Proceedings of the 10th International Conference on Economic Management and Green Development
DOI: 10.54254/2754-1169/2026.38364

Investor Sentiment, Cross-Market Contagion, and Asset Pricing in China: A Sentiment-Extended Fama–French Framework

Yufei Li

University of Manchester, Manchester, UK
yufei.li-5@student.manchester.ac.uk

Abstract. Investor sentiment can transmit non-fundamental price pressure across markets, yet standard asset-pricing models do not explicitly capture the interaction between broad market sentiment, high-frequency attention, and narrative shocks. This paper develops a China-adapted Sentiment-Extended Fama–French Six-Factor model (CASE-FF6) for A-share equities. The framework combines a lagged Baker–Wurgler-style market sentiment component, NLP-derived financial-news tone, and abnormal Baidu search intensity into a hybrid sentiment state variable. After orthogonalization against conventional factors and macroeconomic controls, the residual signal is converted into a tradable sentiment-mimicking factor and added to a China-adapted Fama–French five-factor baseline. The proposed model is evaluated through nested comparisons with FF3, sentiment-augmented FF4, and FF5 specifications, using time-series pricing errors, the Gibbons–Ross–Shanken test, Fama–MacBeth cross-sectional regressions, and out-of-sample performance. A complementary event framework uses high-salience Nvidia-related announcements in 2026 to examine herding, attention-driven buying, and liquidity amplification in Chinese concept stocks. The study therefore distinguishes sentiment measurement from factor pricing and provides a falsifiable design for testing whether sentiment represents a distinct priced return component or a temporary conditioning mechanism in China's equity market.

Keywords: investor sentiment, Fama–French model, Chinese A-share market, cross-market contagion, behavioral finance

1. Introduction

Investor sentiment is widely recognized as a systematic driver of deviations between market prices and fundamental values, particularly for securities whose valuation is subjective and whose mispricing is difficult to arbitrage [1]. In the Chinese A-share market, the potential pricing effect of sentiment is especially relevant because trading remains heavily influenced by individual investors, short-selling is institutionally constrained relative to developed markets, and digitally transmitted narratives can become salient within very short event windows.

Behavioral finance provides two complementary explanations for why such price effects need not be immediately eliminated. The noise-trader-risk model shows that correlated non-fundamental

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.38364

beliefs can themselves create risk for rational arbitrageurs [2], while the limits-to-arbitrage framework explains why delegated capital, short horizons, funding constraints, and the possibility that mispricing worsens before correcting can deter corrective trading [3]. These mechanisms imply that a high-salience external information shock may generate a temporary but economically meaningful return component even when the shock has little direct effect on a domestic firm's near-term cash flows.

A central empirical difficulty is measurement. The Baker–Wurgler (BW) index provides a market-based composite proxy for broad investor sentiment [1], whereas search-volume measures capture shifts in investor attention at substantially higher frequency [4]. Textual measures extracted from financial news further allow researchers to identify the directional tone of information rather than attention alone. A hybrid measurement strategy can therefore combine slow-moving market sentiment, high-frequency semantic tone, and active search attention while reducing the weaknesses of any single proxy.

This paper develops such a framework and embeds it in a Fama–French asset-pricing structure. Importantly, the original draft's proposed "sentiment-extended four-factor model" is theoretically coherent only when the Fama–French three-factor model is used as the baseline. Because the Fama–French five-factor model adds profitability and investment factors [5], and because Chinese-market evidence indicates that profitability can improve model fit while the incremental contribution of investment is more mixed [6], a more rigorous specification is a nested sentiment-extended six-factor model. The paper therefore proposes a China-adapted Sentiment-Extended Fama–French Six-Factor model (CASE-FF6), in which a tradable sentiment-mimicking factor is added to market, size, value, profitability, and investment factors.

