).

The relevance of our instruments is grounded in the underlying adoption mechanisms. Peer adoption pressure and a thicker local ecosystem plausibly increase banks' AI adoption by shifting expected benefits and lowering implementation costs. The identifying assumption is that these lagged shifters affect bank outcomes primarily through bank AI intensity, conditional on fixed effects and controls. The main concern is that province-level technology and innovation measures may also proxy for broader economic development that directly affects banks. We address this issue by progressively conditioning on richer local controls. In Model 1, we include standard bank controls and year fixed effects. In Model 2, we add lagged province fundamentals (income per

capita, industrial structure, inflation, and financial-sector wages) to absorb time-varying local development. In Model 3, we further include lagged measures of broad ICT infrastructure (telecom service volume and broadband access ports) to separate AI-specific ecosystem effects from general digital readiness.

In **Table 11**, the first-stage regression relates the AI-intensity index to the excluded instruments, the full set of controls, and fixed effects. The estimated signs align with our economic intuition. The leave-one-out peer AI measure enters with a negative coefficient, suggesting that, conditional on bank characteristics and local fundamentals, banks tend to differentiate from nearby peers. This pattern is consistent with strategic positioning and potential crowding-out in local markets (e.g., competition for scarce AI talent or vendor capacity). In contrast, the local ecosystem shifters have positive coefficients. Lagged growth in provincial AI firms predicts higher bank AI intensity, consistent with a thicker vendor and talent market lowering implementation frictions. Likewise, stronger innovation activity and faster growth in R&D human-capital intensity are positively associated with bank AI intensity, consistent with greater experimentation capacity and technical labor supply facilitating adoption.

The identification diagnostics indicate that the instrument set is relevant. The underidentification test based on Kleibergen-Paap rk LM statistics rejects the null that the equation is underidentified, and the Kleibergen-Paap rk Wald F statistics exceed conventional thresholds (and the relevant Stock–Yogo critical values), suggesting that weak-instrument concerns are limited. Finally, Hansen’s  $J$  test does not reject the joint validity of the instruments. While overidentification tests cannot establish exclusion, the strength of the first stage and the stability of results across increasingly stringent controls are consistent with the interpretation that the instruments primarily shift banks’ AI intensity through adoption and implementation channels. Overall, the IV estimates reinforce our baseline results, indicating that higher AI intensity is associated with a statistically and economically meaningful decline in banks’ Z-scores.

## **5. Additional Analyses**

### **5.1. AI Intensity Thresholds and Bank Stability**

In this section, we ask whether the baseline relation between AI intensity and stability is

concentrated among banks operating in “high-AI” states, as opposed to reflecting a purely linear effect. **Appendix 4** motivates this state-based perspective and provides additional details on the AI adoption process. The transition matrix shows meaningful persistence in banks’ relative AI positioning (one-year persistence of 51.5%), but also substantial mobility: 44.3% of bank-years in the low tercile move up in the following year, and 38.2% of bank-years in the high tercile move down. These dynamics indicate that AI intensity evolves through gradual re-ranking in contrast to a one-time treatment event with a common adoption date. Accordingly, we complement the baseline continuous specification with a threshold design that uses a high-intensity indicator capturing whether a bank-year falls in the upper tail of the AI-intensity distribution. Formally, we estimate the following fixed-effects specification:

$$\ln(Z - score)_{i,t} = \alpha_i + \delta_t + \beta HighAI_{i,t-1} + \gamma' X_{i,t-1} + \varepsilon_{i,t},$$

where  $\ln(Z - score)_{i,t}$  is the log Z-score,  $\alpha_i$  and  $\delta_t$  denote bank and year fixed effects,  $HighAI_{i,t-1}$  is an indicator equal to one if bank  $i$  is classified in the high AI-intensity group in year  $t-1$  (e.g., top quartile or high tercile), and  $X_{i,t-1}$  is the standard vector of lagged bank controls.

This threshold design tests whether transitions into (and out of) high-AI states are associated with discrete changes in subsequent stability, while preserving the within-bank identification and common-shock adjustment of the baseline framework. This approach is preferable to a conventional DiD setup in our setting because “treatment” is not tied to a single, common adoption date or an exogenous shock across banks. Instead, banks enter and exit high-AI states at different times, and the transition process is gradual.<sup>17</sup>

Results in **Table 12** indicate that banks operating at high levels of AI intensity exhibit an enhanced probability of default. In the top-quartile design, the lagged high-AI indicator is negative and statistically significant in both models, with the larger fixed-effects magnitude consistent with identification coming from within-bank transitions after absorbing time-invariant heterogeneity. The tercile specification extends this pattern beyond the extreme right tail. Relative to the low-AI group, both the middle- and high-tercile groups exhibit lower distance-to-default, with

<sup>17</sup> With a relatively small panel (35 banks) and substantial year-to-year persistence in AI intensity (**Appendix Table A4**), a DiD design based on a sharply defined post period or a single event window would have limited statistical power and require stronger parallel-trends assumptions. The indicator-based specification instead exploits observed transitions into and out of high-AI states, avoids imposing an arbitrary common break date, and allows us to test for economically meaningful threshold effects while maintaining the baseline control structure.

economically and statistically meaningful estimates in the fixed-effects model for the high tercile. This pattern of larger effects in the high tercile than in the middle tercile suggests that the stability effect scales with the intensity of AI engagement.

## 5.2. Moderating Factors: Liquidity, Funding Structure, and Loss-Absorption Capacity

A conceptual question that arises from our framework is whether the AI–stability relation varies across banks’ funding structure, liquidity conditions, and loss-absorption buffers. AI adoption can raise short-run earnings volatility and tail risk through faster balance-sheet adjustment, model/operational risk, and learning-by-doing in credit and risk management. When liquidity buffers are thin or funding is unstable, these volatility shocks are more likely to translate into fragility, consistent with classic bank-run and liquidity-risk mechanisms and with evidence that liquidity and stable funding discipline risk-taking and dampen stress amplification (Diamond & Dybvig, 1983; Holmström & Tirole, 1998; Brunnermeier & Pedersen, 2009; Berger & Bouwman, 2013).

To examine this heterogeneity, in **Table 13**, we augment the baseline fixed-effects specification by interacting lagged AI intensity with a lagged moderating variable  $M$ , introduced one at a time, and estimate the following specification:

$$\begin{aligned} \ln(Z - score)_{i,t} &= \alpha_i + \gamma_t + \beta_1 \times AI - intensity\ index_{i,t-1} + \beta_2 \times M_{i,t-1} \\ &+ \beta_3 \times (AI - intensity\ index_{i,t-1} \times M_{i,t-1}) + \delta' \times X_{i,t-1} + \varepsilon_{i,t}, \end{aligned}$$

where  $\alpha_i$  and  $\gamma_t$  denote bank and year fixed effects, and  $X_{i,t-1}$  is the baseline vector of lagged bank controls. The coefficient  $\beta_3$  captures how the marginal effect of AI adoption on stability varies with  $M$ , since  $\partial \ln(Z - score)_{i,t} / \partial AI - intensity\ index_{i,t-1} = \beta_1 + \beta_3 M_{i,t-1}$ .

The first moderating factor is reliance on interbank deposits, measured as interbank deposits divided by total deposits plus interbank liabilities. In Column (1), the interaction term is negative and statistically significant, indicating that the negative association between AI intensity and distance-to-default becomes more pronounced as banks rely more on runnable wholesale-type funding. This pattern is consistent with classic fragility mechanisms in which funding structures that are more exposed to rollover risk and coordination problems amplify adverse shocks and

accelerate liquidity stress (Diamond & Dybvig, 1983; Gorton & Metrick, 2012). In this setting, AI-driven changes, such as faster balance-sheet adjustment, more elastic credit supply, or more correlated risk-taking, are more likely to transmit rapidly into funding pressures when liabilities are less stable, thereby intensifying the negative effects of AI intensity on bank stability (Huang & Ratnovski, 2011).

We then consider liquidity conditions using the exchange liquidity ratio. In column (2), the interaction is positive and highly significant, indicating that stronger liquidity buffers reduce the negative effect of AI intensity on banks' distance-to-default. One interpretation is that greater balance-sheet liquidity and more marketable positions reduce the scope for fire-sale dynamics and funding spirals, thereby mitigating the stability consequences of technology-driven balance-sheet adjustments. This mechanism aligns with the theories emphasizing the complementarity between funding liquidity and market liquidity in generating (or preventing) liquidity spirals (Brunnermeier & Pedersen, 2009).

In column (3), we focus on deposit growth. The interaction is positive and statistically significant, suggesting that banks with stronger deposit inflows experience a weaker negative AI–stability relation. A natural interpretation is a “deposit franchise” mechanism. Abundant retail deposits provide relatively stable funding and liquidity insurance, dampening the transmission from AI-enabled expansion or risk-taking into funding stress. This interpretation is consistent with evidence that deposits provide insurance against liquidity risk and become particularly valuable when market liquidity tightens (Gatev et al., 2009).

Finally, in column (4), we examine the effects of loss-absorption capacity using the provision coverage ratio (provisions relative to non-performing loans). The interaction term is positive and significant, indicating that stronger provisioning buffers mitigate the negative effects of AI intensity on banks' stability. This pattern is consistent with the view that larger ex ante buffers reduce the likelihood that model error, accelerated credit growth, or weaker screening and monitoring translate into realized distress. More broadly, prior work emphasizes that provisioning regimes shape the cyclicity of bank risk and the timing of loss recognition, with implications for financial resilience and discipline (Laeven & Majnoni, 2003; Bushman & Williams, 2012; Beatty & Liao, 2014).

**Figure 5** displays these interactions by plotting predictive margins from the fixed-effects

specifications. For each moderating factor—interbank funding dependence, exchange liquidity, deposit growth, and provisioning capacity— we trace fitted values of log Z-scores across the empirical distribution of the moderator at different levels of AI intensity, holding other covariates constant and absorbing bank and year fixed effects.

These findings have direct supervisory implications. They suggest that assessments of AI-enabled credit decisions and risk management should be evaluated jointly with banks’ liquidity and funding profiles, instead of being viewed as a uniform shift in risk. The mitigating role of provisioning coverage further highlights the importance of conservative provisioning practices and robust model governance—encompassing validation, monitoring, and override controls—particularly for banks with more intensive AI use. More broadly, incorporating standardized disclosures on material AI applications alongside core liquidity and funding metrics could strengthen early-warning frameworks and help focus supervisory attention on institutions where AI is most likely to amplify fragility.

## **6. Conclusions and policy implications**

In this paper, we provide new evidence that greater intensity adoption of AI-related technologies by banks is associated with lower bank stability in the short to medium run. Using annual-report disclosures to construct text-based measures of banks’ engagement with *agent adoption*, *cloud computing*, and *big data*, we find that increases in AI intensity predict a subsequent decline in banks’ distance-to-default. Decomposing the stability measure shows that the effect operates primarily through higher earnings volatility and weaker profitability, with no detectable change in capitalization. Consistent with a state-dependent fragility mechanism, this negative effect of AI intensity on bank stability is most pronounced among banks with weaker liquidity positions, greater reliance on more runnable funding, and thinner loss-absorption capacity.

Our moderating-factor evidence indicates that the stability implications of AI adoption are highly state-dependent. The adverse association between AI intensity and subsequent stability is most pronounced for banks with greater reliance on interbank and other runnable wholesale funding, consistent with technology-enabled balance-sheet adjustment interacting with fragile liability structures to amplify stress. In contrast, the association is significantly weaker for banks with stronger liquidity buffers, more robust deposit growth, and greater ex ante loss-absorption

capacity through higher provisioning coverage. Overall, the results imply that liquidity backstops, stable funding, and conservative credit-risk buffers meaningfully dampen the downside risks associated with rapid AI-enabled expansion.

These findings highlight the risks associated with scaling AI-related technologies. For banks, our results emphasize the need to complement AI investments with robust internal risk frameworks—strong model governance, independent validation, change control, and continuous performance monitoring—alongside portfolio-level constraints that contain volatility during rollout and learning phases. Absent such safeguards, rapid technological re-optimization can translate into greater earnings volatility and weaker stability.

Our moderating-factor analysis shows that the stability consequences of AI adoption are far from uniform. The negative AI–stability association is strongest for banks that rely more heavily on interbank and other runnable wholesale funding, consistent with the view that faster balance-sheet adjustment interacts with fragile liability structures to amplify stress. By contrast, the effect is significantly attenuated when banks operate with stronger liquidity positions, experience robust deposit growth, or maintain larger ex ante credit risk buffers through higher provisioning coverage. These patterns suggest that liquidity backstops, stable funding, and conservative credit risk frameworks play a disciplining and absorptive role, limiting the downside risks of AI-driven expansion.

