

China Internet

Post 4Q25: Re-evaluating the Sector; Tencent Our Top Pick

The six weeks since Tencent's 4Q25 print have seen a repricing of China internet names along a different axis than the AI-narrative rally that carried the sector into late January. Earnings revisions now dominate share price direction, AI capex is penalized rather than rewarded, and buy-side conversations have converged on a "real multiple plus earnings floor plus dated catalyst" framework for sizing positions. Against that tape, we refresh our sector view and rank Tencent as our top pick. Sector top picks: Tencent, JD, Trip.com, Baidu, Alibaba and Bilibili.

- We apply four filters consistently across our coverage universe: whether the current multiple represents a real earnings floor; whether disclosed AI spend is self-funded rather than dilutive; whether consensus earnings are set too low; and whether a dated catalyst exists in the next three to six months. Tencent is the only name that clears all four criteria. JD and Trip.com clear three on differentiated grounds. Alibaba, Meituan, and PDD clear fewer, for reasons we detail in the body.
- The case for Tencent rests on four pillars the market is, in our view, mispricing. The earnings floor is real: FY25 free cash flow of RMB182.6bn comfortably funds the current AI investment framework, and our 2026E EPS growth estimate of c8% already embeds elevated spend. The AI monetization bridge is narrower in scope and more advanced in execution than consensus frames it; the disclosed RMB18bn of 2025 spend and the "more than double" in 2026 framework refer to new AI products, not full company-wide infrastructure, and monetization is already live in advertising and gaming. We think the gaming and licensing moat is underappreciated because licensing, not model quality, is the scarce asset in China. Lastly, Weixin carries long-duration optionality to which the current multiple ascribes negligible value.
- Why now? Consensus has reset to "real 14x, earnings floor intact, awaiting a catalyst". Positioning is washed out. The April calendar is dense: Hunyuan 3.0 rolls out as the first real test of domestic AI credibility - the "Gemini 3 moment" analog investors explicitly reference - and there are low expectations for the Honor of King World launch. Cloud monetization should accelerate through the year, and a second model or AI application release later in 2026 provides a further potential catalyst.
- Our thesis is not without risk. Hunyuan 3.0 reception is genuinely binary. AI spend could overshoot the current framework. Core advertising or gaming softness would erode the floor the entire setup rests on. ByteDance, a private-company AI leader investors cannot directly own, may structurally cap listed multiples. We view these risks as priced in; the potential upside is not.

Head of China Equity Research and Co-Head of Asia TMT Research

Alex Yao ^{AC}

(86 21) 6106 6505

alex.yao@jpmorgan.com

SAC Registration Number: S1730523020001

Daniel Chen

(86-21) 6106 6205

daniel.q.chen@jpmorgan.com

SAC Registration Number: S1730521040001

Nancy Liu

(86 21) 6106 6343

nancy.liu@jpmorgan.com

SAC Registration Number: S1730524090001

J.P. Morgan Securities (China) Company Limited

The six weeks since the 4Q25 reporting season has seen a repricing of China internet stocks along a different axis than the AI-narrative rally that ran through January 2026. Earnings revisions now dominate share price direction, AI capex is penalized rather than rewarded, and the buy-side has converged on a "real multiple plus earnings floor plus dated catalyst" framework for sizing positions. Against this tape, we refresh our sector view and rank Tencent (0700 HK) as our top pick.

Our conviction rests on four filters we apply consistently across the coverage universe: whether the current multiple represents a real earnings floor; whether AI spend is self-funded rather than dilutive; whether consensus earnings are set too low; and whether a dated catalyst exists in the next three to six months. Tencent is the only name that passes all four. The setup is asymmetric: downside looks capped near 14x 2026E PE with management preserving financial flexibility, and upside is gated by catalyst delivery in April and beyond rather than by a multi-quarter narrative build.

We retain a constructive position on JD and Trip.com on a more selective basis, remain Neutral on Meituan, and view PDD as having a balanced risk/reward post the recent multiple reset. This report focuses on the sector reset and the relative case for Tencent.