The contribution is threefold. First, the paper clarifies the distinction between a sentiment index, which is an observable state variable, and a priced factor, which should ideally be represented by a factor-mimicking return. Second, it proposes a mixed-frequency sentiment construction combining macro sentiment, NLP-derived tone, and search attention, with orthogonalization designed to isolate the non-fundamental component. Third, it links this factor to a micro-structural transmission framework—herding, liquidity amplification, and attention-driven buying—and illustrates how these channels can be evaluated around high-salience Nvidia-related events in 2026 without attributing causality to unverified stock-price anecdotes.

The remainder of the paper is organized as follows. Section 2 reviews the behavioral and asset-pricing literature. Section 3 defines the sentiment measures, factor construction, and econometric design. Section 4 validates the proposed optimization against alternative Fama–French specifications and sets out the empirical tests required to establish incremental pricing power. Section 5 develops the micro-mechanism event framework. Section 6 concludes.

2. Literature review

The Efficient Market Hypothesis (EMH) provides the classical benchmark for interpreting security prices: available information should be reflected in prices, so persistent abnormal returns require either omitted risk, market frictions, or a failure of informational efficiency [7]. The Fama–French three-factor model subsequently showed that market, size, and value factors explain substantial common variation in equity returns beyond the CAPM [8]. The five-factor extension added profitability (RMW) and investment (CMA), improving average-return descriptions in the original U.S. tests while leaving known weaknesses for some small, low-profitability, high-investment stocks [5].

DOI: 10.54254/2754-1169/2026.38364

Behavioral models do not reject factor pricing as an empirical tool; rather, they question whether every persistent factor necessarily represents fundamental covariance risk. De Long et al. formalize noise-trader risk, under which systematic fluctuations in irrational beliefs create an additional source of uncertainty that can deter rational arbitrage [2]. Shleifer and Vishny further show that arbitrage is limited when professional investors face short evaluation horizons and capital withdrawals precisely when mispricing worsens [3]. These two mechanisms are particularly relevant when a visible external narrative causes investors to trade a category of stocks as a group.

The distinction between risk and mispricing is also central to the sentiment-anomaly literature. Stambaugh, Yu, and Yuan show that many anomalies are stronger following high-sentiment periods and connect this pattern to short-sale impediments and asymmetric overpricing [9]. For China, this logic is important because the market's investor composition and institutional features may magnify the persistence of attention-driven overvaluation.

Sentiment is latent, so empirical work depends on proxies. Baker and Wurgler construct a composite index from market-based indicators and remove macroeconomic components before extracting common variation [1]. The strength of this approach is breadth: it captures a market-wide state rather than the idiosyncratic language of one news source. Its limitation for event research is frequency. Monthly or quarterly construction cannot resolve intraday or daily reactions to a sudden announcement.

Textual measures address the directional component of high-frequency information. Tetlock demonstrates that media pessimism contains predictive information for short-run market pressure and subsequent reversal [10]. Modern financial-language models can extend this principle by scoring context-sensitive positive, neutral, and negative tone. However, text tone measures the content of available language, not whether investors are actively paying attention to it.

Search data provide the complementary attention dimension. Da, Engelberg, and Gao introduce Google Search Volume Index as a direct measure of investor attention [4]. Barber and Odean's attention-driven buying hypothesis predicts that individual investors become net buyers of attention-grabbing securities because buying requires choosing from a very large opportunity set, whereas selling is naturally restricted to currently held assets [11]. In a short-sale-constrained market, positive attention shocks may therefore produce an especially asymmetric buying response.

Traditional contagion research emphasizes common fundamentals and information transmission across integrated markets. King and Wadhwani show how price changes can be transmitted when traders infer information from other markets [12], while Kodres and Pritsker model cross-market contagion through portfolio rebalancing and information asymmetries [13]. Digital information networks add a behavioral channel: a narrative can travel across markets even when the affected firms have weak direct cash-flow links.

