# China Internet Sector

## Presentation - China AI and online verticals update



**Equities**

China

Internet Services

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

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

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

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 34.**

Section 1

# Sector overview

# Coverage list

| Analyst | Company name | Ticker | Mkt cap (USD m) | Rating | Price target (LC) | Price (LC) | YTD return | 1M return |
|---|---|---|---|---|---|---|---|---|
| Wei Xiong | Knowledge Atlas Technology | 2513.HK | 29,357 | Buy | 1,900.00 | 513.50 | 342% | -29% |
| | Baidu, Inc. | BIDU.O | 28,193 | Buy | 160.00 | 81.86 | -37% | -11% |
| | Trip.com | TCOM.O | 25,350 | Buy | 60.00 | 38.88 | -46% | -1% |
| | MiniMax | 0100.HK | 8,160 | Buy | 500.00 | 179.50 | 9% | -29% |
| | Iflytek | 002230.SZ | 13,030 | Buy | 55.00 | 36.30 | -28% | -8% |
| | Tencent Music Entertainment | TME.N | 13,042 | Neutral | 9.50 | 8.48 | -50% | 5% |
| | Kanzhun Limited | BZ.O | 7,366 | Buy | 21.00 | 16.35 | -17% | 10% |
| | Tongcheng Travel | 0780.HK | 2,860 | Buy | 18.00 | 9.88 | -55% | -8% |
| | NetEase Cloud Music | 9899.HK | 2,318 | Neutral | 100.00 | 84.85 | -54% | -3% |
| | Kingsoft Cloud | KC.O | 2,398 | Buy | 19.00 | 8.76 | -15% | -11% |
| | Wondershare | 300624.SZ | 1,445 | Neutral | 65.00 | 50.03 | -29% | -1% |
| Kenneth Fong | Tencent Holdings | 0700.HK | 490,121 | Buy | 740.00 | 417.80 | -29% | -3% |
| | Alibaba Group | BABA.N | 258,026 | Buy | 206.00 | 111.17 | -23% | 2% |
| | PDD Holdings Inc | PDD.O | 113,175 | Buy | 146.00 | 81.45 | -28% | 3% |
| | Meituan | 3690.HK | 59,700 | Buy | 132.00 | 75.50 | -27% | 1% |
| | JD.com | JD.O | 37,979 | Buy | 48.00 | 26.68 | -4% | -2% |
| | NetEase | NTES.O | 78,992 | Buy | 165.00 | 123.73 | -9% | 5% |
| | Kuaishou Technology | 1024.HK | 16,247 | Buy | 60.00 | 28.58 | -55% | -7% |
| | Bilibili | BILI.O | 5,965 | Buy | 27.00 | 14.27 | -42% | -9% |
| Sardonna Fong | Focus Media Information Technology | 002027.SZ | 11,288 | Buy | 9.50 | 5.23 | -25% | 10% |
| | Vipshop | VIPS.N | 6,291 | Buy | 18.20 | 12.86 | -24% | 3% |
| | JOYY Inc | JOYY.O | 4,324 | Buy | 95.00 | 82.27 | 38% | 9% |
| Jenny Yuan | New Oriental Education & Technology | EDU.N | 9,486 | Buy | 81.00 | 57.58 | 6% | 1% |
| | TAL Education Group | TAL.N | 6,923 | Buy | 19.70 | 12.51 | 15% | 7% |
| | China Education Group | 0839.HK | 667 | Buy | 3.50 | 1.89 | -39% | 2% |
| | China Literature | 0772.HK | 2,954 | Buy | 29.50 | 22.18 | -33% | 15% |
| | Weibo Corp | WB.O | 1,599 | Buy | 9.50 | 6.51 | -32% | -2% |
| | iQIYI | IQ.O | 918 | Buy | 1.55 | 0.95 | -50% | -4% |
| | Hello Group | MOMO.O | 696 | Buy | 8.30 | 4.71 | -25% | -4% |

Source: Refinitiv Eikon, UBS estimates. Note: Above data as at close of 14 October 2026
Note: The green indicates stock performance outperforming KWEB, while red indicates underperformance vs. KWEB (YTD: -27.4%).

# Our key views on China Internet

**Sector performance has been soft, with KWEB correcting by 27% YTD**. Key concerns include soft macro conditions (particularly consumption), competition, greater AI investment, varied AI impact across verticals, which have led to EPS downgrades.

We see signs of stabilization, given an undemanding 11x 2026E sector P/E, which has priced in various near-term concerns. We expect a return to positive earnings revisions in the next few quarters, and see the evolving AI narrative across competition, investment and monetization as a key driver for share performance.

**Key observations and structural changes for the China Internet sector**

1) **More heavy-asset investment to build moats**, extending beyond AI investments to physical assets, including upstream production (e.g., PDD's 1P brand XinPinMu), distribution infrastructure (e.g., Meituan's 1P InstaMarts and HappyMonkey Mart), and AI vertical integration (Alibaba T-Head).
2) **AI competition is expanding from models to applications/ecosystems**. LLM competition remains intense, with more use cases (including agents and Claw-like products) opening up.
3) **Heightened focus on profitability for non-AI investments**: In addition to reducing food delivery subsidies, we also notice lower To-C platform subsidies for e-commerce. Operators are also seeking to drive traffic monetisation through higher ad revenue.
4) **Unlocking value from underlying AI assets**: Internet companies have been disclosing more information about the revenue generated and value unlocked from their underlying AI assets (e.g., Baidu's Kunlun, Kuaishou's Kling and BABA's T-Head).
5) **AI disruption risk**: especially in online verticals such as online travel, music and recruiting. While we believe the pace of disruption in China is likely to remain gradual, sector-specific headwinds (e.g., regulatory investigations into Trip.com and rising competition in online music) could pressure the near-term performance of vertical stocks.

# Key themes in China AI: model capability, monetization, and token ROI

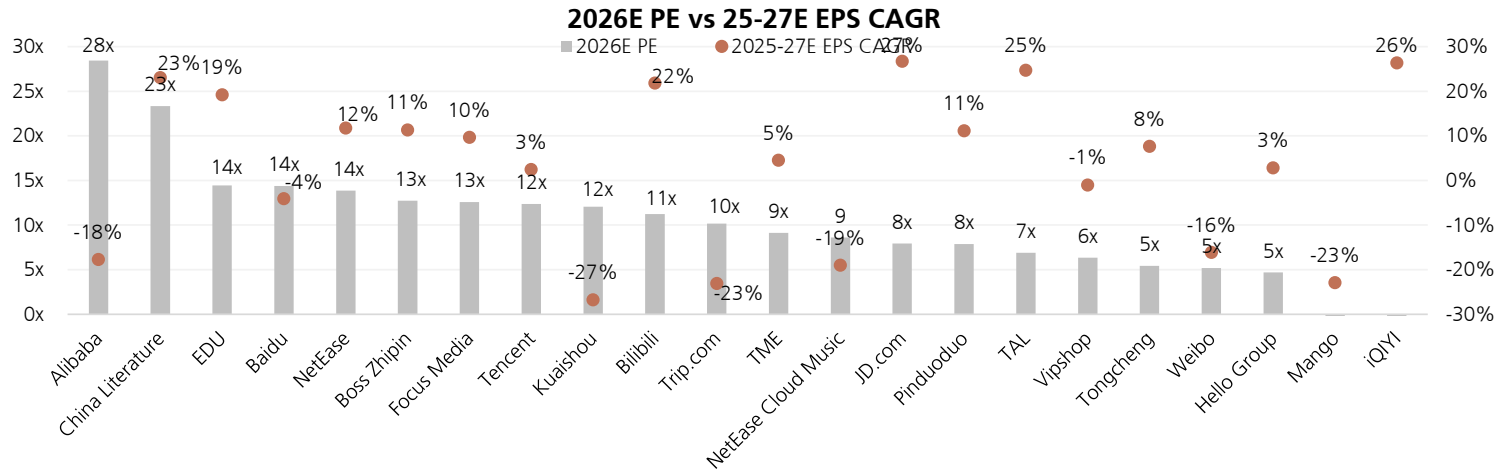
Investor interest in China AI models has been high since the listing of Zhipu and MiniMax in early 2026, with market debates becoming more sophisticated across model iteration, monetization and competition. We set out three key themes to watch in 2H26E:

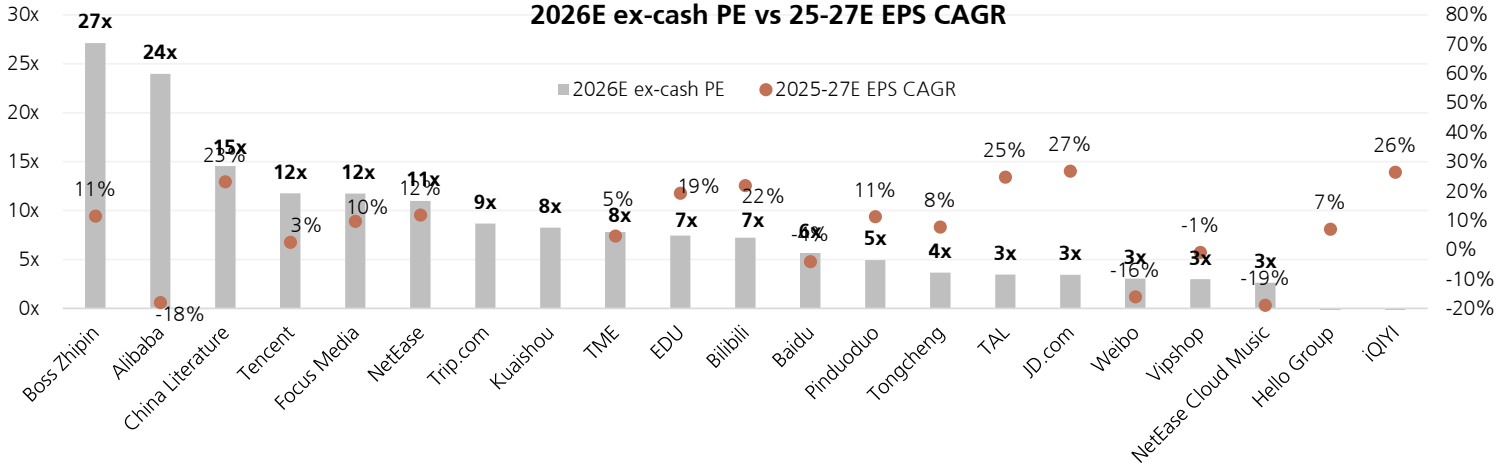
1) **Evolving competition dynamics on model capability and cost efficiency**. While a narrowing capability gap and cost advantages have increased recognition of Chinese models' global competitiveness , we note recent developments that have complicated the competitive outlook and contributed to weak momentum in listed model companies:
2) **We expect model-layer monetization to continue to ramp**, as coding TAM expands beyond developers, while progress in multimodality may be overlooked;
3) **We expect an increasing focus on token ROI**, and the shift of enterprises and users from tokenmaxxing to token optimisation should drive tiered demand and widen divergence in the pricing power of state-of-the-art (SOTA) vs. other models, while Chinese models may benefit from this trend given their structural cost efficiency.

**Upcoming model launches to watch**

- While concerns on intensifying global competition recently drove a sector de-rating, we think Chinese labs' progress in overseas distribution, architectural innovation and open-source R&D synergies remains underappreciated.
- We think continued model launches (e.g., Step 5 Preview, Xiaomi MiMo-V2.6) could help restore investor sentiment by extending the intelligence-cost Pareto frontier, with potential further catalysts ahead including DeepSeek V4.1 Pro, Kimi K3.1, MiniMax M3 Pro, Zhipu GLM-5.5/6.0, etc.