Figure 1: YTD sector share price performance summary

Company	Ticker	Market cap (US\$m)	Price	Stock Return	Relative return vs. MXCN Index	Change in 26E PE	26CY EPS revision	26CY Revenue revision
Tencent	700 HK	604,830	519	-16%	-13%	-12%	-4%	0%
Alibaba	BABA US	323,208	135	-7%	-4%	25%	-2%	-1%
Pinduoduo	PDD US	140,531	99	-12%	-9%	-8%	-4%	0%
NetEase	NTES US	73,260	115	-18%	-15%	-15%	-3%	-2%
Meituan Dianping	3690 HK	68,165	86	-18%	-16%	-183%	-19%	-2%
Zhipu AI	2513 HK	54,089	950	749%	751%	1043%	-2%	77%
Baidu	BIDU US	42,018	123	-6%	-3%	3%	-8%	-3%
JD.com	JD US	41,685	31	7%	9%	21%	-1%	0%
MiniMax	100 HK	35,987	899	462%	465%	315%	36%	9%
Trip.com	TCOM US	35,719	55	-25%	-22%	-20%	-7%	1%
Kuaishou	1024 HK	25,747	46	-29%	-26%	-2%	-2%	-4%
Beike	BEKE US	18,841	16	1%	4%	11%	-9%	-9%
Didi	DIDIY US	17,715	4	-25%	-23%	688%	-9%	0%
TME	TME US	15,141	10	-45%	-43%	-43%	-4%	-2%
Bilibili	BILI US	9,933	23	-7%	-5%	6%	-1%	1%
Alibaba Health	241 HK	9,750	4.7	-9%	-6%	-9%	0%	0%
Vipshop	VIPS US	7,079	14.7	-17%	-15%	-17%	0%	-2%
Kanzhun	BZ US	6,682	14	-33%	-30%	-31%	-4%	-1%
Kingsoft Cloud	KC US	5,253	17.6	67%	69%	46%	14%	9%
Kingsoft	3888 HK	4,327	24.5	-14%	-11%	1%	-1%	-4%
XD	2400 HK	4,062	65	-2%	0%	-3%	1%	-6%
Ping An Healthcare	1833 HK	3,381	12.3	-15%	-12%	-22%	9%	-1%
YY	JOYY US	3,091	60.6	-9%	-6%	-6%	-3%	2%
Weride	WRD US	2,716	8	-7%	-5%	261%	-7%	-14%
Damai Entertainment	1060 HK	2,335	0.6	-37%	-34%	-37%	0%	-1%
AutoHome	ATHM US	2,154	18.7	-16%	-13%	-3%	-1%	-13%
Weibo	WB US	2,113	9	-16%	-14%	-11%	-6%	6%
iQiyi	IQ US	1,274	1.3	-35%	-32%	171%	-7%	-6%
Momo	MOMO US	1,207	6.3	-4%	-1%	-3%	-1%	-1%
Maoyan	1896 HK	893	6.0	-15%	-12%	9%	-2%	-10%
Huya	HUYA US	682	3	11%	14%	4%	7%	3%
Dingdong	DDL US	631	2.7	6%	9%	-1%	7%	1%
KWEB	KWEB US EQUITY	6,605	29.4	-14%	-12%			
MSCI China	MXCN INDEX	12,457,231	81.9	-3%	0%	-11%	10%	9%
Mean				27%	30%	68%	-1%	1%

Source: Bloomberg Finance L.P. (as of Apr 23, 2026).

What Has Changed Since the 4Q25 Earnings Season

The 4Q25 reporting cycle that began on 18 March with Tencent's print did not produce large absolute earnings surprises, but it did reset the framework through which the market is underwriting China internet. In our view, six shifts matter.