The results also have relevant supervisory implications. AI adoption should be assessed in conjunction with banks’ liquidity, funding, and buffer positions, rather than as a standalone risk factor. Supervisors may wish to prioritize AI-intensive institutions in liquidity stress testing and contingency funding assessments, particularly when reliance on interbank funding is high, and to view strong provisioning practices and model-risk governance as complementary instruments for ensuring that AI-related innovation remains compatible with financial resilience.

Finally, our analysis points to several directions for future research as richer data become available. A key constraint in the current evidence base is limited, consistent reporting on AI spending, use-case criticality, third-party exposures (cloud concentration), and model governance outcomes. New supervisory reporting templates, vendor-level datasets, and more granular disclosures—potentially including model inventories, validation events, override rates, and incident logs—would enable sharper tests of mechanisms, separation of adoption from disclosure

intensity, and a better mapping from specific AI applications to stability outcomes across jurisdictions.

### **Data availability statement**

The data used in this study are sourced from proprietary databases and contain confidential bank-level information that cannot be publicly disclosed. Replication code and non-confidential components of the dataset are available from the authors upon request.

### **References**

- Acharya, V. V. (2009). A theory of systemic risk and design of prudential bank regulation. *Journal of Financial Stability*, 5(3), 224-255. <https://doi.org/10.1016/j.jfs.2009.02.001>
- Al Mamun, M., Malik, M., & Mia, M. A. H. (2023). Peer pressure, information technology adoption, and bank performance. *Asian Journal of Business and Accounting*, 161-191. <https://doi.org/10.22452/ajba.vol16no1.6>
- Allen, F., Qian, J., & Qian, M. (2005). Law, finance, and economic growth in China. *Journal of Financial Economics*, 77(1), 57-116. <https://doi.org/10.1016/j.jfineco.2004.06.010>
- Ante, L., & Saggiu, A. (2025). Quantifying a firm's AI engagement: Constructing objective, data-driven, AI stock indices using 10-K filings. *Technological Forecasting and Social Change*, 212, 123965. <https://doi.org/10.1016/j.techfore.2024.123965>
- Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, 103745. <https://doi.org/10.1016/j.jfineco.2023.103745>
- Bailey, W., Huang, W., & Yang, Z. (2011). Bank loans with Chinese characteristics: Some evidence on inside debt in a state-controlled banking system. *Journal of Financial and Quantitative Analysis*, 46(6), 1795-1830. <https://doi.org/10.1017/S0022109011000433>
- Basel Committee on Banking Supervision (BCBS). (2024). *Principles for the sound management of third-party risk*. Bank for International Settlements. <https://www.bis.org/bcbs/publ/d605.pdf>
- Beatty, A., & Liao, S. (2014). Financial accounting in the banking industry: A review of the empirical literature. *Journal of Accounting and Economics*, 58(2-3), 339-383. <https://doi.org/10.1016/j.jacceco.2014.08.009>

- Beccalli, E. (2007). Does IT investment improve bank performance? Evidence from Europe. *Journal of Banking & Finance*, 31(7), 2205-2230. <https://doi.org/10.1016/j.jbankfin.2006.10.022>
- Becerra-Vicario, R., Salas-Compás, B., Valcarce-Ruiz, L., & Sánchez-Serrano, J. R. (2024). The Impact of Artificial Intelligence in The Financial Sector: Opportunities and Challenges. *International Journal of Business & Management Studies*, 5(10), 33-42. [10.56734/ijbms.v5n10a4](https://doi.org/10.56734/ijbms.v5n10a4)
- Beck, T., De Jonghe, O., & Schepens, G. (2013). Bank competition and stability: Cross-country heterogeneity. *Journal of Financial Intermediation*, 22(2), 218-244. <https://doi.org/10.1016/j.jfi.2012.07.001>
- Begenau, J., Farboodi, M., & Veldkamp, L. (2018). Big data in finance and the growth of large firms. *Journal of Monetary Economics*, 97, 71-87. <https://doi.org/10.1016/j.jmoneco.2018.05.013>
- Bellstam, G., Bhagat, S., & Cookson, J.A. (2021). A text-based analysis of corporate innovation. *Management Science*, 67(7), 4004-4031. <https://doi.org/10.1287/mnsc.2020.3682>
- Beltratti, A., & Stulz, R. M. (2012). The credit crisis around the globe: Why did some banks perform better?. *Journal of Financial Economics*, 105(1), 1-17. <https://doi.org/10.1016/j.jfineco.2011.12.005>
- Berg, T., Burg, V., Gombović, A., & Puri, M. (2020). On the rise of Fintechs: Credit scoring using digital footprints. *The Review of Financial Studies*, 33(7), 2845-2897. <https://doi.org/10.1093/rfs/hhz099>
- Berg, T., Fuster, A., & Puri, M. (2022). Fintech lending. *Annual Review of Financial Economics*, 14(1), 187-207. <https://doi.org/10.1146/annurev-financial-101521-112042>
- Berger, A. N., & Bouwman, C. H. (2013). How does capital affect bank performance during financial crises?. *Journal of Financial Economics*, 109(1), 146-176. <https://doi.org/10.1016/j.jfineco.2013.02.008>
- Berger, A. N., Hasan, I., & Zhou, M. (2009). Bank ownership and efficiency in China: What will happen in the world's largest nation?. *Journal of Banking & Finance*, 33(1), 113-130. <https://doi.org/10.1016/j.jbankfin.2007.05.016>
- Berger, A. N., Imbierowicz, B., & Rauch, C. (2016). The roles of corporate governance in bank failures during the recent financial crisis. *Journal of Money, Credit and Banking*, 48(4), 729-770. <https://doi.org/10.1111/jmcb.12316>
- Boyd, J. H., & Runkle, D. E. (1993). Size and performance of banking firms: Testing the predictions of theory. *Journal of Monetary Economics*, 31(1), 47-67. [https://doi.org/10.1016/0304-3932\(93\)90016-9](https://doi.org/10.1016/0304-3932(93)90016-9)
- Bramoullé, Y., Djebbari, H., & Fortin, B. (2009). Identification of peer effects through social networks. *Journal of Econometrics*, 150(1), 41-55. <https://doi.org/10.1016/j.jeconom.2008.12.021>
- Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, 117(1), 339-376. <https://doi.org/10.1162/003355302753399526>

- Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The Review of Financial Studies*, 22(6), 2201-2238. <https://doi.org/10.1093/rfs/hhn098>
- Brynjolfsson, E., & Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4), 23-48. [10.1257/jep.14.4.23](https://doi.org/10.1257/jep.14.4.23)
- Buchak, G., Matvos, G., Piskorski, T., & Seru, A. (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics*, 130(3), 453-483. <https://doi.org/10.1016/j.jfineco.2018.03.011>
- Bushman, R. M., & Williams, C. D. (2012). Accounting discretion, loan loss provisioning, and discipline of banks' risk-taking. *Journal of Accounting and Economics*, 54(1), 1-18. <https://doi.org/10.1016/j.jacceco.2012.04.002>
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128-152.
- Danielsson, J., & Uthemann, A. (2025). Artificial intelligence and financial crises. *Journal of Financial Stability*, 101453. <https://doi.org/10.1016/j.jfs.2025.101453>
- Danielsson, J., Macrae, R., & Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance*, 140, 106290. <https://doi.org/10.1016/j.jbankfin.2021.106290>
- Daoud, S., & Anaya, L. (2025). Implementation of robotic process automation in Jordanian banking sector: Benefits and challenges. *Procedia Computer Science*, 263, 471-480. <https://doi.org/10.1016/j.procs.2025.07.057>
- De Silva, A., Duong, H. N., Nguyen, M., & Nguyen, Y. N. (2023). Bank risk in uncertain times: Do credit rationing and revenue diversification matter?. *Journal of Business Finance & Accounting*, 50(7-8), 1240-1273. <https://doi.org/10.1111/jbfa.12653>
- Demirgüç-Kunt, A., & Huizinga, H. (2010). Bank activity and funding strategies: The impact on risk and returns. *Journal of Financial Economics*, 98(3), 626-650. <https://doi.org/10.1016/j.jfineco.2010.06.004>
- Demirgüç-Kunt, A., & Huizinga, H. (2013). Are banks too big to fail or too big to save? International evidence from equity prices and CDS spreads. *Journal of Banking & Finance*, 37(3), 875-894. <https://doi.org/10.1016/j.jbankfin.2012.10.010>
- Demirgüç-Kunt, A., Peria, M. S. M., & Tressel, T. (2020). The global financial crisis and the capital structure of firms: Was the impact more severe among SMEs and non-listed firms?. *Journal of Corporate Finance*, 60, 101514. <https://doi.org/10.1016/j.jcorpfin.2019.101514>
- DeYoung, R. (2014). Banking in the United States. <https://doi.org/10.1093/oxfordhb/9780199688500.013.0034>
- DeYoung, R., Lang, W. W., & Nolle, D. L. (2007). How the Internet affects output and performance at community banks. *Journal of Banking & Finance*, 31(4), 1033-1060. <https://doi.org/10.1016/j.jbankfin.2006.10.003>
- Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3), 401-419. <https://doi.org/10.1086/261155>

- Egan, M., Lewellen, S., & Sunderam, A. (2022). The cross-section of bank value. *The Review of Financial Studies*, 35(5), 2101-2143. <https://doi.org/10.1093/rfs/hhab089>
- Elekdag, S., Emrullahu, D., & Naceur, S. B. (2025). Does FinTech increase bank risk-taking?. *Journal of Financial Stability*, 76, 101360. <https://doi.org/10.1016/j.jfs.2024.101360>
- European Systemic Risk Board (ESRB). (2025). *Artificial intelligence and systemic risk in the financial system*. European Systemic Risk Board, Advisory Scientific Committee Report. [https://www.esrb.europa.eu/pub/pdf/asc/esrb.ascreport202512\\_AIandsystemicrisk.en.pdf](https://www.esrb.europa.eu/pub/pdf/asc/esrb.ascreport202512_AIandsystemicrisk.en.pdf)
- Financial Stability Board (FSB). (2019). *Third-party dependencies in cloud services: Considerations on financial stability implications*. <https://www.fsb.org/uploads/P091219-2.pdf>
- Financial Stability Board (FSB). (2023). *Cloud adoption in the financial sector and concentration risk*. <https://www.fsb.org/uploads/PIFS.pdf>
- Financial Stability Board (FSB). (2024). *The financial stability implications of artificial intelligence*. Financial Stability Board. <https://www.fsb.org/uploads/P14112024.pdf>
- Financial Stability Board (FSB). (2025). *Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector*. Financial Stability Board. [P101025.pdf](https://www.fsb.org/uploads/P101025.pdf)
- Fiordelisi, F., Marques-Ibanez, D., & Molyneux, P. (2011). Efficiency and risk in European banking. *Journal of Banking & Finance*, 35(5), 1315-1326. <https://doi.org/10.1016/j.jbankfin.2010.10.005>
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance*, 77(1), 5-47. <https://doi.org/10.1111/jofi.13090>
- Fuster, A., Plosser, M., Schnabl, P., & Vickery, J. (2019). The role of technology in mortgage lending. *The Review of Financial Studies*, 32(5), 1854-1899. <https://doi.org/10.1093/rfs/hhz018>
- Gambacorta, L., Huang, Y., Qiu, H., & Wang, J. (2019). How do machine learning and non traditional data affect credit scoring? *New evidence from a Chinese fintech firm* (BIS Working Papers No. 834). Bank for International Settlements. <https://www.bis.org/publ/work834.htm>
- Gatev, E., Schuermann, T., & Strahan, P. E. (2009). Managing bank liquidity risk: How deposit-loan synergies vary with market conditions. *The Review of Financial Studies*, 22(3), 995-1020. <https://doi.org/10.1093/rfs/hhm060>
- Gorton, G., & Metrick, A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3), 425-451. <https://doi.org/10.1016/j.jfineco.2011.03.016>
- Griliches, Z. (1957). Hybrid Corn: An Exploration in the Economics of Technological Change *Econometrica* Vol. 25. <https://doi.org/10.2307/1905380>