# PE/ex-cash PE of China internet companies




Source: Reuters, UBS estimates. As of 14 October 2026

# Coverage valuation comps

| Company name | Ticker | Mkt cap (USD m) | Rating | Price target (LC) | Price (LC) | Implied upside | PS (x) 2026E | PS (x) 2027E | Sales CAGR 2025-27E | PSG (x) FY26E | PE (x) 2026E | PE (x) 2027E | EPS CAGR 2025-27E | PEG FY26E | Stock performance 1M | Stock performance 3M | Stock performance YTD |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **China E-commerce** | | | | | | | | | | | | | | | | | |
| Alibaba Group | BABA.N | 258,026 | Buy | 206.00 | 111.17 | 85% | 1.7 | 1.5 | 11.6% | 0.15x | 16.8 | 11.4 | 55.2% | 0.30x | 2% | -3% | -24% |
| PDD Holdings Inc | PDD.O | 113,175 | Buy | 146.00 | 81.45 | 79% | 1.7 | 1.5 | 7.0% | 0.25x | 8.1 | 6.1 | 11.2% | 0.72x | 4% | -3% | -28% |
| Meituan | 3690.HK | 58,430 | Buy | 132.00 | 75.45 | 75% | 0.9 | 0.8 | 17.0% | 0.06x | n.m. | 14.5 | n.m. | n/a | 1% | -10% | -27% |
| JD.com | JD.O | 37,979 | Buy | 48.00 | 26.68 | 80% | 0.2 | 0.2 | 3.6% | 0.05x | 7.9 | 6.1 | 26.7% | 0.30x | -2% | -10% | -7% |
| Vipshop | VIPS.N | 6,291 | Buy | 18.20 | 12.86 | 42% | 0.4 | 0.4 | -0.3% | NA | 6.6 | 5.2 | (1.0%) | NA | 6% | -11% | -27% |
| **Average** | | | | | | | **1.0** | **0.9** | **7.8%** | **0.13x** | **9.9** | **8.7** | **23.0%** | **0.43x** | **2%** | **-8%** | **-23%** |
| **Online Media and Entertainment** | | | | | | | | | | | | | | | | | |
| Tencent Holdings | 0700.HK | 483,694 | Buy | 740.00 | 417.80 | 77% | 4.0 | 3.6 | 9.1% | 0.44x | 12.5 | 12.7 | (0.1%) | NA | -5% | -9% | -30% |
| NetEase | NTES.O | 78,992 | Buy | 165.00 | 123.73 | 33% | 4.4 | 4.0 | 7.8% | 0.57x | 13.8 | 11.4 | 11.8% | 1.17x | 4% | -6% | -10% |
| Kuaishou Technology | 1024.HK | 15,732 | Buy | 60.00 | 28.58 | 110% | 0.8 | 0.7 | 1.1% | 0.69x | 12.0 | 9.7 | (26.7%) | NA | -4% | -34% | -55% |
| Tencent Music Entertainment | TME.N | 13,042 | Neutral | 9.50 | 8.48 | 12% | 2.5 | 2.3 | 8.0% | 0.31x | 9.1 | 8.4 | 4.6% | 1.98x | 7% | -7% | -52% |
| NetEase Cloud Music | 9899.HK | 2,341 | Neutral | 100.00 | 84.75 | 18% | 2.0 | 1.9 | 2.8% | 0.70x | 8.6 | 8.2 | (18.9%) | NA | -2% | -27% | -54% |
| Bilibili | BILI.O | 5,965 | Buy | 27.00 | 14.27 | 89% | 1.2 | 1.1 | 8.2% | 0.15x | 14.8 | 11.7 | 18.3% | 0.81x | -8% | -21% | -42% |
| iQIYI | IQ.O | 918 | Buy | 1.55 | 0.95 | 63% | 0.2 | 0.2 | -3.5% | NA | n.m. | 13.7 | 26.4% | n/a | -5% | -20% | -50% |
| Weibo Corp | WB.O | 1,599 | Buy | 9.50 | 6.51 | 46% | 0.9 | 0.9 | 0.5% | 1.84x | 5.7 | 5.7 | (16.1%) | NA | -2% | -16% | -36% |
| Hello Group | MOMO.O | 696 | Buy | 8.30 | 4.71 | 76% | 0.5 | 0.5 | -1.4% | NA | 5.1 | 5.1 | 2.8% | 1.80x | -5% | -22% | -28% |
| China Literature | 0772.HK | 2,887 | Buy | 29.50 | 22.18 | 33% | 2.7 | 2.5 | 2.5% | 1.07x | 24.1 | 14.8 | 23.1% | 1.04x | 17% | 14% | -33% |
| Focus Media Information Technology | 002027.SZ | 11,288 | Buy | 9.50 | 5.23 | 82% | 6.0 | 5.7 | 1.9% | 3.17x | 12.5 | 11.9 | 9.7% | 1.29x | 12% | 5% | -29% |
| **Average** | | | | | | | **2.1** | **2.1** | **3.4%** | **0.63x** | **11.8** | **10.1** | **3.5%** | **3.36x** | **1%** | **-13%** | **-39%** |
| **AI, Cloud and Software** | | | | | | | | | | | | | | | | | |
| Knowledge Atlas Technology | 2513.HK | 26,358 | Buy | 1,900.00 | 513.50 | 270% | 30.1 | 12.4 | 343.9% | 0.09x | n.m. | n.m. | 12.3% | n/a | -29% | -54% | 342% |
| MiniMax | 0100.HK | 7,174 | Buy | 500.00 | 179.50 | 179% | 15.2 | 6.8 | 266.6% | 0.06x | n.m. | n.m. | n.m. | n/a | -29% | -17% | 9% |
| Baidu, Inc. | BIDU.O | 28,193 | Buy | 160.00 | 81.86 | 95% | 1.5 | 1.4 | 1.7% | 0.88x | 14.1 | 10.8 | (4.1%) | NA | -9% | -24% | -37% |
| Iflytek | 002230.SZ | 13,031 | Buy | 55.00 | 36.30 | 52% | 2.9 | 2.5 | 12.3% | 0.24x | 86.5 | 67.0 | 23.3% | 3.71x | -7% | -12% | -28% |
| Kingsoft Cloud | KC.O | 2,398 | Buy | 19.00 | 8.76 | 117% | 1.2 | 1.0 | 30.6% | 0.04x | n.m. | n.m. | n.m. | n/a | -8% | -8% | -15% |
| Wondershare | 300624.SZ | 1,446 | Neutral | 65.00 | 50.03 | 30% | 5.7 | 5.1 | 11.7% | 0.49x | n.m. | n.m. | n.m. | n/a | 1% | 9% | -29% |
| **Average** | | | | | | | **3.3** | **2.9** | **18.2%** | **0.18x** | **86.5** | **67.0** | **23.3%** | **3.71x** | **-9%** | **-3%** | **-24%** |
| **Travel and Transactional Platforms** | | | | | | | | | | | | | | | | | |
| Kanzhun Limited | BZ.O | 7,366 | Buy | 21.00 | 16.35 | 28% | 5.3 | 4.7 | 12.2% | 0.43x | 12.8 | 11.3 | 11.4% | 1.13x | 9% | 10% | -20% |
| Trip.com | TCOM.O | 25,350 | Buy | 60.00 | 38.88 | 54% | 2.6 | 2.4 | 7.3% | 0.35x | 10.1 | 9.6 | (23.0%) | NA | -1% | -8% | -46% |
| Tongcheng Travel | 0780.HK | 2,856 | Buy | 18.00 | 9.88 | 82% | 0.9 | 0.8 | 7.7% | 0.12x | 5.4 | 4.9 | 7.7% | 0.70x | -6% | -22% | -56% |
| **Average** | | | | | | | **2.9** | **2.6** | **9.1%** | **0.32x** | **9.4** | **8.6** | **9.6%** | **0.99x** | **1%** | **-7%** | **-41%** |
| **Education** | | | | | | | | | | | | | | | | | |
| TAL Education Group | TAL.N | 6,923 | Buy | 19.70 | 12.51 | 57% | 1.8 | 1.6 | 30.0% | 0.06x | 12.6 | 6.9 | n.m. | n/a | 6% | 22% | 15% |
| New Oriental Education & Technology | EDU.N | 9,486 | Buy | 81.00 | 57.58 | 41% | 1.7 | 1.4 | 16.7% | 0.10x | 16.1 | 13.1 | 17.8% | 0.90x | 2% | 17% | 5% |
| China Education Group | 0839.HK | 654 | Buy | 3.50 | 1.89 | 85% | 0.6 | 0.6 | 2.2% | 0.26x | 2.9 | 2.8 | (13.7%) | NA | 3% | -1% | -39% |
| **Average** | | | | | | | **1.7** | **1.5** | **23.3%** | **0.07x** | **14.4** | **10.0** | **17.8%** | **0.90x** | **4%** | **20%** | **10%** |
| **China Internet average** | | | | | | | **2.0** | **1.8** | **8.0%** | **0.25x** | **14.6** | **11.8** | **7.0%** | **2.09x** | **1%** | **-7%** | **-30%** |

Source: Refinitiv datastream, UBS estimates. Priced as of 14 October 2026.

# China LLM/AI update

# Key themes in 2H26E: Model capability

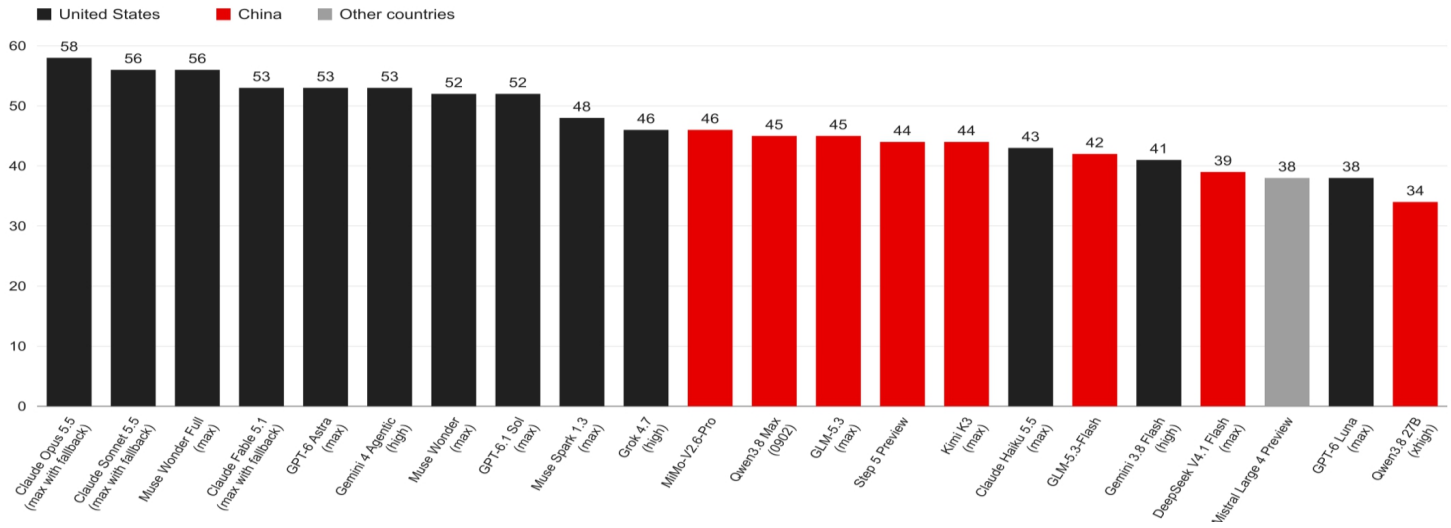
**#1: Evolving competition dynamics on capability and cost efficiency**

While a narrowing capability gap and cost advantages have increased recognition of Chinese models' global competitiveness, we note recent developments that have complicated the competitive outlook and contributed to weak momentum in listed model companies:

1) **Regulatory landscape continues to evolve**, with debate over potential restrictions on open-weight models adding uncertainty to overseas adoption, and tighter anti-distillation measures by global leaders and growing US policy focus on industrial-scale distillation raising questions over the pace of capability catch-up.
2) **Cost-performance advantages are becoming more task-dependent**, particularly in complex, long-horizon tasks, where we observe overseas frontier models partly offsetting their per-token price premium through improved token efficiency and task success rates, while more developers launch smaller/flash models to optimise inference efficiency for volume-driven tasks.

While valuation of China's listed model companies (i.e., Zhipu and MiniMax) has corrected to the low end of the recent range, we await future model launches as catalysts (likely in Oct/Nov) to potentially recover momentum.