Market regime inversion. Through late January 2026, Alibaba and Baidu rallied on an AI narrative despite weak fundamentals. After late January, profit-taking and a lower global risk appetite re-anchored the tape on near-term earnings and free cash flow. Most of the China internet names saw estimate cuts into and out of the 4Q print. Alibaba's US\$100bn five-year cloud target, implying roughly a 47% cloud revenue 5-year CAGR did not support the stock. Long-dated narrative claims now carry a materially lower multiple than they did in 4Q25.

AI capex is being penalized, not rewarded. The US software analog rerating, IGV (iShares expanded Tech-Software Sector ETF) weakness and fears of disruption to core franchises have arrived in China with a lag. Linear extrapolations of AI spend (Tencent RMB36–40bn this year per the company), Alibaba's rapidly scaling of the "all others" line) produce ROI anxiety that the market is translating into multiple compression rather than multiple expansion. The burden of proof has shifted from "show us AI monetization is possible" to "show us AI monetization within the spend framework."

Regulatory wave re-opened, but bounded. The SAMR ghost-kitchen ruling landed on 18 April, with RMB35.97bn in aggregate industry penalties. PDD's RMB1.5bn fine is financially immaterial but carries signaling weight given unusually vivid domestic media coverage. A separate State Council-linked office probe into food delivery subsidy practices remains live. Trip.com's antitrust investigation, launched in early January, should be resolved on a four-to-six-month precedent timetable, implying May to July 2026. In our view, these are clearing events within an already-established framework, not a 2021-style systemic reset, but stocks tied to unresolved cases continue to trade as "dead money" until resolution.

Food-delivery subsidy intensity easing at the margin. Alibaba's FY27–29 loss-narrowing roadmap, roughly halving each year from a base of RMB85–90bn, is now in consensus and largely priced in. The incremental underwriting question has shifted away from food delivery to Alibaba's "all others" AI spending line, where the cadence accelerated from approximately RMB3bn in the September quarter to RMB20bn in the March quarter. Scale remains small relative to food delivery losses, but the acceleration and the opacity are what matter.

Buy-side positioning reset. Late-April counterparty conversations are converging on "14x is real, earnings floor intact, needs a model-quality or gaming catalyst" for Tencent; on "JD has the most FD-easing upside surprise"; on "Trip.com is attractive on 11–12x if the near-term earnings outlook won't be compressed by investment"; and on Alibaba as a multiple-not-EPS story. Expectations are washed out across most of the cohort. Investors are unwilling to size positions ahead of dated events, but are willing to underwrite real earnings floors.

Private-company overhang has intensified. Buy-side framing now treats ByteDance as the China-specific analog of Anthropic: the strongest AI competitor that investors cannot directly own, structurally capping listed multiples. This is a relative-positioning risk, not a disintermediation risk, and the two should not be conflated, in our view.

Revised Sector Framework

We apply four filters to rank names in the current regime. Each filter corresponds to a feature of the 2026 tape rather than a view on any single company.

Filter one: is the current multiple a real floor? A low multiple is only meaningful if earnings downside is capped. We test this through three inputs: management signaling on the earnings trajectory, our own estimate range for 2026 EPS, and free cash flow durability against disclosed investment plans. Tencent and Trip.com both clear this filter. Meituan does not clear it cleanly because normalized food delivery economics still span a wide range (roughly RMB0.2–0.3 per order in the low scenario versus RMB1.0–1.1 per order in the high scenario), depending primarily on Alibaba's strategic posture rather than on policy.

Based on Tencent’s AI investment of Rmb18bn in 2025 (disclosed by management during the 4Q25 earnings call), we estimate that consensus earnings have largely factored in incremental investments in AI in 2026; we estimate even a Rmb60bn AI investment (over 3 times that of 2025’s level) would only imply 6% downside to consensus earnings. This would imply limited room for further earnings cuts, in our view.

Table 1: Tencent new AI product investment analysis and EPS sensitivity

Tencent 2026 new AI product investment (JPMe, Rmb bn)	0	15	30	45	60	75
Delta to 2025 AI investment	-100%	-17%	67%	150%	233%	317%
Tencent 2026E adj NI (Rmb bn) (JPMe)	315	302	290	277	264	251
Delta to consensus	12%	8%	3%	-1%	-6%	-10%

Source: J.P. Morgan estimates, Bloomberg Finance L.P. (as of Apr 23, 2026).