- Gyau, E. B., Appiah, M., Gyamfi, B. A., Achie, T., & Naeem, M. A. (2024). Transforming banking: Examining the role of AI technology innovation in boosting banks financial performance. *International Review of Financial Analysis*, 96, 103700. <https://doi.org/10.1016/j.irfa.2024.103700>
- Hafeez, B., Li, X., Kabir, M. H., & Tripe, D. (2022). Measuring bank risk: Forward-looking z-score. *International Review of Financial Analysis*, 80, 102039. <https://doi.org/10.1016/j.irfa.2022.102039>
- Hernando, I., & Nieto, M. J. (2007). Is the Internet delivery channel changing banks' performance? The case of Spanish banks. *Journal of Banking & Finance*, 31(4), 1083-1099. <https://doi.org/10.1016/j.jbankfin.2006.10.011>
- Hoberg, G., & Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5), 1423-1465. <https://doi.org/10.1086/688176>
- Holmström, B., & Tirole, J. (1998). Private and public supply of liquidity. *Journal of Political Economy*, 106(1), 1-40. <https://doi.org/10.1086/250001>
- Houston, J. F., Lin, C., Lin, P., & Ma, Y. (2010). Creditor rights, information sharing, and bank risk taking. *Journal of Financial Economics*, 96(3), 485-512. <https://doi.org/10.1016/j.jfineco.2010.02.008>
- Huang, R. & Ratnovski, M. L. (2011). The dark side of bank wholesale funding. *Journal of Financial Intermediation* 20(2), 248-26. <https://doi.org/10.1016/j.jfi.2010.06.003>
- Iannotta, G., Nocera, G., & Sironi, A. (2007). Ownership structure, risk and performance in the European banking industry. *Journal of Banking & Finance*, 31(7), 2127-2149. <https://doi.org/10.1016/j.jbankfin.2006.07.013>
- Jaffe, A. B., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics*, 108(3), 577-598. <https://doi.org/10.2307/2118401>
- Jia, X. (2024). FinTech penetration, charter value, and bank risk-taking. *Journal of Banking & Finance*, 161, 107111. <https://doi.org/10.1016/j.jbankfin.2024.107111>
- Jiang, L., Levine, R., & Lin, C. (2016). Competition and bank opacity. *The Review of Financial Studies*, 29(7), 1911-1942. <https://doi.org/10.1093/rfs/hhw016>
- Jondeau, E., Mojon, B., & Sahuc, J. G. (2020). Bank Rollover Risk and Liquidity Supply Regimes. *Swiss Finance Institute Research Paper*, (20-20). <https://dx.doi.org/10.2139/ssrn.3575715>
- Kazinnik, S. (2025). Bank run, interrupted: Modeling deposit withdrawals with generative AI. *Working Paper*. <http://dx.doi.org/10.2139/ssrn.4656722>
- Krishnamurthy, A., Nagel, S., & Orlov, D. (2014). Sizing up repo. *The Journal of Finance*, 69(6), 2381-2417. <https://doi.org/10.1111/jofi.12168>
- Laeven, L., & Levine, R. (2009). Bank governance, regulation and risk taking. *Journal of Financial Economics*, 93(2), 259-275. <https://doi.org/10.1016/j.jfineco.2008.09.003>

- Laeven, L., & Majnoni, G. (2003). Loan loss provisioning and economic slowdowns: too much, too late?. *Journal of Financial Intermediation*, 12(2), 178-197. [https://doi.org/10.1016/S1042-9573\(03\)00016-0](https://doi.org/10.1016/S1042-9573(03)00016-0)
- Leary, M. T., & Roberts, M. R. (2014). Do peer firms affect corporate financial policy? *The Journal of Finance*, 69(1), 139-178. <https://doi.org/10.1111/jofi.12094>
- Li, F. (2010). The information content of forward-looking statements in corporate filings—A naïve Bayesian machine learning approach. *Journal of Accounting Research*, 48(5), 1049-1102. <https://doi.org/10.1111/j.1475-679X.2010.00382.x>
- Li, J. (2025). Artificial intelligence innovation and financial report quality. *International Review of Economics & Finance*, 104832. <https://doi.org/10.1016/j.iref.2025.104832>
- Li, Z. (2025). Does artificial intelligence enhance bank profitability? Evidence from China. *Finance Research Letters*, 109366. <https://doi.org/10.1016/j.frl.2025.109366>
- Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, 66(1), 35-65. <https://doi.org/10.1111/j.1540-6261.2010.01625.x>
- Loughran, T., & McDonald, B. (2014). Measuring readability in financial disclosures. *The Journal of Finance*, 69(4), 1643-1671. <https://doi.org/10.1111/jofi.12162>
- Loughran, T., & McDonald, B. (2015). The use of word lists in textual analysis. *Journal of Behavioral Finance*, 16(1), 1-11. <https://doi.org/10.1080/15427560.2015.1000335>
- Ma, Y., & Li, Z. (2020). Funding liquidity and bank risk-taking: Empirical evidence from China. *China Finance and Economic Review*, 9(1), 64-82. <https://doi.org/10.1515/cfer-2020-090105>
- Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3), 531-542. <https://doi.org/10.2307/2298123>
- Maslej, N., L. Fattorini, R. Perrault, Y. Gil, V. Parli, N. Kariuki, E. Capstick, A. Reuel, E. Brynjolfsson, J. Etchemendy, K. Ligett, T. Lyons, J. Manyika, J. C. Niebles, Y. Shoham, R. Wald, T. Walsh, A. Hamrah, L. Santarlasci, J. B. Lotufo, A. Rome, A. Shi, & Oak, S. (2025). *Artificial Intelligence Index Report 2025*. Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford University. <https://arxiv.org/abs/2504.07139>
- Rauf, A. (2026). Reaching for yield? Artificial intelligence and mortgage credit allocation. *Working Paper*. <http://dx.doi.org/10.2139/ssrn.6095528>
- Stiroh, K. J. (2004). Diversification in banking: Is noninterest income the answer?. *Journal of Money, Credit and Banking*, 853-882. <https://www.jstor.org/stable/3839138>
- Stiroh, K. J., & Rumble, A. (2006). The dark side of diversification: The case of US financial holding companies. *Journal of Banking & Finance*, 30(8), 2131-2161. <https://doi.org/10.1016/j.jbankfin.2005.04.030>
- Vives, X. (2019). Digital disruption in banking. *Annual Review of Financial Economics*, 11(1), 243-272. <https://doi.org/10.1146/annurev-financial-100719-120854>

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT Press.

World Intellectual Property Organization (WIPO) (2019). *WIPO Technology Trends 2019: Artificial Intelligence*. Geneva: World Intellectual Property Organization.  
[https://www.wipo.int/edocs/pubdocs/en/wipo\\_pub\\_1055.pdf](https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf)

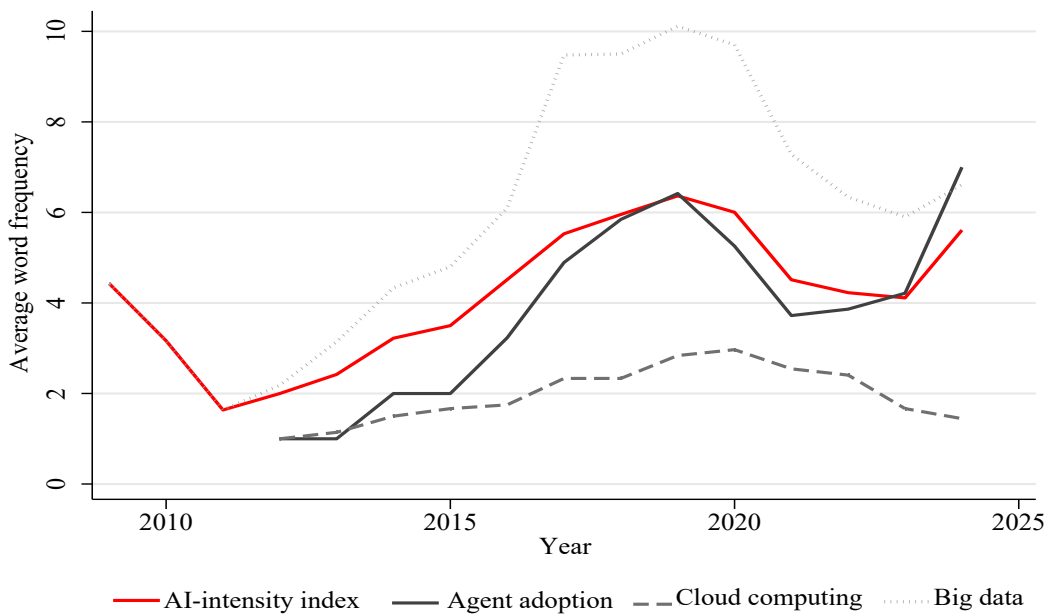
Wu, Y., Yuan, Y., & Song, X. (2025). The impact of AI adoption on R&D productivity: Evidence from Chinese pharmaceutical manufacturing industry. *Journal of Asian Economics*, 97, 101890.  
<https://doi.org/10.1016/j.asieco.2025.101890>

Yildiz, K., & Cil, A. E. (2025). A systematic literature review on applications of explainable artificial intelligence in the financial sector. *Internet of Things*, 101696. <https://doi.org/10.1016/j.iot.2025.101696>

Figure 1

**Trends in AI-related Technologies in Banks**

Figure 1 shows the annual trends in AI-related technologies in banks' disclosures for our baseline estimation sample of 35 banks from China. Lines report yearly mean word-frequency counts for the *AI-intensity index* and its components—agent adoption, cloud computing, and big data—based on keyword dictionaries applied to banks' annual reports, as detailed in Appendix 2. Higher values indicate greater disclosure intensity and stronger emphasis on AI-related technologies across banks in a given year.

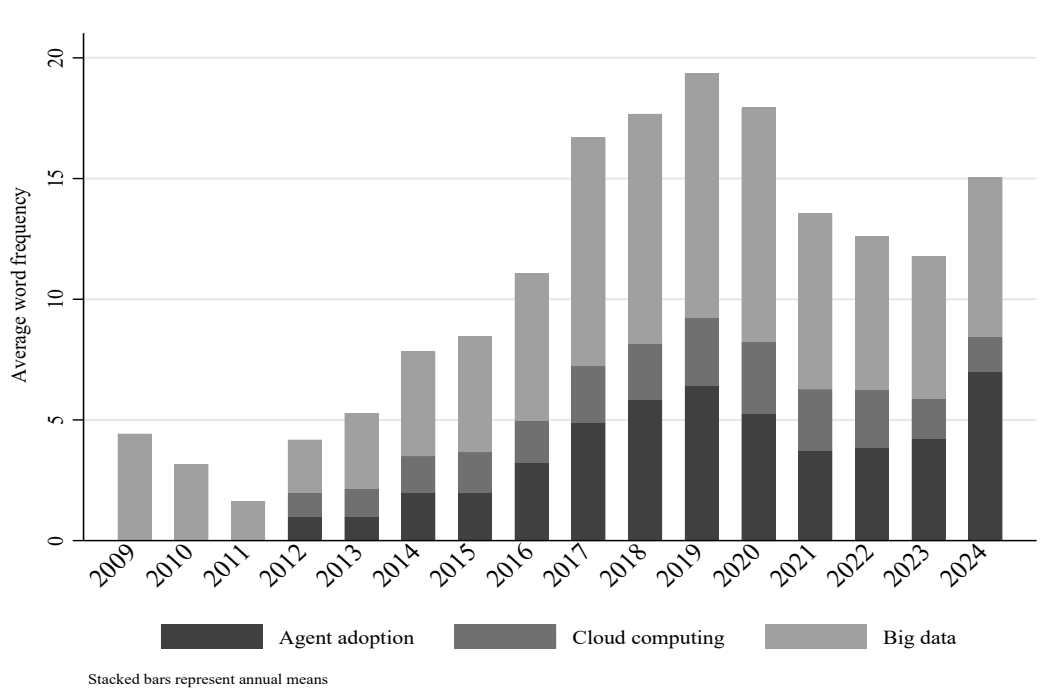


Lines report annual mean word-frequency counts based on banks' disclosures.

Figure 2

**Composition of AI-related Technologies in Banks**

Figure 2 shows the annual composition of AI-related technologies in banks' disclosures for our baseline estimation sample of 35 banks from China. Stacked bars report yearly mean word-frequency counts for agent adoption, cloud computing, and big data, based on keyword dictionaries applied to banks' annual reports, as detailed in Appendix 2. Higher values indicate greater disclosure intensity and a stronger emphasis on AI-related technologies across banks in a given year.



