

Corporate Artificial Intelligence and Stock Price Crash Risk: Evidence from Chinese Listed Firms

Yixuan Li

*School of Finance, Shandong University of Finance and Economics, Jinan, China
finance_sully@163.com*

Abstract. Using a panel of 45,385 firm-year observations of Chinese A-share listed firms over 2010-2025, this paper examines whether firm-level artificial intelligence (AI) development affects stock price crash risk. Two-way fixed effects estimates show that a higher level of corporate AI development significantly reduces both the negative conditional skewness and the down-to-up volatility of firm-specific returns. The finding survives an instrumental variable strategy, the exclusion of turbulent years, alternative measures of the dependent and the explanatory variables, and alternative specifications. Mechanism tests indicate that AI eases financing constraints and raises total factor productivity, while financing constraints amplify and productivity dampens crash risk. Heterogeneity tests indicate that the mitigating effect is concentrated in non-high-tech firms and in firms located in economically developed regions. The study extends the assessment of the economic consequences of AI from real-economy outcomes to extreme downside risk in capital markets.

Keywords: artificial intelligence, stock price crash risk, financing constraints, total factor productivity

1. Introduction

A sharp collapse in equity prices is the most visible form in which accumulated financial risk is released. A crash destroys investor wealth, weakens market confidence, and propagates to the real sector through wealth and collateral channels. Two parallel developments in China make the question of what restrains such collapses timely. Regulatory reform has moved towards compressing the space in which negative information can be concealed, while a growing number of listed firms now embed AI in design, production, risk management and disclosure. Whether firm-level AI development improves the information environment enough to restrain extreme downside price movements is an open empirical question. Research on AI has developed along three lines. Conceptually, its economic essence lies in a drastic reduction in the cost of prediction [1] and in the reshaping of production through automation and the creation of new tasks [2]. Empirically, three measurement strategies dominate, namely robot installation density at the industry or regional level [3], firm-level AI patents and AI-related investment [4], and keyword frequencies extracted from corporate annual reports. On consequences, AI has been shown to raise productivity, product

innovation and firm growth [4], although the accumulation of intangible complements can generate a J-shaped productivity path in the early diffusion phase [5].

The literature on crash risk builds on the opacity hypothesis [6] and the bad-news-hoarding hypothesis [7]. Because of agency conflicts and career concerns managers withhold unfavourable news, and once the accumulated stock exceeds what can be absorbed it is released at once; reporting opacity and tax avoidance [8, 9], accounting conservatism [10] and short-selling constraints [11] are among the documented determinants. Two gaps remain. Work on the consequences of AI concentrates on productivity, innovation and employment, whereas its capital-market consequences, and extreme downside risk in particular, have received limited attention. Studies of crash risk in turn proceed mainly from disclosure quality and governance structure, and rarely trace how a specific technology shapes the accumulation and release of bad news. This paper addresses both gaps with a sample of Chinese A-share listed firms from 2010 to 2025. Three contributions are intended. First, the assessment of AI is extended from real-economy outcomes to extreme risk in the capital market. Second, financing constraints and total factor productivity are examined as transmission channels, which offers a technological interpretation of how crash risk is formed. Third, the effect is examined across industries of different technology intensity and regions of different development levels, which informs the differentiated implementation of technology policy.

2. Theoretical analysis and hypotheses

2.1. Direct effect

Bad-news hoarding theory traces crashes to the selective concealment of negative information, and the ability to conceal depends on the degree of information asymmetry between insiders and outsiders [6, 7]. The core economic function of AI is to collect, clean and predict from large volumes of data at a very low marginal cost [1], which weakens that ability systematically. Once machine learning and natural language processing are deployed, unstructured data scattered across production, sales, supply chains and finance are converted into traceable and verifiable records, operating anomalies are flagged in real time, and management can no longer rely on the lags of hierarchical reporting to postpone the disclosure of bad news. The probability that the board, internal audit or the regulator detects an anomaly rises at the same time, so the expected cost of concealment increases. When bad news is released continuously and in small increments rather than accumulating to a threshold, the probability of an extreme decline falls, which is consistent with the documented association between reporting transparency and crash risk [8, 9]. Technology also narrows the room for earnings manipulation and lowers the cost at which outside investors process firm-specific information, so that more of that information is impounded in prices. A countervailing force should be acknowledged. AI investment is sunk, slow to pay off and uncertain in outcome, so it may inflate optimistic expectations, in which case disappointment triggers a sharper revaluation [5]. On balance the information-improving force is expected to dominate.

H1. Other things equal, a higher level of corporate AI development significantly reduces stock price crash risk.

