China Internet Sector

The market is overpricing the AI bill

Equities

China

Internet Services

Kenneth Fong
Analyst
kenneth-kc.fong@ubs.com
+852-3712 3890

Wei Xiong
Analyst
S1460518100005
wei.xiong@ubs.com
+86-21-3866 8883

Sardonna Fong
Analyst
sardonna.fong@ubs.com
+852-3712 3042

Jenny Yuan
Analyst
S1460525030004
jenny-za.yuan@ubs.com
+86-21-3866 8912

Charles Chen
Analyst
S1460524020001
charles-za.chen@ubs.com
+86-21-3866 8907

Dorothy Chen
Associate Analyst
dorothy.chen@ubs.com
+852-3712 2917

Mycroft He
Associate
S1460126080003
mycroft.he@ubs.com
+86-21-3866 8821

China Internet shares are down 30% YTD, reflecting a 19% reduction in consensus EPS and a 14% valuation de-rating. Weak consumption, rising AI expenditure and uncertainty over future returns have put pressure on both earnings expectations and free cash flow. Yet at 10-11x P/E after incorporating incremental AI investment, we believe selected stocks increasingly price peak spending as a permanent earnings burden, while assigning limited value to eventual cash-flow normalization or successful AI monetization.

Sector: Market may be over-discounting risk, while underpricing optionality

We believe Chinese internet companies are effectively paying an AI insurance premium to preserve their competitive positions amid uncertainty about how AI will transform discovery, content and enterprise workflows. Hyperscalers are allocating approximately one year of core cash generation to infrastructure investment (e.g. Tencent's annual OCF of cRmb300bn vs. capex of Rmb250bn). Application-driven platforms (e.g. JD, Meituan and Bilibili) are investing through R&D opex, which should offer greater flexibility. The valuation debate should therefore differentiate among: 1) temporary investment, including expansion capex, experimentation and product subsidies; 2) recurring participation costs, including inference and maintenance; and 3) permanent franchise changes, including lower monetization, higher capital intensity and traffic loss. Even if AI returns disappoint, reduced procurement and discretionary R&D could support a faster recovery in FCF than in EPS, as depreciation on previously acquired assets persists. If AI proves successful, current valuations provide limited recognition of the associated monetization option, in our view.

Industry: Three structural changes define the industry in the AI era, in our view

1. AI expands supply more readily than demand, in our view, judging from: 1) supply or demand constraints; 2) monetization headroom; and 3) ability of the platform to retain the value created. We see the clearest revenue upside where production costs, content availability or operating complexity constrain growth. With that, we are more positive on online gaming or under-monetized verticals (e.g. online recruitment), neutral on e-com (bound by demand constraints), but cautious on ad (the sharp rise of short drama added significant inventory and lower eCPM). **2. Agentic AI threat to transaction platform likely contained:** Agentic AI could reorganize traffic allocation by searching, comparing and executing across platforms. However, China's integrated mobile ecosystems, walled gardens, fragmented merchant supply and high agent-serving costs raise the hurdle for disruption. **3. AI competition shift from model intelligence to the combination of model capability × inference economics × distribution and monetization × proprietary feedback loops**. Under this backdrop, we see 3 competitive advantages: 1) capital, to fund increasingly expensive model development, secure training and inference capacity, and support working capital; 2) efficiency or specialization, through superior unit economics, architecture or domain expertise; and 3) ecosystem, connecting user context and authorization with execution, monetization and feedback. We think this favors well-funded internet giants platforms with full-stack capabilities and established customer touchpoints and context, while leaving opportunities for independent labs with differentiated technology or specialization.

Financial implications: Stronger moat, but lower margin and asset heavier

AI is pushing internet companies beyond their traditional role as asset-light transaction facilitators, requiring platforms to reinforce their moats and differentiate their offerings to avoid commoditization. Examples include investments in compute (BABA and Tencent's AI infrastructure), supply (PDD's Rmb100bn Xinpinmu, Meituan's InstaMart

This report has been prepared by UBS Securities Asia Limited. **ANALYST CERTIFICATION AND REQUIRED DISCLOSURES, INCLUDING INFORMATION ON THE QUANTITATIVE RESEARCH REVIEW PUBLISHED BY UBS, BEGIN ON PAGE 22.**

and Xiaoxiang Supermarket), fulfilment (BABA and JD's quick commerce), etc. These investments may strengthen competitive barriers, but also weigh on margins, working capital and balance sheets. Despite cost discipline across other non-AI investment (e.g. food delivery & e-com optimizing subsidies), we expect incremental margin pressure into 2027E.

Stock picks: AI beneficiaries, valuation support and bottom up drivers

Our stock-selection framework combines AI upside option with downside resilience from its core cashflow. We identify three opportunities: 1) AI monetization leaders we believe **Alibaba** is entering a EPS upgrade cycle as core e-com stablizing, investment losses narrowing and Cloud growing at an expanding margin and increasingly become a group profit driver. Its full-stack capabilities stands it apart. 2) AI optionality with limited valuation premium; we think **Tencent**'s AI investment can be well supported by its core business (resilient and under-monetized). We expect a -5% EPS decline in 2027 even if new AI-product costs double from Rmb55bn to Rmb110bn as core business earnings growth is offset. Tencent is now trading at a 12x 2027 P/E (on our 9% below cons earnings, after 2% 2027E-28E EPS cut). We believe the market is over-discounting the duration of AI investment pressure. 3) Companies with bottom-up earnings drivers. We like **Meituan** as easing competition (both food delivery and local services) provide a multi-year earnings recovery independent of AI with scope for the multiple to overshoot as recovery confidence builds. **JOYY** offers attractive yield (7-8%) and yet high teens EPS growth profile at a 3x ex-cash P/E. **Boss Zhipin's** bottom-up driver comes from more proactive monetization strategy initiated in 2Q26 to complement user expansion through higher enterprise spending and paying conversion.

Beyond our preferred picks: We see better opportunities for these names in early 2027: We think **JD** stock is attractive at 6x P/E and 7-8% shareholder return, but a potential re-rating driver could come from its strategy on food delivery losses. **Kuaishou** trade close to its net cash value while market is focusing on how its core business turn around and KLing ARR progress after upgrade. **Netease** is defensive, in our view, with Ananta launch in Jan 27 as a potential catalyst, and we think **TCOM** growth could accelerate into 2Q27 if anti-trust investigations end and as comps ease. The market sentiment towards AI Labs **Z.AI** and **Minimax** appears to remain relatively cautious in the near term, awaiting upcoming model iterations as key potential catalysts.

PRICING THE AI RISK, OVERLOOKING THE OPTION

Sector: Market may be over-discounting risk and underpricing optionality

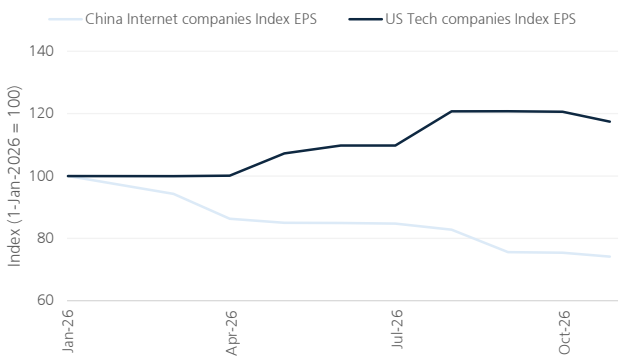
China Internet sector shares have declined 30% YTD, reflecting an 19% EPS downgrade and a 14% valuation multiple de-rating. Fundamental headwinds persist, including subdued macro consumption and AI-related capex pressure on FCF and EPS. This contrasts with US internet names, where resilient core businesses and high-margin cloud growth help offset part of the AI-related depreciation.

Figure 1: HSTech YTD indexed performance



Source: Reuters

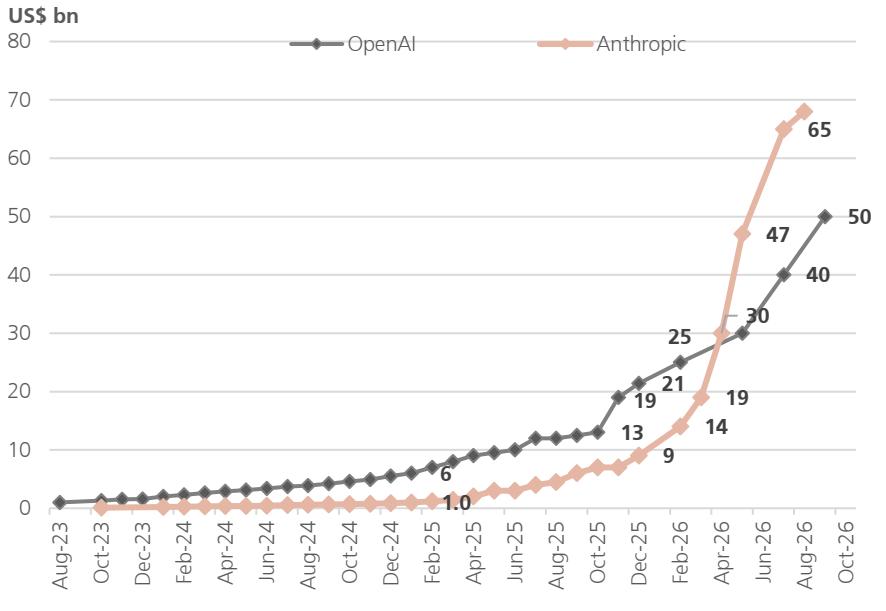
Figure 2: China vs US YTD cons EPS index



Source: Reuters, Company data. Note: China Internet companies included: BABA, Tencent, Baidu, JD, Kuaishou, and Meituan; US Tech companies included: NVIDIA, Microsoft, Meta, Alphabet, Tesla, Apple and Amazon.