Style-investing theory provides a mechanism for this non-fundamental co-movement. Barberis and Shleifer argue that investors categorize assets into styles and shift funds across categories rather than evaluating every security independently [14]. A dominant global technology firm can therefore act as a narrative anchor for a broad "AI hardware," "memory," or "optical interconnect" category. Domestic concept stocks may receive flows because they are cognitively assigned to the same category, not because their expected cash flows change by the same magnitude.

Directly transplanting U.S.-constructed factors into China can be problematic. Liu, Stambaugh, and Yuan show that the smallest Chinese listed firms are strongly affected by shell-value considerations and propose China-specific size and value factor construction that substantially improves explanatory power [15]. Guo et al. find that the Fama–French five-factor model captures strong size, value, and profitability patterns in China, while the investment factor contributes less

DOI: 10.54254/2754-1169/2026.38364

consistently [6]. More recent broad model comparisons likewise show that no single U.S.-style specification uniformly dominates in A-shares and that China-tailored factor construction can materially improve pricing performance [16].

These results validate the direction—but not the exact form—of the draft's proposed Fama–French optimization. Adding sentiment is economically plausible, but simply appending a non-traded sentiment index to FF3 and labeling the result a four-factor asset-pricing model conflates a conditioning variable with a factor return. A stronger design should (i) use a China-adapted baseline, (ii) distinguish slow-moving and high-frequency sentiment information, (iii) purge fundamental and existing-factor components, and (iv) convert the residual sentiment signal into a factor-mimicking portfolio before testing whether its loading earns a cross-sectional premium.

3. Research design and model construction

3.1. Data architecture and sample

The empirical sample should include non-financial A-share firms listed in Shanghai and Shenzhen, with daily returns, market capitalization, book equity, profitability, investment, turnover, trading volume, suspension status, price-limit status, and industry classification obtained from a consistent market database such as CSMAR or WIND. ST/*ST firms, prolonged suspensions, and observations with mechanically constrained returns should be treated separately because price-limit and distress rules can contaminate event-window inference. The risk-free rate should be matched to the return frequency.

Sentiment data should be synchronized to the trading calendar. The low-frequency component consists of a China-adapted BW-style market sentiment index. The high-frequency semantic component is constructed from financial news and relevant technology-market reports using a finance-domain language model. The attention component uses abnormal Baidu search intensity for event-related terms such as "Nvidia," "HBM," "Rubin," "Feynman," and firm-specific names. All source timestamps must be converted to China Standard Time, and information released after the A-share close should enter the next trading day to avoid look-ahead bias.

3.2. Hybrid sentiment state variable

Let BW^{o}_t denote the macro-orthogonalized and standardized BW-style sentiment component, NLP_t the standardized directional text score, and ASVI_t the standardized abnormal search-volume measure. To prevent a monthly BW observation from creating false daily precision, the BW component should be lagged and treated as a slow-moving state, while NLP and ASVI vary at daily frequency. A parsimonious hybrid state variable is:

$$SENT_{raw,t} = w_1 BW^{o}_{t-1} + w_2 NLP_t + w_3 ASVI_t$$

The weights may be estimated by principal component analysis over the aligned standardized series, or fixed ex ante in robustness tests. Because the objective is to identify non-fundamental sentiment rather than macro or conventional-factor information, the raw index is then orthogonalized against the market state, Fama–French factors, and selected macro controls:

$$SENT_{raw,t} = a + b' F_t + c' Macro_t + u_t, \quad SENT^{\perp}_t = u_t$$

DOI: 10.54254/2754-1169/2026.38364

Here F_t contains the contemporaneous conventional factor returns. Orthogonalization should be performed recursively or within the estimation sample rather than using full-sample information, so that the procedure does not leak future information into earlier observations.

3.3. From sentiment index to tradable sentiment factor

A key correction to the original four-factor proposal is that SENT⊥_t is an index, not automatically a traded factor return. To obtain a factor compatible with standard time-series asset-pricing regressions, the study should construct a sentiment-mimicking portfolio (SMF). One practical method is a double sort. At each formation date, firms are first neutralized by size and industry; within those groups, they are sorted by estimated sensitivity to lagged SENT⊥. The SMF return is the value-weighted return of high-sentiment-sensitivity portfolios minus low-sensitivity portfolios. Alternative implementations can use cross-sectional projection weights to create the minimum-variance factor-mimicking portfolio.