2.2. Financing constraint channel

Financing constraints arise from information asymmetry between the firm and outside capital providers. Adverse selection makes external funds more expensive than internal funds, and profitable projects may be abandoned for lack of financing [12]. AI relaxes this constraint in two

ways. It raises the completeness and verifiability of operating data, so that banks and investors can assess the true condition of the firm at a lower cost, which compresses the external financing premium [13]. It also improves cash-flow forecasting and working capital management, so the firm appears a more stable payer in credit screening; the diffusion of digital technologies has indeed been associated with better access to external finance and higher stock liquidity [14]. Looser constraints in turn lower crash risk. Firms under funding pressure face stronger performance pressure and a higher default probability, so management has a strong incentive to conceal unfavourable news in order to preserve financing eligibility and valuation [7]. Constrained firms also cut research and long-term investment and rely on short-term borrowing, which amplifies operating volatility and makes a sudden release of information more likely once liquidity tightens. When the constraint is relaxed the firm neither needs to trade concealment for credit nor lacks the resources to smooth shocks, so both the stock of hoarded bad news and the intensity of its release decline.

H2. Corporate AI development reduces stock price crash risk by easing financing constraints.

2.3. Total factor productivity channel

Total factor productivity measures the improvement in efficiency beyond factor inputs and is a core indicator of firm competitiveness. AI raises it through several routes. By pushing the cost of prediction towards zero it allows data to replace experience in demand forecasting, inventory scheduling, quality inspection and equipment maintenance, which reduces misallocation and redundant input [1]. Automation substitutes for routine tasks while creating new high-skill ones, and the recomposition of the task structure yields a net efficiency gain [2]. Firm-level evidence shows that companies with greater AI investment perform better in product innovation and growth, with productivity improvement as the central transmission link [4]. Higher productivity then restrains crash risk. Efficiency gains stabilise profitability and reduce the frequency of large negative earnings surprises, so fewer unfavourable events need to be concealed in the first place [7]. Productive firms hold more internal resources to absorb shocks and need not resort to earnings management to disguise temporary fluctuations, so accounting quality is higher and bad news accumulates more slowly [10]. A clearer operating outlook also reduces divergence in investor beliefs and therefore the likelihood of a price jump driven by an abrupt revision of expectations.

H3. Corporate AI development reduces stock price crash risk by raising total factor productivity.

3. Research design

3.1. Model specification

The following two-way fixed effects model is used to test the hypotheses above.

$$Crash_{i,t} = \alpha_0 + \alpha_1 AI_{i,t} + \sum_k \gamma_k Control_{k,i,t} + \mu_i + \lambda_t + \varepsilon_{i,t} \quad (1)$$

$Crash_{i,t}$ denotes stock price crash risk, measured by *NCSKEW* and *DUVOL*; $AI_{i,t}$ is the level of corporate AI development; $Control_{k,i,t}$ is a vector of firm-level controls; μ_i and λ_t are firm and year fixed effects; and $\varepsilon_{i,t}$ is the disturbance term. The coefficient α_1 is the parameter of interest and is expected to be negative under H1.

3.2. Variables

Stock price crash risk is measured following Chen et al. [15] and Hutton et al. [8] by the negative conditional return skewness (*NCSKEW*) and the down-to-up volatility (*DUVOL*) of firm-specific weekly returns, both obtained from an extended market model that includes two leads and two lags of the market return.

Corporate AI development (*AI*) is measured from the management discussion and analysis (MD&A) section of the annual report, which carries the forward-looking statements that management chooses to emphasise [16]. Annual reports for 2010-2025 were collected from the CNINFO disclosure platform and the MD&A section was located and extracted by its standard section headings, AI-related words were then counted by dictionary-based frequency counting, and the measure is the natural logarithm of one plus that count.

Ten controls are included: audit by a Big Four auditor (*big4*), the ownership attribute of the controlling shareholder (*soe*), the book-to-market ratio (*bm*), institutional ownership (*Inst*), sales growth (*growth*), board size in logarithm (*lnboard*), the separation of control and cash-flow rights (*sep*), operating cash flow over total assets (*ocf*), the current ratio (*Current*) and the Altman Z-score (*ZScore*).

3.3. Data

The initial sample consists of Chinese A-share listed firms from 2010 to 2025. Financial, market and governance data are drawn from the CSMAR database. The initial sample is then processed as follows: financial firms are excluded; ST, *ST and PT firms and firms that were delisted or suspended are removed; observations with substantial missing values on the main variables are dropped so that all regressions share the same sample; and all continuous variables are winsorised at the 1st and 99th percentiles. The final sample contains 45,385 firm-year observations for 5,056 firms.¹

4. Empirical results

4.1. Baseline regression

Table 1 reports the baseline estimates. With only firm and year fixed effects, the coefficient on *AI* is significantly negative at the 1% level for both crash risk indicators. Adding the full set of controls leaves the sign and the significance unchanged and raises the within R-squared. Among the controls, the book-to-market ratio and institutional ownership enter significantly positive and board size significantly negative, in line with existing evidence on how valuation, institutional trading and governance structure relate to crash risk. H1 is supported.

