

China Artificial Intelligence

After the rally — What's priced in, what could change and how to position?

Zhipu AI and MiniMax have delivered extraordinary post-IPO returns, with both stocks appreciating roughly 5-7x from their Jan 2026 IPO, as the market has rapidly rerated them into the two most visible listed proxies for China's frontier model layer. The central question for investors now is whether the rally has fully discounted the commercial opportunity or **whether further upside exists**, and critically, **what scenarios could drive material re-rating in either direction**. We remain constructive and maintain **Overweight ratings on both names**, and we view **technically driven pullbacks as buying opportunities**. At the same time, the AI foundation model sector still faces meaningful uncertainties around competition, pricing and monetization, so we also **lay out a focused set of fundamental tripwires** that would lead us to take a more cautious view.

- We continue to view both Zhipu and MiniMax as structural beneficiaries of China's AI adoption cycle, supported by strong API demand, improving monetization, and a domestic market that remains early in scaling enterprise AI workloads. Within that constructive view, the two stocks now offer different ways to express theme: **Zhipu offers a clearer monetization framework and a more measurable execution path** towards ~\$1bn year-end ARR, supported by high visibility on API pricing and coding adoption; **MiniMax offers greater ARR upside versus ~\$700mn year-end ARR currently priced in** if its ability to source and deploy compute more flexibly across regions allows it to scale revenue faster while industry remains constrained by tight GPU supply.
- We evaluate both names across four dimensions: (1) what **fundamentals** are already embedded in the share price, (2) **near-term catalysts** that can drive momentum, (3) **fund flow dynamics** including lock-up expiration, financing, Index/Stock Connect inclusion, and quantified 'scarcity premium', and (4) how **US frontier AI leaders'** growth and potential IPO influence valuation anchors. This framework is designed to translate a broad AI theme into stock-specific underwriting. Model cycles such as GLM 5.5 and MiniMax M3 remain the next key event to watch. At the same time, lock-up supply, financing windows and index-related flows will shape the trading path. In this setup, we see **flow-driven volatility as an opportunity to build positions, as long as the underlying execution trajectory remains intact**.
- The **key fundamental tripwires** are relative **model positioning, competitive dynamics, pricing discipline, the China-US capability gap, and ARR conversion from usage growth**. More specifically, we are focusing on signals of (1) loss of relative model leadership in coding, agent and enterprise workloads; (2) internet or cloud majors establishing a clear SOTA lead that weakens the role of independent model providers as premium monetization layers; (3) a return to broad-based API price competition that compresses the link between better models and higher pricing; (4) a renewed widening in the China-US model gap; and (5) weaker ARR conversion despite strong usage growth. **As long as these conditions remain supportive, technically driven pullbacks should be used to add exposure. A deterioration in these conditions would lead us to take a more cautious view.**

Internet

Olivia Xu ^{AC}

(86-21) 6106 6138

olivia.w.xu@jpmorgan.com

SAC Registration Number: S1730525060001

Alex Yao

(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

J.P. Morgan Securities (China) Company Limited

2513.HK, 2513 HK

Overweight

Price: HK\$950.00

21 Apr 2026

Price Target: HK\$950.00

PT End Date: 31 Dec 2026

0100.HK, 100 HK

Overweight

Price: HK\$898.50

21 Apr 2026

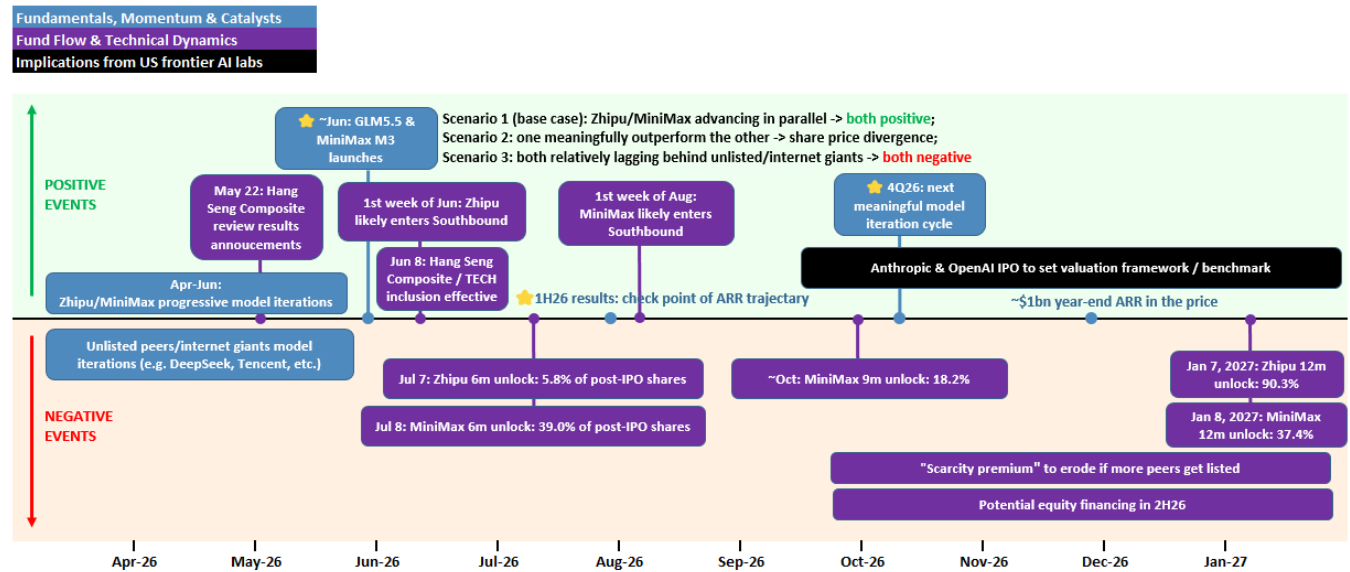
Price Target: HK\$1,100.00

PT End Date: 31 Dec 2026

See page 23 for analyst certification and important disclosures, including non-US analyst disclosures.

J.P. Morgan does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.

Figure 1: Major events timeline for Zhipu AI and MiniMax



Source: Company prospectus (data as of Dec31, 2025), Hang Seng Index website, Reuters, Sina Finance, Sina News, AASTocks, and J.P. Morgan estimates. (data as of Apr20, 2026)

Table 1: Major events timeline, impacts and expected share price direction for Zhipu AI and MiniMax

Timing	Event	Impact	Share price direction
Apr-May 2026	Zhipu / MiniMax ongoing model iterations	Focuses on absolute & relative model progress, pricing power and enterprise adoption cadence	Positive
	DeepSeek V4 launch	Key read-across for China model competition and for domestic GPU-based training/inference validation	Outcome-dependent
	Tencent Hunyuan 3.0 launch	A key test of how quickly scaled platforms can close the capability gap	Outcome-dependent
Jun-26	GLM 5.5 / MiniMax M3	Potential step-function catalyst for ARR; market focus on both absolute capability and relative positioning	Positive, unless relative performance disappoints
	Potential index inclusion	Passive flow catalyst and broader institutional ownership expansion	Positive
	Potential Zhipu Southbound inclusion	Expands Mainland investor access and can improve liquidity / valuation support	Positive
Jul-26	6-month lock-up expiry	Pre-IPO shareholders at 5-7x gains; Zhipu unlocks 5.8% share, MiniMax at 39.0%	Negative
Aug-26	1H26 results	Check point of ARR trajectory, pricing durability and margin trajectory	Outcome-dependent, positive if execution holds
	Potential MiniMax Southbound inclusion	Broadens investor base and may partly offset post-lock-up flow pressure	Positive
Oct-26	9-month lock-up expiry	MiniMax 9m unlock at 18.2%	Negative
2H26	Potential equity financing	Additional financing is highly likely; R&D = 4-5x revenue	Negative
	Potential Anthropic / OpenAI IPO	Sets global AI multiple anchor	Positive
	Potential Kimi / StepFun listing	Validates the sector's investability, but also dilutes scarcity	Mixed to slightly negative
Jan-27	12-month lock-up expiry	Zhipu 12m unlock at 90.3%, MiniMax at 37.4%	Negative

