Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

ASIA PACIFIC INSIGHT

September 14, 2026 09:33 PM GMT

China's AI Path

Consumer AI: Beyond the Chatbot Race

China leads global 2C AI adoption thanks to strong distribution and clear use cases. While standalone AI-native apps capture early mindshare, in-app AI offers better ROI. We see AI accelerating existing mobile-internet monetization and estimate a Rmb300bn China 2C AI monetizable TAM by 2030.

AlphaWise α

China is ahead at the consumer-adoption layer, driven by distribution and use-case fit. Our AlphaWise survey shows 80% weekly personal AI usage in China vs. 54% in the US. Basic access is free in both markets, but China's advantage is that free access is coupled with ecosystem distribution, CNY campaigns and a broader range of everyday use cases. AI is riding an already mature mobile-internet layer and is increasingly embedded into existing apps.

AI-native apps win early consumer mindshare, but in-app AI has better ROI. ByteDance's Doubao remains the clear AI-native app leader, at ~400mn MAU, but our AlphaWise survey also shows that multi-homing is deepening rather than fading. We continue to prefer in-app assistants over AI-native apps, from an economic perspective, because in-app assistants can monetize existing traffic, proprietary data, permissions, merchant supply and payment/fulfillment without rebuilding user habits from scratch. The next competitive metric should increasingly move from MAU to task completion, AI-attributed conversion, and GMV.

We estimate ~Rmb300bn of China 2C AI monetizable TAM by 2030. Unlike traditional internet platforms, which benefit from high fixed costs and low marginal costs, AI products face a structurally different cost curve – every additional query consumes inference compute. We view consumer AI primarily as an accelerant for existing mobile-internet revenue rather than a standalone revenue pool. We size China's 2C AI monetizable TAM of ~Rmb300bn by 2030, with 99% coming from merchant-side intent monetization.

Stock implications – favor distribution and vertical execution; stay disciplined on standalone-app spend. We favor Tencent (WeChat moat looks more likely to be enhanced than displaced by Doubao) and vertical leaders (Meituan, Taobao, OTA, etc.) whose existing products already own demand, data and execution, where incremental acquisition cost is low and conversion can be measured. BABA remains our top pick given its full-stack AI capabilities, but our view specifically on its consumer AI initiatives is more balanced: Taobao/Qwen Shopping Assistant and the broader commerce stack are strategically attractive, while the economics of heavy standalone Qwen-app investment warrant continued discipline. We are more cautious on Baidu, TME and iQIYI because AI can disrupt their core business – search, music and video content creation.

MORGAN STANLEY ASIA LIMITED+

Gary Yu
Equity Analyst
Gary.Yu@morganstanley.com +852 2848-6918

Kathy Zhu
Research Associate
Kathy.Zhu@morganstanley.com +852 3963-2618

Eddy Wang, CFA
Equity Analyst
Eddy.Wang@morganstanley.com +852 2239-7339

Rebecca Xu
Equity Analyst
Rebecca.Xu@morganstanley.com +852 2848-7359

Joanne Lau
Research Associate
Joanne.CY.Lau@morganstanley.com +852 3963-1592

Andy Meng, CFA
Equity Analyst
Andy.Meng@morganstanley.com +852 2239-7689

Yang Liu
Equity Analyst
Yang.Liu@morganstanley.com +852 2239-1911

Tom Tang
Equity Analyst
Tom.Tang@morganstanley.com +852 3963-1860

Lydia Lin
Equity Analyst
Lydia.Lin@morganstanley.com +852 2239-1572

CHINA INTERNET AND OTHER SERVICES
Asia Pacific
Industry View **Attractive**

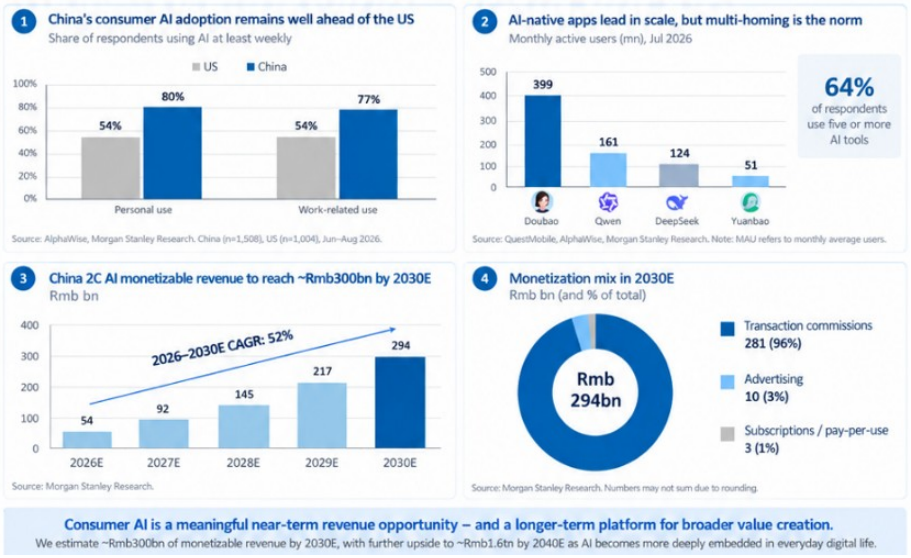
Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only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

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

+= Analysts employed by non-U.S. affiliates are not registered with FINRA, may not be associated persons of the member and may not be subject to FINRA restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account.

Our Thesis in Charts

Exhibit 1: Our Thesis in Charts



Source: Morgan Stanley Research Estimates.

Exhibit 2: Tencent/WeChat ecosystem



Source: Morgan Stanley Research. Tencent. 1. Since 4Q19, 2. In 2024, 3.In 4Q25.

Executive Summary



What the market had assumed…

- AI-native apps will become the dominant consumer AI entry point.
- Consumer AI monetization will mainly resemble US freemium software: paid subscriptions first, with advertising added later.
- The best foundation model should determine the consumer winner.
- AI-native apps are the main disruption threat to incumbents apps.

What our analysis reveals instead…

- **Distribution and use-case fit explain more of China's consumer-adoption lead than frontier-model superiority.** Basic access is free in both China and the US; China has combined it with stronger ecosystem distribution, CNY promotions and a broader set of daily-use scenarios. Our AlphaWise survey shows weekly personal AI use at 80% in China vs. 54% in the US. China personal AI daily-use scenarios remain concentrated in information, creation, decision support and, increasingly, commerce-related tasks.
- **AI has also diffused faster than mobile internet did at a comparable stage.** China's generative-AI user base rose from ~249mn in Dec-2024 to ~602mn in Dec-2025, consistent with AI scaling as a software layer on top of an already mature mobile infrastructure.

- **AI-native apps are efficient at user education and top-of-funnel intent capture, but multi-homing is accelerating rather than fading.** QuestMobile July data show Doubao at ~399mn monthly average users (MAU), Qwen ~161mn, DeepSeek ~124mn and Yuanbao ~51mn; 64% of our AlphaWise survey respondents used five or more tools, which keeps switching costs low and limits winner-take-all economics.
- **In-app AI assistants have the clearest near-term ROI because they start with existing traffic, data and execution rights. Tencent illustrates the distinction**. Our AlphaWise survey shows 83% past-month usage of in-app AI; 75% of respondents find AI search/chatbot features in shopping apps useful, and 57% say these features would make them purchase more often. Tencent illustrates the distinction: Yuanbao remains subscale as a standalone app, while WeChat's social graph, Mini Programs and payments give AI a much stronger path to task completion and monetization inside the core app.
- **We view consumer AI as an accelerator to existing internet monetization.** We estimate ~Rmb300bn of China 2C AI monetizable revenue TAM by 2030E, with ~99% coming from merchant-funded intent monetization through advertising and AI-attributed transaction commissions. Because much of that revenue already exists today, the incremental upside comes from better conversion, monetization and share gains. In-app AI can still earn attractive returns because it leverages an existing user base and distribution.
- **The read-through is therefore more favorable for companies that can embed AI**

into existing commerce and workflows rather than rely on heavy investment in standalone AI-native apps.