**Leading models by AI lab: Artificial Analysis Intelligence Index**



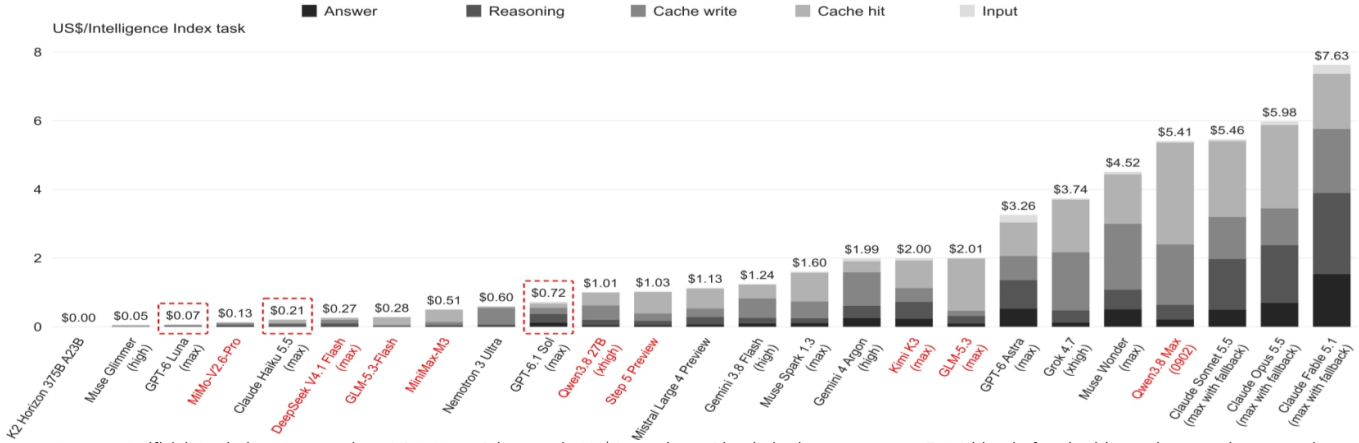

Sources: Artificial Analysis Intelligence Index
Notes: Artificial Analysis Intelligence Index incorporates 7 evaluations:MMLU-Pro, GPQA Diamond, Humanity's Last Exam, LiveCodeBench, SciCode, AIME, MATH-500

# Key themes in 2H26E: Model capability (cont'd)

**AI developers are launching more Flash models, offering near-flagship performance at lower prices**

| 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 |
|---|---|---|---|---|---|---|---|---|---|
| Zhipu | Flash | GLM-5.3-Flash | 26 Aug 2026 | 320B / 18B | 41.8 | $0.15 | $0.03 | $0.50 | $0.10 |
| | Flagship | GLM-5.3 | 18 Aug 2026 | 753B / ~40B | 44.8 | $1.40 | $0.26 | $4.40 | $0.90 |
| Moonshot | Flash | Kimi K2.8 Preview | 11 Sep 2026 | Not disclosed | n.a. | $1.00* | $0.25* | $4.00* | $0.78* |
| | Flagship | Kimi K3 | 16 Jul 2026 | 2.8T / 104B | 43.6 | $3.00 | $0.30 | $15.00 | $2.31 |
| Alibaba | Flash | Qwen3.8-Flash-Next | 26 Aug 2026 | 180B / 6B | 39.8 | $0.15 | $0.016 | $0.47 | $0.09 |
| | Flagship | Qwen3.8 Max | 3 Aug 2026 | Not disclosed | 40.2 | $2.00 | $0.25 | $6.00 | $1.18 |
| Google | Flash | Gemini 3.8 Flash (high) | 2 Sep 2026 | Not disclosed | 40.9 | $0.75 | $0.075 | $3.75 | $0.58 |
| | Flagship | Gemini 3.1 Pro Preview | 19 Feb 2026 | Not disclosed | 29.7 | $2.00 | $0.20 | $12.00 | $1.74 |
| SpaceXAI | Flagship | Grok 4.7 (xhigh) | 21 Sep 2026 | Not disclosed | 46.4 | $2.00 | $0.50 | $6.00 | $1.35 |
| | Flash | Grok Build 0.1 0616 | 16 Jun 2026 | Not disclosed | 27.2 | $1.00 | $0.20 | $2.00 | $0.54 |
| DeepSeek | Flash | DeepSeek V4.1 Flash | 10 Sep 2026 | 552B / 16B | 39.5 | $0.30 | $0.006 | $1.20 | $0.18 |
| | Flagship | DeepSeek V4 Pro 0424 | 24 Apr 2026 | 1.6T / 49B | 30.4 | $0.435 | $0.0036 | $0.87 | $0.18 |
| OpenAI | Balanced | GPT-6.1 Sol | 29 Sep 2026 | Not disclosed | 51.8 | $2.00 | $0.10 | $10.00 | $1.47 |
| | Balanced | GPT-6 Sol | 22 Sep 2026 | Not disclosed | 47.6 | $2.00 | $0.20 | $10.00 | $1.54 |
| | High-volume | GPT-6 Luna | 22 Sep 2026 | Not disclosed | 38.1 | $0.10 | $0.01 | $0.50 | $0.08 |
| | Flagship | GPT-6 Astra | 3 Sep 2026 | Not disclosed | 52.7 | $10.00 | $1.00 | $50.00 | $7.70 |
| Anthropic | Efficient | Claude Haiku 5.5 | 7 Oct 2026 | Not disclosed | 43.4 | $0.10 | $0.01 | $0.50 | $0.08 |
| | Balanced | Claude Sonnet 5.5 | 28 Sep 2026 | Not disclosed | 56.0 | $2.00 | $0.10 | $10.00 | $1.47 |
| | Efficient | Claude Opus 5.5 | 22 Sep 2026 | Not disclosed | 57.6 | $4.00 | $0.20 | $20.00 | $2.94 |
| | Flagship | Claude Fable 5.1 | 1 Sep 2026 | Not disclosed | 53.4 | $10.00 | $0.25 | $50.00 | $7.18 |
| Meta | Flagship | Muse Wonder (max) | 12 Oct 2026 | Not disclosed | 52.4 | $4.00 | $0.20 | $20.00 | $2.94 |
| | Balanced | Muse Spark 1.3 (max) | 2 Sep 2026 | Not disclosed | 48.1 | $1.25 | $0.15 | $4.25 | $0.78 |

**Cost per Intelligence Index Task - New overseas AI models offer competitive task costs versus Chinese peers**



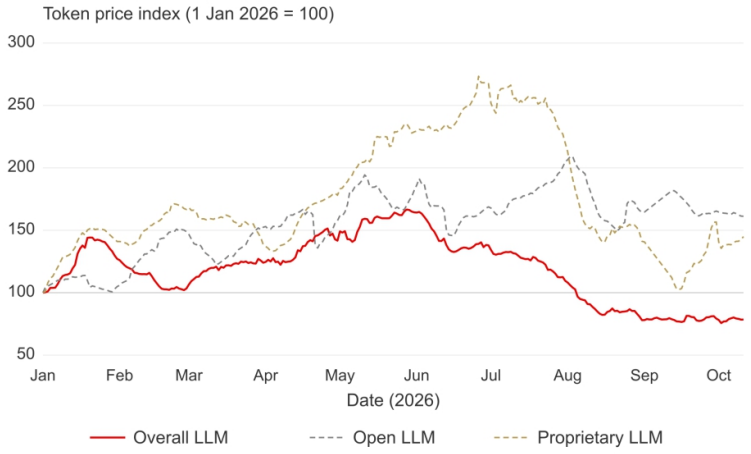

Sources: Artificial Analysis, company data, UBS. Note: Prices are in US$/mn tokens. Blended price represents a 7:2:1 blend of cached-input, input and output token prices. Parameters refer to total/active parameters where disclosed. *Kimi K2.8 Preview pricing is based on a third-party API provider. Data as of 14 October 2026.

# Key themes in 2H26E: Token ROI

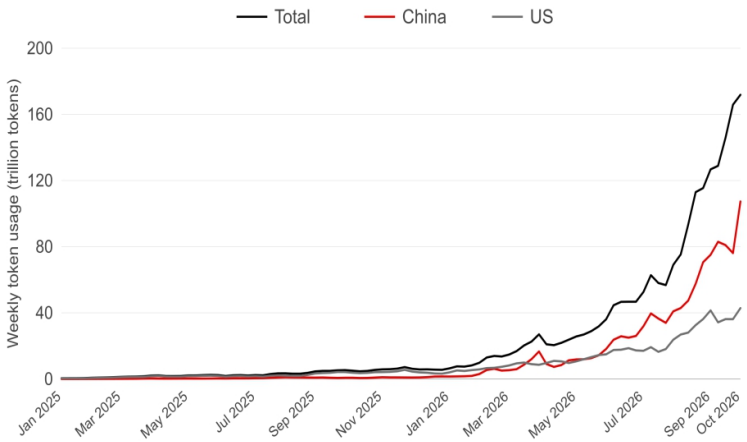
**#2: We expect an increasing focus on token ROI**.

- As AI adoption moves from experimentation to scaled deployment, companies are **shifting from maximising AI token usage (i.e., "tokenmaxxing") towards more ROI driven resource allocation**. In our view, this shift should improve market recognition of Chinese models' cost-performance advantages.
- **We expect the global model market to become increasingly tiered**. While frontier models may retain pricing premiums for complex, high-value tasks, cost-performance models could see broader adoption in high-volume, ROI-sensitive workloads, creating long-term global share-gain opportunities for China's models.

**LLM Token Expenditure Index**




**China vs US AI model token usage trend on OpenRouter**




Source: OpenRouter, SiliconData. Note: Data as of 14 Oct 2026. Normalized blended rate drawn from observations across frontier API providers, open-weight inference platforms, brokered dedicated-instance markets and self-hosted reference deployment

# Key themes in 2H26E: Token ROI

**Overseas apps and platforms are increasingly integrating Chinese AI models**

| App / Platform | Product category | Key functions | Timing | Integration details |
|---|---|---|---|---|
| **Fireworks AI** | AI inference / model hosting platform | **Open-model AI infrastructure** for training / fine-tuning and high-throughput inference, with serverless, on-demand and reserved deployments plus OpenAI / Anthropic-compatible APIs | Kimi / MiniMax / Qwen: 10-12 Jun 2026; GLM-5.2: 16 Jun 2026 | Fireworks announced day-zero **GLM-5.2** availability. Its changelog / model library also lists **Kimi K2.7 Code**, **MiniMax M3, Qwen 3.7 Plus** and other Chinese models on the platform. |
| **Vercel AI** | Developer tool / AI gateway | **Developer AI gateway** providing one API key for hundreds of models, unified billing / observability, provider routing / fallback, spend controls, BYOK and drop-in SDK compatibility | GLM-5.2: 16 Jun 2026; GLM-5.2 Fast via Wafer: 24 Jun 2026 | Vercel added **GLM-5.2** to AI Gateway on 16 Jun under model ID zai/glm-5.2. On 24 Jun, it added **GLM-5.2 Fast** via **Wafer**, a model inference / hosting platform offering serverless and dedicated inference. |
| **Cloudflare Workers AI** | Serverless AI / edge inference platform | **Serverless GPU inference on Cloudflare's global networ**k, with open-source model access through Workers / API, pay-as-you-use pricing and integrations with AI Gateway, Vectorize and Workers | 16 Jun 2026 | Cloudflare announced **GLM-5.2** availability on Workers AI, highlighting agentic coding, function calling and reasoning. The initial context window was 262k, with plans to expand later. |
| **0G Private Computer** | Decentralised AI compute / verifiable inference | **Decentralised AI compute network for router / direct inference**, provider marketplace access, wallet-based settlement, TEE-backed verification and chatbot, image and speech service types | 18 Jun 2026; Kimi added the following day | 0G announced **GLM-5.2** availability on Private Computer, highlighting end-to-end encryption for prompts, code and outputs, as well as verifiable compute. **Qwen-3.7 Plus** was also available, with **Kimi 2.7 Code** added shortly thereafter. |
| **Notion** | Notes / AI workspace | **All-in-one workspace** for docs, wikis / knowledge base, databases, project management, enterprise search, meeting notes and agents | Mid-Jun 2026 | Notion announced **GLM-5.2** availability, particularly for long-horizon tasks. The model is served through Baseten, which also confirmed that Notion is offering **GLM-5.2** via its inference platform. |
| **Perplexity** | AI answer engine / Enterprise AI / APIs | **AI answer engine and enterprise research platform** spanning web research, internal-file / app search, agentic Computer workflows and developer APIs | Mid- to late Jun 2026 | Perplexity's Agent API model list includes **perplexity/glm-5.2** and **perplexity/kimi-k2.7-code**, making both models available through its Agent API. |
| **Devin / Cognition** | AI coding / software engineering agent | **Autonomous AI software engineer** for writing, running and testing code, resolving tickets / bugs, migrations / refactors, PR review, codebase Q&A and engineering support | 24 Jun 2026 | Devin announced that **Kimi K2.7** and **GLM-5.2** are available in Devin Desktop and CLI. |
| **Coinbase** | Crypto exchange / digital asset platform | **Retail and institutional crypto / asset trading platform**, with custody, staking, payments, wallets, developer infrastructure and onchain services | Late Jun 2026; | Coinbase co-founder and CEO Brian Armstrong said the company is experimenting with **GLM-5.2** and **Kimi 2.7** as default open-weight models through its internal LLM gateway, aiming to reduce AI spend without restricting engineers' token usage. |
| **Fello AI** | Multi-model AI assistant / consumer app | **All-in-one AI chatbot app** for Mac, iPhone and iPad, with multi-model chat, web search, file / PDF analysis, image generation / editing, multilingual chat and cross-device sync | 25 Jun 2026 | Fello AI 6.7.0 added **GLM-5.2, MiniMax M3, DeepSeek V4 and Qwen 3.7 Max**, alongside new Podcast and Travel Skills and broader file support. |
| **Poe / Quora** | Multi-model chatbot / bot platform | **Consumer and developer AI aggregation platform** for accessing multiple models / bots, generating text, images, video and audio, web search, multi-bot chat, bot / app creation, workflows and Poe API | Live on platform; no single announcement date found | Poe pages show access to **GLM-5.2**, **DeepSeek V4 Flash / Pro**, **Kimi K2.7 Code, MiniMax M3** and other Chinese models via providers such as Novita AI. |