¹ Descriptive statistics of all variables are omitted here for reasons of space and are available from the authors on request.

Table 1. Baseline regressions

Variable	(1) NCSKEW	(2) DUVOL	(3) NCSKEW	(4) DUVOL
AI	-0.0174*** (0.0064)	-0.0151*** (0.0041)	-0.0173*** (0.0064)	-0.0148*** (0.0042)
big4			-0.0173 (0.0345)	-0.0200 (0.0210)
soe			-0.0289 (0.0191)	-0.0173 (0.0125)
bm			0.3968*** (0.0306)	0.3247*** (0.0202)
Inst			0.0032*** (0.0004)	0.0015*** (0.0003)
growth			-0.0046 (0.0053)	0.0003 (0.0033)
lnboard			-0.0763* (0.0391)	-0.0568** (0.0255)
sep			-0.0003 (0.0011)	0.0004 (0.0007)
ocf			-0.0566 (0.0639)	-0.0621 (0.0419)
Current			0.0052* (0.0029)	0.0056*** (0.0019)
ZScore			0.0020 (0.0013)	0.0010 (0.0009)
Constant	-0.0338* (0.0178)	-0.0336*** (0.0120)	-0.1870** (0.0946)	-0.1288** (0.0616)
Firm / Year FE	Yes	Yes	Yes	Yes
N	45385	45385	45385	45385
R ²	0.0253	0.0287	0.0324	0.0390

Notes: Standard errors clustered at the firm level are in parentheses. ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. The same applies below.

4.2. Endogeneity

The baseline may suffer from reverse causality and omitted variables. Two-stage least squares is therefore applied, using the average AI development level of all other firms in the same province and industry in the same year, excluding the firm itself, as an instrument; leave-one-out industry-region averages of this kind are widely used as instruments in the empirical literature [17]. Firms in the same region and industry face a similar supply of technology, talent and industrial support, so the average adoption of other firms predicts the firm's own adoption intensity, while the adoption decisions of other firms do not directly determine whether a particular firm experiences an extreme price decline. The instrument enters the first stage significantly positive, and the underidentification and weak-instrument tests are comfortably passed. In the second stage the coefficient on AI remains significantly negative for both indicators, so the mitigating effect survives once endogeneity is addressed.²

4.3. Robustness checks

Four sets of checks are performed. First, 2016 and 2020 are dropped in turn, since the circuit-breaker episode and the COVID-19 shock pushed price volatility far from its normal range. Second, the dependent variables are replaced by indicators computed under segmented market capitalisation weighting, which tests whether the result depends on the choice of market benchmark. Third, the explanatory variable is replaced by the share of AI-related words in total MD&A words, with missing values set to zero; the coefficient is larger in absolute value simply because the ratio takes

² Full first- and second-stage results are omitted for reasons of space and are available from the authors on request.

values on a far narrower scale. Fourth, the specification is altered by measuring crash risk in the adjacent period to allow for a lag between adoption and market response, and by adding province and industry fixed effects to absorb regional institutional conditions and industry technological characteristics. In every case the coefficient on AI remains negative and significant, so the conclusion is robust.³

5. Further analysis

5.1. Mechanism tests

The mediator is first regressed on AI and the controls, and is then added to the crash risk equation. Financing constraints (FC) are measured by the absolute value of the SA index [18], with a larger value indicating tighter constraints, and total factor productivity (TFP) is estimated by the Olley-Pakes method [19]. Results are reported in Table 2. For the financing constraint channel, AI enters the constraint equation significantly negative, so technology adoption relaxes the constraint, and the constraint index enters the crash risk equations significantly positive, so more constrained firms are more crash-prone. The two links close a complete chain and H2 is supported. For the productivity channel, AI raises total factor productivity significantly, while productivity enters the crash risk equations significantly negative, so H3 is likewise supported. The coefficient on AI remains significantly negative in all specifications, so a direct effect persists and the two channels do not exhaust the total effect, which matches the theoretical discussion above.

Table 2. Mechanism tests

Variable	(1) FC	(2) NCSKEW	(3) DUVOL	(4) TFP	(5) NCSKEW	(6) DUVOL
AI	-0.0132*** (0.0019)	-0.0155** (0.0064)	-0.0133*** (0.0041)	0.0372*** (0.0066)	-0.0164** (0.0064)	-0.0137*** (0.0041)
FC		0.1359*** (0.0252)	0.1142*** (0.0165)			
TFP					-0.0235** (0.0106)	-0.0290*** (0.0068)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm / Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	45385	45385	45385	45385	45385	45385
R ²	0.3931	0.0332	0.0403	0.3235	0.0326	0.0396

Notes: FC is the absolute value of the SA index; TFP is total factor productivity estimated by the Olley-Pakes method.