- **META's Muse launch provides a live US test of our thesis.** Muse is a free standalone agent spanning shopping, travel, tickets, calendar and email, but its potential advantage depends on integrating Meta's distribution/data and broader service ecosystem. At launch, it is separate from Meta's core apps and does not share data with the ad system, so the key proof points are task completion, conversion, monetizable behavior and ecosystem integration rather than MAU alone (see more in Meta Platforms Inc: Muse Has Arrived; What Matters Next? [9 Sep 2026]).

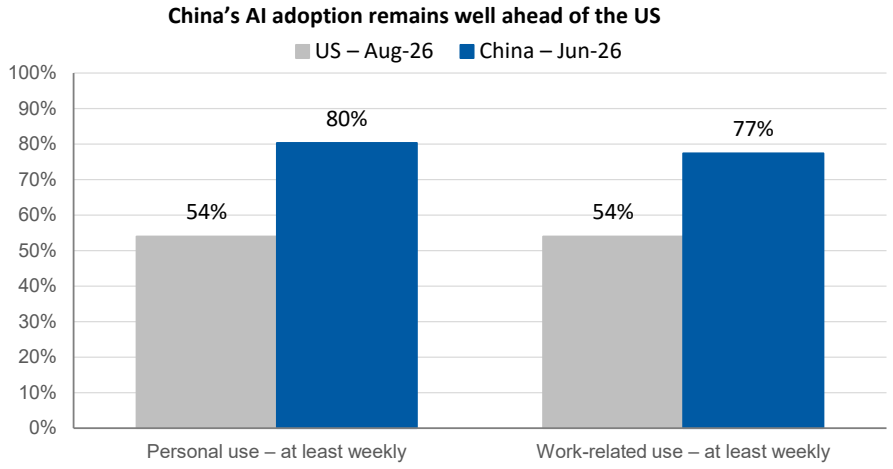
1. China has moved from user education to mass adoption faster than the US because AI is a software layer on top of mature mobile infrastructure.

China's consumer AI penetration is already mainstream. Our latest AlphaWise survey shows China's weekly personal AI usage rising to 80% from 75% in two months, and weekly work-related usage to 77% from 74% in three months. The latest US wave is ~54% for work and personal AI usage. The comparison is not perfectly like-for-like across market structure and survey timing, but the direction is clear: ***China has diffused consumer AI faster despite the US retaining an advantage in frontier compute and several frontier models.***

We believe five factors continue to explain the gap:

- China has stronger consumer distribution channels because AI is being inserted into super apps and vertical apps rather than requiring a new download.
- Free access lowers trial friction in both China and the US. China gains more from it because free AI is pushed through large ecosystems, reinforced by CNY campaigns and attached to a wider set of everyday use cases.
- Many mainstream use cases are low-to-mid token intensity.
- China's platforms have stronger commercial closed loops.
- Daily consumer use cases often require a model that is good enough, fast and accessible, rather than the absolute frontier model.

Exhibit 3: China's AI adoptions remain well ahead of the US



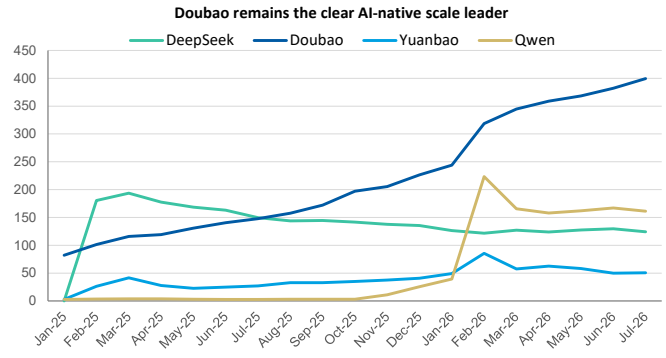
Source: AlphaWise, Morgan Stanley Research. China fieldwork Jun 30–Jul 2, 2026; US Aug 2026 wave.

2. The entry-point race is broadening; AI-native apps win mindshare, but in-app AI wins economics

AI-native apps had been efficient at user education...

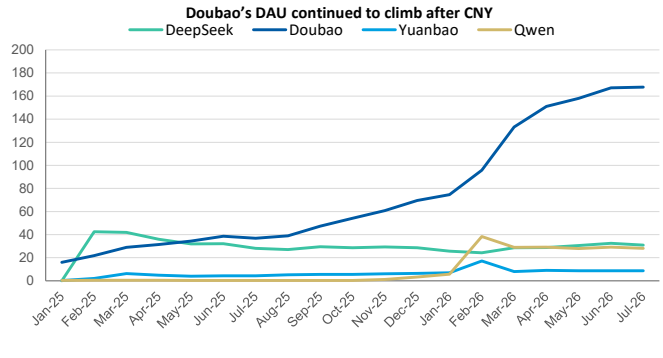
QuestMobile continues to show Doubao as the clear AI-native scale leader. In July, Doubao reached 399mn MAU and 168mn daily average users (DAU), vs. Qwen at 161mn/28mn, DeepSeek at 124mn/31mn and Yuanbao at 51mn/9mn. Doubao also led on DAU/MAU at 42%, while DeepSeek also remained relatively sticky, at 25%.

Exhibit 4: MAU (mn)



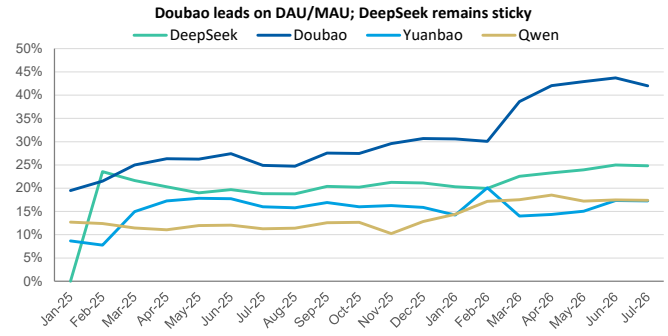
Source: Questmobile, Morgan Stanley Research

Exhibit 5: DAU (mn)



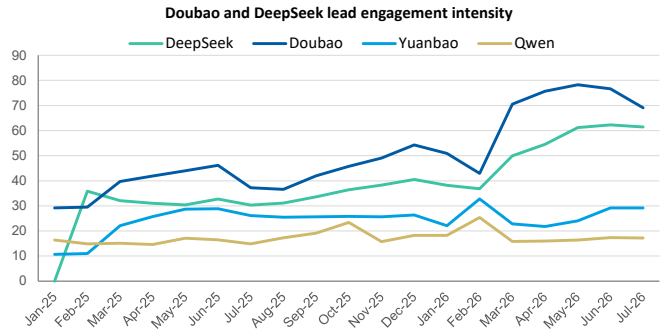
Source: Questmobile, Morgan Stanley Research