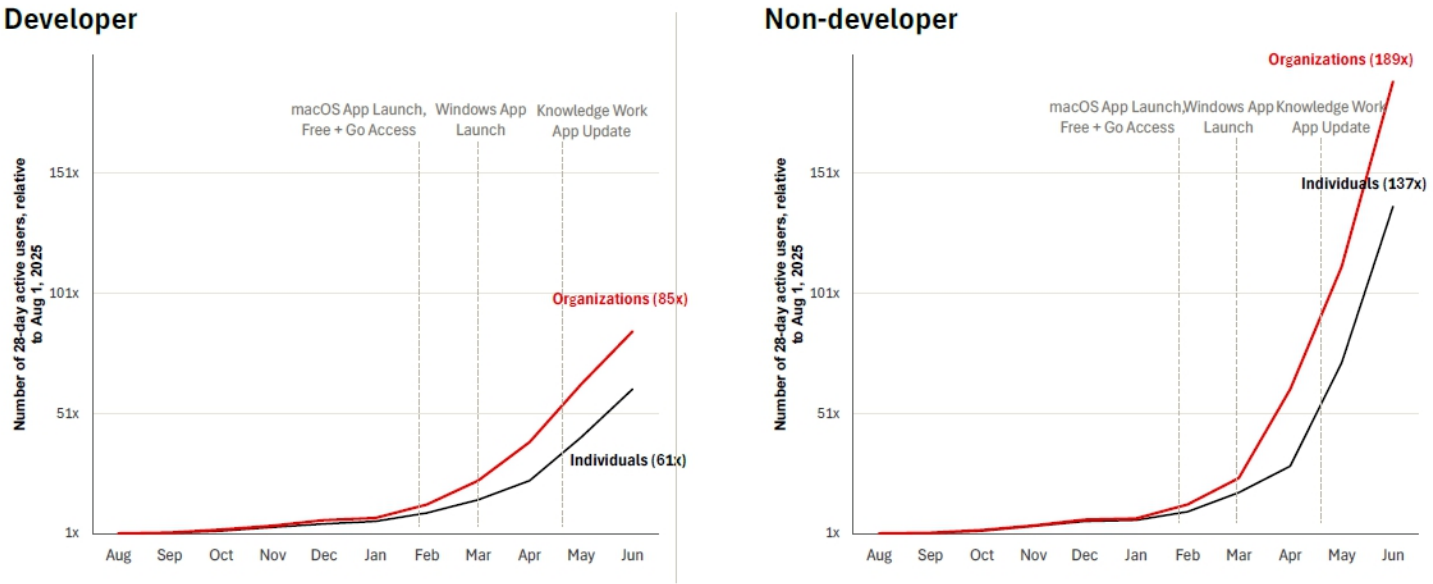
Sources: Company data, UBS

# Key themes in 2H26E: Monetization

**#3: We expect model-layer monetization to continue to ramp up**

- AI coding TAM continues to expand, shifting from developer tools into broader productivity tools for knowledge workers, with feature innovation (e.g. app plugins) to integrate with more workflows and accelerate adoption.
- Progress in multimodality is also worth noting, with Chinese models demonstrating early leadership in video-generation capability and monetization (e.g. US$500m ARR for Kuaishou's Kling in May).
- We continue to observe the model upgrades driving robust token usage and monetization in China, with Zhipu's ARR reaching c.US$1.8bn in September, and MiniMax's ARR exceeding US$800m in Aug (vs c.US$400m in May and US$150m in Feb).

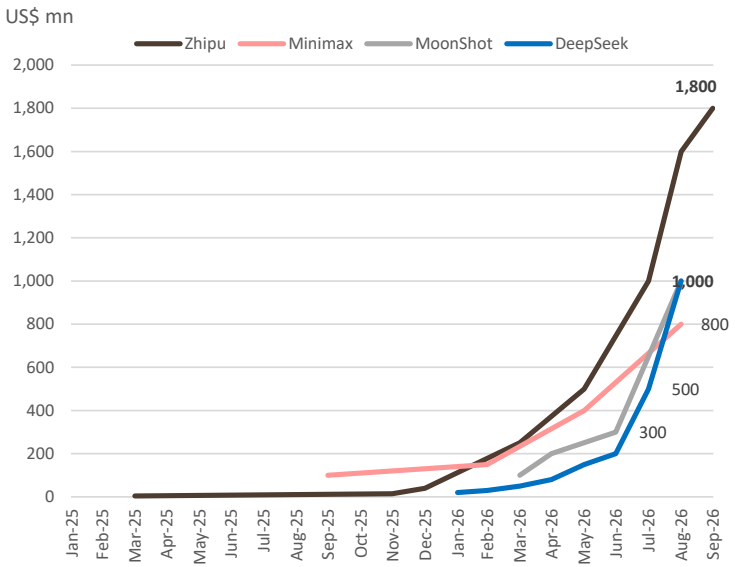
**AI coding agent adoption is broadening beyond developers**



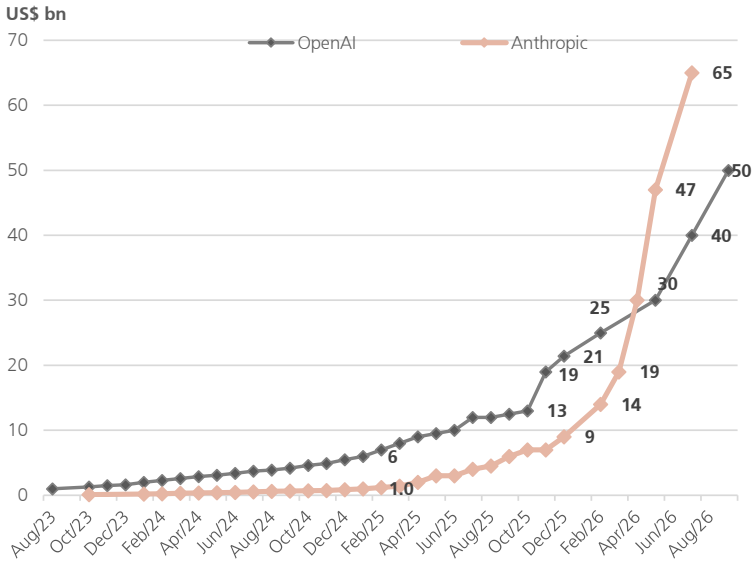

Source: OpenAI (How agents are transforming work). Note: The above data refers to 28-day active users indexed to 1 Aug. 2025

# Key themes in 2H26E: Monetization (cont'd)

**The ARR trend of major China AI labs**




**The ARR trend of major US AI labs**



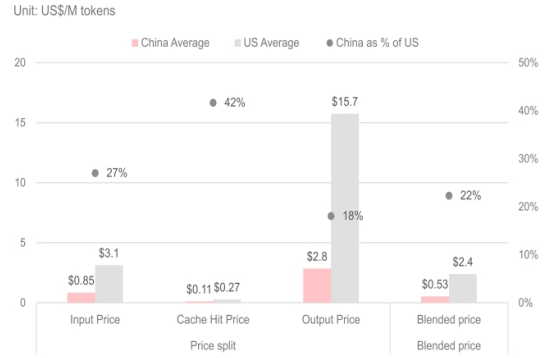

Sources: The Information, company data,

# How do China's AI models compete on cost efficiency?

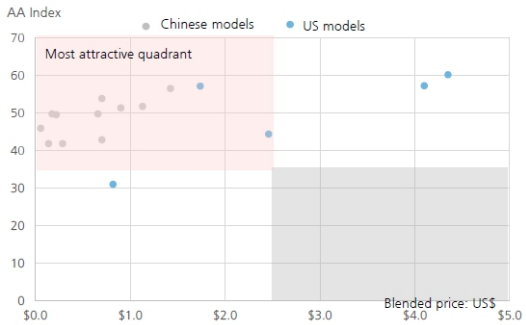
**Assessing the cost leadership of China's AI models**

- We estimate the training costs for China's models are **less than 10%** of those of global leaders, while the average API price of major China models is **~20%** of comparable global peers. However, China's frontier models are closing the capability gap with rapid iteration and improving intelligence, gaining global traction in token usage while preserving healthy gross margins.

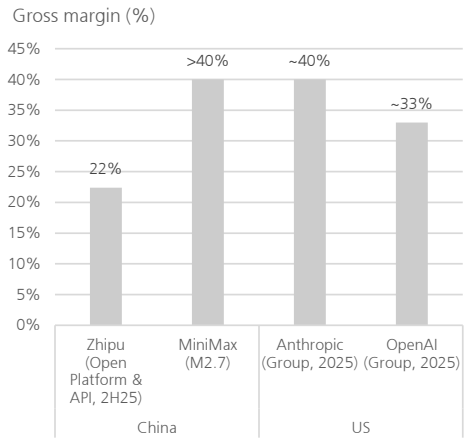
**Price comparison of AI models – China vs. the US**




**China's models have a higher cost-performance ratio**




**Gross margin of API businesses**




Source: Artificial Analysis. Note: Data as of 14 October 2026. The blended price is calculated using cached input, uncached input and output token prices, with a ratio of 7:2:1.

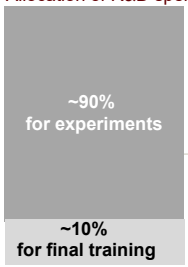
# China's model training efficiency

- **In training**, China's model providers use multiple approaches to reduce workload (via smaller parameters and sparse attention) and improve GPU throughput, helped by targeted R&D and open-source ecosystems.

## #1 STRATEGY FOCUS

**Less trial-and-error spend**

Allocation of R&D spending




**US frontier labs**
More experimentation and path exploration to pursue upper limits of model intelligence.

**China model labs**
More focused paths; smaller experiment share

**Implication: meaningful cost-saving potential from lower trial-and-error compute.**

## #2 TRAINING COMPUTE COST

**Optimize workload, capacity and GPU-hour cost**

**Total training workload**
Smaller parameters (<2trn for leading China models, vs >5trn for frontier US models) + sparse attention to reduce long-context token interactions

**Effective GPU capacity**
Lower precision training (BF16/FP16 to FP8; c.2x theoretical B200 throughput) + higher MFU via scheduling / load balancing

**Cost per GPU hour**
Lower GPU-hour pricing and lower electricity cost (China c.4.4c/kWh vs US c.7.9c/kWh; c.40% saving)

**Training cost ≈ workload / effective GPU capacity x GPU-hour cost**

## #3 OPEN-SOURCE ECOSYSTEM

**Collective R&D leverage**

**Open models allow validated technical ideas to diffuse faster across labs.**

- **Sparse attention**
  Kimi / Zhipu's GLM follow DeepSeek-like latent or sparse attention designs
- **Muon optimizer**
  DeepSeek V4 incorporates Moonshot (Kimi)-related optimizer techniques
- **GRPO-style RL**
  Qwen3 and GLM-4.5 adopt DeepSeek-like RL recipes

**Once one lab validates an innovation, marginal R&D cost can fall for the broader ecosystem in China.**

Source: Epoch AI, SemiAnalysis, company data, UBS

人在草木间
更多资料请加V

# China's model inference efficiency

- **In inference**, key innovations span model architecture (e.g., MoE, attention mechanisms, KV cache and lower precision), engineering optimisation (e.g., P/D separation, batching and scheduling) and infrastructure (incl. domestic chip adaptation).

