China Internet

Cheap is acknowledged; urgency still needs proof

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

Olivia Xu
(86-21) 6106 6138
olivia.w.xu@jpmorgan.com
SAC Registration Number: S1730525060001

Nancy Liu
(86 21) 6106 6343
nancy.liu@jpmorgan.com
SAC Registration Number: S1730524090001

J.P. Morgan Securities (China) Company Limited

Based on our one-week long marketing meetings with US-based investors, we believe US investors remain engaged on China Internet, but the trip did not suggest a broad rotation back into the sector. The consistent message from meetings was that valuation is no longer the central debate. Most investors accept that the large-cap platforms are inexpensive. What they do not yet accept is that cheapness alone is enough to compete for capital against AI hardware, memory, semis, and private or newly listed AI assets where near-term earnings revisions or scarcity narratives remain cleaner.

- That leaves the sector in a selective, proof-driven regime, in our view. Potential marginal buyers are not asking whether China Internet can screen cheap. They are asking where the next observable catalyst is strong enough to change behavior. On this trip, the most actionable debates were Alibaba cloud margin, Tencent AI disclosure, Baidu's Kunlunxin and capital return path, Zhipu relative to MiniMax, and Kuaishou's value capture from Kling.

- Our takeaway is not to press broad sector beta. We would focus marketing and positioning around names where a dated catalyst or measurable proof point can do the work. Alibaba had the clearest tactical traction among the large platforms because cloud margin is a near-term reported variable. Baidu drew event-driven interest around Kunlunxin, capital return, and potential listing-related catalysts. The Zhipu over MiniMax relative preference resonated well, although the debate has shifted from model quality to commercial durability. Tencent remains highly respected but was rarely treated as urgent; investors want AI monetization KPIs before assigning more value to the option. Meituan was viewed as regulation-positive but not necessarily the cleanest expression of food-delivery rationalization, with JD repeatedly raised as an alternative revision vehicle. Kuaishou interest remained concentrated in Kling, with the main question being how much value ultimately accrues to the listed parent.

- We believe US investors are not dismissing China Internet. They are asking for proof before reallocating capital. The sector has investable assets, and valuation support is real. But the market's burden has moved from "is it cheap?" to "what changes revisions, AI monetization, or catalyst density soon enough to matter?"

- In that regime, we would not chase broad beta. We would lean into where proof is the closest: Alibaba cloud margin, Baidu's catalyst path, and the Zhipu over MiniMax relative preference. We would hold Tencent as patient core but not press the entry without AI KPIs. We would keep Meituan and Kuaishou catalyst-dependent, with JD remaining the cleaner food-delivery revision expression if rationalization shows up there first.

The positioning regime: engaged, but still underweight in behavior

US-based investors are not rejecting China Internet. It is deferring the trade, in our view.

See page 7 for analyst certification and important disclosures, including non-US analyst disclosures.

Across the meetings, the broad pattern was consistent. Investors were familiar with the valuation argument and generally accepted it. The pushback was timing. For US allocators, the opportunity cost remains high: AI hardware, memory, semis, and private AI assets still offer cleaner near-term revision stories or scarcity value. Against this backdrop, China Internet is respected but not yet urgent.

In our view, the biggest challenge to China Internet is capital allocation. Investors can agree that Tencent, Alibaba, Baidu, Meituan, JD, Kuaishou, and select AI model labs have investable assets, while still choosing not to fund them today. The question that came back repeatedly was: what changes the near-term revision path, reported AI return profile, or catalyst density enough to pull money out of the existing AI hardware and scarcity trades?

This is why the trip supports a more selective framework. We would separate the sector into three groups:

1. **Names with near-term proof points that investors can underwrite:** Alibaba cloud margin, Baidu's Kunlunxin and capital return path, and the Zhipu over MiniMax relative debate.

2. **High-quality names that need disclosure before urgency returns:** Tencent is the clearest example.

3. **Names where the asset is interesting but the equity bridge remains unproven:** Meituan, Kuaishou, and to a lesser extent any broad beta expression without a catalyst.