Exhibit 6: DAU/MAU ratio



Source: Questmobile, Morgan Stanley Research

Exhibit 7: Monthly sessions per user

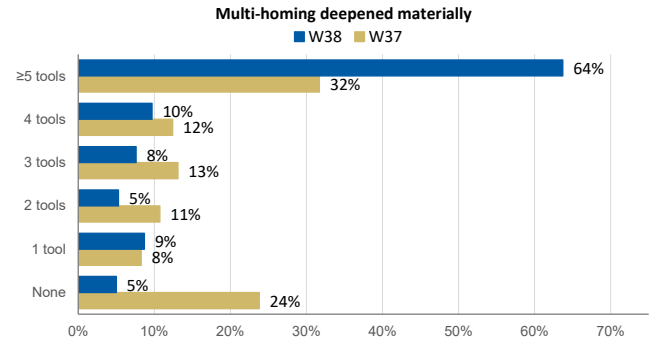


Source: Questmobile, Morgan Stanley Research

...but the entry-point race is broadening

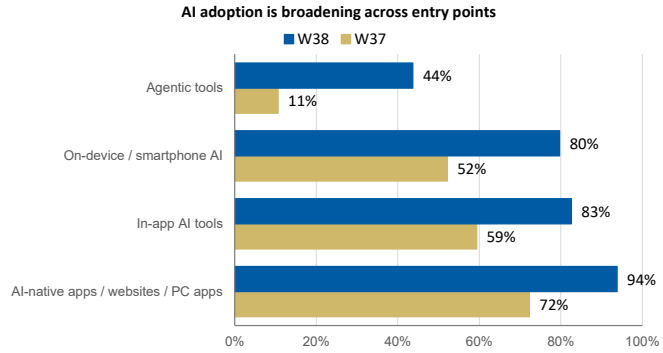
Our AlphaWise survey strengthens the multi-entry-point thesis. Past-month usage penetration reached ~94% for AI-native apps, ~83% for in-app AI, and ~80% for on-device AI, while agentic tools rose to ~44%. The average consumer is therefore not choosing a single AI gateway. The more important question is: which surface can convert intent into a completed task with acceptable customer-acquisition and inference costs?

Exhibit 8: Multi-homing deepened materially: Number of AI tools used in the past month, among total respondents



Source: AlphaWise, Morgan Stanley Research. W38 is conducted in June 2026, W37 is conducted in March 2026

Exhibit 9: AI adoption is broadening across entry points: Past-month use among total respondents

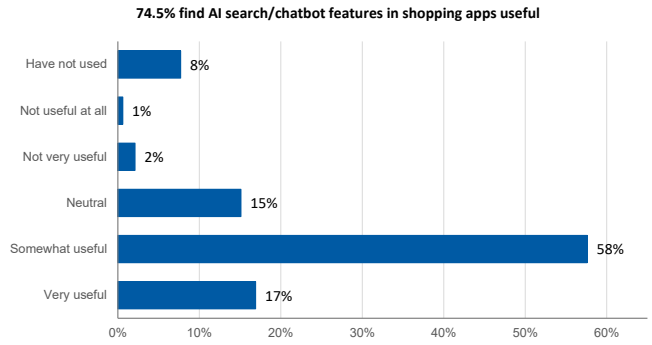


Source: AlphaWise, Morgan Stanley Research. W38 is conducted in June 2026, W37 is conducted in March 2026

...and we believe in-app AI has better ROI

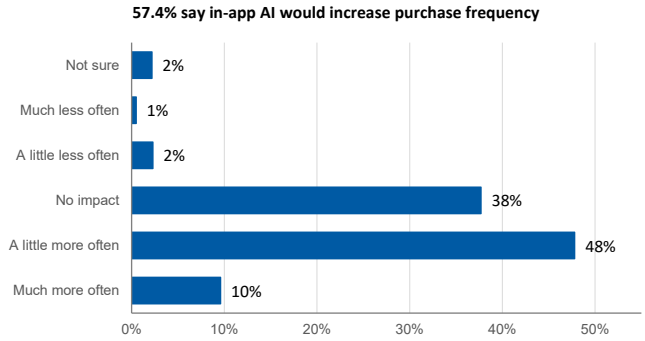
Our AlphaWise survey asked whether AI inside a vertical app changes commercial behavior – 75% of respondents find AI search/chatbot features in shopping apps very or somewhat useful, and 57% say such features would make them purchase much or a little more often. That said, AI creates more economic value when it improves an existing high-frequency workflow, than when it creates a standalone funnel that must be acquired and monetized from scratch.

Exhibit 10: Most consumers find shopping-app AI useful - among total respondents



Source: AlphaWise, Morgan Stanley Research

Exhibit 11: In-app AI can increase purchase frequency



Source: AlphaWise, Morgan Stanley Research

Recent company developments validate the in-app thesis more than the standalone-app thesis.

- **Tencent:** Tencent disclosed a small-scale Xiaowei prototype inside WeChat, powered by a customized WeLM model for privacy, WeChat-specific use cases, and inference efficiency. The agent can use WeChat content, Mini Programs, merchant reach and payments; management explicitly described a future agent-to-agent transaction loop. We view this as a stronger strategic pathway than trying to make Yuanbao a standalone super-app.
- **Alibaba:** Qwen Shopping Assistant launched inside Taobao in May and was described as providing end-to-end assistance from inspiration through after-sales. Management said losses narrowed sequentially as Qwen marketing expenses fell. This makes the ROI distinction between ecosystem AI and standalone-app spend unusually visible.
- **Meituan:** 2Q26 earnings call emphasized moving Xiaotuan from 'answering' to 'getting things done,' with order placement, ride-hailing and reservations, while LongCat 2.0 and merchant/operational agents provide model and workflow support. The direction is consistent with our thesis that vertical platforms should focus on real workflows rather than token volume as an end in itself.

3. Monetization – Freemium at the User Layer, Intent Monetization at the Commercial Layer

Doubao's introduced professional tier formalized a three-step paid structure at Rmb68/200/500 per month while maintaining a free base tier. Qwen introduced paid office-assistant tiers at Rmb19/49/128 per month in August. In both cases, the paid proposition centers on heavier productivity or generation workloads rather than basic chat. This is consistent with our original unit-economics framework: charge the users and tasks that consume more inference, while keeping broad consumer entry points free enough to preserve scale.

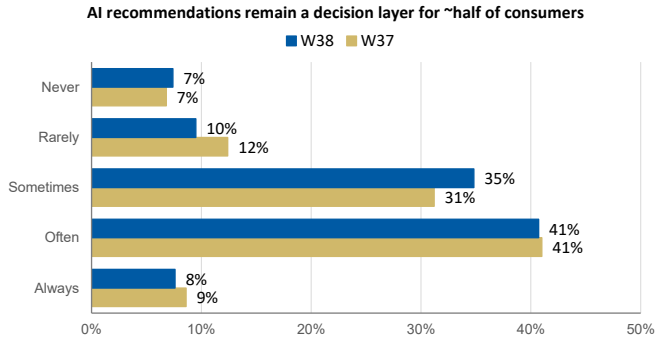
About 64% of our respondents either do not mind ads in a free AI service or can tolerate some ads if an ad-free option exists; only ~17% say they would refuse AI apps that show ads. Ad tolerance has softened slightly in the past three months, but remains high in

absolute terms.

We therefore expect basic AI to stay free for mass-market users, with subscriptions reserved for high-compute features and the larger monetization opportunity coming from sponsored recommendations, merchant advertising/lead generation and transaction commissions tied to commercial intent.

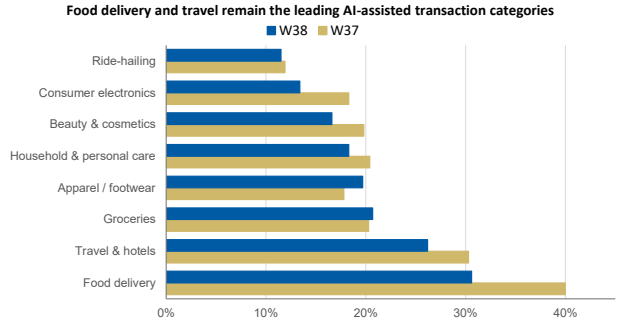
By 2030E, this framework implies ~Rmb294bn (~Rmb300bn) of China 2C AI monetizable revenue TAM, with merchant-funded intent monetization, including advertising and AI-attributed transaction commissions, contributing ~99%. As AI penetrates more consumer workflows, the same framework expands to ~Rmb1.6tn by 2040E. Incremental value continues to come from faster conversion, higher monetization, and share gains versus a no-AI baseline.