## #1 MODEL-SIDE OPTIMISATION

Model architecture design reduces compute and memory cost.

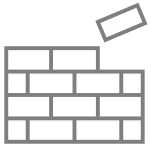


**China vs. US / Global**

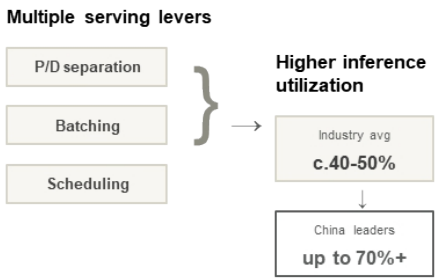
| Dimension | China | US / Global |
|---|---|---|
| MoE | c.3-10% active ratio | c.15-30% active ratio for earlier MoE |
| Attention | More extensive use of sparse attention | Greater use of full attention architectures |
| KV cache | Production-grade KV cache compression (e.g., DeepSeek's HCA and CSA) | More experimental-stage KV cache optimization (e.g., TurboQuant) |
| Precision | Broader adoption of low-precision inference (FP8/INT4) | More BF16/FP16-based deployment |

## #2 SERVING / ENGINEERING

AI infrastructure optimization lifts utilization and throughput.

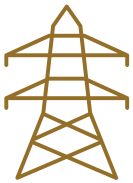


**China vs. US / Global**




## #3 RESOURCE / HARDWARE COST

Utility-cost advantage + accelerating compute localization.



**China vs. US / Global**

**+ Cheaper energy cost**

Lower electricity cost provides a cost buffer

**- AI chip inference efficiency gap**

Domestic chips still require optimization but offer long-term cost-saving potential

Note: HCA stands for Heavily Compressed Attention, and CSA stands for Compressed Sparse Attention

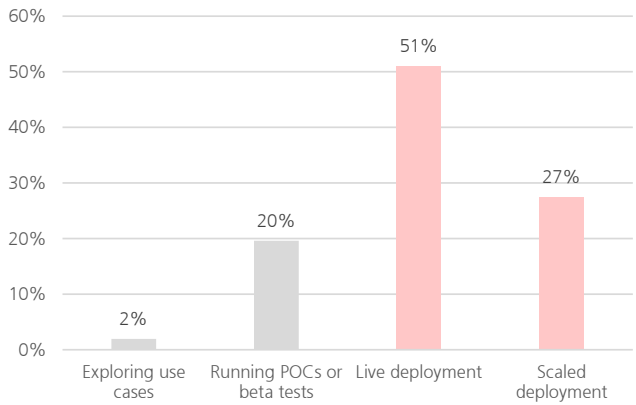
Source: Epoch AI, SemiAnalysis, company data, UBS

# What is the progress of enterprise AI adoption?

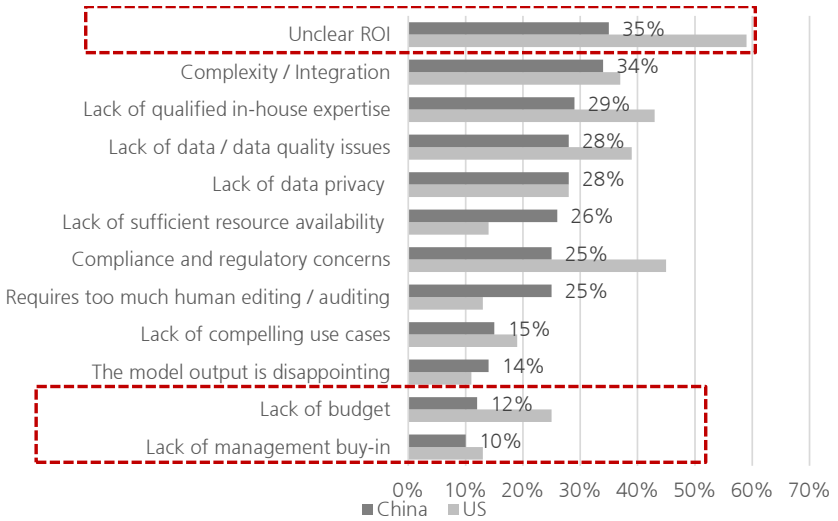
We leverage UBS Evidence Lab's China AI Business Survey of 102 IT executives and data engineers to evaluate AI adoption across organisations. The adoption pace has shown encouraging progress:

- **GenAI deployment has moved beyond the pilot stage with production scaling underway**: 78% of respondents have achieved live deployment (51% in at least one function and 27% at scale across multiple business units); 22% are at the Proof of Concept (PoC) or earlier stage with reaching scale expected to take 9.3 months.
- **AI agents are emerging as the next phase of adoption** with a longer expected timeline of 1.7 years to reach scale; we believe the rapid uptake of OpenClaw-like agent frameworks could accelerate this process. Use cases diverge between China and the US with China oriented to customer-facing scenarios (e.g. visual analytics), while the US is focused on productivity and workflow automation.
- **Uncertain ROI remains the key constraint for AI adoption**, along with IT-related constraints, while budgets are not a major issue and managements are proactive. As such, despite the initial traction, we expect AI adoption in China to deepen further into enterprise productivity use cases as digital infrastructure improves.

**Current state of GenAI deployment**




**Key challenges in GenAI adoption: China vs. US**



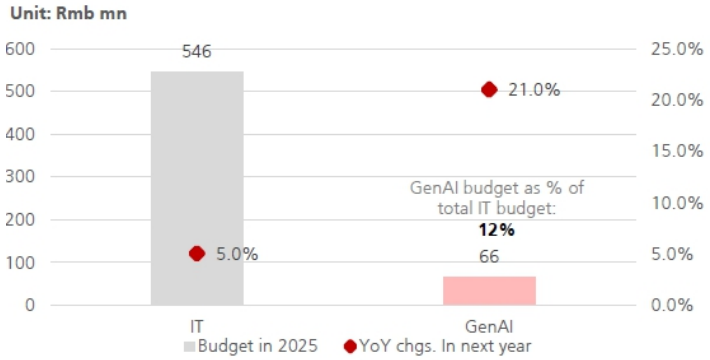

Sources: UBS Evidence Lab

# What is the progress of enterprise AI adoption?

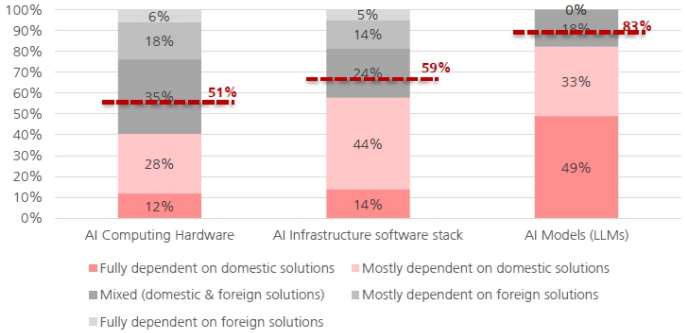
Our survey results also suggest clear commitment from China companies to AI investments:

- **GenAI spending** accounted for 12% of IT budgets in 2025, according to the survey. This ratio will likely increase as AI spending grows significantly faster than overall IT budgets (21.0% vs. 5.0% YoY in 2026), and higher than the US (17.7% vs. 4.6% YoY). **Budget allocation** appears balanced across key cost components with AI-related personnel representing the largest share, alongside similar levels of spending on models and infrastructure, reflecting Chinese companies' efforts to adopt AI on multiple fronts.
- **Localisation** progress continues across the AI value chain with the AI model layer recording faster penetration than hardware and AI infrastructure software, driven by the narrowing capability gap between Chinese and US models, as well as China's cost efficiency advantage, in our view.
- **Performance outweighs cost in cloud and model selection.** According to the survey results, performance and security are the top selection criteria ahead of cost, which underpins recent price increases by major Cloud vendors amid robust demand, in our view.

**IT and GenAI budgets and growth outlook (China)**




**Expected localisation levels across AI stack (2026)**



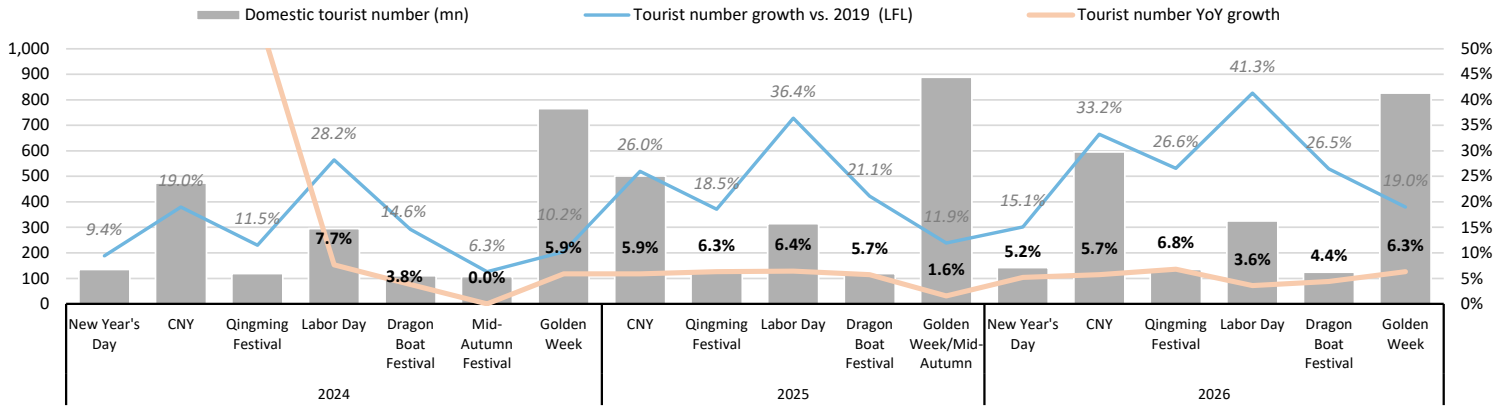

Sources: UBS Evidence Lab. Note: The GenAI budget includes AI models, AI-related personnel, projects, infrastructure and others

Section 2

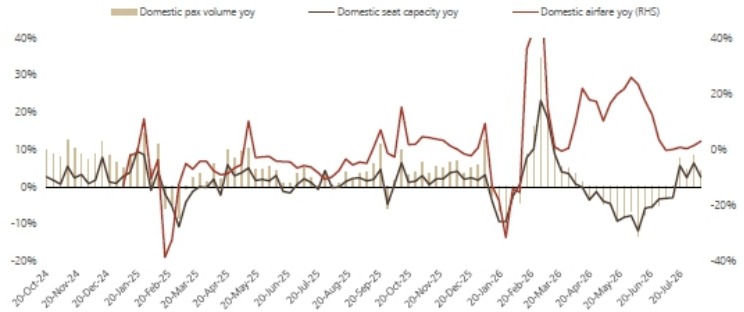
# Subsector updates

# Online travel: Improving pricing, while volume may see pressure from rising oil prices

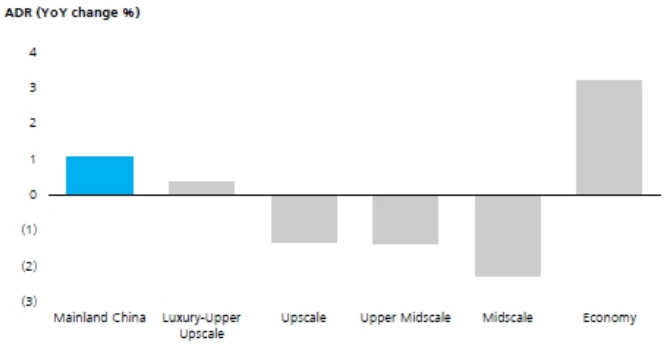
**Number of domestic tourists, YoY growth and the growth vs. 2019 levels during public holidays**