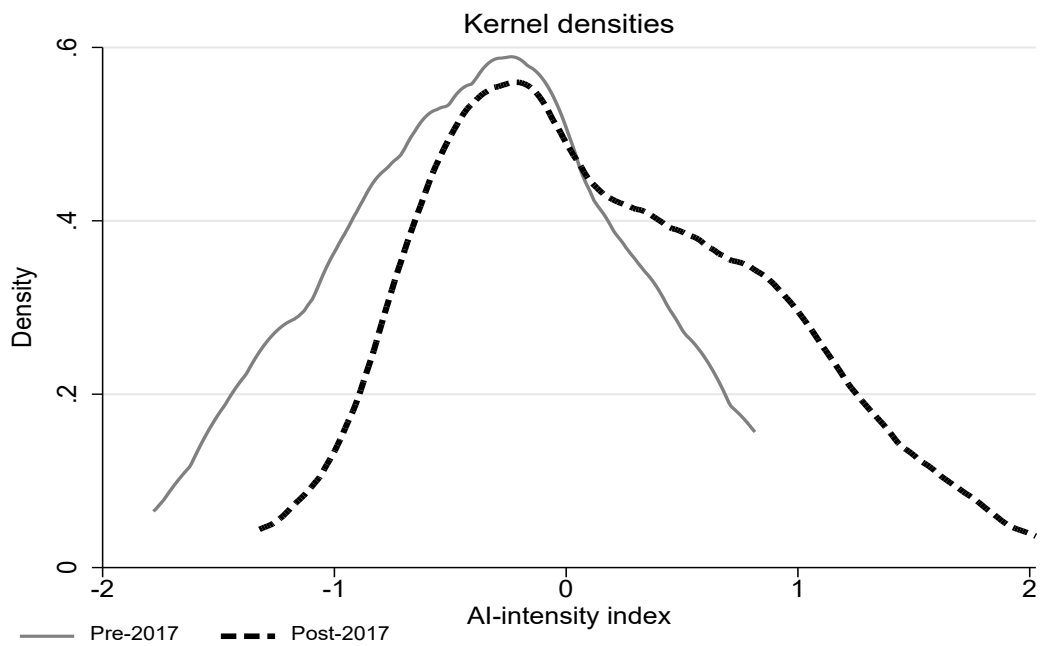
---

Figure 3

**Distribution of the AI-Intensity Index**

---

Figure 3 plots kernel density estimates of the standardized *AI-intensity index* for the periods before and after 2017. AI intensity is measured using standardized word-frequency indices extracted from banks' public disclosures, as detailed in Appendix 2. The sample is restricted to the baseline estimation sample which includes 35 banks from China.



AI intensity is measured by the standardized AI-intensity index. The sample is restricted to the baseline estimation sample.

Figure 4

**Binned Scatter (within-bank): AI-Intensity Index vs Z-score**

Figure 4 presents a binned scatter plot of  $\ln(\text{Z-score})$  against the *AI-intensity index* using within-bank variation. Both variables are residualized with respect to bank and year fixed effects. Each dot represents the mean of the variables within equally sized bins of the *AI-intensity index*. The fitted line depicts the relationship between  $\ln(\text{Z-score})$  and the *AI-intensity index*.

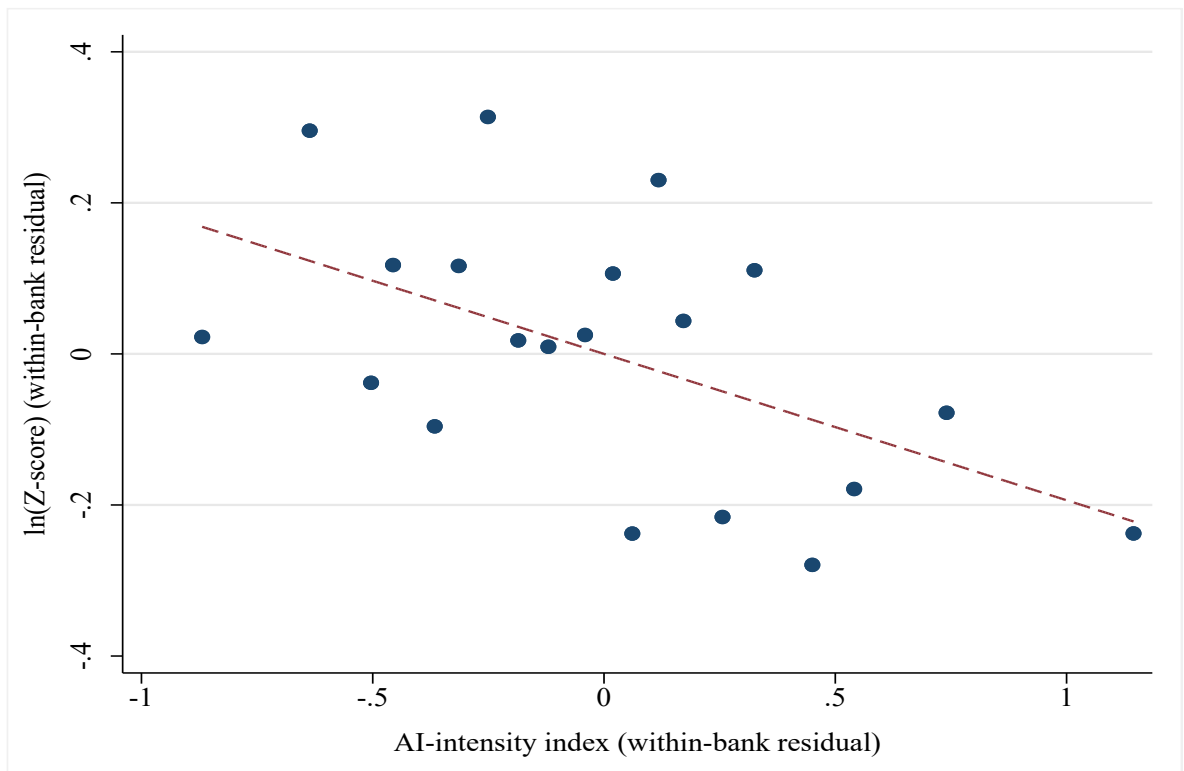


Figure 5

Predictive Margins

Figure 5 illustrates the estimated interaction effects by plotting predictive margins based on the fixed-effects specifications in Table 13. For each moderator - interbank funding dependence, exchange liquidity, deposit growth, and provisioning capacity - the figure traces the fitted values of the predicted  $\ln(Z\text{-score})$  across the moderator's empirical distribution at different levels of AI intensity. Three curves are shown, corresponding to the *AI-intensity index (AI)* at its mean and one standard deviation below/above the mean (*AI Mean - 1 SD*, *AI Mean*, *AI Mean + 1 SD*). All other covariates are held constant and absorbing bank and year fixed effects.

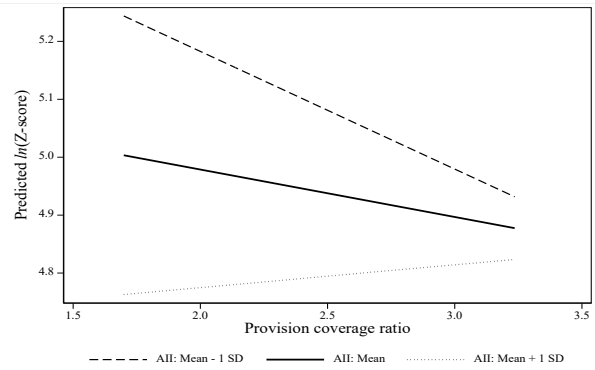
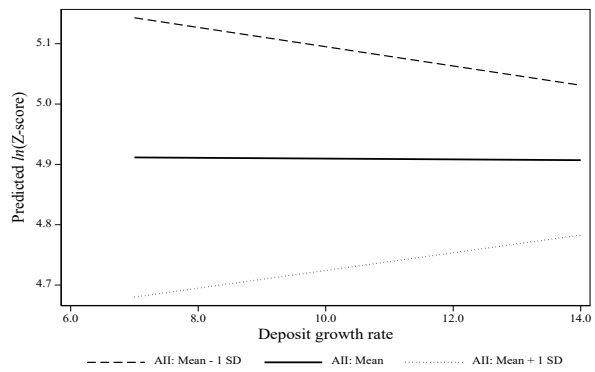
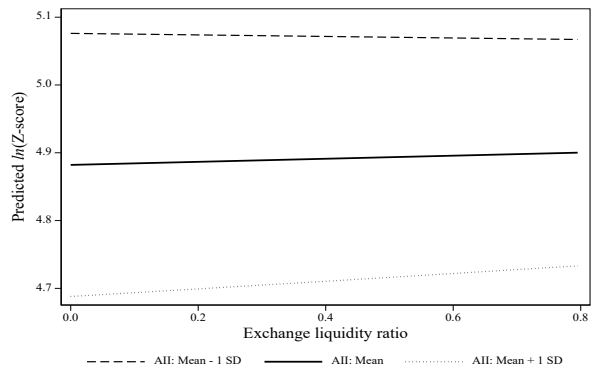
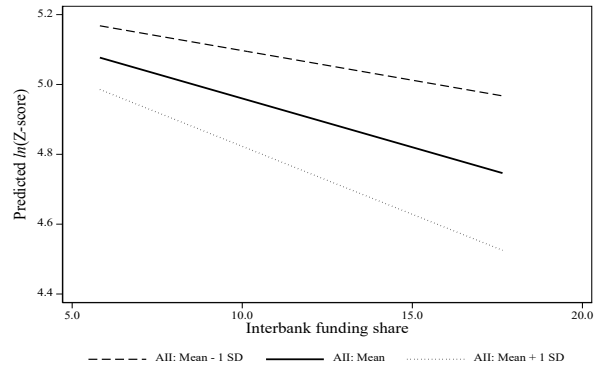


Table 1  
Descriptive Statistics

Variables	Obs	Mean	Std. dev.	Min	Max
<b>A. Bank Stability Measures</b>					
Ln(Z-score)	227	4.96	0.77	3.58	7.33
ROA	227	0.86	0.20	0.29	1.47
Standard deviation of ROA	227	0.07	0.04	0.01	0.21
Leverage ratio	227	6.63	0.89	4.84	9.24
Ln(Z-score) 3 years	222	5.16	0.80	3.63	7.38
Ln(Z-score) 5 years	227	4.83	0.75	3.52	7.40
Ln(Z-score) 8 years	227	4.63	0.74	3.42	7.14
<b>B. AI Adoption and Intensity Indices</b>					
AI-intensity index	227	0.07	0.72	-1.78	2.03
AI-intensity index total	227	0.31	1.97	-4.00	6.08
AI-intensity index PCA	168	-0.02	1.30	-2.99	3.22
AI-intensity index count	227	5.17	3.11	1.00	19.67
AI investments to total assets	212	0.06	0.04	0.01	0.21
Ln(AI investment stock)	212	21.04	1.51	17.30	24.13
Digital infrastructure proxy	227	0.33	0.75	-3.05	1.67
<b>C. AI Disclosure Subcomponents</b>					
Z <sub>Ln_wordfreq_Agent</sub>	198	0.10	0.99	-1.36	3.60
Z <sub>Ln_wordfreq_Cloud</sub>	185	-0.02	0.95	-0.87	3.06
Z <sub>Ln_wordfreq_Data</sub>	226	0.24	0.88	-1.78	2.34
Z <sub>Ln_wordfreq_Blockchain</sub>	42	0.10	1.10	-0.67	3.57
<b>D. Baseline Bank-Level Controls</b>					
Ln(Total assets)	227	28.77	1.48	25.45	31.52
Tier 1 ratio	227	11.31	1.50	8.59	15.64
Net loans to total assets	227	49.77	7.38	29.55	64.48
Non-performing loans ratio	220	1.32	0.36	0.02	2.39
Liquidity to total assets	227	7.64	3.02	3.38	17.88
Deposits to total liabilities	227	77.36	7.92	59.49	97.60
Cost to income ratio	227	29.79	6.12	12.38	64.82
<b>E. Additional Bank-Level Controls</b>					
Risk-weighted assets to earning assets	224	0.75	0.08	0.56	1.00
Loans growth ratio	227	0.13	0.07	-0.02	0.44
Top 5 shareholding concentration	227	59.82	23.61	17.68	97.38
Herfindahl–Hirschman index	227	15.33	15.80	0.64	54.39
Independent directors share	227	38.55	4.95	16.67	55.56
<b>F. Province-Level Controls</b>					
Ln(GDP per capita)	227	11.33	0.38	10.63	12.34
Inflation	227	101.48	0.95	99.70	103.20
Tertiary industry share in GDP	227	54.01	6.28	43.50	85.30
Ln(Average wage financial sector)	227	11.94	0.31	11.44	13.00
Telecommunication services growth ratio	227	0.12	0.49	-0.91	1.62
Internet broadband access ports growth ratio	227	0.07	0.10	-0.04	0.81
<b>G. Moderating Variables</b>					
Interbank deposits ratio	225	11.54	8.69	0.01	41.55
Exchange liquidity ratio	227	0.80	3.25	0.00	34.26
Deposit growth ratio	227	10.33	6.82	-7.00	44.00
Provision coverage ratio	227	2.66	1.16	1.32	5.68
<b>H. Instruments and Threshold Indicators</b>					
AI-intensity peers index	186	0.06	0.52	-1.26	0.97
AI firms growth ratio	197	6.30	37.01	-152.67	110.15
New product development intensity growth ratio	227	2.65	6.96	-37.38	26.01
R&D personnel intensity growth ratio	227	4.67	9.55	-13.47	31.06
AI high tercile	227	0.40	0.49	0.00	1.00
AI middle tercile	227	0.38	0.49	0.00	1.00
AI top quartile	227	0.31	0.46	0.00	1.00

Table 2

**Within- and Between-Bank Variation**

Table 2 reports the decomposition of total variation into within-bank and between-bank components for the main variables used in the empirical analysis. Total, within-bank, and between-bank standard deviations are computed using the panel variance decomposition implied by the fixed-effects estimator. The percentage of within-bank variation is calculated as the ratio of within-bank variance to total variance. The sample consists of 227 bank-year observations for 35 publicly listed Chinese banks and corresponds to the baseline estimation sample.