Exhibit 12: AI recommendations remain a decision layer for ~half of consumers – Frequency of using AI recommendations for purchases



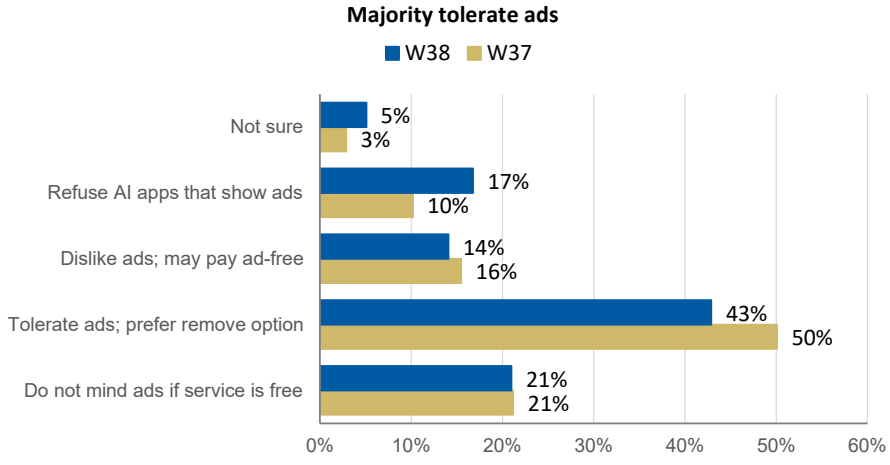
Source: AlphaWise, Morgan Stanley Research. W38 is conducted in June 2026, W37 is conducted in March 2026

Exhibit 13: Food delivery and travel remain the leading AI-assisted transaction categories – Among respondents completing an AI-assisted purchase



Source: AlphaWise, Morgan Stanley Research. W38 is conducted in June 2026, W37 is conducted in March 2026

Exhibit 14: Attitudes toward ads or sponsored results in AI tools – Majority tolerate ads



Source: AlphaWise, Morgan Stanley Research. W38 is conducted in June 2026, W37 is conducted in March 2026

4. Stock implications – favor moats that AI can deepen, not

just traffic that AI can aggregate.

Tencent is the clearest positive read-through in our consumer-AI framework. We do not expect Doubao to displace the WeChat social operating layer; social graph, messaging, Mini Programs, payments and merchant services create high switching costs and a proprietary execution environment. AI can instead make that environment more useful and more monetizable. Importantly, this does not mean every Tencent AI investment has equal ROI: standalone Yuanbao spending should still be evaluated against engagement and commercial conversion.

For **vertical leaders**, the implication is broadly constructive. The best ROI should sit where AI can be embedded into an existing decision and transaction flow – Meituan/local services, Taobao/e-commerce, WeChat/social commerce, OTA/travel, Boss Zhipin/ recruitment and other vertical marketplaces. **BABA** remains our top pick given its full-stack AI capabilities, but our view specially on its consumer AI initiatives is more balanced: we view Qwen Shopping Assistant's integration into Taobao positively, while the economics of heavy standalone Qwen-app investment warrant continued discipline.

Baidu, TME and iQIYI carry less favorable read-throughs because AI directly challenges their core consumer categories: answer engines can pressure Baidu search-query/ad share; AI music creation can fragment TME listening/discovery share; and AI video generation lowers production barriers and expands competing supply, pressuring iQIYI viewing share. The common risk is market-share erosion as AI commoditizes the core product.

AlphaWise Survey Methodology

The report is based on two online surveys conducted during June 30th – July 3rd, 2026 (Wave 38) and March 17th - 20th, 2026 (Wave 37), each covering ~2,000 consumers aged 18-59 residing in Tier 1-4 cities in China across 19 provinces. The survey sample represents urban, more affluent and active Internet consumers relative to the national population. The margin of error for findings based on total sample is ±1.7% at 90% confidence level; higher for subgroups.

In parallel, this report includes data from our monthly US Consumer Pulse Survey of ~2,000 U.S. consumers, which tracks consumer behavior, attitudes, and outlook. Full details to the latest edition conducted from August 20th – 24th, 2026 can be found in Thematic Alpha: US Consumer Pulse Survey: Consumer AI Adoption Continues to Climb (31 Aug 2026).

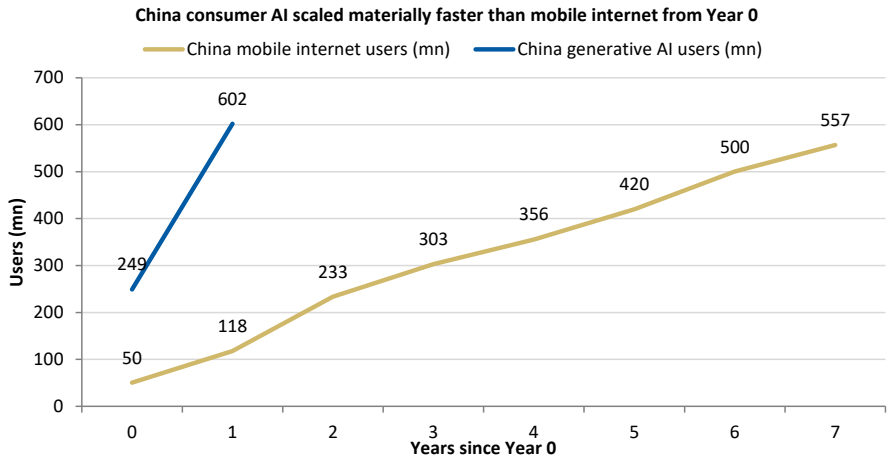
Consumer AI Adoption: China a Global Leader

Why is AI diffusing faster than mobile internet?

China's mobile-internet adoption required a long infrastructure and device cycle: mobile internet users numbered~50mn in 2007 and reached ~557mn by 2014. Consumer AI has scaled much faster, with AI users rising from ~249mn in Dec-2024 to ~602mn in Dec-2025.

The comparison is not perfectly like-for-like, but it shows why AI can diffuse faster: it is a software and interface layer built on top of an already mature mobile internet ecosystem.

Exhibit 15: China consumer AI scaled materially faster than mobile internet from Year 0

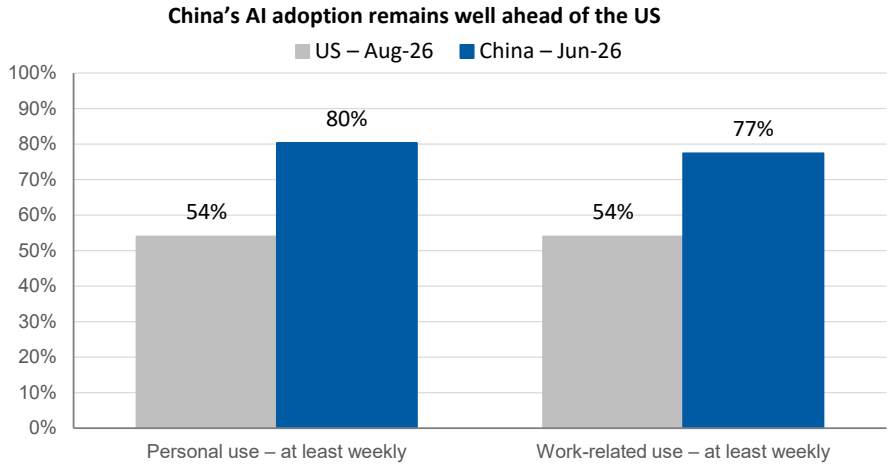


Source: Mobile internet Year 0 = 2007; Consumer AI Year 0 = Dec-2024. The comparison is directional because AI is a software layer distributed over an already mature mobile-internet base. China Internet Network Information Center (CNNIC). Morgan Stanley Research

Why is China's consumer AI penetration higher than that in the US?