Source: Company prospectus (data as of Dec31, 2025), Hang Seng Index website, Reuters, Sina Finance, AASTocks, and J.P. Morgan estimates. (data as of Apr20, 2026)

Our four-dimension framework conclusions:

- **Fundamental: valuations already embed not just the path to ~\$1bn year-end ARR, but also an Anthropic-like post-inflection growth trajectory.** We believe current share prices are anchored to the API/MaaS franchise and imply roughly US\$1bn year-end 2026 ARR for Zhipu and ~US\$700mn for MiniMax. More importantly, by benchmarking against how Anthropic was priced around its ~US\$1.4bn ARR stage in Mar 2025, the market is implicitly underwriting that both companies can sustain high-velocity ARR expansion beyond \$1bn, potentially following a trajectory similar to Anthropic's \$1bn to \$5bn ramp over 6 months – a demanding assumption in a more competitive China market. Even so, we remain structurally positive, as current ARR still likely understates true demand under compute constraints, and both companies retain credible positioning in the highest-value segments of the model market, particularly coding, agent and enterprise workloads.
- **Momentum & catalysts: the next leg depends on whether model upgrades, token growth and pricing continue to translate into better monetization.** The next major operating catalysts remain the upcoming model cycles, particularly GLM 5.5 and MiniMax M3 in June, as each major release has historically driven a visible step-up in adoption and ARR. At the sector level, token consumption remains supportive, though raw usage growth now matters less than whether demand is converting into recognized revenue at higher realized pricing. That conversion is improving through several channels already visible in the market, including direct API repricing at model upgrades, migration from subsidized coding plan toward heavier pay-as-you-go usage, and structural repricing in cache-related economics. Our read remains positive: the demand backdrop is still firm, and sector revenue quality is improving. The next question is whether product progress continues to drive monetization at a pace sufficient to support current expectations.
- **Fund flow & technical dynamics: near-term fund flow headwinds are real, but we view them as technical pressure rather than a break in thesis.** The main flow risks are lock-up expiry, likely post-IPO financing, and a gradual reduction in scarcity premium as more China AI names move toward listing. MiniMax faces the heavier near-term supply calendar (39.0% share lock-up expiration in July), while Zhipu's setup is relatively clean (5.8% share) given a smaller six-month unlock and an earlier path toward Southbound inclusion. These events can create volatility and short-term derating, especially after such strong post-IPO gains. However, we do not view them as reasons to turn structurally cautious. Index inclusion (June) and Stock Connect (June/August) should provide meaningful offsetting demand support, while any technically driven pullback may create better entry points if company-level execution remains intact.
- **US frontier AI: accelerating US ARR growth is supportive for Chinese LLM market validation, while future IPOs could reset the valuation framework.** Continued ARR acceleration at Anthropic and OpenAI reinforces the view that frontier-model monetization can scale to levels far above prior market assumptions, which remains supportive for the long-duration revenue opportunity available to Chinese LLM companies. That external validation is important as it helps sustain the valuation anchor for Zhipu and MiniMax. At the same time, eventual IPOs from Anthropic and OpenAI may shift the sector's benchmark from private-market ARR comparisons toward a more public-market focused matrix, which may create framework risk, especially for names whose equity story still relies heavily on long-duration growth. Even so, our central view remains unchanged: as long as Zhipu and MiniMax continue to defend competitive relevance and convert demand into recognized revenue, US frontier AI progress is more supportive than threatening for Chinese listed LLM names.

Table of Contents

Fundamental: What's priced in & what to watch?	5
ARR ramp-up: the single most important variable — ~\$1bn year-end ARR is in the price, and implicitly expecting Anthropic-like post \$1bn inflection growth trajectory	5
Global competition: US-China frontier gap – currently priced in 9-12 months gap	6
Domestic competition: internet giants vs. independents – base case assumes independents will remain competitive, even superior in coding/agent applications	8
Competition within independents: currently expecting advancing in parallel	8
Momentum & Catalysts	9
• Model launches: still the most visible company-level catalyst	9
• Token consumption remains supportive, structural acceleration underway	9
• Pricing is becoming an increasingly important earnings driver, with several mechanisms now moving in the same direction	10
Fund Flow & Technical Dynamics	12
• -ve: Lock-up expiry: pre-positioning typically begins two months ahead of the unlock, though relative-value support may limit de-rating	12
• +ve: Stock Connect/Index inclusion should provide the meaningful source of flow support	16
• Neutral/-ve: Potential placement should be treated as part of the base-case funding framework rather than as a tail risk	16
• -ve: New listings pipeline: Cambricon case study indicates 25-30% “scarcity premium”	16
US Frontier AI: Implications for Chinese LLM Valuations	18
• Accelerating US frontier ARR: structurally bullish for Chinese LLM valuations..	18
• Anthropic & OpenAI IPO: the valuation anchor event – matter as much for valuation methodology as for valuation level	18
• AI value-chain positioning: China’s listed LLM layer already carries meaningful weight, while the US comparison still reflects private-market marks	19

Fundamental: What's priced in & what to watch?

ARR ramp-up: the single most important variable — ~\$1bn year-end ARR is in the price, and implicitly expecting Anthropic-like post \$1bn inflection growth trajectory

At ~\$40-55bn market cap, we believe the market is pricing Zhipu AI and MiniMax to reach approximately \$1bn and \$700mn in API ARR by year-end 2026 respectively. Current valuation also embeds an external benchmark that investors should state explicitly. At US\$40-55bn market cap, Zhipu/MiniMax are no longer being valued simply as a fast-growing China AI company; it is being benchmarked against Anthropic at a comparable revenue stage. Anthropic's March 2025 Series E valued the company at about US\$61.5bn on roughly US\$1.4bn ARR, implying a trailing multiple near 44x. The implication is that **the market is underwriting an Anthropic-like post-inflection growth trajectory once Zhipu / MiniMax crosses the first US\$1bn ARR threshold**, namely another 5x ARR hike in 6 months. That benchmark clarifies where upside remains available and where risk sits: if ARR keeps compounding rapidly after the next model cycle, current valuation can still be supported; if growth normalizes materially below that path, the stock becomes more exposed to multiple compression.


- The case for sustaining Anthropic-level velocity rests on a stage comparison between China's AI adoption in 2026 and the US in early 2025. China enters the \$1bn ARR phase with **structural tailwinds absent when Anthropic was at the same point**: AI agent frameworks (OpenClaw) have achieved mass developer adoption, token-intensive API calls are embedded into production workflows, and enterprise willingness to pay has been validated by Zhipu's API price hike and Alibaba Cloud's compute price increase — both of which saw demand accelerate.
- **The primary uncertainty is competitive fragmentation**. Anthropic's \$1bn→\$5bn sprint occurred facing limited direct competition. Chinese independents face a fundamentally more crowded field: internet giants like Alibaba, Bytedance and Tencent are intensifying their B2B push: Tencent launched WorkBuddy, Alibaba is scaling Qoder and TRAE, ByteDance is bundling Doubao into enterprise suites. Among independents, the model iterations are all targeting better coding and agent applications: if the \$1bn→\$5bn opportunity is distributed across 6-8 credible players rather than concentrated in 2, per-company velocity would be lower even if aggregate market growth matches the US pace.