Variables	Total SD	Within-bank SD	Between-bank SD	% Within
Ln(Z-score)	0.77	0.60	0.66	44.81
AI-intensity index	0.72	0.60	0.47	61.26
$Z_{Ln\_wordfreq\_Agent}$	0.99	0.74	0.74	50.28
$Z_{Ln\_wordfreq\_Cloud}$	0.95	0.82	0.54	69.62
$Z_{Ln\_wordfreq\_Data}$	0.88	0.71	0.58	60.10
Ln(Total assets)	1.48	0.27	1.48	3.30
Tier 1 ratio	1.50	0.86	1.29	30.92
Net loans to total assets	7.38	3.88	5.91	30.08
Non-performing loans ratio	0.36	0.20	0.30	31.19
Liquidity to total assets	3.02	2.19	2.08	52.67
Deposits to total liabilities	7.92	3.04	8.04	12.53
Cost to income ratio	6.12	2.21	6.91	9.24

Table 3

**Correlation Matrix**

Table 3 reports Pearson pairwise correlations for the main variables used in the baseline regressions. All variables are measured at the bank-year level and are defined in Appendix 1. The sample consists of 227 bank-year observations. \*\*\*, \*\*, and \* denote statistical significance at the 1%, 5%, and 10% levels, respectively. The estimation sample corresponds to the baseline specification.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
(1) Ln(Z-score)	1											
(2) AI-intensity index	-0.02	1										
(3) $Z_{Ln\_wordfreq\_Agent}$	-0.01	0.80*	1									
(4) $Z_{Ln\_wordfreq\_Cloud}$	-0.05	0.70*	0.29*	1								
(5) $Z_{Ln\_wordfreq\_Data}$	-0.04	0.80*	0.48*	0.31*	1							
(6) Ln(Total assets)	0.00	0.38*	0.46*	0.25*	0.20*	1						
(7) Tier 1 ratio	0.34*	0.18*	0.17	0.14	0.08	0.32*	1					
(8) Net loans to total assets	0.03	0.09	0.16	0.12	-0.08	0.34*	0.46*	1				
(9) Non-performing loans ratio	-0.18*	0.21*	0.15	0.08	0.17	0.28*	0.03	0.23*	1			
(10) Liquidity to total assets	-0.16	-0.03	0.06	0.04	-0.07	0.37*	-0.05	-0.08	0.13	1		
(11) Deposits to total liabilities	-0.03	0.06	0.28*	0.13	-0.1	0.52*	0.28*	0.36*	0.09	0.57*	1	
(12) Cost to income ratio	0.15	0.05	0.02	0.07	0.06	-0.03	0.05	0.09	-0.28*	0.04	0.19*	1

Table 4

## Main Results

Table 4 reports panel regressions of bank stability on the lagged *AI-intensity index*. The dependent variable is the  $\ln(Z\text{-score})$  computed using a four-year rolling window for ROA volatility (excluding the current year). The *AI-intensity index* is a standardized index constructed from log word-frequency measures, as detailed in Appendix 2. All right-hand-side variables are lagged one year. Column (1) estimates a bank fixed-effects model including lagged *AI-intensity index* and year fixed effects only. Column (2) adds the baseline bank controls— $\ln(\text{Total assets})$ , *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio*—while retaining bank and year fixed effects. Columns (3)–(4) report random-effects specifications with the same controls and year fixed effects; column (4) additionally includes province fixed effects. The estimation sample consists of 227 bank-year observations for 35 publicly listed Chinese banks over 2013–2024. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	(1)	(2)	(3)	(4)
	FE: AI only Ln(Z-score)	FE: Baseline + controls Ln(Z-score)	RE: Baseline + controls Ln(Z-score)	RE: + province FE Ln(Z-score)
AI-intensity index	-0.264*** (-3.696)	-0.253*** (-4.080)	-0.189*** (-3.131)	-0.221*** (-3.650)
Ln(Total assets)		-0.119 (-0.185)	0.126* (1.705)	-0.028 (-0.323)
Tier 1 ratio		0.187** (2.432)	0.215*** (3.956)	0.206*** (4.284)
Net loans to total assets		-0.061*** (-3.127)	-0.039*** (-3.097)	-0.044*** (-3.350)
Non-performing loans ratio		0.073 (0.356)	-0.018 (-0.101)	-0.008 (-0.043)
Liquidity to total assets		0.031 (0.588)	0.029 (0.629)	0.057 (1.236)
Deposits to total liabilities		-0.051** (-2.669)	-0.030** (-1.995)	-0.025* (-1.806)
Cost to income ratio		0.004 (0.227)	0.029** (2.136)	0.030*** (3.054)
Constant	4.577*** (14.591)	12.418 (0.641)	1.990 (1.215)	6.089** (2.267)
Bank FE	YES	YES	NO	NO
Province FE	NO	NO	NO	YES
Year FE	YES	YES	YES	YES
SE clustered (bank)	YES	YES	YES	YES
Observations	227	227	227	227
Banks	35	35	35	35
R-squared (within)	0.287	0.401	0.375	0.375
Economic effect: $(\exp(\beta_{AI\text{-intensity index}})-1)$	-0.23	-0.22	-0.17	-0.20
Economic effect: $\beta_{AI\text{-intensity index}} \times (Stdev_X / Stdev_Y)$	-0.25	-0.24	-0.18	-0.21

Table 5

## AI Subcomponents and Enabling Technologies

Table 5 reports bank fixed-effects regressions of  $\ln(Z\text{-score})$  on lagged AI-related subcomponents and enabling-technology measures. The dependent variable is the baseline  $\ln(Z\text{-score})$  constructed using a four-year rolling window for *ROA volatility* (excluding the current year). Each column includes one lagged standardized component—agent adoption, cloud computing, big data, blockchain—measured as the standardized log word-frequency of the corresponding technology term (as detailed in Appendix 2). All right-hand-side variables are lagged one year. All specifications include the baseline bank-level controls -  $\ln(\text{Total assets})$ , *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio*, bank fixed effects, and year fixed effects. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. The sample is an unbalanced panel of publicly listed Chinese banks. The number of observations varies across columns because disclosure of specific technology terms is not uniform across banks and years. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	(1)	(2)	(3)	(4)
	Agent adoption Ln(Z-score)	Cloud computing Ln(Z-score)	Big data Ln(Z-score)	Blockchain Ln(Z-score)
$Z_{Ln\_wordfreq\_Agent}$	-0.219*** (-3.376)			
$Z_{Ln\_wordfreq\_Cloud}$		-0.117** (-2.518)		
$Z_{Ln\_wordfreq\_Data}$			-0.134** (-2.202)	
$Z_{Ln\_wordfreq\_Blockchain}$				-0.051 (-0.615)
Constant	25.569 (0.950)	35.311 (0.875)	6.896 (0.338)	443.195*** (3.672)
Initial controls	YES	YES	YES	YES
Bank FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
SE clustered (bank)	YES	YES	YES	YES
Observations	185	187	225	38
Banks	33	32	35	20
R-squared (within)	0.434	0.457	0.377	0.751
<i>Economic effect: <math>(\exp(\beta_{AI\text{-intensity index}})-1)</math></i>	-0.20	-0.11	-0.13	-0.05
<i>Economic effect: <math>\beta_{AI\text{-intensity index}} \times (Stdev_X / Stdev_Y)</math></i>	-0.28	-0.14	-0.15	-0.07

Table 6

**Channels: Z-score Decomposition**

Table 6 examines the effects of AI intensity on each channel by decomposing the *Z-score* into its underlying components and re-estimating the baseline specification at the bank-year level. The dependent variables are: (1) *ROA volatility* (the standard deviation of ROA computed over a four-year rolling window excluding the current year), (2) the *ROA level*, and (3) the *Equity to assets ratio*. The main regressor is the standardized AI-intensity index detailed in Appendix 2. All specifications include the baseline bank-level controls - *Ln(Total assets)*, *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio*, bank fixed effects, and year fixed effects. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. The sample is an unbalanced panel of 35 publicly listed Chinese banks over 2009–2024, yielding 227 bank-year observations in the estimation sample. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	(1) ROA volatility (Standard deviation of ROA)	(2) ROA level	(3) Equity to total assets
AI-intensity index	0.012*** (3.071)	-0.022** (-2.209)	-0.027 (-0.469)
Ln(Total assets)	-0.023 (-0.727)	0.339*** (3.957)	-0.882** (-2.231)
Tier 1 ratio	-0.009** (-2.366)	-0.011 (-1.086)	0.172*** (4.323)
Net loans to total assets	0.003*** (2.946)	0.003 (0.676)	0.044** (2.500)
Non-performing loans ratio	0.000 (0.023)	-0.080 (-1.691)	0.177 (1.382)
Liquidity to total assets	-0.001 (-0.285)	-0.012 (-1.322)	-0.012 (-0.504)
Deposits to total liabilities	0.002* (1.914)	0.004 (1.158)	-0.035*** (-3.063)
Cost to income ratio	0.001 (0.496)	0.003 (0.734)	-0.018 (-1.362)
Constant	0.505 (0.548)	-8.454*** (-3.392)	29.703** (2.674)
Bank FE	YES	YES	YES
Year FE	YES	YES	YES
SE clustered (bank)	YES	YES	YES
Observations	227	227	227
Banks	35	35	35
R-squared (within)	0.409	0.771	0.696

Table 7

**Alternative Z-score windows**

Table 7 reports fixed-effects panel regressions of bank stability on lagged *AI-intensity index* using alternative constructions of the dependent variable. The dependent variable is the *ln(Z-score)* computed using *ROA volatility* estimated over a 3-year, 5-year, or 8-year rolling window (each excluding the current year). AI intensity is measured by the *AI-intensity index*, as detailed in Appendix 2. All specifications include bank fixed effects and year fixed effects. Control variables are lagged one period and include *Ln(Total assets)*, *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio* (definitions in Appendix 1). Standard errors are clustered at the bank level. t-statistics are reported in parentheses. The panel is unbalanced. \*\*\*, \*\*, and \* denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Variables	(1) Ln(Z-score) 3 years	(2) Ln(Z-score) 5 years	(3) Ln(Z-score) 8 years
AI-intensity index	-0.203** (-2.438)	-0.190*** (-3.761)	-0.175*** (-3.762)
Ln(Total assets)	-0.334 (-0.493)	0.276 (0.523)	0.706 (1.684)
Tier 1 ratio	0.123 (1.669)	0.237*** (3.672)	0.185*** (3.329)
Net loans to total assets	-0.052** (-2.405)	-0.052** (-2.498)	-0.027 (-1.195)
Non-performing loans ratio	-0.090 (-0.452)	0.151 (0.684)	0.215 (1.233)
Liquidity to total assets	-0.020 (-0.271)	0.006 (0.154)	0.016 (0.407)
Deposits to total liabilities	-0.036 (-1.422)	-0.036** (-2.388)	-0.023** (-2.198)
Cost to income ratio	0.013 (0.586)	0.008 (0.388)	0.015 (1.103)
Constant	17.956 (0.925)	-0.409 (-0.026)	-14.305 (-1.178)
Year FE	YES	YES	YES
Bank FE	YES	YES	YES
SE clustered (bank)	YES	YES	YES
Observations	222	229	230
Banks	35	35	35
R-squared (within)	0.384	0.356	0.323