This conversion has two advantages. First, the factor has units of return and can be directly compared with MKT, SMB, HML, RMW, and CMA. Second, the sign and economic magnitude of the sentiment premium become interpretable as compensation or mispricing associated with exposure to sentiment-sensitive stocks, rather than as an arbitrary regression coefficient on a standardized index.

3.4. Proposed China-adapted sentiment-extended six-factor model (CASE-FF6)

The optimized specification uses a nested architecture. FF3 + sentiment is retained as a benchmark because it corresponds to the original draft's four-factor idea. FF5 provides the broader conventional benchmark. The preferred specification then adds the tradable sentiment-mimicking factor to a China-adapted five-factor baseline:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{M,i}\mathrm{MKT}_t + \beta_{S,i}\mathrm{SMB}_t^{CH} + \beta_{H,i}\mathrm{HML}_t^{CH}$$

$$+\beta_{R,i}RMW_t + \beta_{C,i}CMA_t + \beta_{Sent,i}SMF_t + \varepsilon_{i,t}$$

SMB^CH and HML^CH denote China-adapted size and value factors, using a construction designed to reduce shell-value contamination in the smallest listed firms [15]. The model is called CASE-FF6 because it contains six return factors: market, size, value, profitability, investment, and sentiment. If CMA proves redundant in the study sample, a pre-specified reduced CASE-FF5 variant may be reported, but factor removal should follow formal spanning and joint-pricing tests rather than ex post significance fishing.

3.5. Econometric tests

Time-series regressions first test whether adding SMF reduces absolute pricing errors and increases adjusted R^2 across test portfolios. Model comparison should include the Gibbons–Ross–Shanken joint alpha test, which evaluates whether intercepts are jointly zero [17], alongside information criteria and out-of-sample error measures. Because sentiment effects may be concentrated in difficult-to-arbitrage firms, test assets should include portfolios sorted by size, idiosyncratic volatility, intangible intensity, turnover, and retail-trading proxies.

Cross-sectional pricing is then examined with the Fama–MacBeth two-pass procedure [18]. First-pass time-series regressions estimate factor exposures; second-pass cross-sectional regressions test

DOI: 10.54254/2754-1169/2026.38364

whether β_Sent commands a statistically significant premium after controlling for conventional factor loadings. Standard errors should be corrected for serial dependence where appropriate. A complementary two-way fixed-effects panel can test whether firm-level returns respond more strongly to SENT⊥ when arbitrage is more difficult or retail attention is higher.

The analysis should distinguish explanatory power from true factor pricing. A significant contemporaneous coefficient on a sentiment state variable does not by itself demonstrate a priced risk factor. Stronger evidence requires (i) significant factor-mimicking returns, (ii) reduced pricing errors relative to nested benchmarks, (iii) a stable cross-sectional premium on sentiment beta, and (iv) robustness to alternative sentiment definitions and event exclusions.

4. Validation of the Fama–French optimization

4.1. What is valid in the original four-factor proposal

The original proposal is valid as a first-stage benchmark: starting from FF3, adding one sentiment variable yields a four-variable return model. This is consistent with the broader literature showing that sentiment is related to cross-sectional mispricing and that its effects are stronger when short-sale constraints or valuation uncertainty impede arbitrage [1, 9]. Therefore, the economic intuition behind the extension is well grounded.

However, three modifications are necessary for a defensible "optimized" model. First, the modern Fama–French benchmark contains five factors, so a sentiment extension is naturally a six-factor model rather than a four-factor model [5]. Second, Chinese factor definitions should account for market-specific distortions, particularly the smallest firms' shell-value characteristics [15]. Third, the sentiment regressor must be separated from conventional factor and macro information and, for a strict factor-pricing interpretation, transformed into a factor-mimicking return.