The ARR trajectory is not linear — it is a step-function tied to model iterations & compute power unlock. Zhipu AI disclosed its API ARR surpassed \$250mn on Mar31, up 60x YoY and 6.4x YTD in 2026, each successive GLM release (4.5 → 4.6 → 4.7 → 5 → 5.1) produced a discrete ARR jump — the curve goes near-vertical after GLM-5. MiniMax shows a similar pattern post the launch of M2.2-2.5-2.7. Such ARR growth pattern means the next model cycle (GLM 5.5, MiniMax M3, both expected ~June) is not just a product event — it is the single most predictable ARR catalyst on the calendar.

Critically, current ARR understates true demand. Every major Chinese LLM provider has reported that inference compute cannot keep pace with the rapidly expanding token demand, and such demand remains rigid despite upward revising pricing actions along the supply chain – at LLM side, Zhipu AI saw API price almost double year-to-date; on cloud side, Alibaba announced AI compute price +34% effective on Apr18. Such compute under-supply situations means that, any capacity expansion, whether domestic/overseas GPU allocations, cloud partnerships, or inference optimization, would directly convert suppressed demand into recognized revenue. **The ARR growth we see today is the floor, not the ceiling**.

The DeepSeek V4 wildcard. DeepSeek's forthcoming V4 model (expected to launch at late April) is [reportedly](#) being built to run natively on Huawei Ascend chips — the first frontier model to bypass Nvidia entirely for inference optimization. This has a second-order implication for the compute bottleneck: if V4 validates Ascend as a viable inference platform at frontier scale, it effectively expands the addressable compute supply for the entire Chinese LLM ecosystem, and more available compute could translate into faster ARR conversion.

Figure 2: Zhipu AI / MiniMax ~\$1bn year-end ARR is in the price, and implicitly expecting Anthropic-like post \$1bn inflection growth trajectory



Source: Company official website, The information, CNBC, Yahoo Finance, Reuters, J.P. Morgan estimates (data as of Apr 20, 2026)

Table 2: Anthropic and OpenAI's historical valuation

Anthropic				
Date	Funding Round	Post-Money Valuation	Reported/Estimated ARR	Implied P/ARR
Feb-26	Series G	\$380 billion	\$14 billion	27.1
Sep-25	Series F	\$183 billion	\$5 billion+	36.6
Mar-25	Series E	\$61.5 billion	~\$1.4 billion	43.9
Jan-24	Series D	\$19.4 billion	~\$200 million	97.0
May-23	Series C	\$4.5 billion	~\$100 million	45.0
OpenAI				
Date	Funding Round	Post-Money Valuation	Reported/Estimated ARR	Implied P/ARR
Apr-26	Series G	\$852 billion	\$25 billion	34.1
Mar-25	SoftBank-led Series F	\$300 billion	~\$10 billion	30.0
Oct-24	Strategic Funding	\$157 billion	\$3.7 billion	42.4
Jan-23	Microsoft Strategic	\$29 billion	~\$200 million	145.0

Source: Company websites, The information, Reuters, CNBC (all data as of Apr 20, 2026)

Global competition: US-China frontier gap – currently prices in 9-12 months gap


We believe the current market pricing assumes Chinese LLM players maintain a 9-12 month technology gap versus US leaders (Anthropic Claude, OpenAI GPT, Google Gemini) and are catching up rapidly. The **downside risk** lies in the scenario where the next generation of US models **delivers a step-function improvement that meaningfully re-widens the gap**. If US next-generation models demonstrate capabilities qualitatively beyond what Chinese players can replicate within 9-12 months, the ‘catch-up’ narrative would be reasonably challenged, and the valuation premium for Chinese independence in AI is likely to erode.

Conversely, **the upside scenario is that the gap continues to narrow** or that Chinese players achieve parity in specific verticals (coding, enterprise agents, multimodal applications), which would validate higher multiples and potentially attract incremental global institutional flow into HK-listed AI names.

Figure 3: US versus China: frontier language model intelligence over time

Frontier Language Model Intelligence By Country, Over Time

Artificial Analysis Intelligence Index v4.0 incorporates 10 evaluations: GDPval-AA, τ^2 -Bench Telecom, Terminal-Bench Hard, SciCode, AA-LCR, AA-Omniscience, IFBench, Humanity's Last Exam, GPQA Diamond, CritPt




Source: Artificial Analysis ([link](#), data as of Apr 20, 2026)

Figure 4: US versus China: model intelligence over time

Performance of top United States vs. Chinese models on the Arena

Source: Arena, 2026 | Chart: 2026 AI Index report



Source: Sanford University Human-Centered Artificial Intelligence ([link](#), data as of Apr 20, 2026)

Domestic competition: internet giants vs. independents – base case assumes independents will remain competitive, even superior in coding/agent applications

The structural investment thesis for Zhipu and MiniMax as independent LLM providers **rests on a core assumption that independent players can remain competitive against the model development efforts of China's internet giants** — Alibaba, ByteDance, and Tencent. This assumption is non-trivial: the incumbents internet giants possess superior distribution, larger captive datasets, deeper compute budgets, and the ability to subsidize AI development through profitable existing businesses.

The bear case here is that if any of these giants launches a model that achieves clear state-of-the-art (SOTA) performance and is meaningfully superior to Zhipu's GLM or MiniMax's offerings, it challenges the fundamental premise that the market needs independent LLM companies. The counterargument, and the basis of the bull case, is that enterprise customers often prefer vendor-neutral AI infrastructure, and that the independent players' singular focus on model quality gives them an edge in frontier capability that diversified tech conglomerates cannot match. **Therefore, we see Tencent's launch of Hunyuan 3.0 a key test of how quickly scaled platforms can close the capability gap.**

Competition within independents – currently expecting advancing in parallel


The current market assumption is that China's leading independent LLM players are advancing roughly in parallel. Zhipu, MiniMax, Moonshot AI (Kimi), and DeepSeek are all seen as credible frontier competitors. If this assumption holds, both listed stocks should maintain correlated momentum.

However, if relative gaps widen — for example, if one player's next-generation model significantly outperforms peers, or if one captures disproportionate enterprise market share — share price momentum would diverge. This creates both a risk and an opportunity: investors must evaluate whether their preferred name is pulling ahead or falling behind within the cohort. If the differentiation persists, we expect a market to shift from 'own any AI exposure' to 'pick the winners'.

Figure 5: Chinese independent LLM players' model intelligence over time

Frontier Language Model Intelligence, Over Time

Artificial Analysis Intelligence Index v4.0 incorporates 10 evaluations: GDPval-AA, r^2 -Bench Telecom, Terminal-Bench Hard, SciCode, AA-LCR, AA-Omniscience, IFBench, Humanity's Last Exam, GPQA Diamond, CritPt



Source: Artificial Analysis ([link](#), data as of Apr 21, 2026)

Momentum & Catalysts: model launches, token consumption and pricing dynamics

Model launches: still the most visible company-level catalyst

Model launches continue to matter because model capability remains the core battlefield in the LLM sector. Pricing, distribution, and product packaging all influence commercialization, but the industry's competitive position still starts from capability. A new model cycle therefore carries significance well beyond benchmark headlines or user attention. It shapes how investors think about long-term monetization potential, short-term market-share trajectory, and whether a company deserves to sustain a premium valuation.