Figure 2: Tencent trades close to 10 year trough



Source: Bloomberg Finance L.P. (as of Apr 23, 2026)

Filter two: is AI spend self-funded or dilutive? We look for three pieces of evidence: explicit scoping of the disclosed spend figure (new products versus total infrastructure), visibility of AI-linked revenue contribution in existing businesses, and a credible monetization sequence. We believe that Tencent’s RMB18bn new AI product investment in 2025 and the roughly RMB36-40bn 2026 framework (per the company) refer to investment in new AI products rather than full company-wide AI infrastructure, with advertising and gaming already benefiting and cloud next in sequence. Alibaba has not provided equivalent scoping on the "all others" line, which we think is the primary reason multiples on BABA trail the underlying operating improvement.

Filter three: is there out-of-consensus earnings upside? We look for gaps between consensus and our numbers that could close on a defined print. JD ranks highest on this filter given a meaningful quarter-on-quarter improvement in food-delivery and new-business losses, while consensus appears to assume roughly flat year-on-year new business losses. Meituan's substantial next-year profit improvement is already embedded in consensus and is therefore priced in.

Filter four: is there a dated catalyst in the next three to six months? We look for binary events that could reopen the narrative on a defined timeline. Tencent's April calendar, Hunyuan 3.0 rollout and Honor of King World launch, is the densest in the cohort. Trip.com's potential antitrust resolution between May and July would be a clearing event. JD's 1Q26 print is its primary earnings catalyst. Alibaba's next inflection point depends on the performance of its "all others" business segment, which is not a single-date event.

Applying these filters, Tencent is the only name that clears all four. JD clears three (floor, out-of-consensus earnings, catalyst). Trip.com clears three (floor, clearing event, manageable near-term earnings risk). Baidu clears three (floor, self-funded AI investment, catalyst). Alibaba clears one. Meituan clears one (catalyst, through FD easing, although already priced). This ranking is the spine of our top-pick call.

Why Tencent Tops the Ranking

Thesis

Tencent trades at roughly 14x 2026E PE based on Bloomberg consensus, offering a credible earnings floor with asymmetric optionality on sequenced, self-funded AI monetization. Downside is capped by fundamentals; upside potential is catalyzed rather than narrative-gated. This is the cleanest structure in the sector, in our view.

Pillar one: the earnings floor is real. FY25 produced revenue of RMB751.8bn, non-IFRS net profit of RMB267.0bn, and free cash flow of RMB182.6bn. Our 2026 EPS growth estimate of c.8% already embeds elevated AI investment. Management has framed 2026 as a year when revenue growth will outpace earnings growth by design, not because cost discipline is deteriorating. That framing is compatible with earnings that continue to grow, even if modestly, and is the foundation of the floors, in our view.

Pillar two: AI spend is self-funded, not dilutive. The disclosed RMB18bn of 2025 spend and the roughly RMB36-40bn 2026 framework refer to new AI products, not Tencent's full company-wide AI infrastructure. Linear extrapolations to RMB70bn or RMB100bn in future years therefore mis-scope the underwriting, in our view. More importantly, monetization is already live in the businesses that investors can underwrite today: advertising, where AI is improving targeting and click-through, and gaming, where AI is accelerating the production of skins, assets, and monetizable content on top of a large licensed installed base. Cloud is the next layer as GPU capacity expands. Token monetization and Weixin agent monetization remain back-end-loaded but carry the largest long-duration optionality.

Pillar three: The gaming and licensing moat is underappreciated, in our view. In China, game licensing remains the scarce asset, not model quality. Tencent's edge in AI-enabled gaming is its orchestration across the production and live-operations stack, deployed onto the largest licensed installed base in the market. This is not a moat that a frontier-model competitor can replicate without also securing license approvals and distribution. Honor of King World launches into this moat.