At the same time, the global AI debate is shifting from demand growth to the pace of demand to justify the capital committed. Yet some prospective investors were reportedly disappointed by the deceleration, illustrating how quickly expectations have risen alongside AI investment.

Figure 3: US models ARR saw a deceleration on MoM growth



Source: The Information, Bloomberg

While the sentiment can change quickly given how rapid the technology is developing, 3 trends are increasingly clear, in our view: 1) AI investment is expanding, 2) global model pricing is becoming more competitive, and 3) the eventual distribution of profits remains

uncertain. **For Chinese internet companies, we believe some view AI investment as strategic insurance against potentially significant changes in customer behavior.** As companies cannot know precisely how AI will reshape discovery, commerce, advertising, content and enterprise workflows, they can't risk under-investment.

We divide the Chinese internet companies into two camps, with **1) the capex-heavy approach**, e.g. Tencent and Alibaba; and **2) the more opex-driven approach**, e.g. JD, Bilibili and Meituan, which are deploying AI to upgrade their platforms and enhance user and merchant experience (e.g. more accurate recommendations and better merchant operating efficiency).

One mitigating factor is that China remains capacity constrained, which naturally limits the pace of infrastructure spending. Even for hyperscalers, planned AI capex remains broadly comparable with annual cash generation. For instance, Tencent's annual OCF is cRmb300bn, Alibaba's core e-com EBITA is cRmb190bn, compared with their 2026E and 2027E annual capex of cRmb250bn. **Essentially, hyperscalers are effectively committing around a year of core cash flows to capture the AI opportunity and insure against the risk of under-investment.** Application-led platforms are scaling up their R&D expenses, which are relatively flexible in nature.

So, we believe investors should distinguish between temporary investment, recurring costs of participation, and permanent changes to the economics of the core franchise. This is because if AI ROI ultimately proves less attractive than expected, Chinese internet companies can scale back investment, potentially driving a rebound FCF for hyperscalers, and margin recovery for application-led businesses.

On valuation, we believe selected Chinese internet names may be over-discounting the duration of AI investment pressure, while assigning insufficient value to potential AI monetization. Following the sharp YTD share price decline and multiple de-rating, the sector is trading at an undemanding valuation of 10-11x, even after factoring in incremental AI investments. Conversely, if AI proves less attractive than expected, investment intensity would be likely to moderate, making valuations on normalized underlying earnings and FCF even more attractive, in our view.

Industry: Three structural changes define the industry in the AI era, in our view

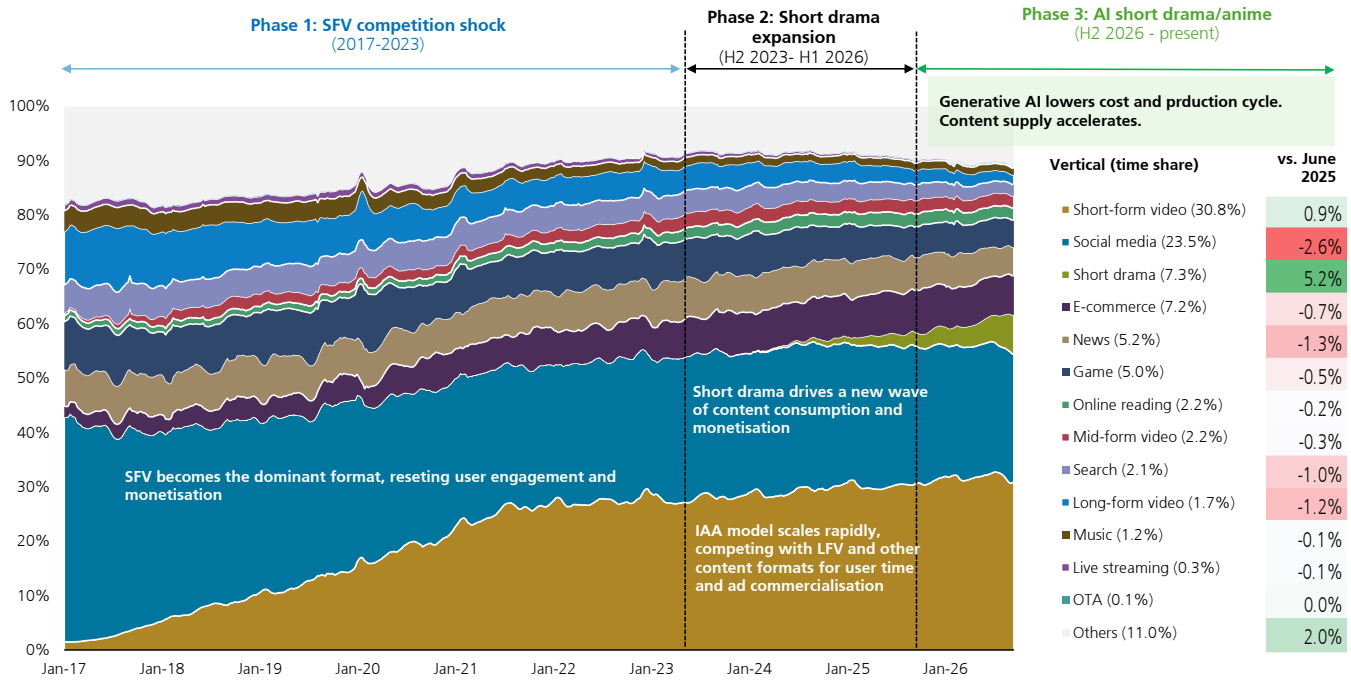
1) Monetization and traffic: AI expands supply more readily than demand

AI expands supply faster than demand, creating clear winners and losers. We assess beneficiaries through three lenses: does AI remove the binding constraint, is there monetization headroom, and who captures the value created? Where production cost or content availability constrains growth, as in short dramas and games, AI can translate more directly into revenue. Where the constraint is consumer spending or advertising budgets, as in e-commerce and parts of advertising, AI should improve efficiency and redistribute share rather than materially expand the profit pool.

Advertising: AI short drama redistributes attention and ad budgets

AI short drama is emerging as a major beneficiary of the shift toward bite-sized entertainment. The evolution of short dramas illustrates a broader shift in content formats, from short-form video (SFV; 2017–H1 2023) to live-action short dramas (H2 2023–H1 2026), and more recently to AI-generated short dramas and anime (H2 2026–present). The format combines structurally lower content-production costs with rising user demand: AI production costs are approximately one-tenth of live action, with production cycles shortened to 2–3 weeks, while more than 50k AI-generated short dramas and anime are being released monthly on Douyin.

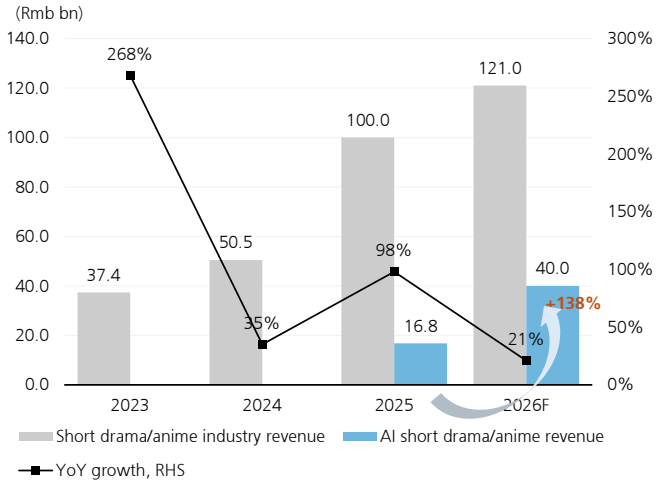
Figure 4: From SFV to AI-driven content - a three-phase shift in user time spent



Source: UBS Evidence Lab (> Access Dataset), UBS. Note: Each vertical includes major apps/developers in that vertical.