There are two dimensions investors should focus on when assessing a new model release:

- **The first is absolute capability improvement.** This refers to whether the new model meaningfully advances coding performance, reasoning quality, agent reliability, long-context handling, and multimodal capability relative to the company's own prior generation and relative to the global frontier. This dimension matters because the long-term revenue opportunity for listed LLM companies depends on whether they can continue participating in the highest-value workloads, rather than being pushed into lower-priced or narrower use cases. A model that improves commercialization while falling behind the global frontier would still face a lower long-duration valuation ceiling.
- **The second dimension is relative capability change within the domestic peer group.** This relative ranking matters because the Chinese LLM market is crowded, switching costs at the model layer remain manageable, and the market is likely to reward the companies that keep moving closer to the top of the domestic capability stack. Even if the whole sector continues to improve, share-price performance can diverge once one or two names begin to pull ahead on coding, agent performance, or enterprise-grade reliability.

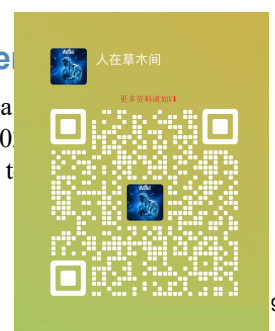
This is why model launches remain the most visible near-term catalyst. **A successful release can do three things at once: improve the company's position against the global frontier, strengthen its relative standing against domestic peers, and support stronger monetization through better usage growth and higher realized pricing.** That combination is what turns a technical event into an equity catalyst.

The bar has also become higher after the sharp post-IPO rerating. Earlier in the cycle, a new release could support the shares largely through narrative momentum, benchmark visibility, or renewed investor interest in the theme. At current valuations, the market is likely to apply a more demanding standard. Investors will still pay attention to launch headlines, though the more important question is whether the new model improves the company's competitive position in a way that can translate into share gain, pricing power, and another step-up in recurring API spend.

For that reason, **upcoming model cycles should be evaluated less as isolated product events and more as competitive checkpoints.** In a sector where model capability remains the real battlefield, each release provides new evidence on which companies are moving up the stack, which companies are holding ground, and which companies are starting to lose relative position. That assessment is likely to play a central role in the next phase of stock differentiation across China's listed LLM names.

Token consumption remains supportive, structural acceleration under


The strongest near-term support for China's LLM sector still comes from the demand curve: China usage exceeded 140tn in March 2026, up from 100bn at the beginning of 2024 and 100tn at end-2023 (Daily), which remains an unusually steep demand acceleration and lays the foundation for market t ARR assumptions for listed frontier-model names.



OpenRouter data offers a useful but incomplete read-through. The platform is directionally useful because it is highly sensitive to developer behavior, coding activity, and agent experimentation, all of which tend to lead broader commercialization trends. At the same time, OpenRouter should not be treated as a complete proxy for industry demand, because it underrepresents direct enterprise API traffic and model usage flowing through large cloud platforms.

By April, token growth on its own has less incremental value in the valuation debate than it did earlier in the cycle. Investors have already accepted that demand is scaling quickly. The larger question now is how much of that demand is being converted into recognized revenue, and at what realized price. That leaves first-party monetization signals, product pricing, and model-level ARR step-ups carrying more weight than raw volume headlines alone. Anthropic's early-April policy change around third-party harnesses illustrates the point: usage that had previously sat inside subscription coverage is now moving toward API billing or separate prepaid spend, which can lower visible third-party volume while improving monetization quality.

Figure 6: Global token consumption captured by OpenRouter



Source: OpenRouter (data as of Apr 13, 2026)

Pricing is becoming an increasingly important earnings driver, with several mechanisms now moving in the same direction

Pricing has become one of the more important changes in the sector year to date. Through much of 2024 and 2025, the concern was that rapid demand growth would be competed away through aggressive API cuts. The 2026 setup looks different: pricing is improving through multiple channels, which is supporting revenue conversion across the model layer.

- Mechanism 1: direct API repricing around model upgrades.** As model quality improves and demand remains strong, LLM providers have become more willing to charge meaningfully higher prices for newer models. Zhipu, specifically, has its model API pricing moving materially higher through the GLM-4.7 to GLM-5 to GLM-5.1 cycle, with effective token price almost doubled year-to-date. The significance of such price hike is not just the higher list price; it is that demand appears to have remained strong alongside higher pricing – suggesting a much firmer pricing environment.
- Mechanism 2: subscription seeding followed by migration toward higher-value usage.** Deeply discounted plans help reduce adoption friction, build developer habits, and encourage users to integrate model usage into their daily work. Once those workflows become more embedded and token intensity rises, the commercial logic shifts toward higher-tier subscription products or pay-as-you-go billing. Effective Mar 20, the Alibaba Cloud coding plan's Lite tier stopped accepting new subscriptions, leaving the Pro tier as the main active entry point. Zhipu AI

uses a similar separation: Coding Plan quota can only be used in supported coding tools, API calls are billed separately; both Alibaba and Zhipu AI are encouraging users to adopt a “pay as you go” mechanism. That product design matters because discounted or quota-based plans can reduce adoption friction, while the larger revenue opportunity remains tied to heavier usage and separately billed API consumption once customers move beyond bundled workflows. For investors, this mechanism matters because it tends to lift revenue quality over time. Entry-level plans can be used to acquire usage, but the larger economic value sits in customers whose consumption expands beyond the original bundle.

- Mechanism 3: structural repricing through KV cache economics, affecting realized monetization even when headline input/output rates appear unchanged.** MiniMax’s M2.7 versus M2.5 demonstrates a good example: the standard input and output pricing does not fully capture the change in monetization. M2.5 and M2.7 carry the same headline input/output pricing at \$0.3 / \$1.2 per million tokens, while prompt-caching read pricing rises from \$0.03 per million tokens on M2.5 to \$0.06 on M2.7. In coding and agent workloads, where context is reused repeatedly, billing is influenced heavily by the cached-token component. Under those conditions, higher cache-read pricing can raise the effective realized price even when the headline model comparison looks stable – we estimate that the effective price on M2.7 is 30+% higher than M2.5.

Table 3: Three mechanisms to improve token/API price

Mechanism	Example	Effect on realized pricing	Earnings implication
Direct API repricing	Zhipu GLM-5.1 versus GLM-5	Visible uplift in official input/output pricing	Supports higher ASP if demand remains firm
Migration from coding plan to “pay as you go”	Alibaba Coding Plan; Z.ai GLM Coding Plan	Lower-friction entry, separate monetization for heavier API usage	Helps acquisition while preserving upside to higher-value consumption
KV-cache repricing	MiniMax M2.7 versus M2.5	Cache-read price doubles while headline input/output rates stay flat	Can raise effective price materially in coding and agent workloads

Source: Company official website, OpenRouter, J.P. Morgan estimates (data as of Apr 20, 2026)

Fund Flow & Technical Dynamics: lock-up expiration, Stock Connect/Index inclusion, potential placement and eroding 'scarcity premium' on peers' IPO pipeline