We believe the feedback signal is the strongest at the sector level: cheapness is acknowledged, but opportunity cost is still winning. It is somewhat less precise at the name level, because individual investor preferences varied by mandate, holding period, and exposure to AI hardware. Still, the same proof burden repeated often enough to be useful.

What investors pressed on

The trip was not about whether China Internet looks cheap on screens. It was about whether reported proof can arrive quickly enough to compete with the current AI trade.

Four debates ran through the week.

First, what breaks the long hardware / underweight internet regime?

Investors recognize that AI hardware and memory are crowded. But most did not want to fade that trade without a trigger. A valuation gap by itself was not considered enough. A rotation would likely require either disappointment in hardware revisions or positive revision breadth in China Internet across multiple large caps

Second, can AI investment show up in reported numbers?

Investors were less interested in broad AI narratives than in reported evidence. For Tencent, the focus was Hunyuan, WeChat agents, WorkBuddy, cloud APIs, advertising monetization, and whether any of those surfaces can become visible in revenue or margin. For Alibaba, the debate was more concrete: whether cloud margin can expand as AI demand scales, and whether losses in consumer AI or local services offset the cloud benefit.

Third, does China's lower-cost AI path create durable economics?

Investors accepted that Chinese models may operate on a different cost curve. The question was whether that translates into commercial advantage, pricing power, and retention, or whether the model layer remains structurally competitive with limited rent capture. This was especially relevant to Zhipu and MiniMax.

Fourth, does food-delivery regulation reset economics or simply create a tradeable window?

Investors were broadly willing to treat the SAMR rules as regulation-positive for the sector. But they were more divided on equity expression. Meituan was not rejected, but several investors framed JD as the cleaner revision surprise if rationalization holds.

The common thread is that investors wanted proof with dates, not just frameworks. Cloud margin, AI KPIs, unit economics, listing events, lock-up clearance, capital return, and revenue disclosure all mattered more than broad sector valuation.

US versus Asia: same regime, different lens

Our Asia marketing ran from early April through late June, while the US feedback reflects one week of meetings from June 15–22. The two periods are not perfectly comparable. The tape, hardware trade, and several company-specific catalysts moved during the Asia window. We therefore read the comparison as directional rather than precise.

With that caveat, both cohorts broadly agreed on the central regime: China Internet is inexpensive, but opportunity cost against AI hardware, memory, domestic chips, and model labs remains high.

The difference was in how the debate was framed. Asia investors tended to spend more time inside the operating mechanics: Alibaba cloud margin, Tencent Hunyuan and agents, Baidu Kunlunxin, food-delivery anti-involution, and AI cloud returns. US investors were more direct about the allocation problem: why own China Internet now when the AI hardware and scarcity trades still have cleaner revisions?

That distinction matters. Asia could become the earlier marginal bid if reported proof lands first in Alibaba cloud, Tencent AI disclosure, Baidu's event path, or Kuaishou/Kling value discovery. But Asia may also simply be expressing the same caution with more name-level detail. We would not overstate a regional divergence. The real divide is not US versus Asia. It is proof versus narrative.

Feedback by name

Tencent: respected, but still waiting for AI disclosure

Tencent remains the highest-quality large-cap platform in the group, but respect did not convert into urgency on this trip.

We carry Overweight on Tencent based on compounding earnings, bounded AI investment risk, and a trough multiple. Investors generally accepted the quality argument. The pushback was not that Tencent lacks assets. It was that the AI revenue path still feels too long-dated relative to what investors can own elsewhere today. For many investors, Tencent is not being rejected; it is being deferred.

The key debate is whether Tencent is a long-duration AI winner or simply a high-quality funding source for more immediate AI hardware exposure. The Hy3 release helped narrow the execution discount, but the market still wants disclosure rather than narrative. Investors asked for evidence around Hunyuan at scale, WeChat agents, WorkBuddy, cloud APIs, and advertising monetization.

The implication for our call is clear. Tencent remains a patient core position, not a name we would press purely on valuation. To shift from respected compounder to urgent buy, Tencent needs AI KPIs that can be connected to reported revenue, margin, or engagement over the next few quarters. Without that, the AI option remains outside base earnings in investors' minds.