4.2. Nested model comparison

The optimization should be evaluated through nested models rather than by assuming that more factors are automatically better. Model A is FF3; Model B is FF3 + SMF; Model C is FF5; Model D is CASE-FF6. The sentiment factor is supported only if Model B improves materially on Model A and Model D improves on Model C after accounting for the cost of an additional regressor. This design directly tests the incremental information content of sentiment.

Table 1. Nested comparison of Fama–French and sentiment-extended asset-pricing specifications

Model	Specification	Factors	Purpose
Model A	FF3	MKT, SMB, HML	Classical benchmark
Model B	Sentiment-FF4	MKT, SMB, HML, SMF	Tests the original draft idea
Model C	FF5	MKT, SMB, HML, RMW, CMA	Modern conventional benchmark
Model D	CASE-FF6	MKT, SMB^CH, HML^CH, RMW, CMA, SMF	Preferred China-adapted sentiment model

4.3. Validation criteria and falsifiable hypotheses

The proposed optimization should be accepted only if it survives falsifiable tests. H1 predicts that SMF has non-zero average returns and is not spanned by the conventional factors. H2 predicts that CASE-FF6 produces smaller absolute alphas than FF5 across sentiment-sensitive test assets. H3

DOI: 10.54254/2754-1169/2026.38364

predicts a significant cross-sectional premium on β_Sent in Fama–MacBeth regressions. H4 predicts stronger sentiment exposure among smaller, more volatile, more intangible-intensive, and more retail-attention-sensitive firms. H5 predicts partial mean reversion after extreme positive sentiment shocks, consistent with temporary price pressure rather than permanent cash-flow news.

If SMF is fully spanned by existing factors, if its premium is unstable across subsamples, or if pricing errors do not improve out of sample, the correct conclusion would be that sentiment is a useful conditioning state variable but not a distinct priced factor. This distinction prevents the model from overstating what the empirical evidence can establish.

5. Micro-mechanism case framework: deconstructing the "Jensen Huang effect"

The Chinese outline identifies three micro-channels—herding, liquidity spirals, and attention-driven price pressure. These channels can be retained, but the event claims should be anchored to independently verifiable announcements and then tested with market and search data rather than treated as established causal facts. The 2026 event set is suitable because Nvidia announcements generated unusually high global salience and contained information directly related to AI computing, memory, and optical interconnects.

At CES 2026, Nvidia presented the Rubin platform as its next generation of AI computing infrastructure and emphasized system-level co-design across computing, memory, networking, and storage [19]. The event therefore provides a clean positive-salience shock for the "AI hardware/memory" category. The mechanism hypothesized in the Chinese outline is herding: domestic investors may buy a broad basket of memory- and substrate-related concept stocks without firm-specific evidence that the announcement changes their contracted shipments or near-term cash flows.

The empirical test should not infer herding from a price increase alone. It should combine abnormal returns, abnormal turnover, cross-sectional return correlation within the concept-stock basket, and synchronized retail-attention measures. A stronger herding interpretation would require unusually high within-basket co-movement after controlling for market and semiconductor-sector returns, particularly among firms with weak direct supply-chain exposure.

At GTC 2026, Nvidia officially identified Feynman as the architecture beyond Vera Rubin and described a platform involving next-generation compute, memory, storage, networking, copper connections, and co-packaged optical scale-up [20]. This verifies the technological premise of the draft's second event, but it does not by itself verify the draft's stronger claim that a particular Chinese optical-stock selloff was caused by a "misreading" of Feynman or by forced liquidation. Those causal statements require transaction-level or margin-financing evidence.

The appropriate mechanism test is therefore conditional. If a high-valuation optical-interconnect portfolio exhibits a sharp negative abnormal return after a narrative reversal, the liquidity-spiral hypothesis predicts that the decline should be amplified where pre-event margin financing, turnover, and valuation are high. Evidence of simultaneous deleveraging or margin-balance contraction would strengthen the mechanism; absent such evidence, the event should be described as a liquidity-amplification candidate rather than a proven liquidation spiral.