Table 4: Summary of key fund flow & technical dynamics for Zhipu AI and MiniMax

Topic	Zhipu	MiniMax	Equity implication
Public float at IPO	3.90%	5.40%	Both names started with tight float
6-month unlock	Early Jul-26: 5.8%	Early Jul-26: 39.0%	MiniMax faces the larger near-term supply event
9-month unlock	—	Around Oct: 18.2% conditional unlock	Adds a second 2H26 supply event for MiniMax
12-month unlock	Jan-27: 90.3%	Jan-27: 37.4%	Zhipu's larger supply event is later
Hang Seng Composite / Hang Seng TECH review	22 May 2026 results; 8 Jun 2026 effective	22 May 2026 results; 8 Jun 2026 effective	First likely index catalyst for both names
Expected Stock Connect inclusion	First week of Jun-26	First week of Aug-26	Zhipu likely accesses mainland demand earlier; MiniMax delayed by WVR seasoning
Potential placement	Likely over time	Likely over time	Funding remains part of the sector base case

Source: Company prospectus (data as of Dec31, 2025), Hang Seng Index website, Reuters, Sina Finance, AASTocks, (data as of Apr 20, 2026) and J.P. Morgan estimates. (data as of Apr 20, 2026)

-ve: Lock-up expiry: pre-positioning typically begins two months ahead of the unlock, though relative-value support may limit de-rating

Lock-up expiry will become the first meaningful technical headwind in 2H26, with MiniMax facing the heavier near-term supply event. Based on the lock-up timetable calculations, Zhipu's schedule is relatively back-end loaded: 3.9% of post-IPO shares were in free float at listing, another 5.8% becomes tradable in early July 2026, and the remaining 90.3% unlocks in January 2027. MiniMax's schedule is more front-loaded: 5.4% was in free float at listing, 39.0% becomes tradable in early July 2026, a further 18.2% may unlock around October 2026 depending on Stock Connect inclusion timeline, and the final 37.4% unlocks in January 2027. In practical terms, Zhipu faces a modest six-month unlock followed by a much larger January event, while MiniMax is exposed to a larger increase in tradable supply already in 2H26.

For Zhipu and MiniMax, the setup is more nuanced because that **technical window will overlap with the next model cycle around June**. This means share-price performance into the unlock is unlikely to be driven by flow alone. Strong model releases that improve capability, relative positioning, and monetization could offset part of the pre-unlock pressure, while a less decisive model cycle would leave the stocks more exposed to the market's usual tendency to front-run supply expansion.

MiniMax's larger unlock points to heavier near-term technical pressure, though the downside may be partly cushioned by relative-value support because the market still tends to treat MiniMax and Zhipu as a paired China frontier-LLM trade. The scarcity of listed pure-play model names and the overlap in investor positioning mean MiniMax is still often assessed against Zhipu rather than on a fully standalone basis. As a result, if MiniMax were to underperform primarily because of lock-up-related flow pressure rather than weaker fundamentals, relative-value investors are likely to benchmark the stock increasingly against Zhipu and reassess whether any widening in the valuation gap is justified by operating differences. That dynamic could help cushion the downside if the gap moves materially beyond what differences in model capability, monetization, or ARR trajectory would support. It would not remove unlock-related volatility, though it does suggest that a technically driven de-rating may be less durable if the fundamental narrative remains broadly intact.

Table 5: Lock-up schedule summary for Zhipu AI

Lock-Up Expiry Date	Shareholder Category	Specific Shareholders	Number of Shares to be Unlocked	Percentage of Post-IPO Shares	Cumulative % of Post-IPO shares	Unlock Status & Notes
At IPO	Public free float		17,350,800	3.9%	3.9%	
July 7, 2026 (6 Months Post-IPO)	IPO Cornerstone Investors	11 cornerstone investors	25,681,600	5.8%		Mandatory lock-up expired, 100% tradable
Subtotal (6 Months)			25,681,600	5.8%	9.7%	
	Controlling Shareholders & Core Executives		133,044,230	29.8%		Mandatory lock-up expired, 100% tradable
January 7, 2027 (12 Months Post-IPO)	Pathfinder Senior Independent Investors (SIIs)	Legend Capital SIIs, Meituan SII	44,326,430	9.9%		Mandatory lock-up expired, 100% tradable
	All Remaining Pre-IPO Investors		225,440,030	50.6%		Mandatory lock-up expired, 100% tradable
Subtotal (12 Months)			402,810,690	90.3%	100.0%	
Total Unlocked Shares			445,843,090	100.0%		

Source: Company IPO filings (as of Dec31, 2025), J.P. Morgan estimates (data as of Apr 20, 2026)

Table 6: Lock-up schedule summary for MiniMax

Lock-Up Expiry Date	Shareholder Category	Specific Shareholders	Number of Shares to be Unlocked	Percentage of Post-IPO Shares	Cumulative % of Post-IPO shares	Unlock Status & Notes
At IPO	Public free float		17,073,200	5.4%	5.4%	
July 8, 2026 (6 Months Post-IPO)	Certain group of Pre-IPO Investors	Specified investors under Prospectus 48.92% pre-IPO shares	105,834,624	33.7%		Mandatory lock-up expired, 100% tradable
	IPO Cornerstone Investors	15 cornerstone investors	16,504,040	5.3%		Mandatory lock-up expired, 100% tradable
Subtotal (6 Months)			122,338,664	39.0%	44.5%	
Which is earlier of: (1) October 8, 2026 (9 Months Post-IPO) or (2) 20th trading day starting from the date on which the share is included as an eligible stock of Stock Connect						
	Certain group of Pre-IPO Investors	Specified investors under Prospectus 20.34% pre-IPO shares	56,963,811	18.2%		Lock-up expired, 100% tradable
Subtotal (9 Months)			56,963,811	18.2%	62.6%	
	Controlling Shareholders & Core Executives	Yan Junjie, Yun Yeyi & other key management	91,494,505	29.2%		Mandatory lock-up expired, 100% tradable
January 8, 2027 (12 Months Post-IPO)	Pathfinder Senior Independent Investors (SIIs)	miHoYo, IDG Capital	25,765,128	8.2%		Mandatory lock-up expired on July 8, 2026; voluntary lock-up expired on January 8, 2027
Subtotal (12 Months)			117,259,633	37.4%	100.0%	
Total Unlocked Shares			313,635,308	100.0%		

Source: Company IPO filings (as of Dec31, 2025), J.P. Morgan estimates (data as of Apr 20, 2026)

Compared to Kuaishou, SenseTime and Horizon Robotics, what is the read-across to Zhipu and MiniMax?

The stock price movements observed during the IPO lock-up expiries of Kuaishou (1024.HK), SenseTime (0020.HK), and Horizon Robotics (9660.HK) highlight the **impact of liquidity and the scale of unlocked shares**. Kuaishou experienced a sharp price drop following its lock-up expiry in 2021, with its stock falling from HKD 250 to HKD 120, although it later stabilized. Similarly, SenseTime saw a significant dip post-expiry but gradually recovered over the following months. Horizon Robotics, on the other hand, had a relatively smooth transition, with a modest decline in its share price before rebounding. We believe turnover plays a key role; Kuaishou, for example, with lock-up period average turnover of HK2+bn, had sufficient liquidity to absorb the increased float. These dynamics underscore that while the initial release of shares can lead to volatility, markets with higher liquidity tend to stabilize over time as trading volumes pick up.


In contrast, the situation for Zhipu AI (2513.HK) and MiniMax (0100.HK) appears more favorable. Zhipu AI's 6M release of 6% of total shares and MiniMax's 39% release are considerably smaller compared to the previous cases. Furthermore, both companies exhibit significantly higher daily turnover (~HKD 1.2-1.5bn), suggesting that the market has ample liquidity to absorb these share releases without substantial disruption. MiniMax, despite the larger proportion of shares unlocked, benefits from a turnover ratio that will likely mitigate any major price swings. Given these liquidity advantages and relatively moderate float expansion, we expect the lock-up expiries for Zhipu and MiniMax to have a limited impact on their stock prices, with the market effectively absorbing the new shares without significant volatility.