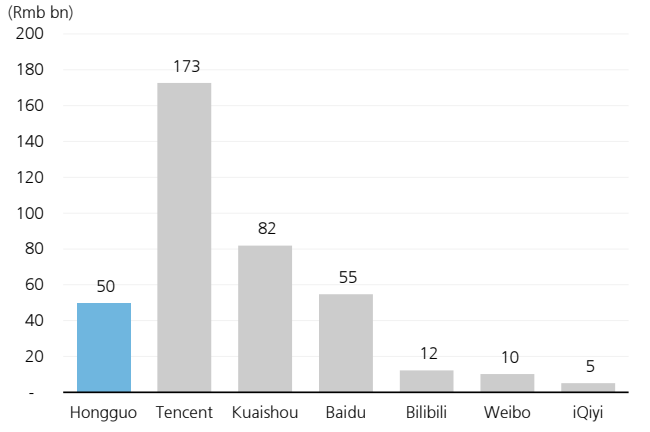
Short drama has become a meaningful new pool of user attention. DataEye estimates China's AI short-drama/anime market at RMB40bn in 2026, +138% YoY. Based on our industry data, short drama has become the second-largest entertainment format by total time spent; Hongguo reached 203mn DAU in September 2026, +127% YoY, with average daily usage of 166 minutes, while its dedicated short-anime app reached 79mn MAU.

Figure 5: AI short drama/anime market revenue is expected to increase 138% YoY to Rmb40bn in 2026



Source: DataEye

Figure 6: Hongguo ad revenue vs. major ad platforms, 2026E



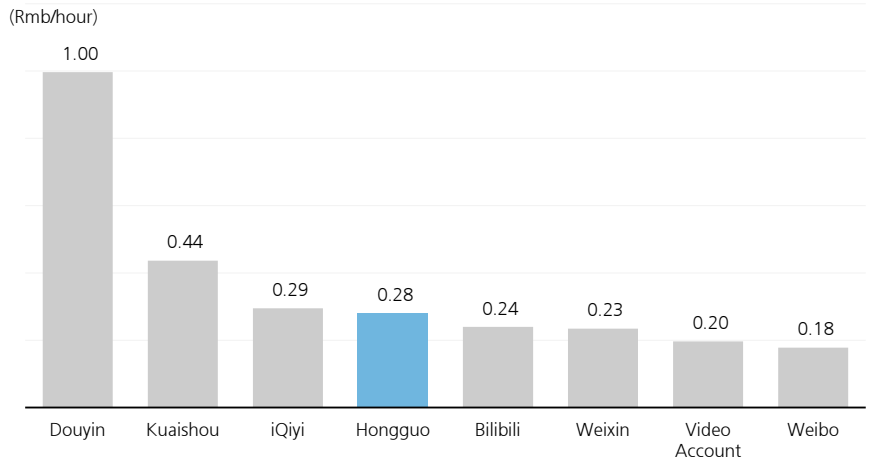
Source: QQ news, UBSe

Ad budgets are following attention, increasingly shifting monetization from subscription-led IAP toward advertising-led IAA. Our channel checks (information gathered from UBS-authorized expert calls or third-party data base) indicate Hongguo short drama generated approximately RMB50bn of advertising revenue, equivalent to roughly 30%/60% of our 2026E advertising revenue estimates for Tencent/Kuaishou, respectively. With 70%+ of short-drama revenue spent on traffic acquisition, economics are increasingly accruing to the platforms that control distribution rather than content

production alone.

However, we see this primarily as ad-budget redistribution, rather than pure market expansion. Rapid growth in IAA-supported inventory should intensify competition for finite advertising budgets and pressure industry eCPMs. Importantly, we estimate Hongguo ad revenue per user-minute remains less than half that of Douyin, suggesting substantial monetization headroom even without assuming equivalent monetization intensity. The investment implication is therefore two-sided: AI expands content supply and engagement, but scarcity shifts from content to user attention, strengthening the economics of scaled distribution platforms.

Figure 7: Ad revenue per user time spent comparison, Q326E



Source: Company data, UBS estimates. *2026 for Hongguo.

Online games: Lower production costs meet untapped monetization

We see online games as a direct AI beneficiary, as AI addresses both content supply and monetization. On supply, AI can shorten development cycles, accelerate gameplay testing and increase the frequency and variety of live-service content, particularly cosmetic items tailored to different player preferences. The impact should vary by genre: we see greater upside in strategy-led titles such as SLG, where gameplay depth, content density and iteration matter more than visual differentiation alone. On demand, we estimate payer penetration by MAU for Tencent's evergreen HoK and PUBG titles at below 20%, with ARPPU of around RMB200, leaving meaningful monetization headroom if better content lifts engagement and payer conversion. AI therefore offers a dual benefit: lower the cost and time to create content, while expanding the amount of content that can be monetized across an under-penetrated paying user base.

E-commerce: AI improves efficiency, but demand remains the binding constraint

We see more limited near-term revenue upside from AI in e-commerce, where the constraint is consumer demand rather than product discovery or content supply. AI can improve targeting, CTR and conversion, but higher efficiency does not necessarily translate into proportionately higher consumer spending or merchant ad budgets. China's already sophisticated recommendation engines further reduce the incremental uplift from better discovery. Our channel checks illustrate this constraint: Douyin e-commerce GMV growth slowed to low-single digits YoY in 3Q26 from approximately 25% in 1H26, despite continued improvements in recommendation and content tools.

The more immediate value lies in lowering merchants' cost to operate. AI can automate product listings, creative generation, customer service and campaign management, reducing participation costs and improving ROI. **We therefore expect AI to drive merchant efficiency, conversion and market-share redistribution before it materially expands the e-commerce profit pool**. A larger GMV opportunity should

emerge only when stronger consumer demand allows these productivity gains to translate into incremental spending.

2) Platform base business: How would agentic AI disrupt transactional platforms?

While Meta Muse has expanded investor expectations for personal agents, we are cautious about extrapolating the US opportunity directly to China. For an agent to displace an incumbent platform, the incremental value from better search, comparison, pricing or execution must materially exceed the switching cost from an already optimized user experience. We see four structural hurdles:

a) Integrated mobile ecosystems raise the adoption hurdle. We estimate approximately 85% of online transactions occur on mobile in China versus 50% in the US, leaving less friction for agents to remove from fragmented PC workflows. China's major transaction verticals, including e-commerce, OTA and food delivery, are already concentrated around a handful of apps integrating discovery, recommendation, payment and fulfillment. Purchasing habits are also entrenched among important spending cohorts born in the 1980s and 1990s, while for some discretionary categories, browsing itself is part of the consumer experience. Agents therefore need to create substantially more value than "fewer clicks" to change behavior. The hurdle rises further as incumbents embed their own assistants inside existing apps, reducing the need for users to migrate to a third-party interface.

b) Walled gardens turn agent access into an economic negotiation. Independent agents require access to account data, merchant inventory and transaction functions, but technical ability does not guarantee commercial permission. Amazon's reported blocking of Muse illustrates the incentive for incumbent platforms to protect customer ownership. In China, where supply and transactions are concentrated within a small number of large ecosystems, agent penetration may depend as much on negotiated access, authorization and revenue sharing as on model capability. This potentially slows cross-platform aggregation and preserves bargaining power for incumbent platforms.

c) Fragmented merchants strengthen incumbents' execution moat. China's consumer-facing platforms do more than aggregate supply. Serving a fragmented merchant base requires local BD, onboarding and ongoing support, while leading platforms increasingly participate in pricing, inventory, fulfillment and after-sales operations. An agent can aggregate information more readily than it can replicate the operating infrastructure required to complete a transaction. We therefore expect AI to make discovery more contestable before it disintermediates platforms such as Meituan, Alibaba or JD from execution.

d) Agent economics remain challenging. Industry estimates put Muse's service cost at approximately US$10–20 per user per month, potentially higher when dedicated virtual-computing resources are required. This creates an unusually high monetization hurdle for a consumer Internet product: agents must generate enough incremental consumer value or transaction economics to cover inference, compute and execution costs. Lower model costs should help, but cheaper inference could also stimulate greater agent usage, limiting the reduction in aggregate serving costs.

Figure 8: Summary for Agentic AI Opportunity in E-commerce: US vs. China

	US	China
Online ecosystem	Predominantly PC-based; c50% mobile-based transactions	Mobile-first; c85% mobile-based transactions
Platform structure	More open ecosystem; Amazon (c50% share), alongside a fragmented long tail of independent DTC sites	Mostly walled-garden super apps; c90% share by platforms (e.g. BABA, PDD, JD, Douyin)
Platform offerings	More fragmented, with services spread across multiple platforms	More integrated; with platforms typically offering one-stop services within each vertical
Merchant structure	Higher chain store penetration	More fragmented with higher SME mix
Agent economics	More favourable, with higher ASP, and greater value addition from time savings from cross-site search, helping offset higher service costs	Less favourable; limited ToC monetization potential given lower incremental value creation

Source: UBS estimates

3) AI competition: From model leadership to system economics

Model intelligence remains important, but benchmark leadership alone is becoming less defensible as iteration cycles shorten. We see competition shifting toward a broader equation: model capability × inference economics × distribution and monetization × proprietary feedback loops.