The third event is the most directly verifiable. On 7 June 2026 in Seoul, Jensen Huang stated that shortages across the industry supply chain—from wafers and packaging to silicon photonics—would persist for several years because demand remained exceptionally high [21]. The statement provides a high-salience authoritative shock that can be matched to the next A-share trading session and to Baidu search intensity.

DOI: 10.54254/2754-1169/2026.38364

The attention-pressure hypothesis predicts three linked outcomes. First, event-related ASVI should jump relative to its pre-event baseline. Second, abnormal turnover and buy-side pressure should be concentrated in domestic stocks cognitively associated with memory, packaging, and semiconductor design, even when direct Nvidia supply-chain links are weak. Third, if the initial price response reflects temporary attention rather than revised cash flows, cumulative abnormal returns should partially reverse over subsequent days. An event-study design over [−5,+10] trading days, with narrower [0,+2] windows for immediate pressure, can distinguish rapid repricing from slower mean reversion.

The three mechanisms are not mutually exclusive. A high-salience announcement first changes the information set; media repetition and search activity then determine which domestic securities become focal; category-based investing synchronizes demand across concept stocks; and leverage or thin liquidity can amplify the resulting price movement. In the proposed model, these mechanisms should appear as cross-sectional variation in β_Sent and as event-time spikes in the underlying sentiment state. The combination of factor pricing and event microstructure therefore links broad return comovement to a testable behavioral transmission process.

6. Conclusion

This paper develops a behavioral extension of Fama–French asset pricing for the Chinese A-share market and clarifies the conditions under which sentiment can be treated as more than an auxiliary explanatory variable. The literature supports the core intuition of the original draft: investor sentiment can affect securities that are difficult to value and arbitrage, attention can generate asymmetric buying by individuals, and market frictions can allow non-fundamental price pressure to persist [1-4, 9, 11]. The original "FF3 + sentiment" four-factor structure is therefore a valid benchmark, but it is not the strongest final specification.

The proposed optimization is CASE-FF6, a China-adapted sentiment-extended six-factor model containing market, size, value, profitability, investment, and a tradable sentiment-mimicking factor. Its principal methodological improvement is to separate the measurement of sentiment from the pricing of sentiment. A hybrid state variable combines a lagged macro sentiment component, high-frequency NLP tone, and abnormal search attention; this signal is orthogonalized against conventional factors and macro controls; and a factor-mimicking portfolio converts the residual signal into return units suitable for time-series and cross-sectional asset-pricing tests. China-adapted size and value construction further reduces distortions associated with the smallest listed firms [15].

The 2026 Nvidia events provide a useful micro-structural setting for examining the model's behavioral mechanisms. CES 2026 supplies a positive category-salience event [19], GTC 2026 verifies the Feynman roadmap and the growing role of optical interconnection [20], and Huang's June 2026 Seoul remarks provide a clearly documented memory-shortage shock [21]. These events should be used as test windows rather than as substitutes for empirical evidence: claims of herding, forced liquidation, or mean reversion require confirmation with returns, turnover, margin data, and search intensity.

The final empirical judgment must therefore remain conditional on the data. CASE-FF6 is validated only if the sentiment-mimicking factor is not spanned by existing factors, reduces pricing errors relative to FF5, earns a stable cross-sectional premium, and improves out-of-sample performance. If those tests fail, the correct interpretation is that investor sentiment acts as a conditioning state variable or transient mispricing mechanism rather than a distinct priced factor. This falsifiable distinction is essential for converting a plausible behavioral narrative into a rigorous asset-pricing contribution.