Pillar four: Weixin carries long-duration optionality the market is not paying for. Management has framed the core consumer AI opportunity as improving task completion and deepening transaction behavior inside Weixin, payments, mini-shops, commerce, and related advertising, rather than as standalone token sales. This is consistent with the public framing from Shunyu Yao - that Tencent's consumer AI edge is augmenting user context rather than pure model scaling. We do not underwrite this pillar in the 2026 numbers, but we note that the current multiple ascribes negligible value to it. Therefore, the Google re-rating mechanism (model quality improvement, distribution surface, token monetization, multiple expansion path) offers potential upside to the share price, in our view.

Why now?

Consensus has converged on "real 14x, earnings floor intact, needs a catalyst." Positioning is washed out. The April catalyst calendar is loaded: Hunyuan 3.0 rolls out as the first real test of domestic AI credibility - the "Gemini 3 moment" analog the buy-side explicitly references - and Honor of King World launches into a low-expectation tape. Cloud monetization should accelerate as GPU capacity expands, and a second major model or AI application release later in 2026 would provide a second bite. Google's rerating from the mid-teens to the high-20s on Gemini 3 and product integration is the most relevant template, in our view. We do not claim Tencent will trace the same path, but we observe that the mechanism, a credible model inflection combined with a distribution surface, is the same.

Variant perception

The market is pricing Tencent as an AI-cost-burden story. We think the causality is backwards in two specific ways. First, the disclosed spend frame is narrow, not total. Second, the ByteDance overhang is a relative-positioning concern, not a

disintermediation concern; if Tencent solves for "good enough" model quality, Weixin distribution handles the conversion to adoption.

Valuation

We anchor Tencent on a 12x 2026E PE floor. Our base case is 16–18x on catalyst delivery and partial narrative repair. A bull case towards 21x requires both a credible Hunyuan inflection and visible AI-linked revenue disclosure. Under current market conditions, we view 21x as the bull rather than the base.

Risks to the Sector and to Our Top Pick

Sector risks. Three developments would invalidate the sector framework. Regulatory re-escalation beyond current cases, particularly any ruling that changes the economic substance of platform take rates rather than just the form, would reset the earnings floor logic across multiple names. A resumption of the food delivery subsidy war, triggered by a change in Alibaba's strategic posture, would widen normalized earnings ranges across Meituan, Alibaba, and JD and break the "FD-easing" leg of the tape. A deeper global risk-off episode or a sustained absence of China inflows would cap multiple expansion even where fundamentals deliver; several late-April counterparty conversations explicitly identified inflows, not fundamentals, as the binding constraint.

Tencent-specific risks. Four risks matter for our top-pick call. Hunyuan 3.0 reception is genuinely binary: a "still behind" verdict from the market would re-open the capability-lag debate and extend the narrative-repair timeline. AI spend overshoot above the current RMB36–40bn framework would re-trigger the capex penalty regime, particularly if disclosed without matching revenue evidence. Softness in core advertising or gaming, either consumption-driven or from short-video share shift, would erode the earnings floor that the entire setup rests on. ByteDance's launch of a consumer-grade agentic product would not disintermediate Tencent but would likely compress the relative multiple the market is willing to pay.

What Would Change Our View

We would reassess the top-pick call on any of the following.

On Tencent, a combination of a mixed-to-negative Hunyuan 3.0 reception, and a 1Q26 print showing advertising growth decelerating below mid-teens digits would, in aggregate, weaken the thesis meaningfully. Conversely, a Hunyuan 3.0 release judged to be in the GLM 5.1 class, visible AI-linked revenue disclosure alongside the 1Q26 print, and Honor of King World achieving a top-three grossing position in its launch month would move us from the 14x floor toward the 21x bull case sooner than our base case anticipates.

On the sector framework, we would revisit the "earnings floor plus catalyst" template if two or more 1Q26 prints break below consensus on the core earnings line rather than only on below-the-line items. That pattern would indicate the regime has moved further, from capex-penalty toward outright earnings-quality anxiety, and would tighten the filters we apply.