This raises the importance of three competitive advantages: **1) capital**, to fund increasingly expensive model development, secure training and inference capacity, and support working capital; **2) efficiency or specialization**, through superior unit economics, architecture or domain expertise; and **3) ecosystem**, connecting user context and authorization with execution, monetization and feedback. The shift should favor cash-generative, full-stack Internet platforms, while independent labs increasingly need differentiated technology or vertical specialization to justify their funding burden.

Capital: the cost of competing extends well beyond model R&D

AI is becoming more capital-intensive across three fronts. **First, model development costs continue to rise** as model scale increases, access to frontier-model outputs tightens and high-quality differentiated data becomes more valuable. Our examples include Kimi K3's 2.8tn-total-parameter architecture and Alibaba's planned larger 5T Qwen models.

Second, compute is needed not only to train models, but increasingly to monetize them. Our channel checks suggest strong Kimi K3 demand has encountered GPU constraints, while Z.AI's GLM 5.2 has been able to serve only approximately 10% of demand directly, with the balance served through third-party capacity. Hardware inflation compounds the issue: our checks indicate NVIDIA B300 pricing in China has approximately doubled YTD, while Huawei Ascend 950 and Cambricon 690 quotations have risen approximately 20–40%.

Third, working capital is becoming an increasingly important part of the AI funding equation. Tencent's infrastructure investment has included both substantial cash capex and compute prepayments. At Z.AI, trade and bills receivables increased from RMB303mn to RMB783mn, while current compute-service prepayments rose from RMB116mn to RMB1.24bn; higher compute-service payables, from RMB727mn to RMB1.43bn, provided only a partial offset. Scaling AI therefore requires capital not only to build models, but to secure capacity and fund commercialization before the associated cash is collected.

Investment implication: funding capacity becomes part of the moat. Cash-generative hyperscalers can absorb model losses, reserve scarce compute resources and fund inference ahead of monetization. Independent labs can still differentiate through architecture, specialization and infrastructure partnerships, but face a tighter trade-off between capability, capacity, margins and future fundraising.

Figure 9: Both Internet Company and LLM labs actively seeking funding supporting AI investment

China internet company	Date	Capital type	Amount (USD bn)	Detail
Tencent	14-Oct	Debt – Bonds	5.5	USD 1.8bn + RMB 26bn bonds
BiliBlli	4-Sep	Debt – Convertible Note	0.7	USD 700m convertible-note package
Alibaba	24-Aug	Equity – Placement	10.3	HKD 80bn equity placement
Tencent	10-Jun	Debt – Bonds	4.6	USD 2.45bn + RMB 15bn bonds
JD	10-Apr	Debt – Senior Notes	1.4	RMB 10bn senior unsecured notes
Kuaishou	16-Jan	Debt – Senior Notes	2.0	USD 1.5bn + RMB 3.5bn senior notes
China internet company total			**24.5**	
LLM labs				
Zhipu	13-Sep	Mixed – Equity & Convertible Bonds	5.0	USD 5bn combined financing
Moonshot AI / Kimi	29-Jul	Equity – Series F Round	3.5	Over USD 3.5bn, Series F round
MiniMax	14-Jul–16-Jul	Mixed – Equity & Convertible Bonds	2.1	HKD 16.0bn combined financing
Zhipu	13-Jul	Equity – Placement	4.0	HKD 31.4bn equity placement
Kling AI / Kuaishou subsidiary	2-Jul	Equity – Financing	3.0	Up to USD 3bn equity financing
DeepSeek	Late Jun	Equity – External Round	7.3	RMB 51bn first external round
Moonshot AI / Kimi	May	Equity – Series D Round	2.0	USD 2bn, Series D round
Moonshot AI / Kimi	Feb	Equity – Series C+ Round	0.7	Over USD 700m, Series C+ round
MiniMax	9-Jan	Equity – IPO	0.7	HKD 5.5bn gross IPO proceeds
Zhipu	8-Jan	Equity – IPO	0.6	HKD 5bn gross IPO proceeds
LLM labs total			**28.9**	
Total			**53.4**	

Source: Company data, Bloomberg, Forbes, UBS estimates.

Efficiency or specialization: superior unit economics, architecture or domain expertise

Low pricing alone is becoming less defensible as global frontier models compete on both intelligence and economics. In agentic workflows, repeated model calls, tool use and retries make cost per successfully completed task, incorporating reliability, latency and human intervention, more relevant than token price alone. OpenAI's GPT-6 Luna illustrates the pressure: input/output pricing fell to US$0.10/ US$0.50 per million tokens from US$0.20/US$1.20 for GPT-5.6 Luna. For Chinese providers, sustaining a pricing advantage increasingly requires structural cost reduction, not margin compression.

We see three routes to lower inference costs: 1) intelligent routing, matching model capability to task complexity; 2) architecture efficiency, including MoE and KDA-type approaches that reduce computational intensity; and 3) hardware/software co-optimization, lowering serving costs across the stack. As frontier-model pricing falls, engineering efficiency increasingly becomes as important as model research in determining commercial competitiveness.

Full-stack cost leadership, Alibaba: Alibaba combines Qwen's open-source ecosystem with optimization across models, serving software, cloud infrastructure and proprietary silicon. Qwen's 460+ open-source models have generated 300,000+ derivatives, supporting task-specific selection rather than deploying the most expensive flagship model for every workload. We estimate approximately 10% of Alibaba's compute capacity runs on proprietary T-Head accelerators, with our channel checks positioning the 810E and M890 against NVIDIA's H20 and H200, respectively. With GPU compute representing an estimated 50–60% of inference costs in our checks, higher proprietary-chip penetration could provide a meaningful structural cost lever.

Software-engineering efficiency, DeepSeek: Our industry checks indicate DeepSeek achieves an approximately 10–20% inference-margin advantage versus peers despite competitive pricing. We see this as evidence that architecture, routing and serving optimization can create a cost moat without owning the full hardware stack.

Vertical specialization offers the alternative to cost leadership. Providers do not necessarily need to be the cheapest if they generate materially better outcomes in high-value workflows. Z.AI's coding focus illustrates the opportunity to combine model capability with domain knowledge and proprietary feedback loops.

We therefore expect the market to bifurcate: cost-efficient general-purpose providers competing for volume, and differentiated specialists competing on task value. Flagship models can retain premium pricing where incremental intelligence matters, while smaller and optimized models increasingly serve high-volume workloads.

Figure 10: Beyond flagship models, AI developers are expanding Flash offerings for cost-efficient options

Developer	Tier	Model	Release date	Parameters (total / active)	AA Index	Input US$/mn tokens	Cached input US$/mn tokens	Output US$/mn tokens	Blended price 7:2:1 blend	Cost per task US$/task
Zhipu	Flash	GLM-5.3-Flash	26 Aug 2026	320B / 18B	41.8	$0.15	$0.03	$0.50	$0.10	$0.28
	Flagship	GLM-5.3	18 Aug 2026	753B / ~40B	44.8	$1.40	$0.26	$4.40	$0.90	$2.01
Moonshot	Flagship	Kimi K3.1 Pro	23 Oct 2026	2.8T / 104B	55.1	$2.00	$0.20	$10.00	$1.54	$2.11
	Flash	Kimi K2.8 Preview	11 Sep 2026	Not disclosed	n.a.	$1.00*	$0.25*	$4.00*	$0.78*	n.a.
	Flagship	Kimi K3	16 Jul 2026	2.8T / 104B	43.6	$3.00	$0.30	$15.00	$2.31	$2.00
Alibaba	Flash	Qwen3.8-Flash-Next	26 Aug 2026	180B / 6B	39.8	$0.15	$0.016	$0.47	$0.09	$0.37
	Flagship	Qwen3.8 Max	3 Aug 2026	Not disclosed	40.2	$2.00	$0.25	$6.00	$1.18	$2.67
Google	Flash	Gemini 3.8 Flash (high)	2 Sep 2026	Not disclosed	40.9	$0.75	$0.075	$3.75	$0.58	$1.24
	Flagship	Gemini 3.1 Pro Preview	19 Feb 2026	Not disclosed	29.7	$2.00	$0.20	$12.00	$1.74	$1.30
SpaceXAI	Flagship	Grok 4.7 (xhigh)	21 Sep 2026	Not disclosed	46.4	$2.00	$0.50	$6.00	$1.35	$3.74
	Flash	Grok Build 0.1 0616	16 Jun 2026	Not disclosed	27.2	$1.00	$0.20	$2.00	$0.54	n.a.
DeepSeek	Flash	DeepSeek V4.1 Flash	10 Sep 2026	552B / 16B	39.5	$0.295	$0.006	$1.20	$0.18	$0.48
	Flagship	DeepSeek V4 Pro 0424	24 Apr 2026	1.6T / 49B	30.4	$0.435	$0.0036	$0.87	$0.18	$0.12
OpenAI	Balanced	GPT-6.1 Sol	29 Sep 2026	Not disclosed	51.8	$2.00	$0.10	$10.00	$1.47	$0.72
	Balanced	GPT-6 Sol	22 Sep 2026	Not disclosed	47.6	$2.00	$0.20	$10.00	$1.54	$1.04
	High-volume	GPT-6 Luna	22 Sep 2026	Not disclosed	38.1	$0.10	$0.01	$0.50	$0.08	$0.07
	Flagship	GPT-6 Astra	3 Sep 2026	Not disclosed	52.7	$10.00	$1.00	$50.00	$7.70	$3.26
Anthropic	Efficient	Claude Haiku 5.5	7 Oct 2026	Not disclosed	43.4	$0.10	$0.01	$0.50	$0.08	$0.21
	Balanced	Claude Sonnet 5.5	28 Sep 2026	Not disclosed	56.0	$2.00	$0.10	$10.00	$1.47	$5.46
	Efficient	Claude Opus 5.5	22 Sep 2026	Not disclosed	57.6	$4.00	$0.20	$20.00	$2.94	$5.98
	Flagship	Claude Fable 5.1	1 Sep 2026	Not disclosed	53.4	$10.00	$0.25	$50.00	$7.18	$7.63
Meta	Flagship	Muse Wonder (max)	12 Oct 2026	Not disclosed	52.4	$4.00	$0.20	$20.00	$2.94	$4.54
	Balanced	Muse Spark 1.3 (max)	2 Sep 2026	Not disclosed	48.1	$1.25	$0.15	$4.25	$0.78	$1.60