Table 7: Compared to Kuaishou, SenseTime and Horizon Robotics, what is the read-across to Zhipu and MiniMax?

Metric	Kuaishou (1024.HK)	SenseTime (0020.HK)	Horizon Robotics (9660.HK)	Zhipu AI (2513.HK)	MiniMax (0100.HK)	Note
Listing Date	2021-02-05	2021-12-30	2024-10-24	2026-01-08	2026-01-09	
6M Lock-up Expiry Date	2021-08-05	2022-06-30	2025-04-24	2026-07-08	2026-07-09	
Price Performance						
IPO Price (HKD)	115	3.85	3.99	116.2	165	
Price at 6M Expiry (HKD)	84.35	2.54	5.81	NA	NA	
6M Post Expiry Price (HKD)	89.06	2.18	8.88	NA	NA	
Change from IPO to 6M Expiry	-27%	-34%	46%	NA	NA	
6M Post Expiry Return	6%	-14%	53%	NA	NA	
Share Structure						
Total Shares Outstanding (mn)	4,100	33,282	13,030	446	314	
IPO Shares Issued	365	1,725	1,355	37	29	
IPO Issued as % of total	9%	5%	10%	8%	9%	
Public Shares at 6M (Float)	91	1,725	1,355	17	17	
Pre-6M Float (% of total)	2.2%	5.2%	10.4%	3.9%	5.4%	
Shares Released at 6M (mn shares)	3,882	23,300	9,540	26	122	
6M Release as % of total	95%	70%	73%	6%	39%	Zhipu & MiniMax are much lower than studied cases
Post-6M Float (mn shares)	3,973	25,025	10,895	43	139	
Turnover						
Pre-6M Avg Price (HKD)	235.9	6.0	5.4	580.0	720.0	
Total Market Cap Pre-6M (HKD bn)	967	199	71	259	226	
Float Market Cap Pre-6M (HKD bn)	21.5	10.3	7.3	10.1	12.3	
Daily Turnover Pre-6M (HKD bn)	2.1	0.3	0.4	1.5	1.2	Zhipu & MiniMax already have sizable trading value
Turnover/Float Ratio Pre-6M (%)	10%	3%	5%	15%	10%	
Float Expansion at 6M	43x	14x	7x	1x	7x	
6M Expiry Day Performance	-2%	-47%	-7%	NA	NA	
6M Expiry Day Volume (mn shares)	120	873	337	NA	NA	
Daily Turnover Post-6M (HKD bn)	2.0	0.4	1.8	NA	NA	
Daily Turnover Post-6M vs Pre-6M	1.0x	1.6x	5.2x	NA	NA	No meaningful changes in daily turnover at Kuaishou's case
Post-6M Float Market Cap (HKD bn)	342	50	85	NA	NA	
Post-6M Turnover Ratio (%)	1%	1%	2%	NA	NA	


Source: Company filings (data as of Feb 2021 for Kuaishou, Dec 2021 for SenseTime, Oct 2024 for Horizon Robotics, Dec 2025 for Zhipu and MiniMax), Bloomberg Finance L.P. Yahoo Finance, J.P. Morgan estimates (data as of Apr 20, 2026). Note: Past performance is not an indicator of future results

Figure 7: 1-year share price performance of Horizon Robotics post IPO



Source: Bloomberg Finance L.P. (data as of Apr 20, 2026). Note: Past performance is not an indicator of future results

Figure 8: 1-year share price performance of Kuaishou post IPO



Source: Bloomberg Finance L.P. (data as of Apr 20, 2026). Note: Past performance is not an indicator of future results

Figure 9: 1-year share price performance of SenseTime post IPO



Source: Bloomberg Finance L.P. (data as of Apr 20, 2026). Note: Past performance is not an indicator of future results

+ve: Stock Connect/Index inclusion should provide the meaningful source of flow support

Hang Seng Composite & Hang Seng TECH inclusion should provide the first meaningful source of passive and index-related support. The Hang Seng Index has announced that 1Q26 review results will be released on 22 May 2026 and become effective on 8 June 2026. Under the published methodology, newly listed first-quarter names that satisfy the relevant size and liquidity criteria may be added at the June rebalance. On that basis, both Zhipu and MiniMax are candidates for Hang Seng Composite inclusion effective 8 June 2026, subject to the usual review outcome. **Stock Connect Inclusion / Southbound timing should diverge. We expect Zhipu to enter earlier (1st week of June) and MiniMax to follow later (1st week of Aug),** due to its weighted voting rights structure.

Neutral/-ve: Potential placement should be treated as part of the base-case funding framework rather than as a tail risk

Potential placement should be treated as part of the base-case funding framework rather than as a tail risk. For both Zhipu AI and MiniMax, we see post-IPO financing within 12 months are near-certain given R&D spend exceeds revenue by 4-5x. The timing of any future placement will matter at least as much as the fact of the placement itself. A financing executed after another visible step-up in ARR and model capability would be easier for the market to absorb, particularly if accompanied by index-related and Southbound flows. A financing attempted into a large unlock window, or against softer operating momentum, would carry greater valuation risk because it would add supply at a point when the market is already digesting a broader free-float expansion.

The near-term technical setup therefore looks cleaner for Zhipu than for MiniMax. Zhipu benefits from a smaller six-month unlock and an earlier path into Southbound demand, while MiniMax faces a denser 2H26 flow calendar because of its larger July unlock, later Southbound access, and another meaningful potential supply event around October. For both stocks, technical support from Hang Seng Composite & TECH inclusion remains relevant, though the supply calendar is becoming less forgiving than it was immediately after IPO.

-ve: New listings pipeline — Cambricon case study indicates 25-30% ‘scarcity premium’


Zhipu and MiniMax are currently the only listed pure-play frontier AI model companies globally, and such scarcity commands a premium, in our view. The Kimi and StepFun listing pipeline are likely to dilute it — and the Chinese GPU sector provides a directly relevant, quantifiable precedent.

Cambricon (688256.SH), the sole listed AI chip pure-play since 2020, surged to ~RMB 1,500/share by late November 2025 (market cap ~Rmb630bn). During Dec 2025, three additional domestic GPU players listed: Moore Threads (listed on Dec 5, 2025), MetaX Integrated Circuits (listed on Dec 17, 2025) and Biren Technology (listed on Jan 2, 2026) — and the combined sector market cap exceeded Rmb1.1tn — **the pie expanded significantly.**

However, with more peers becoming investor-accessible on the public market, **the Cambricon’s ~25-30% ‘scarcity premium’ is gradually digested.** By combining the fundamental changes on Cambricon, here we quantify the ‘scarcity premium’ on the name: on a year-to-date basis, even with 12m forward revenue / earnings are raised by 17% / 26%, Cambricon’s share price is still down 2% versus the SHSZ300 index +3% — the valuation multiple has fallen by 25-30%.


Implications for Zhipu and MiniMax: new LLM listings (Kimi, StepFun) will likely be **net positive for sector enthusiasm short-term**, as the IPO event generates capital inflows. However, the scarcity premium per name declines structurally. The window during which Zhipu and MiniMax benefit from being “the only way to play China AI” is likely 6-12 months.