5.2. Heterogeneity

Industry technology intensity determines the initial level of a firm's digital infrastructure and information processing capacity, and therefore the marginal improvement a new technology can deliver. By this dimension, the coefficient on AI is insignificant among high-tech firms and significantly negative among non-high-tech firms for both indicators. High-tech firms already possess a solid digital base and strong information processing capacity, so the marginal

³ Robustness results are omitted for reasons of space and are available from the authors on request.

improvement brought by AI is limited; their valuations also embed a larger growth component, so price movements are driven more by revisions of expectations, which dilutes the effect on downside risk [5]. Non-high-tech firms start from a lower level of information infrastructure and operating efficiency, so the same technology delivers a larger improvement and a more visible risk-governance effect.

Regional development shapes both the digital infrastructure and talent supply on which adoption depends and the investor structure and market maturity that condition price stability. By this dimension, the coefficient is significantly negative in developed regions, comprising first-tier, new first-tier and second-tier cities, for both indicators, whereas in less developed regions it is insignificant for NCSKEW and only marginally significant for DUVOL. Developed regions offer better digital infrastructure, a deeper pool of technical talent and a more mature market for data services, so AI investment is converted into effective governance capacity more efficiently and is more likely to pass the threshold at which it begins to matter. Where complementary conditions are weaker, disclosed AI content may not correspond to comparable implementation, and the effect is correspondingly less stable.⁴

6. Conclusions

Using 45,385 firm-year observations of Chinese A-share listed firms from 2010 to 2025, this paper finds that corporate AI development significantly reduces stock price crash risk. The mitigating effect still holds after an instrumental variable strategy, the exclusion of turbulent years, alternative measures of the dependent and the explanatory variables and alternative specifications. Easing financing constraints and raising total factor productivity are two identifiable transmission channels, and the effect differs markedly across firm types, being stronger among non-high-tech firms and firms in developed regions. Two limitations remain. AI development is captured by annual-report text, which cannot fully separate strategic announcement from actual investment, and the sample is restricted to listed firms, so the evidence does not speak to unlisted enterprises.

⁴ Heterogeneity results are omitted for reasons of space and are available from the authors on request.

References

- [1] Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press. Boston.
- [2] Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542.
- [3] Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244.
- [4] Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, art. 103745.
- [5] Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372.
- [6] Jin, L., & Myers, S. C. (2006). R2 around the world: New theory and new tests. *Journal of Financial Economics*, 79(2), 257–292.
- [7] Kothari, S. P., Shu, S., & Wysocki, P. D. (2009). Do managers withhold bad news? *Journal of Accounting Research*, 47(1), 241–276.
- [8] Hutton, A. P., Marcus, A. J., & Tehranian, H. (2009). Opaque financial reports, R2, and crash risk. *Journal of Financial Economics*, 94(1), 67–86.
- [9] Kim, J. B., Li, Y., & Zhang, L. (2011). Corporate tax avoidance and stock price crash risk: Firm-level analysis. *Journal of Financial Economics*, 100(3), 639–662.
- [10] Kim, J. B., & Zhang, L. (2016). Accounting conservatism and stock price crash risk: Firm-level evidence. *Contemporary Accounting Research*, 33(1), 412–441.
- [11] Callen, J. L., & Fang, X. (2015). Short interest and stock price crash risk. *Journal of Banking & Finance*, 60, 181–194.
- [12] Myers, S. C., & Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics*, 13(2), 187–221.
- [13] Bharath, S. T., Pasquariello, P., & Wu, G. (2009). Does asymmetric information drive capital structure decisions? *Review of Financial Studies*, 22(8), 3211–3243.
- [14] Goldstein, I., Jiang, W., & Karolyi, G. A. (2019). To FinTech and beyond. *Review of Financial Studies*, 32(5), 1647–1661.
- [15] Chen, J., Hong, H., & Stein, J. C. (2001). Forecasting crashes: Trading volume, past returns, and conditional skewness in stock prices. *Journal of Financial Economics*, 61(3), 345–381.
- [16] Li, F. (2010). The information content of forward-looking statements in corporate filings: A naive Bayesian machine learning approach. *Journal of Accounting Research*, 48(5), 1049–1102.
- [17] Fisman, R., & Svensson, J. (2007). Are corruption and taxation really harmful to growth? Firm level evidence. *Journal of Development Economics*, 83(1), 63–75.
- [18] Hadlock, C. J., & Pierce, J. R. (2010). New evidence on measuring financial constraints: Moving beyond the KZ index. *Review of Financial Studies*, 23(5), 1909–1940.
- [19] Olley, G. S., & Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6), 1263–1297.