**China weekly domestic air seat capacity, pax volume and airfare YoY**



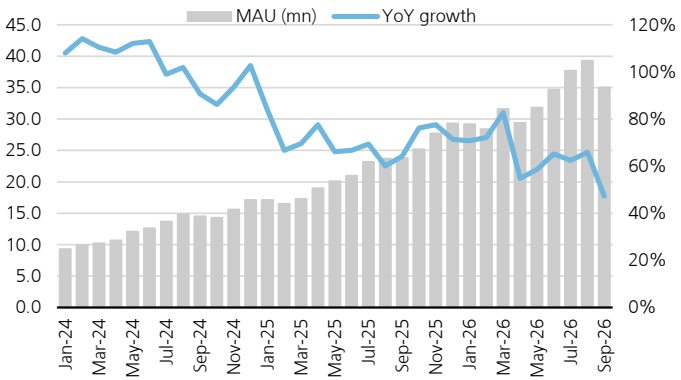

**ADR YoY performance by segment (Aug 2026)**



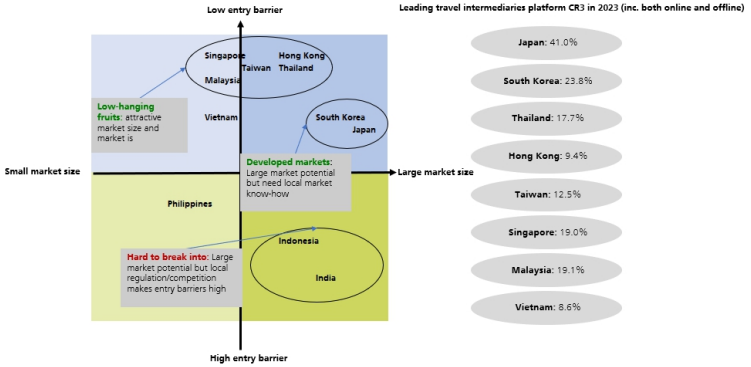

Sources: Flight Master, NBS, STR

# Online travel: Global expansion as a long-term growth driver

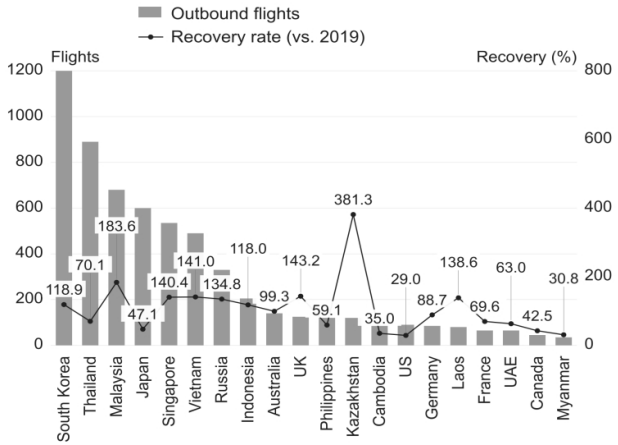
**Trip.com global MAU (m) – Trip.com**




**Analysis of Trip.com's key Asian markets**




**International flights—planned flight volume and recovery rate (vs. 2019) during National Day holiday period in 2026**




Sources: Flight Master, SensorTower, Euromonitor, UBS estimates

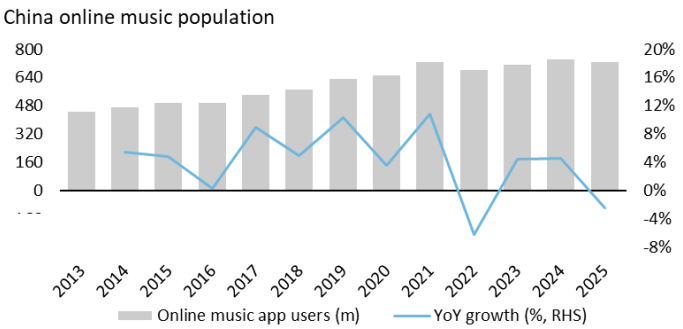
**Favorable policies for inbound travel**

| Timeline | Visa-free transit policy |
|---|---|
| 12/17/2024 | China further relaxed its visa-free transit policy on Dec. 17, extending stay periods to 240 hours: 1) transit stays were extended from 72/144 hours to 240 hours (for 54 countries); 2) 21 new ports were added as eligible exit-entry ports (for a total of 60); 3) the permitted stay area was expanded from 19 to 24 provinces. |
| 7/30/2024 | China granted 144-hour visa-free entry to foreigners visiting Hainan in tour groups via Hong Kong or Macao. |
| 7/16/2024 | The 144-hour visa-free transit policy was expanded to three more entry ports, taking the number of ports covered by the policy to 37. |
| 11/17/2023 | China expanded the scope of its 72/144-hour transit visa-free policies to 54 countries. |

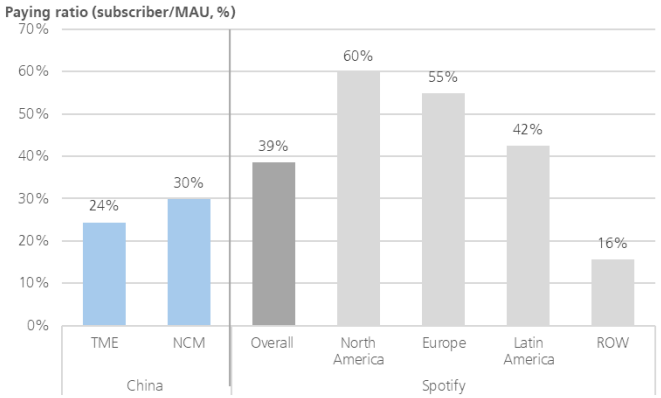
| Timeline | Instant VAT refund policy for foreign visitors |
|---|---|
| 4/8/2025 | The pilot policy lets eligible foreign visitors (staying <183 days) receive up to an 11% VAT refund at designated stores on purchases over RMB500. It covers 324 product types, excluding food, alcohol, tobacco, and autos. |
| 4/26/2025 | Threshold lowered to RMB200; cash cap raised to RMB20,000; store/product coverage expanded; refund process improved. |

# Online music: Stable user base as the market enters a mature stage

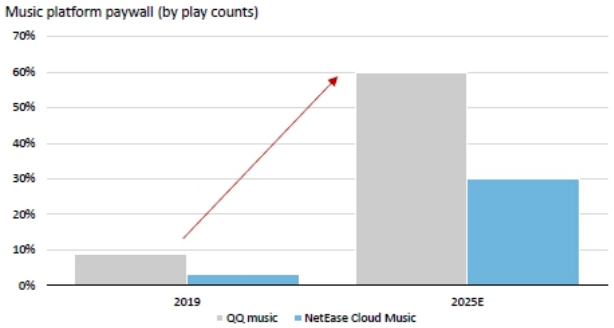
**Online music population in China and YoY growth rate**



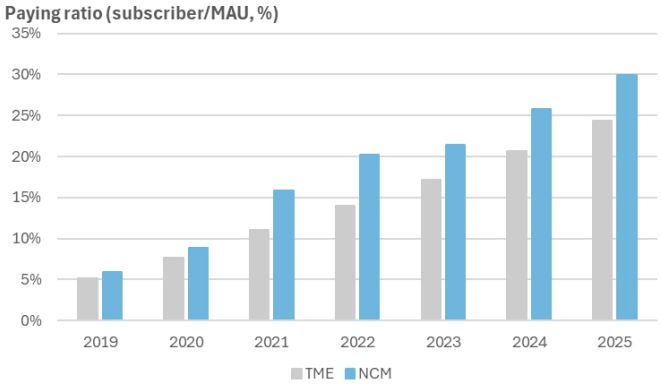

**Online music paying user ratio—TME, NCM and Spotify (2025)**




**Rising paywall in China's music streaming platforms**



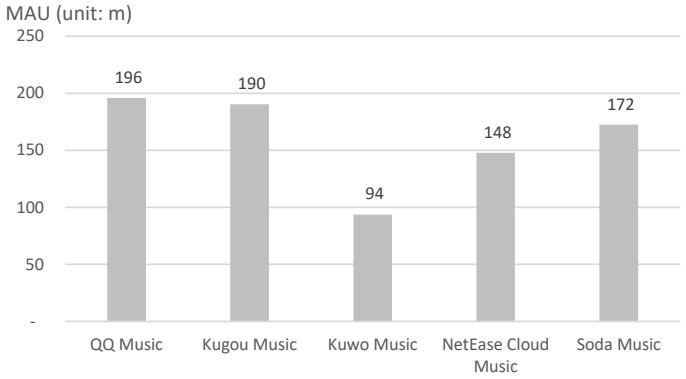

**Online paying user ratio increased notably in the same period**




Sources: Company data, UBS estimates

# Online music: Competition overhang remains, diverging strategies among key players

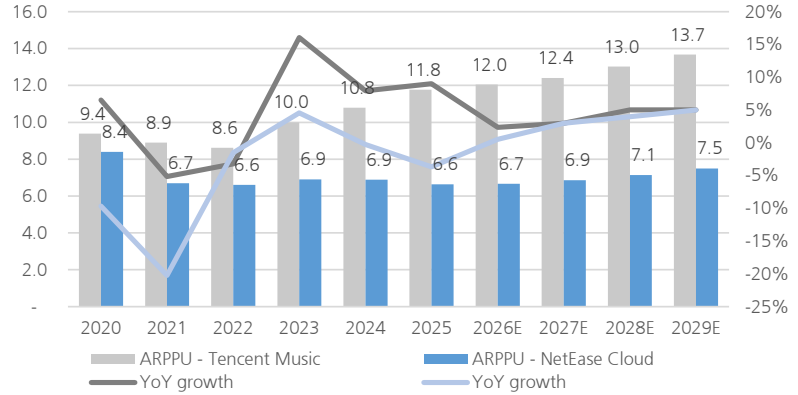
**The MAU of China music streaming platforms (as of Aug 2026)**



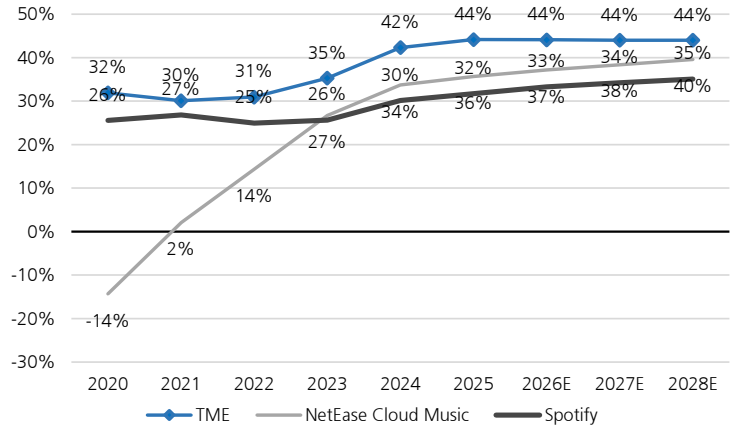

**TME's multi-tiered membership**

| Membership type | Explanation |
|---|---|
| **SVIP** | In addition to standard music-related benefits, it includes high-fidelity audio options, ecosystem benefits across long-form audio and online karaoke, as well as offline and fan-related benefits such as priority ticket booking, exclusive artist merchandise, and fan-artist interaction events. |
| **Standard VIP** | The standard paid membership tier focused on core music benefits. Key benefits typically include access to pay-for-streaming content, a fixed number of monthly downloads, faster streaming speeds, personalized app themes or avatar designs, as well as selected early-access rights for new song releases. |
| **Ads Membership (Freemium)** | A reward-based membership tier targeting freemium users. Users watch advertisements in exchange for specific benefits, such as trial subscriptions, digital albums, or access to select music content that typically requires payment. |

**TME and Cloud Village ARPPU in 2010-2029E**



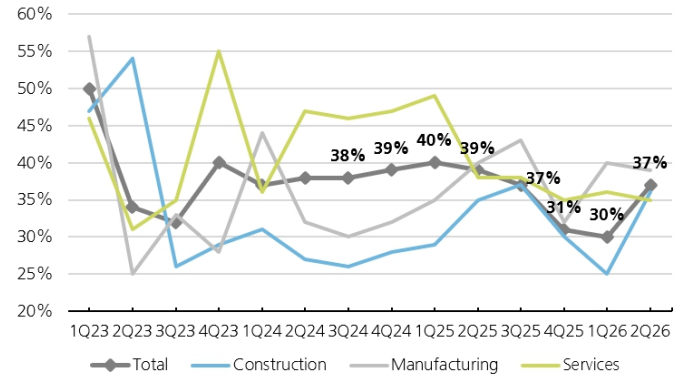

**The gross margins of TME, NetEase Cloud Music and Spotify**



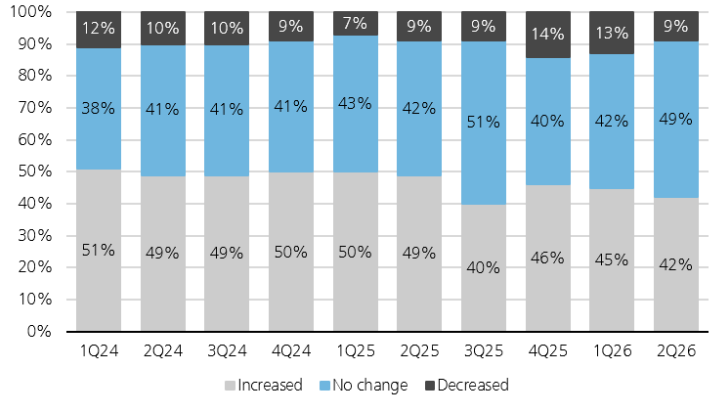

Sources: QuestMobile, Company data, UBS estimates.