Our AlphaWise survey shows 80% weekly personal AI usage in China (vs. 54% in the US), and 77% weekly work-related AI usage (vs. 54% in the US).

Exhibit 16: China's AI adoption remains well ahead of the US



Source: AlphaWise, Morgan Stanley Research. China fieldwork Jun 30–Jul 2, 2026; US Aug 2026 wave.

China's higher 2C AI penetration is driven more by distribution and use-case fit than by frontier-model capability. Although the US still leads in frontier models and compute, China has been faster in pushing AI into mass-market consumer scenarios through super-app ecosystems, low trial friction, low-cost use cases and stronger commercial closed loops.

- **1) China has stronger consumer distribution channels.** AI adoption is not limited to standalone chatbot downloads. It is increasingly embedded into high-frequency apps such as WeChat, Douyin, Taobao, Alipay, Amap, Meituan, and Ctrip, etc. We believe the next phase of adoption is likely to happen inside existing apps, where users already search, consume content, shop, pay, book travel and order local services.
- **2) Free access lowers trial friction but is not unique to China.** What is more distinctive is how Chinese platforms combine free usage with ecosystem traffic, CNY promotions and a wider range of daily-life functions. During the CNY period, Doubao, Yuanbao and Qwen together acquired around 130mn new users, while Qwen downloads reached 170mn in February 2026. This helped convert AI from a novelty product into a broader consumer habit (see more in China Internet: China Consumer AI – 2026 CNY Campaign Review (2 Mar 2026)).
- **3) Most current 2C AI use cases are not very token-intensive.** Search, Q&A, short-form writing, shopping comparison and travel planning usually require only short prompts and relatively short outputs (see more for consumer AI use cases here: **2. Competitive Landscape and use cases**). These use cases are much cheaper than long-document analysis, deep research, coding agents or multi-step autonomous agents. As a result, platforms can theoretically support free or low-cost usage and monetize indirectly through ads, traffic routing or transaction conversion (Exhibit 17).
- **4) China's platforms have stronger commercial closed loops.** Chinese super apps combine content, search, merchants, payment, fulfillment, customer service and advertising in one ecosystem. This allows AI to move directly from answering questions to recommending products, booking services, placing orders or handling after-sales. In the US, individual verticals, such as search, social and e-commerce,

are also highly concentrated, but they are owned by separate platforms. End-to-end agent execution therefore depends on more cross-platform integrations, permissions and commercial agreements.

- **5) Daily consumer use cases do not always require the best frontier model.** For mass-market consumer scenarios, such as search, summarization, short-form writing, shopping guidance and travel planning, the decisive factor is often not whether the model is globally best-in-class, but whether it is good enough, free/low-cost, easy to access. This explains why China can lead in consumer AI penetration even if the US still leads in frontier model capability and compute infrastructure.

Exhibit 17: Token usage budget by task (assumes pricing as DeepSeek V4 Flash - Rmb1.5 per million input tokens, Rmb4.5 per million output tokens)

Use case	Complexity	Input tokens	Output tokens	Cost / run	Runs per Rmb1
Simple Q&A / short answer without tools	Easy	100–2k	100–2k	Rmb 0.0006–0.012	83.33–1666.7
Short-form content generation: email, ad copy, social post, bullet summary	Easy	300–3k	200–1.5k	Rmb 0.0014–0.011	88.89–740.7
Single-search factual lookup / current fact answer	Easy	2k–128k	300–8k	Rmb 0.0043–0.228	4.39–229.9
RAG / knowledge-base Q&A / file-search answer	Medium	2k–25k	200–2k	Rmb 0.0039–0.047	21.51–256.4
Image understanding / vision Q&A	Medium	1k–15k	200–3k	Rmb 0.0024–0.036	27.78–416.7
Medium- to long-form content generation: blog, landing page, proposal	Medium	2k–30k	2k–15k	Rmb 0.0120–0.113	8.89–83.3
Long-text rewrite / translation / polishing	Medium	1k–200k	800–260k	Rmb 0.0051–1.470	0.68–196.1
Shopping recommendation / product comparison	Medium	10k–100k	1k–5k	Rmb 0.0195–0.173	5.80–51.3
PPT / report writing from a short brief	Medium	2k–20k	5k–30k	Rmb 0.0255–0.165	6.06–39.2
Data analysis / code-interpreter / notebook-style analysis	Medium	5k–150k	2k–30k	Rmb 0.0165–0.360	2.78–60.6
Long-document summarization	Hard	20k–1M	1k–20k	Rmb 0.0345–1.590	0.63–29.0
PPT / report writing from extensive source materials	Hard	50k–500k	10k–60k	Rmb 0.1200–1.020	0.98–8.3
Coding agent: small bug fix / explanation / patch	Hard	10k–80k	5k–30k	Rmb 0.0375–0.255	3.92–26.7
Coding agent: medium feature / refactor / test generation	Hard	80k–800k	30k–300k	Rmb 0.2550–2.550	0.39–3.9
Deep research / multi-step research agent	Hard	100k–1M	20k–200k	Rmb 0.2400–2.400	0.42–4.2
Large software agent / long-running multi-agent session	Hard	500k–10M	100k–2M	Rmb 1.2000–24.000	0.04–0.8

Source: DeepSeek official website. Morgan Stanley Research

Entry Points: AI-Native Apps Win Mindshare, In-App AI Wins ROI

China's consumer AI race is no longer only about who builds the best chatbot. It is becoming a race for AI entry points: **AI-native apps for mindshare, in-app AI for monetization, and OS-level AI for long-term orchestration.** We prefer in-app AI assistants from an economics perspective because they can monetize existing traffic, proprietary data and closed-loop transactions without rebuilding user habits from scratch.

AI-native apps have been the fastest way to educate users and capture top-of-funnel intent, but multi-homing and low switching costs limit defensibility. While Doubao, Qwen, DeepSeek and Yuanbao have captured early mindshare, their long-term defensibility varies by product iteration, traffic access, ecosystem depth, user stickiness and monetization path.

- **Doubao** clearly leads on user scale through ByteDance's traffic engine and fast product iteration;
- **Qwen** has the strongest transaction-loop advantage through Alibaba's ecosystem;
- **Yuanbao** remains behind as a standalone app, but WeChat Agent could become Tencent's real AI entry point;
- **DeepSeek** remains a model-first outlier, driven by technical capability rather than user acquisition.

In-app AI assistants have the clearest near-term ROI because they sit inside existing user habits and transaction flows, with proprietary data, merchant supply, payment, fulfillment and attribution.

OS-level AI assistants may become the long-term command layer, especially for local context and cross-app tasks, but execution still depends on app permissions, commercial agreements and cloud inference for complex workflows.

Dimensions	AI Native Apps	In-app AI Assistants	OS-level AI Assistants
Key players	Doubao, Qwen, DeepSeek, Yuanbao, Kimi, Baidu Wenxin...total 440mn+ users in March	Baidu/Douyin/WeChat/Amap/Tencent news/Toutiao/Taobao/Meituan AI search, WeChat AI agent (not released), WeChat Yuanbao friend, Douyin Doubao friend...	Oppo's Xiaobu, Huawei's Xiaoyi, Xiaomi's Xiaoai, Vivo's BlueLM, Honor's YOYO, Apple Intelligence; Xiaomi MiClaw (not released)
Core logic	Build a new super entry	Embed AI into existing apps for a seamless user experience upgrade	Control the system-level command layer

Competitive advantages	High product freedom, can polish pure AI interaction experiences	No extra downloads needed; user preferences/data already accumulated on the host app; stable agentic execution given high permissions; naturally form a transaction closed loop	Zero user acquisition cost; on-device data privacy protection; offline availability; potential cross-app operations
Competitive disadvantages	High CAC, needs to build up user habits and brand awareness from scratch; agentic execution depends on app-vendor negotiations.	Limited by host apps' positioning (no matter how good the AI in a shopping app is, users won't go there just to "chat", except for WeChat)	Agentic execution depends on app-vendor negotiations.