Figure 10: Cambricon's share price performance & 12-month forward revenue forecasts during the period of peers' listing at end-2025



Source: Bloomberg Finance L.P., J.P. Morgan estimates (data as of Apr 20, 2026). Note: Past performance is not an indicator of future results

Figure 11: Cambricon's share price performance & SHSZ300 index price during the period of peers' listing at end-2025



Source: Bloomberg Finance L.P., J.P. Morgan estimates (data as of Apr 20, 2026). Note: Past performance is not an indicator of future results

US Frontier AI: Implications for Chinese LLM Valuations

Accelerating US frontier ARR: structurally bullish for Chinese LLM valuations

On April 6, Anthropic's run-rate revenue surpassed \$30bn, up from \$9bn at end-2025 — a 3.3x increase in ~3.5 months, with over 1,000 enterprise customers spending above \$1mn annually (source: Bloomberg), while OpenAI is on track for ~\$29bn in 2026 revenue (source: Bloomberg). The combined US frontier ARR run-rate now exceeds \$59bn.

Read-across to Chinese LLMs:

- **TAM validation.** If a single US player can sustain 10x+ annual growth at \$30bn ARR with no deceleration, the global LLM TAM is structurally larger than what consensus assumes. China's token consumption (100bn daily in early 2024 → 140tn by March 2026) mirrors this with a lag. Zhipu's \$1bn year-end target may prove conservative if China's adoption curve follows even a fraction of the US trajectory.
- **Relative discount deepens.** At constant P/ARR, the combined US valuation could reach \$1.5+tn by year-end based on Anthropic and OpenAI's latest ARR trajectory. Zhipu + MiniMax at \$80bn would represent ~5% of US peers' valuation — down from ~10% today. Chinese LLM stocks mechanically become cheaper on a relative basis even with no change in their own fundamentals.

The fundamental risk is not the valuation gap itself — it is whether Chinese independents can capture enough domestic enterprise share to participate in the TAM expansion — especially considering the aggressive B-end focus among Alibaba, ByteDance and Tencent. The gap between China's aggregate token growth and Zhipu/MiniMax's individual share is where investment risk concentrates.

This is why US frontier LLM revenue acceleration should be seen as supportive, though not sufficient on its own. It helps validate the size of the market and keeps investor appetite for the theme intact. The next step still depends on whether Chinese listed names can hold their technological position, convert demand into recognized revenue, and defend share in a market where the competitive field is broader and the ecosystem advantages of larger platforms remain meaningful.

Anthropic & OpenAI IPO: the valuation anchor event – matter as much for valuation methodology as for valuation level

According to news, both Anthropic and OpenAI are targeting 2H26 IPOs, where news reported Anthropic recently draws VC interest at valuation as high as \$800bn on its recently reported \$30bn ARR (*Reuters, Apr 2026*), and OpenAI is eyeing Nasdaq at \$850bn-\$1tn (*Reuters, Oct 2025*).

So far, most discussion around frontier AI valuation has centered on price-to-ARR. That has been a practical framework in a sector where revenue growth is exceptionally fast, business models are still taking shape, and near-term profitability is being suppressed by heavy investment. Private investors have generally been willing to underwrite several years of expansion before demanding conventional earnings discipline. Public markets may place the emphasis differently. Once Anthropic or OpenAI begins trading, investor focus could broaden toward gross profit conversion, operating leverage, free cash flow potential, capital intensity, and revenue quality, instead of relying predominantly on ARR multiples.

That distinction carries valuation implications. Price-to-ARR tends to favor companies that are still early in monetization but showing rapid growth. **A framework that places greater weight on profitability or free cash flow would likely reduce the premium investors are willing to pay for revenue growth alone**, especially if compute, traffic acquisition, and model development costs remain elevated. In that setting, public investors may start drawing a sharper line between model companies with strong top-line momentum and AI businesses with a clearer path toward durable earnings power.

For Chinese LLM companies, the read-across could be meaningful. If Anthropic and OpenAI list successfully and continue to be valued mainly on forward revenue, that would reinforce the current benchmarking approach for Zhipu and MiniMax.

If US public markets begin to give more weight to gross margin, EBIT potential, or free cash flow generation, investors would likely revisit how much of Chinese listed LLM valuations can still be anchored by ARR acceleration alone. This would matter most for companies whose current equity stories still rely primarily on growth and market positioning, while medium-term profitability remains less visible.

The range of outcomes remains fairly wide. A successful IPO accompanied by strong revenue growth and improving margin visibility would broaden the sector's valuation framework and help investors underwrite long-duration earnings power with greater confidence. A weaker aftermarket performance, especially if public investors focus more on cash burn, compute intensity, or limited near-term free cash flow, would put more pressure on revenue-based multiples across the sector. In that scenario, Chinese LLM stocks could face multiple compression even if their own operating momentum remains solid.

Table 8: Compare US vs China LLM's valuation under different scenarios

Scenarios	US combined LLM valuation	China \$90bn as % of US
Latest private round (Anthropic's \$380bn & OpenAI's \$850bn)	~\$1.2tn	7.5%
News reported at Anthropic's \$800+bn valuation that draws VC interests	~\$1.6tn	5.6%
Post-IPO: if 50-100% upside after IPO	~\$2.4-3.2tn	2.8-3.8%

Source: Company official websites, WSJ, Reuters (data as of Apr 20, 2026), J.P. Morgan estimates. (data as of Apr 20, 2026)

AI value-chain positioning — China's listed LLM layer already carries meaningful weight, while the US comparison still reflects private-market marks

We estimate that China's investable AI stack is roughly divided into chips at ~US\$350bn, cloud and infrastructure at ~US\$200bn, and listed pure-play LLMs at ~US\$90bn. On that basis, the model layer already accounts for around 14% of the domestic investable AI universe. That proportion is not low, particularly when several important China model players, including DeepSeek and Moonshot, remain unlisted, while Qwen and Doubao sit inside much larger platform companies.

The US split appears lower in the same framework, with the model layer at around 8% of the AI value chain, versus roughly 51% for chips and 34% for cloud and infrastructure. That number, however, is still built largely on private-market valuations and reported funding marks. OpenAI and Anthropic are not yet listed, while Google DeepMind remains embedded within Alphabet, so the US model layer has not yet gone through the same public-market price discovery process as the chip and cloud layers.

On current numbers, China's listed model layer appears to take a larger share of its domestic AI value chain than the US model layer does in the comparable US framework. Part of that reflects **scarcity and investability** in China, where listed pure-play model exposure is limited. Part of it also reflects the **measurement gap in the US**, where the most important model assets are still referenced through private valuations rather than live secondary-market trading.

If US frontier model companies list and still see meaningful upside in the aftermarket, the model layer's share of the US AI value chain would move higher and the current gap versus China would narrow. In that setting, the difference between China's ~14% and the US's ~8% would look less like a structural divergence and more like a timing issue in the valuation cycle. China LLM valuations therefore need to be judged against where the US model layer may settle after IPO, rather than against today's static private-market snapshot alone.