Source: Artificial Analysis, company data, UBS. Note: Prices are in US$/mn tokens. Blended price represents a 7:2:1 blend of cached-input, input and output token prices. Parameters refer to total/active parameters where disclosed. *Kimi K2.8 Preview pricing is based on a third-party API provider. Data as of 26 October 2026.

Ecosystem: workflow ownership matters more than model scale alone

As model capability converges, competitive advantage increasingly shifts from owning the highest-ranked model to owning the workflow around it. A durable closed loop combines context + authorization + execution + distribution + monetization + feedback, allowing real-world usage to improve both the product and underlying models.

ByteDance: content-to-monetization loop. Fanqie provides content/IP, Seedance lowers creation costs, Hongguo provides distribution, Volcano Engine supplies infrastructure, and advertising closes the monetization loop. The advantage is not Seedance's model ranking alone, but ByteDance's ability to connect creation, consumption and monetization within one ecosystem.

Muse: context and product experience can outweigh incremental model intelligence. Once underlying capability clears the threshold required for a use case, personal context, memory, authorization and reliable execution become increasingly important to adoption. This supports our view that the competitive battleground shifts from "who has the smartest model?" toward "who understands the user and can reliably complete the task?"

Tencent: harness and distribution become the moat. WorkBuddy's early traction supports an execution-led thesis: broad model access, harness engineering and its Skills ecosystem allow Tencent to orchestrate models around real workflows. Weixin potentially extends this advantage to consumer agents, combining user context with Mini Programs, payments and services.

Investment implication: model capability gets users through the door; workflow ownership determines who captures the economics. We therefore favor platforms that can connect AI with proprietary context, distribution and transaction execution, rather than model developers competing primarily on benchmark leadership.

Financial implications: Stronger moat, lower margins and heavier assets

AI raises the cost of staying competitive. Beyond hyperscaler capex, application platforms must absorb higher R&D, product-development and inference costs to protect traffic and user experience, often ahead of meaningful AI monetization. For example, JD's R&D increased c40% in 1H26 and its R&D ratio rose to 2.1% of revenue, while we estimate incremental AI-related P&L costs of cRMB2.0bn/cRMB0.5bn for Meituan/ Bilibili. We therefore expect sector margins to compress in 2026E, with more modest pressure in 2027E as investment growth normalizes and efficiency gains catch up.

Cost discipline should partially fund the AI bill. As capital is redirected toward higher-ROI AI initiatives, platforms are becoming more disciplined on subsidies and other traffic-acquisition spending. Our channel checks point to lower food-delivery subsidies and improving UE into September; greater ROI discipline at Douyin local services; lower advertising-agency rebates as AI improves placement efficiency; and reduced platform-funded Singles' Day subsidies. The implication is not that AI investment is costless, but that part of the incremental burden can be self-funded through lower legacy customer-acquisition and operating costs.

More importantly, China Internet is shifting from facilitating transactions to controlling the infrastructure behind them. AI makes discovery and traffic more contestable, increasing the value of assets and capabilities that are harder to replicate: compute, proprietary supply, merchant operations and physical fulfillment. PDD's RMB100bn three-year Xinpinmu commitment deepens participation in brands and upstream supply chains, while Meituan's Xiaoxiang expansion combines stores, front warehouses and instant delivery. These investments strengthen execution moats against third-party agents, but structurally increase capital intensity through infrastructure, inventory and operating commitments.

Figure 11: Increasing investments among China internet companies in the AI era

Infrastructure	Assets	Financial implications
Compute	Chips, data centers	Capex, depreciation, working capital (pre-payments), R&D expense
Supply	1P and private labels	Inventory, and supply chain investment
Fulfilment	Warehouses, delivery network	Lease, working capital
Merchant operations	Local BD	G&A expense

Source: UBS

Investment implication: stronger moats may come at the expense of capital efficiency as platform move from asset light to asset heavy. We expect the Internet model in the AI era to become less asset-light, with higher depreciation and working-capital requirements, lower asset turnover and potentially lower normalized ROE/ROIC. Peak AI spending can moderate and FCF can recover, but the sector may not fully return to its pre-AI margin and capital-return structure.

Figure 12: China internet companies' lower ROE in the AI era

	ROE						YoY chg (ppt)				
	2023	2024	2025	2026E	2027E	2028E	2024	2025	2026E	2027E	2028E
NetEase	26.1%	24.7%	25.7%	28.0%	26.7%	24.9%	-1.4%	0.9%	2.4%	-1.3%	-1.8%
Tencent	19.0%	23.1%	22.6%	19.7%	16.9%	16.6%	4.1%	-0.5%	-3.0%	-2.8%	-0.4%
Bilibili	-22.8%	-0.4%	16.5%	17.3%	19.6%	20.9%	22.4%	16.9%	0.9%	2.2%	1.3%
Kuaishou	22.5%	31.8%	32.9%	13.5%	13.3%	14.4%	9.3%	1.1%	-19.4%	-0.2%	1.1%
JD	12.3%	15.7%	8.9%	10.6%	12.9%	13.9%	3.3%	-6.8%	1.7%	2.3%	1.0%
Baidu	11.9%	10.9%	10.4%	10.3%	9.9%	9.4%	-1.0%	-0.5%	-0.2%	-0.4%	-0.5%
Alibaba	14.2%	14.3%	5.9%	9.3%	12.4%	13.9%	0.2%	-8.5%	3.5%	3.1%	1.5%
Meituan	30.1%	54.9%	-23.1%	3.1%	35.4%	43.7%	24.8%	-78.0%	26.2%	32.4%	8.2%
Kingsoft Cloud	-16.5%	-13.7%	-7.7%	-4.0%	0.6%	3.7%	2.8%	6.0%	3.7%	4.6%	3.1%
Z.AI	-28.3%	-40.5%	-29.9%	-47.0%	-36.8%	-14.0%	-12.2%	10.6%	-17.1%	10.2%	22.8%

Source: Company data, UBS estimates

Assessing the impact of a potential pullback in AI investment. We conduct sensitivity analyses on the impact of AI investment on earnings, and of capex on OCF among Chinese internet companies. Overall, application-led companies, where AI investment is primarily reflected in R&D expenses, may have greater flexibility to adjust

spending, while hyperscalers could see a faster FCF recovery as capex moderates.

Figure 13: Earnings sensitivity on AI opex

(Rmb m)	Adj. OP		AI opex		As % of adj. OP	
	2026E	2027E	2026E	2027E	2026E	2027E
Kingsoft Cloud	295	623	1,397	2,160	473%	347%
Z.AI	(4,753)	(4,736)	6,286	10,961	132%	231%
Minimax	(4,927)	(6,662)	5,414	8,663	110%	130%
Baidu	12,081	13,608	10,664	13,444	88%	99%
Kuaishou	10,813	13,075	5,200	7,062	48%	54%
Tencent	288,764	270,064	55,000	110,000	19%	41%
JD	23,316	39,415	9,117	12,506	39%	32%
BABA	132,535	179,391	44,861	50,000	34%	28%
Bilibili	2,896	3,767	500	829	17%	22%
Meituan	(2,071)	31,256	2,000	2,500	97%	8%

Source: Company data, UBS estimates. Note: BABA's adj. earnings for 2026E/27E refer to adj. EBITA for FY27E/28E, and AI opex refers to EBITA loss in AI labs and applications.