DOI: 10.54254/2754-1169/2026.38364

Future research can extend the model in three directions. First, mixed-frequency state-space or MIDAS methods can replace the simple lagged alignment of monthly BW and daily high-frequency indicators. Second, transaction-level data can identify whether attention shocks are driven primarily by retail order imbalance or institutional sector rotation. Third, network-based text methods can map how a global technology narrative moves through Chinese media, search behavior, and concept-stock classifications before appearing in prices. Such extensions would allow the sentiment factor to be connected more directly to the information-transmission network that generates it.

References

[1] Baker, M., & Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. *Journal of Finance, 61*(4), 1645–1680. doi:10.1111/j.1540-6261.2006.00885.x.

[2] De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy, 98*(4), 703–738. doi:10.1086/261703.

[3] Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance, 52*(1), 35–55. doi:10.1111/j.1540-6261.1997.tb03807.x.

[4] Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. *Journal of Finance, 66*(5), 1461–1499. doi:10.1111/j.1540-6261.2011.01679.x.

[5] Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics, 116*(1), 1–22. doi:10.1016/j.jfineco.2014.10.010.

[6] Guo, B., Zhang, W., Zhang, Y., & Zhang, H. (2017). The five-factor asset pricing model tests for the Chinese stock market. *Pacific-Basin Finance Journal, 43*, 84–106. doi:10.1016/j.pacfin.2017.02.001.

[7] Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance, 25*(2), 383-417. doi:10.1111/j.1540-6261.1970.tb00518.x.

[8] Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics, 33*(1), 3-56. doi:10.1016/0304-405X(93)90023-5.

[9] Stambaugh, R. F., Yu, J., & Yuan, Y. (2012). The short of it: Investor sentiment and anomalies. *Journal of Financial Economics, 104*(2), 288-302. doi:10.1016/j.jfineco.2011.12.001.

[10] Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance, 62*(3), 1139-1168. doi:10.1111/j.1540-6261.2007.01232.x.

[11] Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies, 21*(2), 785-818. doi:10.1093/rfs/hhm079.

[12] King, M. A., & Wadhwani, S. (1990). Transmission of volatility between stock markets. *Review of Financial Studies, 3*(1), 5-33. doi:10.1093/rfs/3.1.5.

[13] Kodres, L. E., & Pritsker, M. (2002). A rational expectations model of financial contagion. *Journal of Finance, 57*(2), 769–799. doi:10.1111/1540-6261.00441.

[14] Barberis, N., & Shleifer, A. (2003). Style investing. *Journal of Financial Economics, 68*(2), 161–199. doi:10.1016/S0304-405X(03)00064-3.

[15] Liu, J., Stambaugh, R. F., & Yuan, Y. (2019). Size and value in China. *Journal of Financial Economics, 134*(1), 48–69. doi:10.1016/j.jfineco.2019.03.008.

[16] Hanauer, M. X., Jansen, M., Swinkels, L., & Zhou, W. (2024). Factor models for Chinese A-shares. *International Review of Financial Analysis, 91*, 102975. doi:10.1016/j.irfa.2023.102975.

[17] Gibbons, M. R., Ross, S. A., & Shanken, J. (1989). A test of the efficiency of a given portfolio. *Econometrica, 57*(5), 1121–1152. doi:10.2307/1913625.

[18] Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy, 81*(3), 607–636. doi:10.1086/260061.

[19] NVIDIA. (2026, January 5). NVIDIA Rubin platform, open models, autonomous driving: CES 2026 special presentation. *NVIDIA Blog*. https://blogs.nvidia.com/blog/2026-ces-special-presentation/

[20] NVIDIA. (2026, March 19). NVIDIA GTC 2026: Live updates on what's next in AI. *NVIDIA Blog*. https://blogs.nvidia.com/blog/gtc-2026-news/

[21] Reuters. (2026, June 7). Nvidia, SK to detail cooperation plan as Huang flags prolonged chip shortage. *Reuters*. https://www.reuters.com/business/media-telecom/nvidia-ceo-sk-chairman-announce-cooperation-plan-monday-report-2026-06-07/