Dimensions/Scores (1-5)	**AI Native Apps**	**In-app AI Assistants**	**OS-level AI Assistants**
User scale	4 (Doubao alone has ~400mn MAU as of July 2026)	5 (leverage/constrained by the app's user base, but WeChat is the biggest variable, and adoption rate of AI features varies)	4 (high penetration via OS pre-installation, but activation rate is a variable)
User stickiness	3 (extremely low switching costs, ~10min average daily time spent per user)	5 (relies on app's established user habits)	4 (deep OS integration creates dependence)
Scenario coverage	3 (high product freedom, but limited by group's ecosystem)	4 (confined to app's specific domain, e.g., shopping/travel, etc., but WeChat Agent is a separate case)	5 (system-level access allows cross app scenarios)
Scenario depth	3 (depends on tasks)	5 (deepest vertical capability with proprietary closed-loop data and high permission)	3 (depends on tasks)
Monetization potential	3 (not yet clear, potentially subscriptions + ads + commissions)	4 (natural monetization within the existing transaction loop)	3 (hardware premium, cloud-service bundles, etc.)
User acquisition efficiency	3 (high CAC, needs to build up user habits and brand awareness from scratch; now may have passed the user education phase especially for Doubao)	5 (can be exposed to the existing user base, but activation depends on entry design, user permission and habit formation)	4 (can be exposed to the existing user base, but activation depends on entry design, user permission and habit formation. Smartphone brands are fragmented in China)
Data/technical moat	3 (model capabilities are converging; lack of proprietary data flywheel)	5 (domain-specific data moat)	4 (hardware/OS)
Total Scores	22	33	27

Footnote: 1 = weak, 5 = strong. Based on Morgan Stanley Research assessment of the variables.

AI-Native Apps

AI-native apps have been the fastest way to educate users and capture top-of-funnel intent.

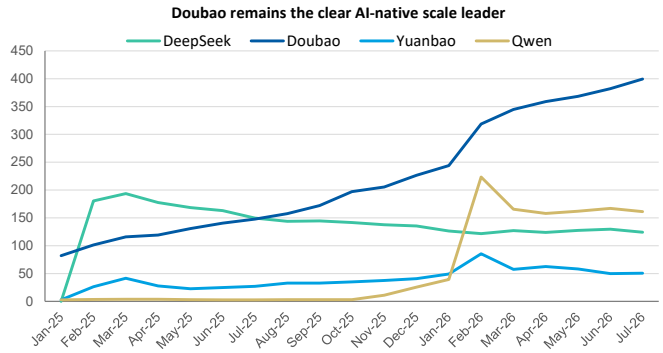
- **Doubao** is the only clear scale and engagement leader, supported by ByteDance traffic, product iteration and mass-market positioning.
- **Qwen** has weaker post-campaign engagement, but it has the best path from question to transaction because it plugs into Alibaba's commerce, local service, payment and travel assets.
- **DeepSeek** maintains strength in perceptions of its reliability and accuracy, and its user engagement, but its consumer distribution and monetization remain limited.
- **Yuanbao** underperforms as a standalone app, while Tencent's real consumer AI opportunity is likely to sit inside WeChat rather than the Yuanbao app itself.

However, AI-native apps face high costs for user acquisition, and they are more likely to aggregate intent than fully monetize it. We would track retention, commercial-query conversion, AI-attributed gross merchandise value (GMV) and paid-feature adoption rather than MAU/DAU alone.

AI Native Apps		Doubao	Qwen	DeepSeek	Yuanbao
User scale (July 2026)	**MAU** (mn)	399	161	124	51
	DAU (mn)	168	28	31	9
	DAU/MAU ratio	42%	17%	25%	17%
User engagement (July 2026)	**Average daily time spent per user** (min/day)	9.9	4.0	16.5	11.2
	Monthly sessions per user	69	17	61	29
Product positioning		AI companion, scale and engagement leader	AI assistant that gets things done	Pure chatbot with frontier reasoning	Yuanbao standalone positioning remains unclear but we believe Yuanbao sits inside Tencent's ecosystem and data flywheel with Hy LLM. WeChat AI agent is real opportunity
Core advantages		ByteDance traffic + product iteration	Alibaba commerce / local / travel ecosystem	Technical credibility and brand mindshare	Social graph, Mini Programs, payment
Main weaknesses		Still needs stronger transaction closure	Post-subsidy engagement weaker	Weak consumer distribution / monetization	Product rollout slower; UX discipline may slow monetization

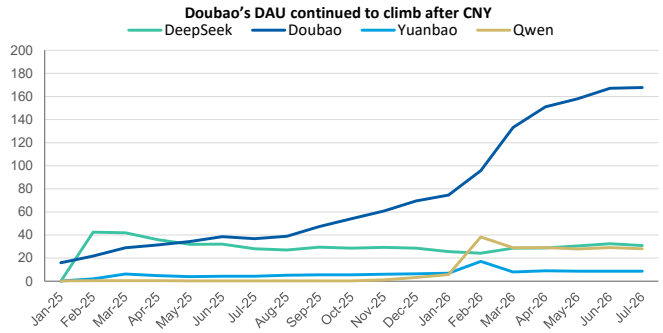
Monetization potential	Ads, Douyin Mall commission, prosumer paid features	Transaction commission, ads, ecosystem GMV uplift, payment	API	WeChat AI agent, transaction commission, ads, payment
User acquisition efficiency/distribution channel	High efficiency given traffic redirection from Douyin/Toutiao and viral content marketing	Acquired sizable users through CNY consumer subsidies	No user acquisition efforts, purely brand awareness and word-of-mouth	Yuanbao as a WeChat friend allows it to reach WeChat's 1bn+ users without any additional user acquisition costs

Exhibit 18: MAU (mn)



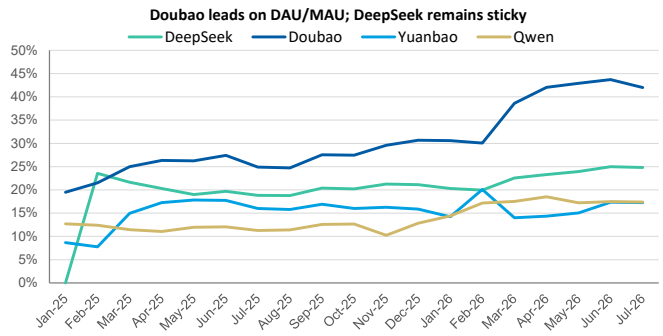
Source: Questmobile, Morgan Stanley Research

Exhibit 19: DAU (mn)



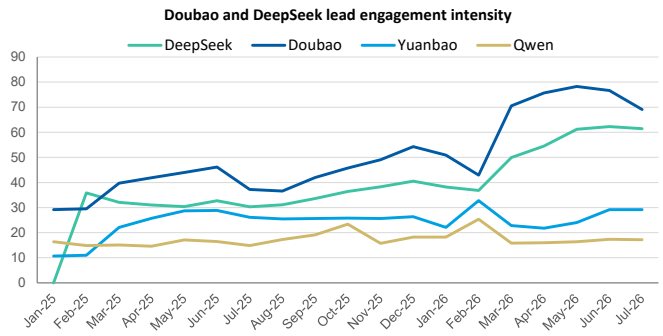
Source: Questmobile, Morgan Stanley Research

Exhibit 20: DAU/MAU ratio



Source: Questmobile, Morgan Stanley Research

Exhibit 21: Monthly sessions per user



Source: Questmobile, Morgan Stanley Research

Doubao (ByteDance): Move early, spend big, iterate fast

In the mobile era, search (Baidu), social (Tencent), and feeds (ByteDance) were the

gateways. Now, ByteDance is trying to make Doubao the new gateway in the AI era. By locking in user mindshare through aggressive marketing spend and massive DAU, Doubao aims to become the default first choice for users seeking information and solving problems.