Figure 14: OCF sensitivity on capex

(Rmb m)	OCF		Capex		As % of OCF	
	2026E	2027E	2026E	2027E	2026E	2027E
Kingsoft Cloud	4,767	6,003	7,160	7,339	150%	122%
BABA	110,710	216,247	250,277	256,521	226%	119%
Kuaishou	18,489	22,559	24,995	22,638	135%	100%
Baidu	34,596	28,460	32,972	27,806	95%	98%
Tencent	325,260	401,167	252,993	247,360	78%	62%
Meituan	16,293	42,333	14,598	16,058	90%	38%
JD	48,039	44,162	13,299	9,309	28%	21%
Bilibili	4,894	7,222	1,026	987	21%	14%
Z.AI	(2,273)	(831)	536	928	n.a.	n.a.
Minimax	(4,791)	(3,909)	233	219	n.a.	n.a.

Source: Company data, UBS estimates. Note: BABA's financials for 2026E/27E refer to FY27E/28E.

STOCK PICKING

Stocks: AI beneficiaries, valuation support and bottomup drivers

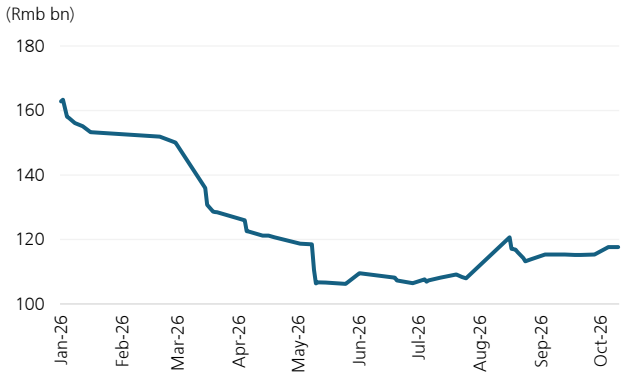
We identify three potential opportunities: **1) AI monetization leaders**, where AI investment is translating into increasingly visible revenue and earnings; **2) AI optionality at limited valuation premium**, where resilient core cash generation and spending flexibility provide downside support while AI offers incremental upside; and **3) bottom-up earnings recovery independent of AI**, driven by under-monetized assets, rationalizing competition or identifiable profit-recovery catalysts.

1) AI monetization leaders

BABA: Cloud monetization, full-stack advantage and an approaching EPS inflection

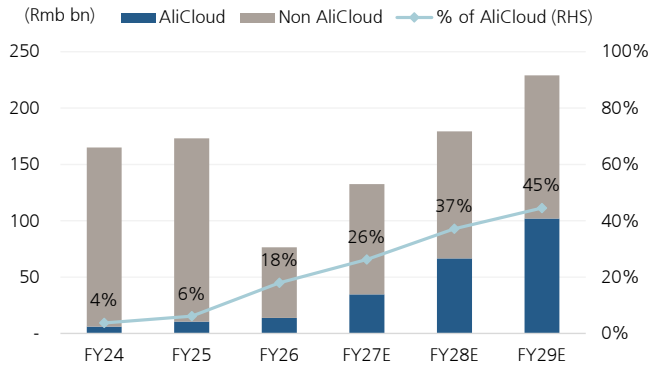
We believe an EPS estimate upgrade cycle is approaching. After several quarters of downgrades driven by weak e-commerce and elevated AI/quick-commerce investment, we see three potential inflection points: 1) core e-commerce stabilization, supporting cRMB190bn of EBITA; 2) accelerating cloud growth and margin expansion, lifting our estimated cloud contribution to group EBITA to 26%/37% in FY27/28E; and 3) narrowing investment losses, particularly in quick commerce. We expect Dec-quarter cloud EBITA to cover AI-lab losses, an important milestone: AI infrastructure earnings begin to fund model-development expenditure.

Figure 15: Cons FY27E adj. EBITA revision YTD



Source: Visible Alpha

Figure 16: AliCloud's increasing EBITA contribution to BABA



Source: Company data, UBS estimates

Cloud margin upside is increasingly the earnings story. We estimate AI/non-AI cloud EBITA margins at 33–35%/5–6%, respectively. A rising AI mix could lift overall AliCloud margin toward 20% by mid-2027, versus consensus FY27/28E margins of 13%/14% and our 14%/18%. As cloud becomes a larger contributor, we expect the valuation debate to shift from "how much will Alibaba spend on AI?" toward "what returns can AI generate?"

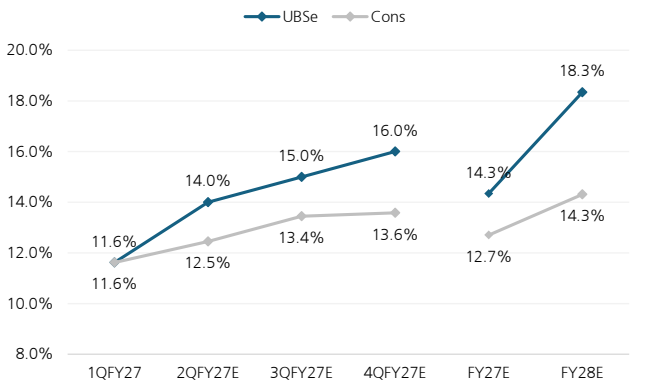
Full-stack economics strengthen Alibaba's AI position. We believe Alibaba can optimize cost across Qwen, serving software, cloud infrastructure and proprietary T-Head silicon, an increasingly important advantage as agentic workloads multiply inference calls. Qwen's 460+ open-source models and 300,000+ derivatives also enable task-specific deployment rather than using expensive flagship models for every workload. We see the combination of model breadth, proprietary compute and cloud scale as a credible route toward low-cost leadership.

Figure 17: AliCloud's EBITA margin should expand from 12% to 14%/18% in FY27E/28E

	1QFY27	2QFY27E	3QFY27E	4QFY27E	FY27E	FY28E
AliCloud revenue	**48,437**	**60,931**	**67,090**	**65,769**	**242,227**	**363,341**
AI	12,376	18,041	22,045	25,080	77,542	154,348
Non-AI and internal	36,061	42,890	45,046	40,689	164,685	208,993
Revenue mix						
AI	26%	30%	33%	38%	32%	42%
Non-AI and internal	74%	70%	67%	62%	68%	58%
AliCloud EBITA	**5,628**	**8,530**	**10,064**	**10,523**	**34,745**	**66,651**
AI	3,713	5,773	7,275	8,277	25,037	54,022
Non-AI and internal	1,915	2,757	2,789	2,247	9,708	12,629
AliCloud EBITA margin	**11.6%**	**14.0%**	**15.0%**	**16.0%**	**14.3%**	**18.3%**
AI	30.0%	32.0%	33.0%	33.0%	34.0%	35.0%
Non-AI and internal	5.3%	6.4%	6.2%	5.5%	5.9%	6.0%

Source: Company data, UBS estimates

Figure 18: AliCloud EBITA margin: UBSe vs cons

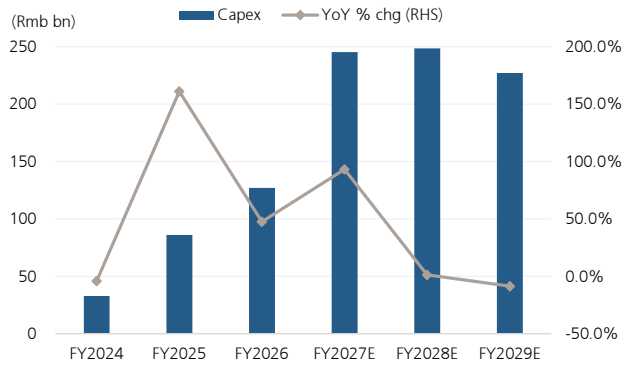


Source: Visible Alpha, UBS estimates

Valuation still gives limited credit to AI monetization. Applying 7–8x P/E to e-commerce and 5x P/S to cloud implies US$100–110/share. At a normalized 20% cloud margin, 5x P/S equates to c25x operating profit, which we consider reasonable against c50% growth and improving profitability. Shares traded around the US$80 range before the 2025 DeepSeek inflection; the key change versus the pre-DeepSeek period is that investors are no longer buying AI optionality alone: cloud revenue growth and margin expansion are increasingly visible.