Alibaba: the strongest large-cap tactical set-up

Alibaba had the most tactical traction among the large platforms because the debate is tied to a dated and measurable proof point: cloud margin.

Investors were no longer debating whether Alibaba has an AI narrative. They were debating whether the cloud margin path is real, how quickly it can scale, and how much of the benefit leaks into consumer AI or local services losses. Our most recent review raised our 12-month conviction after external cloud growth near 40% year on year, the first MaaS ARR disclosure, and a more visible path for margin expansion.

The important nuance is that the cloud-margin thesis is no longer undiscovered. Many investors already understand the bull case: external cloud margin rising from the high-single-digit to low-double-digit range toward a higher long-term level. The investable edge is not in repeating that AI cloud is important. It is in assessing the slope of margin expansion and the degree of reinvestment leakage elsewhere.

That makes Alibaba the cleanest tactical large-cap expression of the current regime, in our view. The catalyst is measurable, near-term, and tied to reported numbers. If cloud margin expands sequentially while losses in local services and consumer AI remain controlled, Alibaba can shift from tactical AI beneficiary to a more durable compounder. If the margin path stalls or is offset by losses elsewhere, the current traction likely fades.

Baidu: catalyst-dense, but still more event path than platform call

Baidu drew event-driven interest because Kunlunxin and capital return provide a denser catalyst path than broad sector beta.

Our most recent published view frames cloud inference momentum and Kunlunxin demand as the swing variable. Investors were interested in several potential catalysts: Kunlunxin IPO progress, a Hong Kong primary listing, Stock Connect inclusion, the first dividend, and a possible special distribution. Those are concrete events, and they help Baidu stand out in a sector where many investors still ask what the near-term trigger is.

The pushback was durability. Investors questioned whether Baidu is an asset-monetization trade or a sustainable AI platform. Legacy search remains under pressure, capex is below the largest platforms, and communication has not always helped investors underwrite the transition.

In our view, Baidu's event path can work, but long-only sizing likely requires a clearer capital return framework, more confidence in parent earnings quality, and better evidence that AI cloud revenue is durable rather than episodic.

Meituan: regulation-positive, but not the cleanest expression

Meituan was viewed as a beneficiary of food-delivery rationalization, but the feedback did not support upgrading on regulation alone.

We are Neutral on Meituan with a fair value of HKD88. Investors generally accepted that SAMR subsidy rules are directionally positive for industry economics. The harder question was whether Meituan is the best way to express that view. Several investors asked whether the revision surprise is cleaner in JD.

For Meituan, the core debate remains normalized unit economics. Investors focused on the gap between management's roughly CNY1 per order ambition and bear cases closer to CNY0.20–0.30. The question is whether regulation protects Meituan's profit pool or exposes lower normalized economics once rational competitors regain share.

Our thesis continues to rest on durable unit economics and high-value order share. The trip reinforced our reluctance to upgrade on regulation alone. We need 2Q26 and 3Q26 to show that the improvement is owned rather than rented. If the sector rationalizes but JD captures the cleaner revision surprise, then JD may remain the better expression of the same theme.

JD: the unprompted alternative in food-delivery rationalization

JD was not a lead marketed name on this trip, but it repeatedly surfaced as the preferred alternative delivery rationalization trade.

That is important feedback. Where investors were cautious on upgrading Meituan, several framed J n vehicle. If the thesis is rationalization plus positive earnings surprise, JD may be perceived as havir consensus and a more straightforward path to earnings revision.



We carry Overweight on JD with estimates above consensus, so this feedback aligns with our existing positioning. Because JD was not systematically tested across all meetings, we would treat the signal as medium confidence rather than definitive. But the fact that investors volunteered the comparison is meaningful. It raises the bar for Meituan: Meituan needs to win on its own unit economics, not simply benefit from the same regulatory tailwind that investors may prefer to express through JD.