Our take on key factors behind Doubao's success

- **Ecosystem/traffic - traffic leverage from Douyin:** Doubao's strong growth relies heavily on organic traffic from Douyin. through internal traffic diversion, splash screen ads, and seamless integration into its content ecosystem – Influencers create viral content interacting with Doubao, such as having it supervise children's homework or assist with fashion styling, which drives millions of organic downloads daily. This speaks to ByteDance's mature product strategy.
- **Product design - differentiated product positioning, agile product-market fit (PMT) testing:** While other players focus on long-form text generation, Doubao took a differentiated path – emotional connection & mass-market accessibility. Doubao is the only mainstream AI chatbot with a personified persona, focusing on a "friendly and approachable" style that supports **voice, video, and dialect-based interactions**. By intentionally lowering the barrier to entry for complex queries, catering more effectively to the usage habits of the general public and the demand for "companionship" and "efficiency", including the elderly and children. According to Questmobile, 50%+ of Doubao's March MAU was from Tier-3 and below cities.
- **Robust foundation model support:** The Doubao LLM family ranks among the top tier in China, with its image/video generation model Seedance 2.0 being internationally competitive. Daily average usage of Doubao models exceeded 120tn in March (double the level from three months ago).
- **Organizational commitment:** As mentioned in the above timeline, ByteDance has iterated the organizational structure of its AI business, allowing the tech giant to maintain entrepreneurial agility at scale. Since late 2024, founder Yiming Zhang participated monthly in Seed's review and strategy sessions, ensuring AI strategy aligns with the company's high-level design.

Qwen (Alibaba): From Question to Action

Our take on Qwen's key advantages:

- **Qwen's most irreplaceable advantage - ecosystem:** Qwen's market entry is distinct given it is not just a chatbot, but an AI agent designed to execute real-world tasks by plugging directly into Alibaba's ecosystem – Taobao, Taobao Instashopping, Fliggy, Amap, Alipay, Freshippo, Damai, etc. – so that users can order food, shop, book flights and hotels, plan a trip, take a taxi, buy movie tickets, completing a closed-loop form, expressing intent to fulfillment and delivery. Being close to transaction is probably the most

valuable commercial unit (an actual order) also provides Qwen with stronger monetization potential.

- **Top-tier model capabilities:** The Qwen model family is a global leader, especially in open-source. This provides the Qwen app with technical reliability and should help ensure better user experience.

Our take on Qwen's challenges:

- **Post-subsidy retention/engagement:** Qwen's DAU peaked at 74mn on Feb 7, 2026, and has now stablized at 29mn (March 2026), which is still well-above pre-campaign levels. However, Qwen's user engagement metrics significantly lag peers – average daily time spent per user was 4min (vs. peers at 10-17min), monthly sessions was 17 (vs. Doubao at 69, Deepseek at 61, Yuanbao at 29), and DAU to MAU at 17% (vs. Deepseek at 25% and Doubao at 42%).

Yuanbao (Tencent): Standalone scale trails, but its strategic value lies in Hy's product–model flywheel

Yuanbao still trails Doubao, Qwen and DeepSeek in standalone app scale, while its consumer positioning remains less distinctive. Tencent is increasingly positioning Yuanbao as a broad consumer feedback and capability-incubation layer for Hy. Its conversational and agent use cases generate real-world feedback, while functionalities developed within Yuanbao can become reusable "atomic capabilities" across WorkBuddy, CodeBuddy, Weixin and QQ Browser. Standalone MAU therefore remains relevant, but increasingly as an input into the flywheel rather than the sole measure of Yuanbao's strategic value.

The more important opportunity is **WeChat**. Tencent owns one of China's strongest consumer operating environments: social graph, messaging, Mini Programs, Official Accounts, Channels, payments and merchant services, and its Weixin AI Agent Xiaowei is set to connect user intent toward completing tasks with the ecosystem capabilities. **This makes it difficult for AI-native apps to disrupt the WeChat ecosystem, while positioning the platform to gain further strength from AI enhancement.**

Yuanbao benefited from the DeepSeek R1 wave and CNY marketing, but it still lags Doubao, Qwen and DeepSeek in scale, engagement, product clarity and feature depth. Its brand identity also remains heavily associated with DeepSeek rather than Tencent's own Hunyuan model. As a standalone AI-native app, Yuanbao does not yet have a sharp enough user proposition.

The more important opportunity is **WeChat**. Tencent owns one of China's

strongest consumer operating environments: social graph, messaging, Mini Programs, Official Accounts, Channels, payments and merchant services. **This makes it difficult for AI-native apps to disrupt the WeChat ecosystem, while positioning the platform to gain further strength from AI enhancement.** If WeChat Agent can connect user intent with Mini Program execution, Tencent could turn WeChat from a communication and service platform into a broad AI task-completion layer.

Our take on why Yuanbao is falling behind

- **Yuanbao's standalone positioning remains unclear.** For many users, Yuanbao's key appeal has been stable access to DeepSeek-style reasoning, which makes its own brand identity less differentiated. Compared with Doubao's "AI companion" positioning, and Qwen's "AI assistant that gets things done" positioning, Yuanbao still lacks a memorable consumer label and a distinctive daily-use scenario.
- **Feature iteration lags:** Doubao and Qwen both offer comprehensive tool suites, including deep research, AI tutor, AI writing assistants, translation, video/phone calls, etc., which Yuanbao doesn't offer.

In-app AI Assistants

In-app AI is the most attractive commercialization layer because it improves existing behavior rather than requiring users to form new habits. We believe in-app AI assistants benefit on context, customer reach, data, execution, and monetization.

- Leverage massive, pre-existing user base with near-zero incremental acquisition costs and education costs - it enhances existing user behavior without changing user paths;
- Powered by proprietary data and user behavior data (purchase history, payment records, social graphs, travel trails) that are inaccessible to external AI apps;
- Deeply integrated into app ecosystems, enabling high-privilege, agentic task execution;
- Establish an immediate monetization loop - not by charging for the AI itself, but by boosting apps user engagement, retention, ad conversions and GMV.

We expect in-app AI monetization to come mainly from engagement uplift, ad ROI improvement, sponsored recommendations, merchant tools, GMV conversion and take-rate / commission. The constraint is that most in-app AI assistants are domain-specific – Taobao for shopping, Amap / Ctrip for travel, Meituan for local services – while WeChat is the notable exception because Mini Programs provide much broader scenario coverage.

ByteDance

Apps	MAU (Jul 2026)	DAU (Jul 2026)	AI feature	Description
Douyin	1.03bn	714mn	AI search	Natural language searches that generate answers from video content
			Doubao friend	Doubao is embedded as a default friend within Douyin, offering functionality largely comparable to the standalone Doubao app.
			@Doubao	Users can tag @Doubao in the comments section to summarize content or ask questions
Toutiao	373mn	170mn	AI chatbot	Users can chat and ask questions, convert news into podcasts
Fanqie Novel	266mn	144mn	AI character	Readers can have AI-powered conversations with book characters while reading

Alibaba

Apps	MAU (Jul 2026)	DAU (Jul 2026)	AI feature	Description
Taobao	967mn	388mn	AI omni search	Embedded in search box. Natural language product searches with AI recommendations. Released in August 2025
			AI try-on	AI virtual try-on feature that helps users preview how clothes will look on them
			Place in my home	Home furnishing AI. After users upload a floor plan, AI will provide home layout suggestions and preview how home furnishings look on their home.
			AI photo search (Pailitao)	Snap a product, AI will help find similar products, summarize comments, and recommend other products you might need.
Alipay	947mn	352mn	AI agent "Ah Bao"	A native conversational interface within the Alipay app, It acts as a personal assistant, allowing users to coordinate more than 10,000 everyday services – such as ride-hailing, food delivery, and EV charging – through natural voice commands
Amap	937mn	184mn	Teacher Xiaogao (AI travel agent)	Map-based AI agent, supporting multi-modal interactions including voice, text, and vision, with reasoning, planning, and execution capabilities. Released in Aug 2025.
Idle Fish	241mn	62mn	AI camera	Snap to list products, AI auto-fills category, generates title & descriptions, and suggests fair reference price based on platform transaction data.
Quark	117mn	33mn	Quark AI search	Quark AI search page is now functionally similar to the Qwen app.