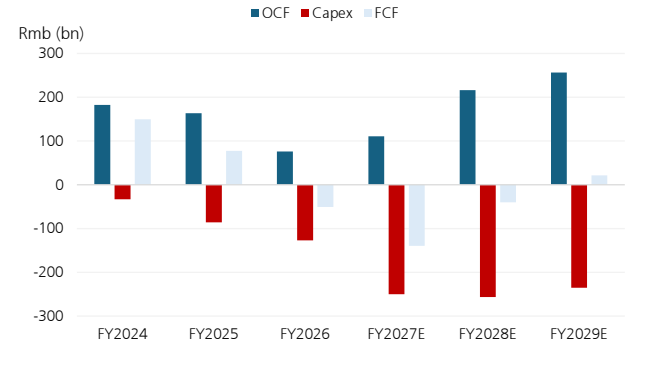
Improving profitability should reduce equity-funding risk. We view the previous placement primarily as replenishing offshore liquidity after substantial chip-related prepayments and capex. Stabilizing e-commerce profits, narrowing quick-commerce losses and improving cloud earnings should strengthen funding capacity. We estimate annual capex of Rmb200–250bn against Alibaba's ambition to reach 20GW of cloud capacity by 2032 through owned and leased infrastructure. Our estimated 30%+ annual cash payback and 25%+ ROIC compare favourably with an approximately 10% cost of equity for the earlier equity placement. **Medium-term funding flexibility should improve as procurement urgency eases and earnings expand.** Less immediate pressure to secure scarce GPUs could allow greater use of debt and other funding channels. Stronger profitability and balance-sheet capacity should also support financing flexibility, although offshore funding remains subject to liquidity, leverage and regulatory requirements.

Figure 19: BABA's capex estimated at Rmb200-250bn in FY27E-29E



Source: Company data, UBS estimates

Figure 20: BABA's FCF is expected to turn +ve in FY29E



Source: Company data, UBS estimates

2) AI optionality at limited valuation premium

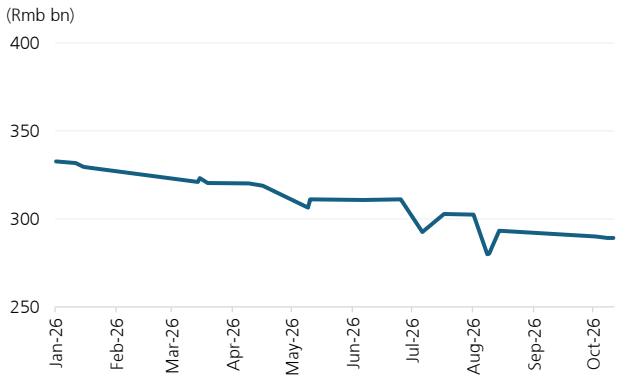
Tencent: Core earnings fund the AI option; FCF recovery protects the downside

Strong core earnings support a manageable AI investment cycle. We estimate cRMB300bn of annual operating cash flow against cRMB250bn of capex, while 2026E adjusted operating profit excluding new AI products grows c15%. Importantly, the core franchise still has monetization runway: we estimate payer penetration below 20% for key evergreen games and Video Accounts ad load at only c5%. Games and advertising therefore provide both the funding source for AI investment and bottom-up earnings growth independent of AI monetization.

AI-related earnings pressure should be less severe than feared. We estimate 2026E core operating profit excluding new AI investment at cRMB345bn. Assuming 10% growth in 2027E generates cRMB35bn of incremental profit; even if new-AI-product costs double from Rmb55bn to Rmb110bn, this offsets most of the incremental burden, leaving cRmb20bn of net OP pressure, or c7% of the 2026E adjusted OP base. We forecast adjusted net profit growth of –8%/–5% in 4Q26E/2027E versus consensus –2%/+3%. Further cuts may pressure sentiment, but at c12x P/E, spending duration and normalized FCF matter more than a few percentage points of near-term EPS. **We expect fuller recognition of the investment drag over the next one to two quarters to reduce the downgrade overhang and support re-rating.**

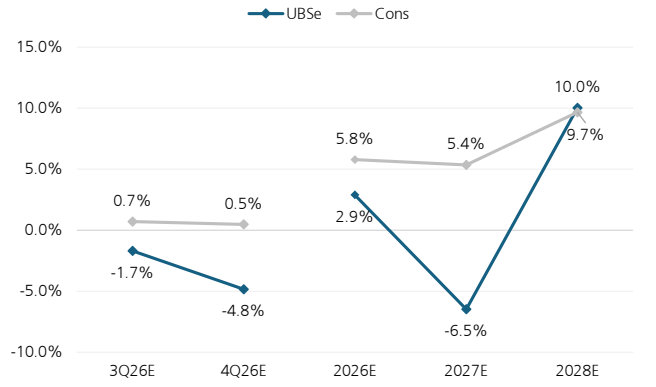
FCF can recover ahead of EPS if the investment cycle moderates. Heavy infrastructure investment creates a depreciation tail, so lower capex does not immediately restore reported profit. However, if industry AI returns disappoint, lower incremental procurement and prepayments can release cash materially faster. This creates an important valuation asymmetry: Tencent can participate aggressively while the opportunity is uncertain, but still retain a route to FCF normalization if expected returns deteriorate, in our view.

Figure 21: Cons 2027E adj. NP revision YTD



Source: Visible Alpha

Figure 22: Adj. OP YoY growth: UBSe vs cons

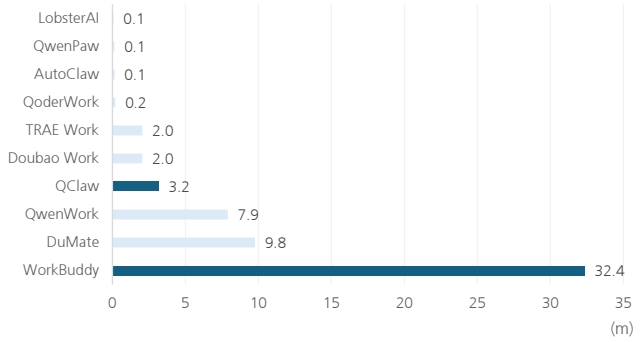


Source: Visible Alpha, UBS estimates

WorkBuddy validates Tencent's application capabilities, but subscriptions alone are unlikely to move group earnings. Strong early adoption supports the relevance of its harness engineering, model choice and Skills ecosystem, while deeper enterprise integration could reinforce distribution and generate useful task-execution feedback. Yet even an illustrative 10mn paying users × Rmb1,000 annual subscription × 50% margin produces cRmb5bn of profit, only c2% of group earnings. The larger opportunity is therefore workflow automation and capturing part of customers' labor-productivity gains, rather than seat subscriptions alone.

Xiaowei is longer-duration optionality, not a 2027 earnings thesis, in our view. Weixin potentially offers a differentiated combination of user context, Mini Programs, services and payments, but we think three hurdles remain: 1) economic sharing with merchants; 2) reliable multi-step execution at acceptable cost; and 3) incremental monetization of transactions that already occur within Weixin. Scaling before these economics are proven could deepen losses. Our base case is gradual feature expansion with limited 2027 monetization, leaving meaningful consumer-agent upside a potential option value.

Figure 23: Monthly interactions of PC-based AI-native office agents in China



Source: Analysys. Data as of Sep 2026.

Figure 24: Quantifying WorkBuddy's earnings impact to Tencent

China white collar population (m)	170
User penetration	50%
Paid conversion	20%
Tencent's market share in AI productivity	60%
Tencent's AI productivity paying users (m)	**10.2**
Annual subscription (Rmb)	1,000
Revenue from AI productivity (Rmb m)	**10,200**
Estimated NPM	50%
NP from AI productivity (Rmb m)	**5,100**
2027E adj. NP	248,270
% of 2027E adj. NP	2.1%

Source: CIC, UBS estimates

3) Bottom-up earnings recovery independent of AI

Meituan: Competition normalization unlocks a multi-year earnings recovery

Competition is rationalizing, improving visibility on a multi-year earnings recovery. Management's aggressive subsidy response following 2Q raised concerns that food-delivery unit economics could remain structurally impaired. We disagree with these concerns. We see competition shifting from subsidy-led share acquisition toward ROI, service quality and execution, supported by regulatory scrutiny and competitors' growing need to allocate capital toward AI. Our channel checks support this inflection: food-delivery UE improved to Rmb0.4–0.5/order in September from Rmb0.1–0.2 in July–August, while we expect 3Q in-store margins to remain around 25% despite competitive pressure. We therefore see the recovery as structural normalization, rather than a temporary subsidy pause.

The competitive cycle has also validated Meituan's moat. Our checks indicate approximately 70% GTV share in food-delivery orders above Rmb30, demonstrating resilience in the economically valuable customer segment. Meanwhile, PinHaoFan, at c15–20% of food-delivery volume, provides a broadly UE-breakeven route to serve price-sensitive users. The combination allows Meituan to defend both ends of the market without applying the same subsidy intensity across its user base. Merchant tools, satellite stores, Raccoon Kitchen and instant retail further deepen the consumer–merchant–fulfillment loop, raising the hurdle for competitors as the basis of competition shifts from subsidies toward execution.