Table 8

**Alternative Methodologies for the AI-intensity Index**

Table 8 reports fixed-effects panel regressions of bank stability on alternative methodologies for the *AI-intensity index*. The dependent variable is  $\ln(Z\text{-score})$  computed using the baseline rolling window reported in Table 4. The key explanatory variable is lagged AI intensity, measured alternatively as: (1) a standardized total-count index (*AI-intensity index total*), (2) a principal component-based index (*AI-intensity index PCA*), and (3) a raw word-count index (*AI-intensity index count*), as detailed in Appendix 2. Each specification includes the same set of lagged bank controls as in Table 4 ( $\ln(\text{Total assets})$ , *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio*) and absorbs bank fixed effects and year fixed effects. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. The panel is unbalanced, and the number of observations may vary across columns if an index requires additional data availability (e.g., PCA construction). \*\*\*, \*\*, and \* denote statistical significance at the 1%, 5%, and 10% levels respectively

Variables	(1) Total (std)	(2) PCA	(3) Raw count
AI-intensity index total	-0.099*** (-4.316)		
AI-intensity index PCA		-0.162*** (-3.043)	
AI-intensity index count			-0.057*** (-3.306)
Ln(Total assets)	-0.188 (-0.299)	-0.757 (-0.533)	-0.134 (-0.206)
Tier 1 ratio	0.189** (2.563)	0.101 (1.281)	0.187** (2.342)
Net loans to total assets	-0.062*** (-3.270)	-0.075*** (-2.912)	-0.061*** (-3.060)
Non-performing loans ratio	0.087 (0.429)	-0.246 (-0.778)	0.075 (0.359)
Liquidity to total assets	0.024 (0.462)	0.091 (1.130)	0.034 (0.569)
Deposits to total liabilities	-0.046** (-2.349)	-0.008 (-0.288)	-0.049** (-2.516)
Cost to income ratio	0.004 (0.207)	-0.004 (-0.178)	0.003 (0.146)
Constant	14.221 (0.752)	27.791 (0.655)	13.088 (0.670)
Year FE	YES	YES	YES
Bank FE	YES	YES	YES
SE clustered (bank)	YES	YES	YES
Observations	227	164	227
Banks	35	31	35
R-squared (within)	0.408	0.509	0.391

Table 9

## Alternative AI proxies

Table 9 reports bank fixed-effects regressions of next-year bank stability on alternative proxies for AI adoption and digitalization. The dependent variable is  $\ln(Z\text{-score})$ . Column (1) uses *AI investments to total assets*. Column (2) uses the *log of AI investment stock*. Column (3) uses a *Digital infrastructure proxy*, defined as the standardized value of  $\ln(\text{wordfreq\_digital})$  for positive counts. All main regressors are lagged one year. All specifications include bank and year fixed effects and control for lagged bank characteristics: *Ln(Total assets)*, *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio*. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. The sample is restricted to observations with `estimation_sample = 1`. \*\*\*, \*\*, and \* denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Variables	(1)	(2)	(3)
	AI inv. intensity (%) Ln(Z-score)	AI inv. stock (ln) Ln(Z-score)	Digital proxy (z) Ln(Z-score)
AI investments to total assets	-4.209** (-2.047)		
Ln(AI investment stock)		-0.260** (-2.335)	
Digital infrastructure proxy			-0.107 (-1.195)
Ln(Total assets)	0.143 (0.213)	0.302 (0.440)	0.116 (0.176)
Tier 1 ratio	0.229*** (2.953)	0.222*** (2.910)	0.210** (2.569)
Net loans to total assets	-0.069*** (-3.502)	-0.067*** (-3.329)	-0.068*** (-3.264)
Non-performing loans ratio	-0.019 (-0.084)	-0.012 (-0.056)	0.003 (0.013)
Liquidity to total assets	0.122 (1.553)	0.115 (1.521)	0.026 (0.420)
Deposits to total liabilities	-0.056** (-2.511)	-0.052** (-2.324)	-0.048** (-2.384)
Cost to income ratio	0.009 (0.386)	0.009 (0.384)	0.007 (0.346)
Constant	4.823 (0.241)	5.039 (0.249)	6.019 (0.305)
Year FE	YES	YES	YES
Bank FE	YES	YES	YES
SE clustered (bank)	YES	YES	YES
Observations	211	211	227
Banks	33	33	35
R-squared (within)	0.370	0.374	0.370

Table 10

## Additional Controls

Table 10 reports fixed-effects regressions of bank stability on AI intensity augmented with extra controls to pre-empt omitted-variables concerns. The dependent variable is *ln(Z-score)* computed over a four-year rolling window. The key regressor is the *AI-intensity index*, a standardized disclosure-based AI intensity measure, detailed in the Appendix 2. The sample comprises 35 publicly listed Chinese banks over 2009–2024. All regressors are lagged one year, and all specifications include year fixed effects. Columns (1)–(7) add one additional control at a time to the baseline specification: (1) *Risk-weight-assets-to-earning-assets* to capture shifts in regulatory risk weights and portfolio risk intensity; (2) *Loan growth ratio* to proxy rapid balance-sheet expansion and risk-taking; (3) *Top-5 ownership concentration* to capture monitoring incentives and controlling-shareholder influence; (4) *Herfindahl–Hirschman index* based on the top five institutions to measure local market concentration/competition; (5) *Independent directors share* as a corporate-governance proxy; (6) regional *ln(GDP per capita)* to capture local economic conditions; and (7) *inflation (CPI)* to capture macroeconomic pressure and business-cycle risk. Column (8) includes all selected controls jointly. Column (9) replaces bank fixed effects with bank-nature fixed effects (e.g., state-owned/joint-stock/city/rural) to absorb time-invariant differences across broad bank types while preserving additional cross-sectional variation. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Risk-weighted assets to earning assets	Loans growth ratio	Top 5 shareholding concentration	Herfindahl–Hirschman index	Independent directors share	Ln(GDP per capita)	Inflation	All selected	Bank nature FE
Variables	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)
AI-intensity index	-0.260*** (-3.896)	-0.256*** (-4.380)	-0.266*** (-4.081)	-0.262*** (-4.372)	-0.253*** (-4.035)	-0.263*** (-3.962)	-0.246*** (-3.935)	-0.279*** (-4.134)	-0.194*** (-3.280)
Risk-weighted assets to earning assets	0.387 (0.389)							0.167 (0.161)	
Loans growth ratio		-0.352 (-0.409)						0.015 (0.016)	
Top 5 shareholding concentration			0.033** (2.494)					0.029 (1.611)	
Herfindahl–Hirschman index				0.034 (1.307)				0.015 (0.516)	
Independent directors share					-0.002 (-0.130)			-0.004 (-0.278)	
Ln(GDP per capita)						1.761 (0.650)		1.762 (0.666)	
Inflation							0.160 (0.914)	0.107 (0.581)	
Constant	10.185 (0.467)	11.407 (0.581)	11.800 (0.632)	13.820 (0.726)	12.526 (0.631)	-10.724 (-0.265)	-3.238 (-0.116)	-24.843 (-0.519)	1.286 (0.286)
Initial controls	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Bank FE	YES	YES	YES	YES	YES	YES	YES	YES	NO
Banknature FE	NO	NO	NO	NO	NO	NO	NO	NO	YES
SE clustered (bank)	YES	YES	YES	YES	YES	YES	YES	YES	YES
Observations	224	227	227	227	227	227	227	224	227
Banks	35	35	35	35	35	35	35	35	35
R-squared (within)	0.409	0.401	0.414	0.407	0.401	0.403	0.403	0.427	0.385

Table 11

## Endogeneity Analysis: IV2SLS

Table 11 reports panel IV2SLS (fixed-effects) estimates of bank stability on AI intensity. The dependent variable is  $\ln(Z\text{-score})$  over a four-year rolling window, and the endogenous regressor is the one-year-lagged standardized *AI-intensity index* (as detailed in Appendix 2). AI intensity is instrumented with one-year-lagged peer AI disclosure and province-level AI innovation/supply growth proxies (*AI-intensity peers index*, *AI firms growth ratio*, *New product development intensity growth ratio*, *R&D personnel intensity growth ratio*). The table reports both first-stage and second-stage results for 35 listed Chinese banks (2009–2024), with bank and year fixed effects and robust bank-level inference. Columns (1)–(3) progressively expand controls: (1) baseline bank controls (*Ln(Total assets)*, *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio*); (2) adds local macro/structure controls (*Ln(GDP per capita)*, *tertiary industry share*, *inflation*, and *Ln(Average wage financial sector)*); and (3) further adds local digital-infrastructure growth controls (*Telecommunication services growth ratio* and *Internet broadband access ports growth ratio*). Standard errors are clustered at the bank level. *t*-statistics are reported in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	(1)		(2)		(3)	
	1 <sup>st</sup> stage	2 <sup>nd</sup> stage	1 <sup>st</sup> stage	2 <sup>nd</sup> stage	1 <sup>st</sup> stage	2 <sup>nd</sup> stage
	AI-intensity index	Ln(Z-score)	AI-intensity index	Ln(Z-score)	AI-intensity index	Ln(Z-score)
AI-intensity peers index	-0.912*** (-4.214)		-0.977*** (-4.322)		-0.992*** (-4.485)	
AI firms growth ratio	0.301*** (2.739)		0.313*** (2.778)		0.321*** (2.889)	
New product development intensity growth ratio	0.054*** (3.784)		0.049*** (3.116)		0.052*** (3.193)	
R&D personnel intensity growth ratio	0.021*** (4.023)		0.019*** (3.454)		0.020*** (3.471)	
AI-intensity index		-0.322** (-2.109)		-0.335** (-2.098)		-0.526*** (-3.408)
Ln(Total assets)	-0.795 (-1.250)	-1.763** (-2.104)	0.061 (0.063)	-1.568 (-1.410)	0.025 (0.025)	-1.417 (-1.255)
Tier 1 ratio	-0.107 (-1.268)	0.150* (1.785)	-0.098 (-1.149)	0.140* (1.670)	-0.092 (-1.055)	0.098 (1.105)
Net loans to total assets	0.004 (0.178)	-0.023 (-0.860)	0.009 (0.398)	-0.016 (-0.548)	0.007 (0.321)	-0.010 (-0.347)
Non-performing loans ratio	0.604*** (3.935)	0.112 (0.604)	0.583*** (3.703)	0.107 (0.574)	0.591*** (3.827)	0.258 (1.228)
Liquidity to total assets	-0.019 (-0.357)	0.019 (0.354)	-0.008 (-0.145)	0.029 (0.497)	-0.006 (-0.107)	0.015 (0.266)
Deposits to total liabilities	0.005 (0.313)	-0.064*** (-3.552)	0.008 (0.475)	-0.060*** (-3.215)	0.006 (0.362)	-0.056*** (-2.957)
Cost to income ratio	-0.007 (-0.437)	-0.016 (-0.710)	-0.004 (-0.198)	-0.016 (-0.743)	-0.001 (-0.070)	-0.023 (-1.094)
Ln(GDP per capita)			4.391 (0.871)	2.360 (0.584)	5.512 (0.992)	-0.911 (-0.207)
Tertiary industry share in GDP			-0.082 (-0.465)	0.127 (0.895)	-0.053 (-0.277)	0.016 (0.110)
Inflation			0.034 (0.139)	0.192 (0.901)	0.027 (0.107)	0.271 (1.167)
Ln(Average wage financial sector)			-1.345 (-0.885)	-0.930 (-0.718)	-1.468 (-1.000)	-0.831 (-0.621)
Telecommunication services growth ratio					-0.579 (-0.528)	1.992** (2.252)
Internet broadband access ports growth ratio					-0.032 (-0.035)	-1.079 (-1.180)
Year FE	YES	YES	YES	YES	YES	YES
Bank FE	YES	YES	YES	YES	YES	YES
SE clustered (bank)	YES	YES	YES	YES	YES	YES
Observations		178		178		178
Banks		28		28		28
R-squared		0.430		0.434		0.416
<i>Underid LM Stat</i>		23.55		24.41		22.30
<i>Underid p-val</i>		0.00		0.00		0.00
<i>Weak Id F Stat</i>		10.85		10.99		11.91
<i>SY 10% LIML CV</i>		5.44		5.44		5.44
<i>Hansen J Stat</i>		3.30		2.60		2.48
<i>J p-val</i>		0.35		0.46		0.48