Table 9: Value chain distribution of US AI supply chain

Company	Mkt Cap / valuation (approx.)	Category	Approx. % of value chain
Nvidia	~\$4.6T	GPU / AI Accelerator	
Broadcom	~\$1.8T	Custom Silicon / Networking	
AMD	~\$350B	GPU / CPU	
Marvell	~\$94B	Custom Silicon / Networking	
Micron	~\$425B	Memory (HBM)	
Intel	~\$90B	GPU / Foundry	
Chips total	~\$8T		51%
	Mkt Cap / valuation (approx.)	Cloud Mkt Share	Approx. % of value chain
Amazon (AWS)	~\$1.7–2.1T	~31–32%	
Microsoft (Azure)	~\$1.5–1.8T	~22–24%	
Alphabet (Google Cloud)	~\$1.1–1.4T	~11–12%	
Oracle	~\$300–400B (est.)	~3–4%	
CoreWeave	~\$40–50B (public mkt cap)	N/A (GPU-native niche)	
Cloud & Infra total	~\$5T		34%
	Mkt Cap / valuation (approx.)	Latest ARR	Approx. % of value chain
OpenAI	~\$852B (post-money Mar 2026) ~\$800B per Reuter's Apr15 2026 report;	~\$25B ARR (Feb 2026)	
Anthropic	~\$380B (Series G, Feb 2026)	~\$30B ARR (Apr 2026)	
Google DeepMind	Embedded in ~\$3.8T Alphabet	Embedded in GCP/Search	
LLM total (only counted OpenAI and Anthropic)	~\$1.6T		8%

Source: Company official websites, Bloomberg, WSJ, Reuters, J.P. Morgan estimates. (data as of Apr 20, 2026); market cap based on Apr 20 prices

Table 10: Value chain distribution of China AI supply chain

Company	Mkt Cap / valuation (approx.)	Category	Approx. % of value chain
Huawei / HiSilicon	N/A (unlisted)	GPU / AI Accelerator	
Cambricon	~\$83B	AI Inference ASIC	
Moore Threads	~\$43B	GPU (Training + Inference)	
MetaX Integrated Circuits	~\$36B	GPU (Training + Inference)	
Biren Technology	~\$12B	GPU (High-End Training)	
Chips total	~\$350B		27%
	Mkt Cap / valuation (approx.)	Parent Mkt Cap	
Alibaba Group	~\$90–110B (est.)	~\$300B	
ByteDance (unlisted)	Embedded in ~\$550B+ ByteDance	~\$550B+ (private est.)	
Tencent Holdings	~\$50–60B (est.)	~\$570B	
Baidu	~\$10–15B (est.)	~\$37B	
Huawei (unlisted)	Unlisted / Not tradeable	Unlisted	
Cloud & Infra total	~\$200B		25%
	Mkt Cap / valuation (approx.)	Latest ARR	Approx. % of value chain
Zhipu AI	~\$55B	~250mn ARR as of Mar31, 2026	
MiniMax	~\$38B	~150mn ARR in Feb 2026	
DeepSeek	Unlisted / Not tradeable		
Moonshot AI (Kimi)	Unlisted / Not tradeable		
Alibaba (Qwen / Tongyi)	Embedded in ~\$300B Alibaba		
ByteDance (Doubao / Seed)	Embedded in ~\$550B+ ByteDance		
LLM total (only counted Zhipu AI and MiniMax)	~\$90B		14%

Source: Company official websites, Bloomberg, Sina News, Reuters, J.P. Morgan estimates. (data as of Apr 20, 2026); market cap based on Apr 20 prices

Knowledge Atlas Technology Joint Stock Co. Ltd. (Overweight; Price Target: HK\$950.00)

Investment Thesis

We rate Knowledge Atlas (Zhipu AI) Overweight. Our investment view is anchored in a structural assessment of how value is created and sustained in the foundation-model industry. We believe that long-term economic outcomes are determined primarily by a company's ability to **remain within the global first tier of model capability across multiple technology cycles**, while business model form, deployment modality, and near-term margin structure are largely downstream expressions of that capability.

Zhipu has reached an important inflection point, especially for its global API business. The release of GLM-4.5/4.6/4.7 to GLM5, together with a visible strategic shift toward agentic systems, tool-augmented reasoning, and developer-facing infrastructure, indicates that the company is aligning its technical roadmap with the dimensions of capability that increasingly define the global frontier — particularly production-grade coding, long-context reasoning, and multi-step execution stability. These attributes are commercially important because they enable genuine workflow substitution rather than marginal assistance, unlocking materially larger enterprise and developer budgets.

The company's business architecture reinforces this trajectory. Zhipu has established a substantial on-premise deployment base in regulated Chinese industries, which we view as a **structurally durable demand pool**. As foundation models iterate, this installed base has the potential to evolve into upgrade-driven, recurring economics. In parallel, cloud-based APIs represent a scalable growth engine, and we expect adoption to accelerate as GLM gains recognition among developers, particularly in coding-centric workflows where willingness to pay and usage intensity are higher.

Valuation

Our Dec-26 PT of HK\$950 is derived from 30x 2030E P/E, when Zhipu AI will generate normalized earnings, and discounted back using a 15% WACC. The 30x multiple represents a valuation premium over tier-1 Chinese internet players, mostly to reflect the company's 100+% revenue CAGR in 2026-30E.

Price target calculation

2030E Revenue (Rmb mn)	98,832
2030E Adjusted net profit (Rmb mn)	20,360
2030E Adjusted EPS (Rmb/share)	43
Target multiple	30
WACC	15%
Dec-26 PT (HK\$)	950

Source: J.P. Morgan estimates.

Risks to Rating and Price Target

Risks include export controls, geopolitical risk, and entity list designation; intensified competition; high and ongoing R&D investment creates execution risk and pressure on profitability; commercialization and customer adoption remain subject to uncertainty; dependence on computing infrastructure and external suppliers exposes Zhipu to cost and availability risks.

MiniMax Group Inc - H (Overweight; Price Target: HK\$1,100.00)

Investment Thesis

We rate MiniMax OW, for its rare combination of technical strength, multi-modal commercialization potential, and global scalability in the AI foundation model space. The company has established a strong technical track record, with its models delivering strong performance across key benchmarks, while its full-spectrum portfolio lays the groundwork for dual B2B/B2C monetization. Its global footprint further enhances scaling potential and profitability. Our investment theses lie in:

- Model capability will increasingly dominate competitive outcomes in the AI model landscape, where MiniMax has delivered a strong technical track record. MiniMax's models deliver leading performance in global benchmark, its models of text, speech and video generation all ranked among tops globally. The group has achieved high ROI on R&D, cumulatively \$450mn R&D investment as of Sep 2025, supporting bi-monthly model iterations. MiniMax's historical behavior — characterized by consistent investment in model training, regular generation upgrades, and early commitment to multimodality — supports the case that it is well-positioned to participate in this AI competition dynamic.
- Full-spectrum models enabling dual B2B/B2C commercialization. The coexistence of text, video, and audio models enables MiniMax to address enterprise and consumer use cases that are increasingly multimodal by design, reducing dependence on any single monetization path.
- Global orientation fueling scaling and profitability. MiniMax's early global orientation, visible in both customer targeting and infrastructure strategy, introduces economic flexibility that is uncommon among peers at a comparable stage.

Valuation

Our Dec-26 PT of HK\$1,100 is derived from 30x 2030E P/E, when MiniMax will generate normalized earnings, and discounted back using 15% WACC. The 30x multiple represents a valuation premium over tier-1 Chinese internet players, mostly to reflect the company's 100+% revenue CAGR in 2026-30E.

Price target calculation

2030E Revenue (US\$mn)	9,136
2030E Adjusted net profit (US\$mn)	2,322
2030E Adjusted EPS (\$/share)	7.2
Target multiple	30
WACC	15%
Dec-26 PT (HK\$)	1,100

Source: J.P. Morgan estimates.

Risks to Rating and Price Target

Risks include MiniMax's litigation proceedings with US studios; intensified competition; high and ongoing R&D investment creates execution risk and pressure on profitability; commercialization and customer adoption remain subject to uncertainty; dependence on computing infrastructure and external suppliers exposes MiniMax to cost and availability risks.