Earnings recovery plus multiple normalization creates dual upside. We expect adjusted net profit to recover to approximately Rmb28bn/Rmb42bn in 2027E/2028E, without requiring margins to return to historical peaks. At approximately 9x normalized P/E, we believe the valuation underprices the visibility and duration of the recovery. The stock can therefore work through earnings upgrades first and multiple expansion second, with Keeta, Xiaoxiang Supermarket and potential portfolio monetization providing additional optionality rather than being required for our base case.

JOYY: Best of both worlds; Visible solid growth and attractive yield

JOYY offers an unusual combination of visible earnings growth and high cash returns. The stock is still largely perceived as a yield story, but we forecast high-teens 2026E–28E earnings growth alongside a 7–8% dividend yield. Based on our investor conversations, recognition has improved, but positioning remains centered on shareholder returns rather than the emerging growth contribution from Bigo Ads and Shopline. We see scope for re-rating as the narrative shifts from "yield with limited growth" toward "growth plus yield."

Multiple drivers for growth ahead: i) Bigo Ad: Following Bigo Ad's official collaboration with Google AdMob, we expect Bigo Ad's topline growth to maintain healthy momentum. Our checks also suggest Bigo Ad's algorithm upgrade and improving ROAS to support its c50%+ growth in 3P advertising revenue in 3Q and over the long term. ii) Shopline: The likely broader adoption of agentic AI for transaction execution should support the decentralization trend and migration toward independent websites among merchants. We believe this is favourable to Shopline's topline growth and loss narrowing ahead.

What could be a catalyst for further re-rating after the strong performance YTD? At approximately 11x 2027E P/E, or c3x ex-cash, we forecast c20%+ EPS growth, combining high-teens adjusted net-profit growth with an approximately 5% reduction in share count, alongside a c7% dividend yield. The next leg of re-rating should require the market to move toward an SOTP framework, as higher-growth ToB businesses contribute a larger share of earnings. Delivery on revenue growth and margin expansion is therefore more important than further yield compression alone.

Figure 25: JOYY's SOTP valuation

	Business segments	Revenue	OP margin	Target P/E (x)	Valuation (US$ m)	Per share (US$)	Mix (%)	Valuation methodology
1)	Live streaming	1,666	10%	5	634	12.4	14%	5x 2027E P/E
2)	Advertising	857	12%	10	1,141	22.3	26%	10x 2027E P/E
3)	Net cash				2,662	52.1	60%	2027E
	Target valuation				**4,437**	**86.8**		
	Valuation (US$)	86.8	end-2026E					
	Price target (US$)	**95.0**	end-Oct 2027E					

Source: Company data, UBS estimates

Boss Zhipin: Monetization runway, blue-collar growth and cash returns

Monetization, rather than user growth alone, should increasingly drive earnings. We expect the more proactive strategy initiated in 2Q26 to lift paying conversion and enterprise ARPPU, supported by differentiated services and improving employer ROI. The key test is whether higher monetization can be achieved without compromising retention and repurchase.

AI can strengthen the value proposition. The company has introduced AI across candidate screening, outreach, communication and interview organization. We believe deeper workflow integration could improve recruitment outcomes, customer stickiness and willingness to pay, allowing the platform to capture a greater share of hiring budgets. Moreover, we see stable competition and limited regulatory risk provide a favorable operating backdrop, and easing market concerns on AI disruption to job market.

Valuation offers growth plus cash returns. At c12x P/E, versus our 15% EPS CAGR and c10% dividend-plus-buyback yield, we see attractive risk/reward. Paying conversion, ARPPU and blue-collar expansion provide visible bottom-up earnings drivers without requiring an AI windfall.

Beyond preferred picks: Key debates and catalysts

JD: Loss reduction is the re-rating catalyst. At c6x 2027E P/E and a 7–8% dividend-plus-buyback yield, the debate is about the speed of earnings normalization, not the value of the core franchise. We forecast FY26 adjusted earnings of cRmb32bn, including cRmb20bn of food-delivery losses, versus cRmb48bn in 2024. Three factors determine the recovery: **1) food-delivery loss reduction**, with consensus assuming a 40–50% decline in 2027E and our expectation of greater focus on higher-value customers and higher-tier cities; **2) JD Retail resilience**, where efficiency and a higher 3P mix should offset softer demand and rising AI expenditure; and **3) discipline in Jingxi, overseas and other new businesses.** The proposed acquisition of Ceconomy is likely to offer incremental earnings optionality, subject to regulatory approval and execution. A credible ceiling on food-delivery losses is the clearest potential catalyst for re-rating.

NetEase: Defensive earnings with a pipeline catalyst. We view NetEase as a defensive non-AI opportunity supported by evergreen franchises and continued gross-margin discipline. H2 growth should remain subdued against tougher comparisons and a limited pipeline, but durable existing titles provide earnings visibility while the company waits for the next product cycle. At 14x 2026E P/E, or c13x excluding investment losses, valuation remains reasonable for the earnings resilience. Ananta's planned January 15, 2027 global launch provides the next meaningful catalyst.

Kuaishou: Asset-value support, but catalysts are still needed. We think the current valuation gives limited credit to either Kling or a recovery in the core business. Our SOTP comprises cRmb99bn net cash, Rmb84bn attributable to Kling based on our 68% stake/US$18bn valuation assumptions, and Rmb56bn for the core business at 6x normalized EBIT. Unlocking this value would require two developments, in our view: Kling ARR acceleration following the 4.0 upgrade, demonstrating commercial differentiation versus Seedance 2.5; and a core earnings trough, as we forecast 2H26 revenue down c8% and adjusted profit down c80% on operating deleverage. We think the stock is inexpensive on assets, but needs commercial AI traction or core stabilization to turn valuation support into a catalyst.

Bilibili: Under-monetized traffic plus a stronger game cycle. Bilibili combines structural advertising share gains with an under-monetized, highly engaged user base. Despite macro softness, we expect advertising to continue outperforming the industry as ad-tech improves, while disciplined AI spending should preserve the margin trajectory. The next leg increasingly depends on games: San Wang in late December and the 2027 pipeline including RO3, Three Kingdoms: The Ravages of Time and Guild Wars Card Game could return game revenue to growth and support further operating leverage. The debate is less about traffic and more about converting engagement into higher advertising and gaming monetization.

Trip.com: Strong franchise, but wait for the growth trough. We remain constructive on the long-term franchise but see limited urgency in the near term. Soft travel demand, fuel prices, geopolitical uncertainty and weather disruptions weigh on near-term revenue, while investors continue to debate domestic accommodation take rates and regulation impact. Our base case is low-single-digit revenue growth through 2H26 and potentially 1Q27, before acceleration from 2Q27. Domestic leadership and international optionality remain intact, but we prefer to wait for better visibility on the growth inflection.

TME: Membership competition clouds otherwise improving monetization. Competition remains the key constraint on the online music thesis. Emerging platforms such as Soda Music could continue to weigh on membership revenue growth even after the Ximalaya consolidation. We are more constructive on non-subscription and IP-led premium content, which can diversify monetization and strengthen the broader music ecosystem. The re-rating debate is whether these new revenue pools can grow fast enough to offset membership pressure without sacrificing the margin-expansion trajectory.

China foundation models (Z.AI, MiniMax): Rapidly evolving capability, monetization and competition landscape

Recent sentiment for AI labs including **Z.AI** and **Minimax** remains cautious, as investors concern about their future capability catch-up given lack of advanced compute yet larger training workload, and the sustainability of cost efficiency advantage in light of improvement of latest US model iteration (note). While valuation has corrected to the low end of the recent range, we await future model launches as catalyst (likely in Oct/Nov) to potentially recover momentum. Key debates:

1) Model capability: We note increasing market concerns around the capability gap between China and US, and expect the upcoming domestic launches as key focus. Recent US model launches, including GPT-6/6.1 and the Muse Wonder series, have pushed the frontier of model intelligence further, widening the capability gap between the latest US and Chinese foundation models. In our view, meaningful capability gains in the next generation of domestic models will be critical to restoring confidence in their

competitive positioning.

2) Token pricing and cost efficiency: We expect ongoing token price deflation, driven by tiered pricing and broader flash-model offerings, while overseas competition intensifies through direct price cuts, free-token quota resets and KV-cache optimization. At the same time, Chinese models' cost advantages are becoming less clear-cut across tasks. For complex, long-horizon tasks, overseas frontier models may offer stronger performance and lower cost per task despite higher per-token prices, and the market is monitoring whether Chinese models could deliver similar cost per task improvement in future iterations.

3) Compute expansion: The availability of advanced compute remains a key concern for future model training and monetisation outlook. Meanwhile, MiniMax and Z.AI continue to expand their compute capacity, resulting in higher AI investments and increased funding requirements. The increasing hardware cost, higher depreciation over time, and continued R&D investment could also pressure future profitability trends.