Table 12

**AI Intensity Thresholds and Bank Stability**

Table 12 reports panel regressions of bank stability on variables capturing AI-intensity thresholds (as detailed in Appendix 4). The dependent variable is the  $\ln(\text{Z-score})$ . Columns (1)–(2) use tercile classifications: AI high tercile equals one if the bank-year observation falls in the top tercile of the pooled AI-intensity distribution, and AI mid tercile equals one if it falls in the middle tercile (the omitted category is the bottom tercile). Columns (3)–(4) use a top-quartile indicator: AI top quartile equals one if the observation falls in the highest quartile (omitted category: the bottom three quartiles). All specifications include the baseline lagged control set— $\ln(\text{Total assets})$ ,  $\text{Tier 1 ratio}$ ,  $\text{Net loans to total assets}$ ,  $\text{Non-performing loans ratio}$ ,  $\text{Liquidity to total assets}$ ,  $\text{Deposits to total liabilities}$ ,  $\text{Cost to income ratio}$ —and year fixed effects. Columns (1) and (3) report random-effects estimates; columns (2) and (4) report bank fixed-effects estimates. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. The panel is unbalanced and covers 35 Chinese commercial banks over 2009–2024. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
	Terciles (RE)	Terciles (FE)	Top quartile (RE)	Top quartile (FE)
Variables	$\ln(\text{Z-score})$	$\ln(\text{Z-score})$	$\ln(\text{Z-score})$	$\ln(\text{Z-score})$
AI high tercile	-0.261** (-2.448)	-0.351*** (-3.187)		
AI middle tercile	-0.229** (-2.326)	-0.246** (-2.120)		
AI top quartile			-0.209 (-1.640)	-0.256** (-2.104)
Initial controls	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Bank FE	YES	YES	YES	YES
SE clustered (bank)	YES	YES	YES	YES
Observations	227	227	227	227
Banks	35	35	35	35
R-squared within	0.367	0.388	0.358	0.383

Table 13

**Moderating Factors**

Table 13 reports fixed-effects regressions that test moderating factors by interacting the one-year-lagged *AI-intensity index* with one-year-lagged moderator variables (*AI-intensity index* \* *Moderator factor*). The dependent variable is  $\ln(Z\text{-score})$  computed over a four-year rolling window. The sample is restricted to the baseline estimation sample of 35 banks from China. Columns (1)–(4) use different moderators: (1) *Interbank deposits ratio*, (2) *Exchange liquidity ratio*, (3) *Deposit growth rate*, and (4) *Provision coverage ratio*. All specifications include lagged baseline bank controls, bank fixed effects, and year fixed effects; standard errors are clustered by bank. t-statistics are reported in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
	Interbank deposits ratio	Exchange liquidity ratio	Deposit growth ratio	Provision coverage ratio
Variables	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)	Ln(Z-score)
AI-intensity index	-0.112 (-1.131)	-0.286*** (-4.503)	-0.486*** (-4.608)	-0.583*** (-4.495)
Interbank deposits ratio	-0.018 (-0.970)			
AI-intensity index * Interbank deposits ratio	-0.011** (-2.176)			
Exchange liquidity ratio		-0.002 (-0.112)		
AI-intensity index * Exchange liquidity ratio		0.082*** (3.009)		
Deposit growth ratio			0.002 (0.291)	
AI-intensity index * Deposit growth ratio			0.021** (2.712)	
Provision coverage ratio				-0.079 (-1.025)
AI-intensity index * Provision coverage ratio				0.146*** (2.988)
Initial controls	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Bank FE	YES	YES	YES	YES
SE clustered (bank)	YES	YES	YES	YES
Observations	226	227	227	227
Banks	35	35	35	35
R-squared (within)	0.407	0.408	0.424	0.423

## Appendix 1

## Variable Definitions

Note: AC denotes the Authors' Computation. CSMAR denotes the China Stock Market & Accounting Research Database. CSMAR-CNBank denotes the CSMAR China Bank Research Dataset; CSMAR-FinTech denotes the CSMAR FinTech Research Dataset; CSMAR-Eco denotes the CSMAR China Regional Economy Research Dataset; NBS denotes the National Bureau of Statistics of China.

Variable	Unit	Definition	Source
<b>A. Bank Stability Measures</b>			
Ln(Z-score)	Log points	Natural log of the Z-score computed as $Z\text{-score} = (ROA + Equity/Assets) / \sigma(ROA)$ , where $\sigma(ROA)$ is the rolling standard deviation of ROA over years $t-4$ to $t-1$ .	AC using CSMAR-CNBank
ROA	Percent	Return on assets.	CSMAR-CNBank
Standard deviation of ROA	Percent	Rolling standard deviation of ROA over years $t-4$ to $t-1$ .	AC using CSMAR-CNBank
Leverage ratio	Percent	Equity-to-assets ratio (capital buffer).	CSMAR-CNBank
Ln(Z-score) 3 years	Log points	Natural log of the Z-score using $\sigma(ROA)$ computed over years $t-3$ to $t-1$ .	AC using CSMAR-CNBank
Ln(Z-score) 5 years	Log points	Natural log of the Z-score using $\sigma(ROA)$ computed over years $t-5$ to $t-1$ .	AC using CSMAR-CNBank
Ln(Z-score) 8 years	Log points	Natural log of the Z-score using $\sigma(ROA)$ computed over years $t-8$ to $t-1$ .	AC using CSMAR-CNBank
<b>B. AI Adoption and Intensity Indices</b>			
AI-intensity index	Log points	Mean-based AI-intensity index: average of standardized $\ln(1+\text{count})$ disclosure measures for agent adoption, cloud computing, and big data.	AC using CSMAR-FinTech
AI-intensity index total	Log points	Total-count AI-intensity index: sum of standardized $\ln(1+\text{count})$ disclosure measures for agent adoption, cloud computing, and big data.	AC using CSMAR-FinTech
AI-intensity index PCA	Principal component score	First principal component of standardized $\ln(1+\text{count})$ disclosure measures for agent adoption, cloud computing, and big data.	AC using CSMAR-FinTech
AI-intensity index count	Counts (mean)	Mean of raw disclosure word-frequency counts for agent adoption, cloud computing, and big data.	AC using CSMAR-FinTech
AI investments to total assets	Percent	AI investment intensity, measured as $100 \times (FirmAIinvestlevel)$ , where $FirmAIinvestlevel$ is the ratio of AI-related assets to total assets ( $AI\ Assets / Total\ Assets$ ).	CSMAR-FinTech
Ln(AI investment stock)	Log points	Natural log of the stock of AI investment: $\ln(FirmAIinvesttotal)$ , where $FirmAIinvesttotal$ is the total capital invested in AI technologies (sum of AI-related intangible assets and fixed assets).	AC using CSMAR-FinTech
Digital infrastructure proxy	Log points	Standardized $\ln(1+\text{count})$ of digital keywords occurrences in bank disclosures. The dictionary includes the following terms: "digital currency, smart contract, distributed computing, decentralization, bitcoin, alliance chain, differential privacy technology, and consensus mechanism."	AC using CSMAR-FinTech
<b>C. AI Disclosure Subcomponents</b>			
$Z_{La\_wordfreq\_Agent}$	Log points	Standardized $\ln(1+\text{count})$ of AI keyword occurrences in bank disclosures. The dictionary includes the following terms: "artificial intelligence, business intelligence, image understanding, investment decision support systems, intelligent data analysis, intelligent robots, machine learning, deep learning, semantic search, biometrics, face recognition, speech recognition, identity verification, autonomous driving, and natural language processing."	AC using CSMAR-FinTech
$Z_{La\_wordfreq\_Cloud}$	Log points	Standardized $\ln(1+\text{count})$ of cloud-computing keyword occurrences in bank disclosures. The dictionary includes the following terms: "cloud computing, stream computing, graph computing, memory computing, multi-party secure computing, brain-inspired computing, green computing, cognitive computing, fusion architecture, billion-level concurrency, EB-level storage, internet of things, cyber-physical systems."	AC using CSMAR-FinTech
$Z_{La\_wordfreq\_Data}$	Log points	Standardized $\ln(1+\text{count})$ of big-data keyword occurrences in bank disclosures. The dictionary includes the following terms: "big data, data mining, text mining, data visualization, heterogeneous data, credit investigation, augmented reality, mixed reality, virtual reality."	AC using CSMAR-FinTech
$Z_{La\_wordfreq\_Blockchain}$	Log points	Standardized $\ln(1+\text{count})$ of blockchain keyword occurrences in bank disclosures. The dictionary includes the following terms: "digital currency, smart contract, distributed computing, decentralization, bitcoin, alliance chain, differential privacy technology, and consensus mechanism."	AC using CSMAR-FinTech
<b>D. Baseline Bank-Level Controls</b>			
Ln(Total assets)	Log points	Natural log of total assets.	CSMAR-CNBank
Tier 1 ratio	Percent	Tier 1 capital adequacy ratio.	CSMAR-CNBank
Net loans to total assets	Percent	Net loans and advances divided by total assets.	AC using CSMAR-CNBank
Non-performing loans ratio	Percent	Nonperforming loan (NPL) balance divided by total loans.	AC using CSMAR-CNBank
Liquidity to total assets	Percent	Cash and balances with the central bank divided by total assets (liquidity buffer).	AC using CSMAR-CNBank
Deposits to total liabilities	Percent	Deposits and due to banks and other financial institutions divided by total liabilities.	AC using CSMAR-CNBank
Cost to income ratio	Percent	Cost-to-income ratio.	CSMAR-CNBank

Appendix 1  
Variable Definitions (Cont.)

Note: AC denotes the Authors' Computation. CSMAR denotes the China Stock Market & Accounting Research Database. CSMAR-CNBank denotes the CSMAR China Bank Research Dataset; CSMAR-FinTech denotes the CSMAR FinTech Research Dataset; CSMAR-Economy denotes the CSMAR China Regional Economy Research Dataset; NBS denotes the National Bureau of Statistics of China.

Variable	Unit	Definition	Source
<b>E. Additional Bank-Level Controls</b>			
Risk-weighted assets to earning assets	Percent	Risk-weighted assets divided by interest-earning assets (risk-weight density).	CSMAR-CNBank
Loans growth ratio	Percent	Loan growth ratio.	CSMAR-CNBank
Top 5 shareholding concentration	Percent	Ownership concentration: sum of equity shares held by the top five shareholders.	AC using CSMAR-CNBank
Herfindahl–Hirschman index	Index	Market concentration index based on the top five institutions' shares (Herfindahl-type measure).	AC using CSMAR-CNBank
Independent directors share	Percent	Share of independent directors on the board (number of independent directors divided by board size).	AC using CSMAR-CNBank
Bank nature	Category	Bank type classification (policy bank; state-owned; joint-stock; urban; rural; foreign-funded; other; rural cooperative; rural credit cooperative).	CSMAR-CNBank
<b>F. Province-Level Controls</b>			
GRDP per capita	RMB per person	Gross regional domestic product (GRDP) per capita.	CSMAR-Economy
CPI	Index (last year=100)	Consumer price index.	CSMAR-Economy
Tertiary industry share in GDP	Percent	Share of the tertiary industry in provincial GDP.	CSMAR-Economy
Average wage financial sector	RMB (yuan)	Average annual wage of employees in the financial sector at the provincial level.	CSMAR-Economy
Telecommunication services growth ratio	Percent	Growth rate of telecommunications service volume (province-year).	AC using NBS
Internet broadband access ports growth ratio	Percent	Growth rate of broadband access ports (province-year).	AC using NBS
<b>G. Moderating Variables</b>			
Interbank deposits ratio	Ratio	Interbank deposits received divided by total deposits and interbank liabilities.	AC using CSMAR-CNBank
Exchange liquidity ratio	Ratio	Foreign-currency liquidity ratio (exchange liquidity ratio).	CSMAR-CNBank
Deposit growth ratio	Percent	Deposit growth ratio.	CSMAR-CNBank
Provision coverage ratio	Percent	Provision coverage ratio: (general + specific + special provisions) divided by (substandard + doubtful + loss loans).	CSMAR-CNBank
<b>H. Instruments and Threshold Indicators</b>			
AI-intensity peers index	Log points	Leave-one-out average AI intensity of other banks in the same province-year. We first sum the AI-intensity index across all banks with non-missing values in a given province-year, then subtract bank $i$ 's own value and divide by the number of remaining banks in that cell: $Total(AI-intensity\ index) - AI-intensity\ index\ i,t / (AI-intensity\ peers\ index\ number\ p,t-1)$ . The measure is set to missing when fewer than two banks report non-missing AI intensity in the province-year.	AC using CSMAR-FinTech
AI firm growth ratio	Percent	Province-level growth in AI enterprise activity, measured as the percent log change in the number of registered AI firms: $100 \times [\ln(AI\ firm\ number_{p,t}) - \ln(AI\ firm\ number_{p,t-1})]$ .	AC using CSMAR-Economy
New product development intensity growth ratio	Percent	Growth in new product development (NPD) intensity at the province level. NPD intensity is defined as NPD expenditure of above-scale industrial enterprises scaled by provincial GRDP, and its growth is measured as $100 \times [\ln(1 + NPD\ intensity_{p,t}) - \ln(1 + NPD\ intensity_{p,t-1})]$ .	AC using NBS
R&D personnel intensity growth ratio	Percent	Growth in province-level R&D personnel intensity. R&D personnel intensity is measured as full-time-equivalent R&D personnel in above-scale industrial enterprises scaled by province-year population, and its growth is computed as $100 \times [\ln(1 + R\&D\ personnel\ intensity_{p,t}) - \ln(1 + R\&D\ personnel\ intensity_{p,t-1})]$ .	AC using NBS
AI high tercile	Indicator	Equals 1 if AI-intensity index is in the top tercile of the pooled distribution; 0 otherwise.	AC using CSMAR-FinTech
AI middle tercile	Indicator	Equals 1 if AI-intensity index is in the middle tercile of the pooled distribution; 0 otherwise (bottom tercile is omitted group).	AC using CSMAR-FinTech
AI top quartile	Indicator	Equals 1 if AI-intensity index is in the top quartile of the pooled distribution; 0 otherwise.	AC using CSMAR-FinTech