# Online recruitment: Continued online shift, AI disruption concern overdone

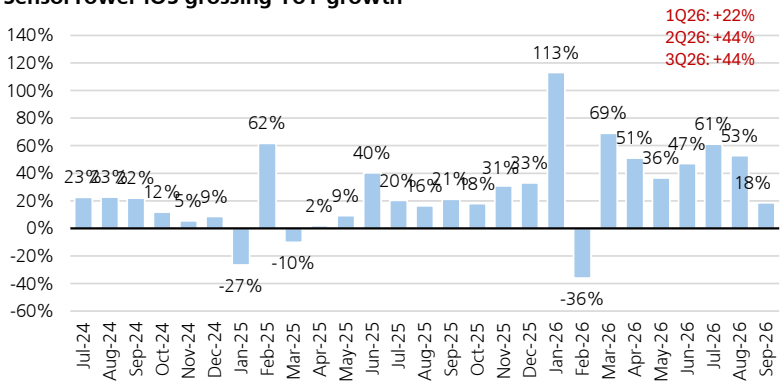
**Expected QoQ net change in staff hiring**



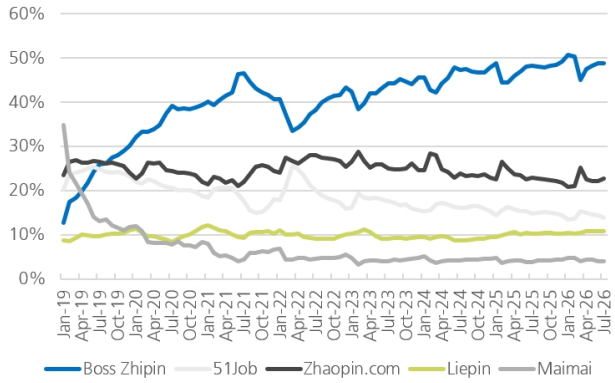

**Expected QoQ net change in hiring budget**




**SensorTower iOS grossing YoY growth**




**MAU share among major recruitment platforms**




Sources: SensorTower, UBS Evidence Lab

Section 3

# Key stocks

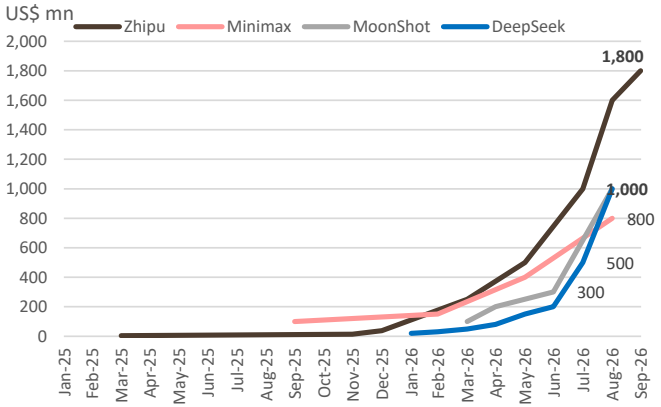
# Z.AI (Zhipu): Accelerated scaling in model development and monetization

- **End-2026 ARR guidance raised to US$3bn with three key drivers.** Management shared that its latest ARR reached c.US$1.8bn and raised end-2026 ARR guidance to US$3.0bn (vs US$2.4bn previously), mainly driven by:
  1) **Coding remains the core growth engine**: continued API and Coding Plan growth, supported by easing compute constraints, alongside the rapid ramp-up of coding harness ZCode (2m+ users);
  2) **Cowork is gaining traction:** GLM-5.3 entered production deployment in cybersecurity, with Cowork orders exceeding Rmb1bn within 1 month of launch
  3) **CSP partnerships expand distribution and overseas monetization:** Zhipu has signed API hosting and revenue sharing agreements with leading domestic and overseas CSPs, leveraging their inference infrastructure, enterprise customers and sales channels to accelerate global distribution.
- **New financing round to support compute expansion:** With the recent US$5bn financing, management expects to invest in compute expansion, with 40%/60% of the added resources allocated to model training/inference, making the company one of the leading labs in terms of compute ownership and supporting further acceleration in inference revenues.
- **In model development,** we see Zhipu emphasizing "right scaling" by optimizing compute allocation across data, base-model parameters, post-training and infra to deliver stronger intelligence at lower cost.
- **Stock view:** We are encouraged by the solid ARR growth, yet recent sector derating amid rising concerns over global competition and uncertainty around the model release pipeline has weighed on share performance. Looking ahead, we expect Zhipu's next major model upgrade to demonstrate a step change in model parameters and capabilities, which will likely come out in 4Q26, while A-share listing and lock-up expiry are major catalysts to watch out for in 1Q27.

**Zhipu's key text model portfolio launched in 2026**

| Tier | Model | Release date | Parameters (total / active) | AA Index | Input ($/mn tokens) | Cached input ($/mn tokens) | Output ($/mn tokens) | Blended price (7:2:1 blend) |
|---|---|---|---|---|---|---|---|---|
| Flash (high-speed) | GLM-5.3-FlashX | n.a. | 320B / 18B | n.a. | $0.37 | $0.075 | $1.25 | $0.25 |
| Flash | GLM-5.3-Flash | 26 Aug 2026 | 320B / 18B | 41.8 | $0.15 | $0.03 | $0.50 | $0.10 |
| Flagship | GLM-5.3 | 18 Aug 2026 | 753B / 40B | 44.8 | $1.40 | $0.26 | $4.40 | $0.90 |
| Flagship | GLM-5.2 | 16 Jun 2026 | 753B / 40B | 33.7 | $1.40 | $0.26 | $4.40 | $0.90 |
| Flagship | GLM-5.1 | 7 Apr 2026 | 744B / 40B | 26.1 | $1.40 | $0.26 | $4.40 | $0.90 |
| Flagship | GLM-5 | 12 Feb 2026 | 744B / 40B | 27.9 (est.) | $1.00 | $0.20 | $3.20 | $0.66 |

**The ARR trend of major China AI labs**

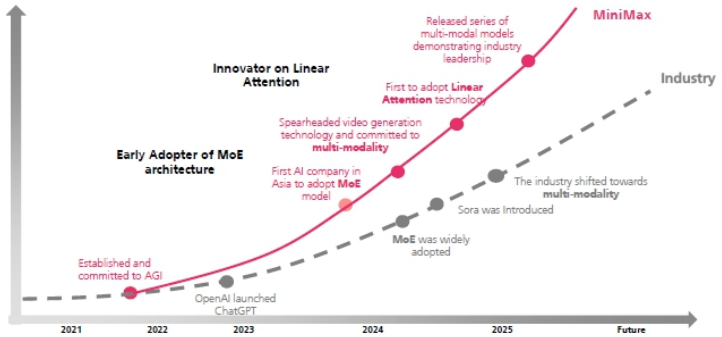



Sources: The Information, company data, UBS estimates. Note: The above models may not be exhaustive.

# MiniMax: Strong ARR trends driven by model improvement and secular demand

- **1H26 results**: MiniMax delivered accelerating topline growth with a narrowing loss ratio in its 1H26 results. We continue to observe the model upgrades driving robust token usage and monetization, with platform token consumption in Jul-26 reaching 20x the Jan-26 level and 2Q revenue growing >80% QoQ.
- **ARR trend**: We are encouraged by MiniMax's ARR ramp-up, exceeding US$800m in Aug (vs c.US$400m in May and US$150m in Feb), which we believe reflects 1) the usage ramp-up of its flagship text model M3, following inference-performance optimization and service refinements; 2) MiniMax's advantage in compute resources to better capture rapidly growing token demand from expanding use cases (e.g., coding, agents and office harness products from leading internet companies); and 3) the better-than-expected performance of its recently launched multimodal model H3. We forecast 2026E year-end ARR to be c.US$1.2bn.
- **Margins**: That said, 1H26 gross margin was lower than the 4Q25/FY25 levels, partly due to one-off impacts from M3's initial launch, as well as a structural revenue mix shift towards lower-margin text models (vs. multimodal models) and Token Plans (vs. API). Looking ahead, we still see room for improvement, with further inference optimization and potential price increases for next-generation models to drive margin accretion.
- **Stock view:** Similar to Zhipu, MiniMax's share price was affected by cautious market sentiment in light of intensifying global competition. Its new flagship text model M3 Pro and an upgrade to its video generation model H3.1 are key catalysts to watch for in the upcoming months. Longer term, we like MiniMax's differentiated positioning with multimodal models and applications, as well as a proactive globalization strategy.

**Multi-modal strategy pioneer's proven model algorithm innovation**




**Latest ranking of MiniMax's frontier models, by modality**

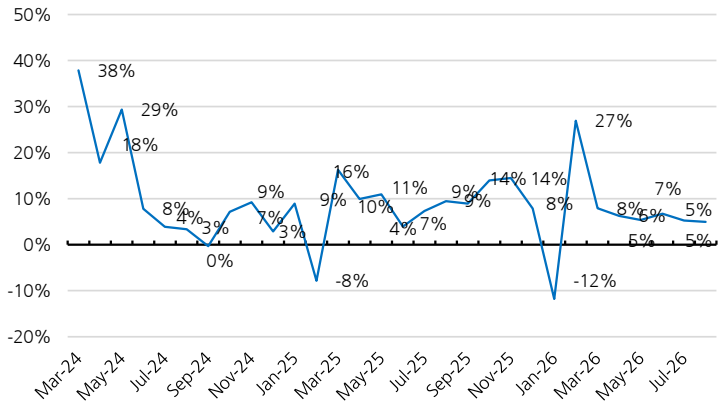
| Modality | Model | Release date | AA ranking |
|---|---|---|---|
| Text (latest preview) | MiniMax-M3.1-Flash-Preview | Sep 2026 | Not yet ranked |
| Text (all models) | MiniMax-M3 | Jun 2026 | 36 |
| Text (open weights) | MiniMax-M3 | Jun 2026 | 13 |
| Video | MiniMax H3 (768p) | Jul 2026 | T2V: 4 / I2V: 1 |
| Speech | Speech 2.8 HD | Feb 2026 | 21 |
| Music | Music 3.0 | Jul 2026 | Not verified on AA |

Sources: OpenRouter, Artificial Analysis. Note: The above data is as of 14 Oct 2026.

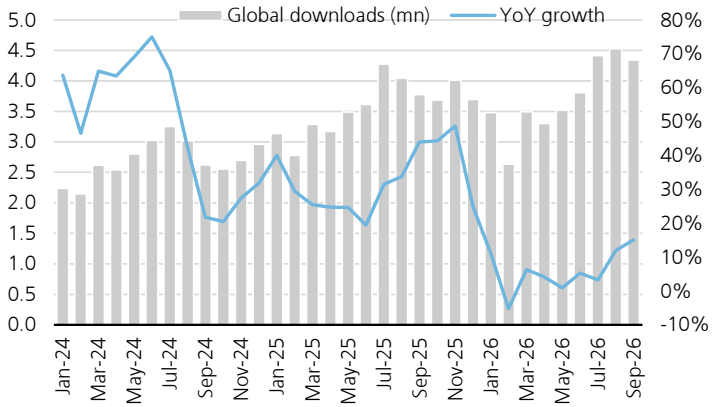
# Trip.com: Near-term earnings better than feared; lower expectations for 2027

- **Rectification transition and industry headwinds to weigh on growth**
  - Following the conclusion of antitrust investigation, we expect **Trip.com to proactively refine its operating model with hotel partners**, which could weigh on hotel take rate from 3Q onwards (UBSe of single digit decline YoY).
  - **Industry headwinds**, including extreme weather and elevated oil prices, also continue to pressure domestic and outbound travel demand in accommodation and transportation.
  - While we believe consumers' travel intent remains resilient, we expect these external headwinds to lead to low single digit total revenue growth in 3Q and 4Q, and potentially 1Q27 given tough comparable base, while gradually ease from 2Q27 onwards.
- **International businesses as structural growth drivers long-term:** On Trip.com overseas, we expect hotel and transportation revenue to remain robust, sustaining c50% YoY in 2Q and 3Q, while improving scale and marketing efficiency should support further loss narrowing. On inbound travel, we expect broadening demand across source markets and destinations to keep the company on track to achieve its long-term target of serving 200m inbound travellers over the next five years.
- **Margins:** we expect Trip.com to maintain cost discipline and narrow losses in international markets, yet the top line weakness and revenue mix shift towards Trip.com overseas are likely to limit the margin expansion pace**.**
- **Stock view:** We remain positive for TCOM's long-term positioning as a domestic leader with increasing overseas exposure, and believe the ongoing headwinds have largely been priced in at 10x 2026 cons PE, while the stock may remain range bound in the near term as investors wait for better earnings visibility in 2027.