Zhipu: preferred listed model-lab exposure, but the burden has moved to commercial proof

Zhipu was one of the most positively received relative calls on the trip. Investors generally preferred Zhipu to MiniMax as listed model-lab exposure, citing model quality, pricing power, capital access, and the fundraising flywheel.

Our long case on ship should not rest only on generic base-model leadership. It needs to rest on coding as a defended workflow vertical where Zhipu can show paid demand, pricing power, retention, and enterprise value capture.

That distinction matters because the model-lab scarcity premium may compress as more assets list or become available to public-market investors. Kimi, DeepSeek, and others may all affect how investors price scarcity. In that environment, benchmark progress is necessary but not sufficient. The commercial test is more important: can Zhipu defend economics inside the coding workflow, retain users, and capture enterprise value despite low switching costs?

The feedback supports the relative preference for Zhipu over MiniMax. It does not remove the need for commercial proof.

MiniMax: least well received; scarcity alone is not enough

MiniMax drew the most negative relative feedback among the names discussed.

Investors focused on the failed price increase, the pricing floor set by DeepSeek, Zhipu's capital advantage, and the lock-up that clears around July 8. The recurring question was whether scarcity can continue to support the stock if relative fundamentals remain weak.

We are cautious on names whose thesis depends mainly on base-model differentiation. MiniMax is the clearest case. The next model cycle, pricing power, customer retention, and technical selling pressure around the lock-up matter more than headline parameter counts. A clean lock-up clearance could support the stock tactically. But without evidence of pricing power and retention, scarcity is a fragile foundation.

Kuaishou: Kling is interesting; parent value capture is the debate

Kuaishou interest was concentrated almost entirely in Kling rather than the core business.

Investors cited a private valuation around US$16-18bn and a possible IPO valuation of US$25bn–30bn within six to twelve months. We treat those figures as investor-stated and unverified, not as our valuation target. The point is not the exact number. The point is that investors see Kling as a potential hidden asset, while remaining unsure how much value reaches Kuaishou shareholders.

The equity bridge is the entire debate, in our view. Investors need clarity on ownership, dilution, listing structure, holdco discount, revenue trajectory, model quality, and the degree to which any Kling value can offset pressure in the core short-video business. Without that bridge, Kling remains an interesting asset rather than a clean equity thesis.

Where the trip sharpened our thesis

The trip did not change the sector call. It sharpened where we should demand proof.

For the sector, broad cheapness is not enough. China Internet needs catalyst density, revision stabilization, and observable AI payback to compete for capital. The feedback argues against chasing broad beta without proof.

For Tencent, the call stands, but the burden shifts further toward AI KPIs and disclosure. Tencent remains a patient core Overweight rather than a valuation-led urgency call.

For Alibaba, the feedback strengthens the tactical case. Cloud margin is the cleanest large-cap proof point, and this is where reported traction can most directly change positioning.

For Baidu, the event path is more compelling than the platform narrative. Kunlunxin and capital return can work, but sustainable rerating needs parent earnings and AI cloud quality.

For Meituan, the trip reinforces our Neutral stance. Regulation is supportive, but the market wants unit economics proof. JD remains a cleaner expression if rationalization shows up first in revisions.

For Zhipu and MiniMax, the relative call strengthened. Zhipu is preferred, but the long thesis has to be defended by workflow economics. MiniMax needs more than scarcity to stabilize investor confidence.

For Kuaishou, the debate moves from “is Kling valuable?” to “how much of Kling’s value belongs to Kuaishou shareholders?”

What would change behavior?

The next test is not whether investors continue to agree that China Internet is cheap. The test is whether dated catalysts change allocation behavior.

Sector and flow indicators

By the end of 3Q26, we would want to see positive EPS revision breadth across at least two large-cap China Internet names. If revisions turn, the funding-source discount should begin to compress. If they do not, we would keep broad beta light and remain name-specific.

We would also watch whether AI hardware, memory, and semi’s earnings disappoint enough to force rotation. A genuine disappointment in the current AI hardware trade remains the cleanest external catalyst for China Internet to stop being treated as a funding source. Without that, China Internet needs to earn allocation through name-specific proof.