## Appendix 2: Construction of AI-Intensity Indices

### A2.1 AI-intensity index subcomponents (Agent, Cloud, Data)

This appendix describes the construction of the disclosure-based AI intensity measures used in the main analysis. We focus on three disclosure subcomponents—agent adoption, cloud computing, and big data—derived from bank annual-report text using the CSMAR China Financial Technology Research Database. We then aggregate these subcomponents into composite AI-intensity indices using mean, total, a principal-components approach, and a raw-count index.

For each bank  $i$  and year  $t$ , CSMAR provides word-frequency counts for three technology categories, based on keyword dictionaries applied to banks' annual-report disclosures: *wordfreq\_Agent*, *wordfreq\_Cloud*, and *wordfreq\_Data*. These variables measure the number of occurrences of category-specific terms in a bank's report in year  $t$ . Higher counts indicate a stronger emphasis on the respective technology in disclosed strategy, operations, or initiatives.

CSMAR constructs these disclosure measures using a taxonomy that maps technology-related terminology into three mutually exclusive indicator classes. Agent adoption (Agent) captures references to artificial-intelligence capabilities and applications—e.g., “*artificial intelligence, business intelligence, image understanding, investment decision support systems, intelligent data analysis, intelligent robots, machine learning, deep learning, semantic search, biometrics, face recognition, speech recognition, identity verification, autonomous driving, and natural language processing.*” Cloud computing (Cloud) captures infrastructure and computing-architecture concepts—e.g., “*cloud computing, stream computing, graph computing, memory computing, multi-party secure computing, brain-inspired computing, green computing, cognitive computing, fusion architecture, billion-level concurrency, EB-level storage, Internet of Things, and cyber-physical systems.*” Big data (Data) captures data-intensive analytics and related technologies—e.g., “*big data, data mining, text mining, data visualization, heterogeneous data, credit investigation, augmented reality, mixed reality, and virtual reality.*” This CSMAR classification links each word-frequency series to a clearly

defined technological domain, so that variation in  $wordfreq\_Agent$ ,  $wordfreq\_Cloud$ , and  $wordfreq\_Data$ , reflects differences in the intensity of banks' disclosures about agent adoption, cloud infrastructure, and data-driven analytics in their annual-report narratives.

Because word-frequency counts are right-skewed and may include zeros, we apply a log transformation,  $\ln(1 + x)$ , to each count series to reduce skewness and limit the influence of extreme observations:

$$\ln\_wordfreq_{k,it} = \ln(1 + wordfreq_{k,it}), k \in \{Agent, Cloud, Data\}.$$

Adding one in the log transformation ensures that zero-count disclosures remain in the sample.

To make coefficients comparable across disclosure dimensions ( $k$ ), we standardize each log-transformed series using its sample mean and standard deviation (computed over the estimation sample):

$$z_{\ln\_wordfreq_{k,it}} = \frac{\ln\_wordfreq_{k,it} - \mu_k}{\sigma_k}, k \in \{Agent, Cloud, Data\},$$

where  $\mu_k$  and  $\sigma_k$  denote the mean and standard deviation of  $\ln\_wordfreq_{k,it}$ .

## A2.2 Composite AI-intensity indices

To summarize banks' disclosure-based adoption of AI-related technologies, we aggregate the three standardized subcomponents, agent adoption, cloud computing, and big data, into composite indices using three complementary approaches. All three methods combine information from the same underlying standardized series:  $z_{\ln\_wordfreq\_Agent,i,t}$ ,  $z_{\ln\_wordfreq\_Cloud,i,t}$ , and  $z_{\ln\_wordfreq\_Data,i,t}$ , where each component is the standardized value of  $\ln(1 + wordfreq_{k,it})$  keyword occurrences in bank  $i$ 's disclosure in year  $t$ . Standardization places all components on a common scale, preventing any single technology from dominating the composite index due to differences in frequency or dispersion.

*Mean-based index.* Our baseline composite measure is the simple average of the three standardized components, which treats each technology dimension symmetrically

and yields an interpretable unit: a one-standard-deviation increase in the average disclosure intensity across agent adoption, cloud, and big data, as follows:

$$AI\text{-intensity index}_{i,t} = \frac{1}{3} (Z_{\ln\_wordfreq\_Agent,i,t} + Z_{\ln\_wordfreq\_Cloud,i,t} + Z_{\ln\_wordfreq\_Data,i,t})$$

This construction is robust to idiosyncratic noise in any single subcomponent and captures a broad-based AI-related technological orientation.

*Total-count index.* As an alternative that emphasizes the overall “mass” of AI-related disclosure, we also compute the corresponding additive index:

$$AI\text{-intensity index total}_{i,t} = Z_{\ln\_wordfreq\_Agent,i,t} + Z_{\ln\_wordfreq\_Cloud,i,t} + Z_{\ln\_wordfreq\_Data,i,t}$$

Relative to the mean-based index, this measure preserves the same ranking but scales the composite by the number of components, which can be useful when interpreting the aggregate level of disclosure intensity across the three domains.

*Principal-components index.* To let the data determine the most informative combination of subcomponents, we construct a principal-components index by taking the first principal component of the three standardized series and using its score as an alternative AI-intensity measure:

$$AI\text{-intensity index PCA}_{i,t} = w_1 Z_{\ln\_wordfreq\_Agent,i,t} + w_2 Z_{\ln\_wordfreq\_Cloud,i,t} + w_3 Z_{\ln\_wordfreq\_Data,i,t}$$

The weights ( $w_1, w_2, w_3$ ) correspond to the eigenvector associated with the largest eigenvalue of the correlation matrix of  $(Z_{\ln\_wordfreq\_Agent,i,t}, Z_{\ln\_wordfreq\_Cloud,i,t}, Z_{\ln\_wordfreq\_Data,i,t})$ . By construction, this index places greater weight on the components that co-move most strongly with the latent common factor underlying AI-related disclosure, and thus provides a parsimonious measure of the shared adoption intensity across the three technologies.

*AI-intensity index count.* Finally, we employ a raw-count index, defined as the mean of the untransformed disclosure frequencies for the three components.

Across specifications, we use the mean-based index as the main proxy for AI intensity, and employ the total-count, PCA, and row-count indices as robust alternatives that vary in their weighting schemes and scaling assumptions.

Appendix 3

**Main Results (winsorized)**

Appendix 3 reports panel regressions of bank stability on lagged AI intensity for an unbalanced panel of 35 publicly listed Chinese commercial banks over 2009–2024. The dependent variable is the *log Z-score* computed using a four-year rolling window for *ROA volatility* (excluding the current year). AI intensity is measured by the standardized mean-based *AI-intensity index* (as detailed in Appendix 2). Column (1) estimates a bank fixed-effects model with year fixed effects and the lagged AI-intensity index only. Column (2) adds the baseline bank-level controls (*Ln(Total assets)*, *Tier 1 ratio*, *Net loans to total assets*, *Non-performing loans ratio*, *Liquidity to total assets*, *Deposits to total liabilities*, *Cost to income ratio*). Column (3) reports the corresponding random-effects specification. Column (4) augments the random-effects model with province fixed effects. All financial variables entering the regressions are winsorized at the 1st and 99th percentiles. Standard errors are clustered at the bank level. t-statistics are reported in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	(1)	(2)	(3)	(4)
	FE: AI only Ln(Z-score)	FE: Baseline + controls Ln(Z-score)	RE: Baseline + controls Ln(Z-score)	RE: + province FE Ln(Z-score)
AI-intensity index	-0.226*** (-3.167)	-0.238*** (-3.494)	-0.181*** (-2.664)	-0.225*** (-3.408)
Ln(Total assets)		-0.001 (-0.001)	0.183*** (2.880)	0.079 (0.646)
Tier 1 ratio		0.072 (0.896)	0.165** (2.499)	0.135** (2.478)
Net loans to total assets		-0.034 (-1.668)	-0.027** (-2.330)	-0.034** (-2.167)
Non-performing loans ratio		-0.088 (-0.245)	-0.230 (-1.080)	-0.169 (-0.637)
Liquidity to total assets		0.081 (0.994)	0.033 (0.620)	0.063 (0.878)
Deposits to total liabilities		-0.037 (-1.443)	-0.030* (-1.818)	-0.028 (-1.234)
Cost to income ratio		-0.000 (-0.001)	0.027* (1.955)	0.024* (1.858)
Constant	4.355*** (36.350)	7.366 (0.349)	0.480 (0.371)	3.902 (0.902)
Bank FE	YES	YES	NO	NO
Year FE	YES	YES	YES	YES
Province FE	NO	NO	NO	YES
SE clustered (bank)	YES	YES	YES	YES
Observations	184	184	184	184
Banks	35	35	35	35
R-squared (within)	0.401	0.441	0.422	0.432

---

Appendix 4

**Transitions in AI Intensity Terciles and Quartiles**

---

Appendix 4 reports year-to-year transition matrices for banks' AI intensity groups. AI intensity is measured using the standardized mean-based *AI-intensity index* constructed from textual disclosures. Panel A reports transition counts, while Panel B reports row percentages, interpreted as transition probabilities from tercile  $t-1$  (rows) to tercile  $t$  (columns). Panel C reports transition counts, and Panel D reports row percentages(transition probabilities) for AI intensity quartiles. Groups are defined on the pooled distribution of the *AI-intensity index* across all bank-year observations. The sample includes transitions for 35 publicly listed Chinese banks over the period 2009–2024.

*Panel A: Counts (from  $t-1$  rows, to  $t$  columns)*

Transitions	To T1	To T2	To T3
From T1	39	20	8
From T2	17	18	18
From T3	7	18	39

*Panel B: Row percentages (transition probabilities, %)*

Transitions	To T1	To T2	To T3
From T1	58.209	29.851	11.940
From T2	32.075	33.962	33.962
From T3	10.938	28.125	60.938
Persistence Pr(same tercile   $t-1$ )	0.522		

*Panel C: Counts (from  $t-1$  rows, to  $t$  columns)*

Transitions	To Q1	To Q2	To Q3	To Q4
From Q1	22	16	6	5
From Q2	13	16	8	7
From Q3	4	11	18	11
From Q4	3	6	12	26

*Panel D: Row percentages (transition probabilities, %)*

Transitions	To Q1	To Q2	To Q3	To Q4
From Q1	44.898	32.653	12.245	10.204
From Q2	29.545	36.364	18.182	15.909
From Q3	9.091	25.000	40.909	25.000
From Q4	6.383	12.766	25.532	55.319
Persistence Pr(same quartile   $t-1$ )	0.446			

---

## Swiss Finance Institute

Swiss Finance Institute (SFI) is the national center for fundamental research, doctoral training, knowledge exchange, and continuing education in the fields of banking and finance. SFI's mission is to grow knowledge capital for the Swiss financial marketplace. Created in 2006 as a public-private partnership, SFI is a common initiative of the Swiss finance industry, leading Swiss universities, and the Swiss Confederation.