**Number of Orders YoY % - Ctrip**




**Trip.com global downloads (m)**

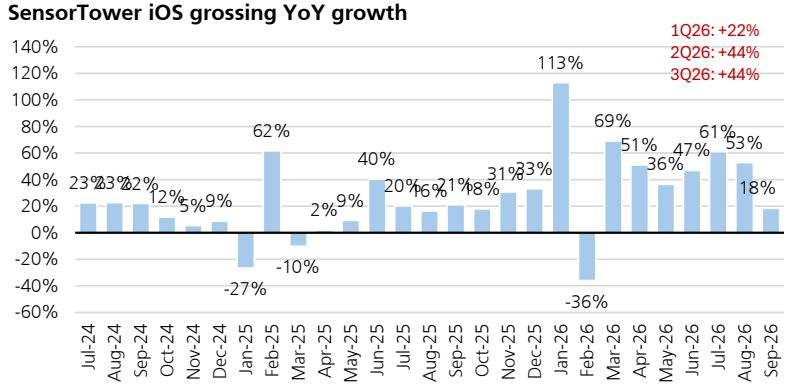



Source: UBS Evidence Lab, QuestMobile, SensorTower

# Kanzhun: Pivoting monetization strategy towards paying ratio and ARPPU

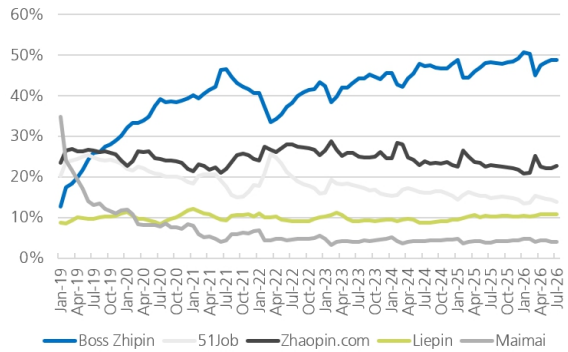
- **More proactive approach towards platform monetization.** We are encouraged that management has become more proactive towards platform monetization amid macro weakness by:
  1) **Laying out its next five-year growth plan**, with more balanced contributions from both user growth and ARPPU (e.g., increasing spend from free/low-paying enterprises);
  2) **Adopting differentiated strategies by market maturity**, gradually increasing monetization intensity in more mature tier-1 and selected tier-2 cities, while continuing to prioritize user penetration in lower-tier cities;
  3) **Recognizing the gap between the value offered and current pricing**, as its unit pricing remains low compared with that of overseas platforms, while new AI-powered products/features could drive incremental revenue opportunities.
- **Shareholder return:** YTD dividends and buybacks totalled US$530m and exceeded 100% of FY25 adjusted net income, implying a c9% total yield for FY2026, by our estimates. We expect management to remain committed to shareholder returns in 2H26 and beyond.
- **Margin**: Despite higher World Cup marketing and AI-related investment, we expect operating leverage and improving efficiency to support expansion in adjusted OPM in 3Q26 and FY26.
- **Overseas expansion as a long-term growth driver**: Offer Today targets US$100-150m in annual revenue in Hong Kong in five years, and may consider entering other selected overseas cities, which we think provides another long-term growth opportunity but takes time to materialize.
- **Stock view:** We expect Boss to increase its paying ratio/ARPPU at a gradual pace (e.g., starting with selected verticals in higher-tier cities), and are not overly concerned about the potential impact on its market share or user experience, yet we acknowledge the execution results and enterprise customer feedback are worth monitoring. Our DCF-based price target of US$21.0 implies 15x 2027E P/E on a low-teens earnings CAGR and a solid shareholder return commitment.

**SensorTower iOS grossing YoY growth**




Source: UBS Evidence Lab, Sensor Tower, QuestMobile

**MAU share among major recruitment platforms**



# Baidu: Stepping up AI investment to strengthen full-stack capabilities

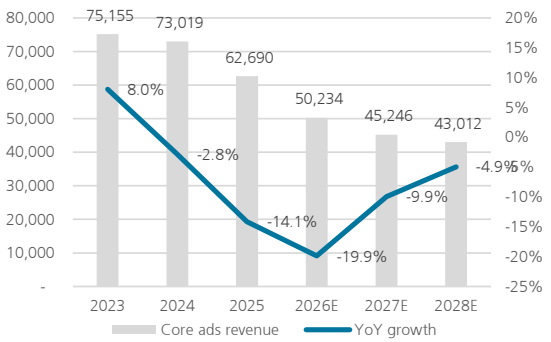
- **AI Cloud Infra remains healthy; higher AI investments in 2H**
  - **AI Cloud Infra** growth moderated to 50% YoY in 2Q (vs. 79% in 1Q), which we think was mainly due to timing of revenue recognition, while compute demand from AI training and inference remained healthy. We expect this segment to reaccelerate to 60%+ growth in 2H26, driven by both higher spending from existing clients and an expanding client base (e.g., autonomous driving, embodied AI).
  - The continued strong growth of higher-margin GPU Cloud business (revenue +283% YoY in 2Q26) could also benefit its Cloud and overall margins.
  - Baidu is committed to **increasing AI investment to pursue better ERNIE model capabilities**, in order to drive incremental growth opportunities in cloud and applications, and strengthen full-stack AI capabilities and cost efficiency over the long term, yet this is likely to weigh on OP margins and cap the stock upside in the near term.
- **Kunlunxin:** Management indicated the listing progress remains on track. We believe the ongoing tight GPU supply, together with Kunlunxin's advantages in performance and compatibility, could support solid revenue growth outlook in 2026-27.
- **Scaling up compute capacity**: Given robust internal and external demand, Baidu expects to potentially double its compute capacity by end-2027, with diverse funding channels to manage cash flow impacts.
- **Shareholder returns**: Buyback execution slowed temporarily in 2Q due to higher AI investment and market volatility, but we believe management remains committed to improving shareholder returns, with dividend plan on track by end-2026..
- **Stock view:** Our SOTP-based PT of US$160 assumes a 3x PE for Core ads, 2x P/S for Cloud, US$56/share for Kunlunxin, US$7/share for Apollo Go (benchmarking Pony.AI's market cap) and net cash and long-term investments. Our PT for 9888.HK is HK$155.

**SOTP valuation of Baidu**

| | Baidu stake % | Investment discount | Value (US$m) | Value Per share (US$) | Value % | Valuation methodology |
|---|---|---|---|---|---|---|
| Search and feed | 100% | | 3,472 | 10 | 6% | 3x P/E on 2026E earnings for a mature business with flattish earnings CAGR in 2026-28E |
| Cloud | 100% | | 8,654 | 25 | 16% | 2x P/S on 2026E sales (ex-Kunlunxin), below industry peers (i.e., Alibaba) |
| Apollo | 100% | 30% | 2,407 | 7 | 4% | Benchmark market cap of Pony.AI (US$3.4bn, as of 17 Aug 2026) |
| Kunlunxin | 59% | 30% | 19,244 | 56 | 35% | Apply 40x PS on Rmb8bn revenue in 2026E |
| Cash and equity investments | | | | | | |
| Net cash and investments | | 30% | 20,641 | 60 | 38% | Apply 30% holding discount on US29.5bn. Based on latest earnings release |
| iQIYI | 45% | 30% | 555 | 2 | 1% | UBS PT based on DCF (WACC: 9.4%, "g": 1.3%) |
| **Total valuation (US$m)** | | | **54,974** | **160** | | |
| Share outstanding (m) | | | 343 | | | |

Sources: UBS estimates, company data

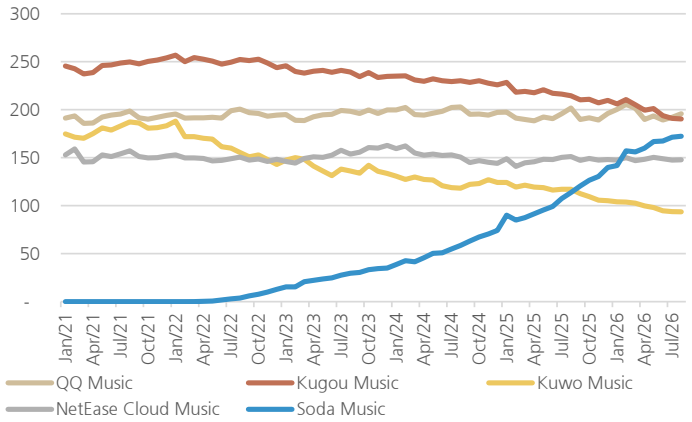
**Baidu Core ads revenue (in Rmb m) and YoY growth**



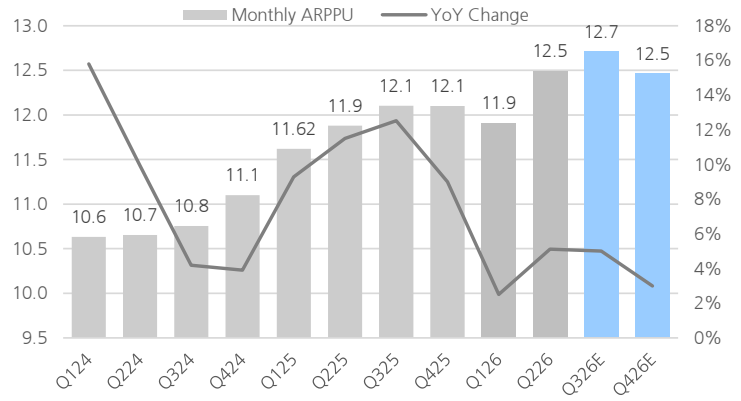
# Tencent Music: Advancing an IP-centric strategy amid ongoing competitive pressure

- **While TME delivered a revenue/earnings beat vs cons in Q2,** we believe it was mainly attributable to the Ximalaya consolidation, which contributed Rmb407m to revenue (recorded in music-related revenue) and is gross margin accretive.
- **Membership services**: Excluding the consolidation, we estimate only limited YoY growth in music membership services in Q2, suggesting lingering industry competition. We expect this competitive pressure to persist in the near term and keep organic membership revenue growth subdued in H2.
- Following Ximalaya's consolidation, we think **TME is transitioning to become a comprehensive music and audio services platform**, with an increasing focus on IP-centric strategies to strengthen its premium content ecosystem and unlock greater value through its integrated platform.
- **Shareholder returns**: we are encouraged by management's execution, with US$400m of share repurchases in Q226, on track to fully utilise its existing US$1bn buyback programme (expiring in March 2027). Meanwhile, its solid cash and investment balance (US$6.5bn as of Q2, or 49% of its market cap) should support further shareholder return initiatives thereafter, in our view.
- **Stock view:** The increasing focus on IP-centric strategies should also benefit its membership business through enriched privileges and support a differentiated positioning versus peers, although we think it will still take time to bear fruit. We thus remain Neutral on TME as we monitor the evolving competitive landscape and its transition.

**Online music peers MAU (mn)**




**TME's ARPPU (Rmb) and YoY growth**




Sources: UBS Evidence Lab, QuestMobile, company data, UBS estimates

**Valuation Method and Risk Statement**

We believe key sector risks include: 1) an evolving competitive landscape and intensifying competition; 2) fast-moving trends in technology, as well as internet users' needs and preferences; 3) uncertain monetisation; 4) the rising cost of traffic acquisition, content and brand promotions; 5) the upkeep of IT systems; 6) expansion into international markets; 7) adverse changes in market sentiment; and 8) regulatory risks. Valuation methods in the sector include SOTP analysis.