Tencent

Apps	MAU (Jul 2026)	DAU (Jul 2026)	AI feature	Description
WeChat	1.16bn	931mn	Yuanbao friend	Add Yuanbao as WeChat friend, chat for Q&A, content creation, image generation. Forward articles from official account, files/messages from other chats, or videos from channels to Yuanbao friend and ask questions. Launched in Apr 2025.
			@Yuanbao	User can tag @Yuanbao in official account/video channels comments to summarize content or ask questions
			WeChat AI Search	Natural language searches that generate answers from official accounts, channels, etc.
			WeChat Agent Xiaowei	Grey testing since June 2026
QQ	635mn	198mn	Yuanbao group chat summary	AI auto-summarizes hot topics, key info, filters content "relevant to me" to generate a summary; clicking on summary jumps to original message
			QQ agent (Xiao Q)	functions similar to Yuanbao app
QQ music	192mn	48mn	Smart music butler (AI music agent)	Natural language dialogues to find songs, set timed playback, etc. Released in late 2025.
			AI listen along	Voice-activated AI interprets song emotions, creation, stories, and era markings.
			AI MV	One-click generation of personal music videos
Tencent News	276mn	130mn	AI news summary	AI-powered news summaries enable users to efficiently capture key takeaways from articles

Weixin's social moat should be durable in the AI era

Tencent's social ecosystem encompasses ~1.9bn users, with WeChat serving as an operating-system-level entry point - messaging, payments, Mini Programs, Channels, and Official Accounts. The cost of migration is exceptionally high for WeChat users. The proprietary, closed-loop data and relationship chains that WeChat possesses are, in the AI era, the scarcest and most valuable "fuel".

Even if Doubao/Qwen were to lead in AI capabilities, they are unlikely to replace WeChat as the infrastructure layer for social relationships. Most AI native apps are mostly task-oriented - users open them to complete a a specific , then close them - but WeChat is not merely a tool but a high-frequency operating environment

Why WeChat is well-positioned for agentic execution

Tencent stated in April that it plans to "[embed] AI across the WeChat ecosystem to enhance user experience, spanning social interactions, information retrieval and analysis, content consumption, merchandise recommendation and customer service", "building next-gen agentic services in WeChat to connect users with our extensive Mini Programs content, social and payment ecosystem", and "adding agents to WeChat boosts ecosystem activity and therefore generates revenue itself".

Mini program infrastructure almost looks tailor-made for AI agents: Mini Programs are, at their core, highly standardized interfaces (API). When an AI agent enters WeChat, it can directly invoke Mini Program APIs rather than navigate or replicate complex graphical user interfaces.

Breadth of scenario coverage: While Qwen's integration with its ecosystem creates a closed loop, the breadth and diversity of the Mini Program ecosystem remains difficult to replicate – millions of mini programs cover all life-service scenarios from e-commerce, housing, transportation, healthcare, entertainment, to social networking, etc.

Latest Progress on WeChat Agents

Tencent has started opening WeChat's AI ecosystem to Mini Program developers, with both automatic and developer integration modes, while the program remains in internal testing and requires developer authorization. Early trial partners include **JD** for shopping, food delivery and logistics; **Meituan** for food delivery and local services; **Trip.com/Tongcheng** for flight booking, hotel reservations and itinerary planning; and **KFC** China for natural-language ordering. Honor has also completed integration through an agent-to-agent protocol, with Huawei, Xiaomi, Oppo and Vivo expected to follow.

Weixin's AI Agent Xiaowei was officially introduced at the August 2Q26 results, described as "the agentic AI inside Weixin". Xiaowei applied WeLM - a dedicated model developed by the Weixin team and with significant cost efficiency. The recent WeLM-80B has 80bn total parameters but activates only 3bn per token, and supports chat and search, native Weixin functions and Mini Program service calls.

Gray-testing underway and expanding: Xiaowei started in June with a small-scale user grey test. Initial testing covered conversational assistance, file and chat summarization, selected native Weixin actions and Mini Program-based services; key actions such as sending messages and making payments retained explicit user confirmation, while final social publishing remained user-controlled. The test scope has since expanded to Moments "AI writing" and "AI comment" functions in August, and Xiaowei-to-Xiaowei social interaction in September, indicating a staged progression from information assistance toward social and transactional execution.

Others

Apps	MAU (Jul 2026)	DAU (Jul 2026)	AI feature	Description
PDD	711mn	390mn	AI summary	AI summary of product features and comments
RedNote	255mn	127mn	Diandian	Natural language searches that generates answers from RedNote content
			@Diandian	Users can tag @Diandian in comments to summarize content or ask questions
Meituan	529mn	181mn	Xiaotuan	Natural language search for restaurants, local services, sightseeing places, etc. On Apr 27, 2026, entry for Xiaotuan was moved to Meituan app homepage.
Ctrip	104mn	15mn	Ctrip Wendao	Travel AI agent. Natural language search for hotels, book flights, itinerary planning, travel consultation.
JD	629mn	151mn	AI search	Embedded in search box. Natural language search for products
Baidu	644mn	167mn	AI search	Embedded in search box.
Kuaishou	453mn	198mn	AI-enabled search/ marketing/merchant tools	AI search / OneSearch; AI-agent shopping and live-commerce assistance; AI-enabled marketing/content workflows

OS-level AI Assistants

OS-level AI should be treated as a hybrid architecture, not pure on-device AI. Light, private and latency-sensitive tasks can run locally, while complex reasoning, long-context analysis, multimodal generation and agentic workflows will still rely on cloud inference. This creates both a strategic advantage and a monetization limitation.

- **The advantage is default access:** OS assistants can use local context, notifications, screen, keyboard, voice and device data, subject to user permission. Over time, they may control the cross-app command layer.
- **The limitation is execution and monetization:** OS agents need app permissions, commercial agreements and APIs to complete transactions, while the economics of orders, ads and commissions often remain inside apps.

We see OS-level AI as a strategic control point, with monetization from device premium, cloud-service bundles, traffic routing and ecosystem lock-in rather than immediate transaction take rate.

Advantages:

- Default entry point: Users can invoke AI from the phone, voice, screen, keyboard, notification layer
- Low user acquisition cost: Can be exposed to the full existing user base, but actual activation depends on entry-point design, user permission and habit formation
- Local context/data: OS can access the most private/authentic behavioral data on the phone (location trails, contacts, phone albums, health metrics, app usage behaviors etc.) based on user permissions, which can improve personalization
- Privacy/latency: Simple tasks can be processed locally with lower latency and lower marginal cloud cost
- Cross-app orchestration: OS can theoretically call multiple apps to complete tasks
- Monetization through hardware price premium

Disadvantages:

- On-device model constraints: Memory, battery, thermal and context-window constraints limit reasoning depth
- Cloud still needed: Complex agentic tasks require larger models and server inference
- App permission bottleneck: OS cannot complete transactions unless apps expose APIs / permissions
- Monetization less direct: OS controls routing, but transaction economics often remain inside apps
- Hardware cycle: Full experience may require newer devices, slowing penetration

What is the next phase of consumer AI?

We expect China's consumer AI penetration to continue rising from broad trial to habitual usage. While total AI-assisted consumer touch points should be meaningfully higher once in-app assistants and OS-level agents are included.

The next phase of consumer AI is likely to shift from answer generation to task completion. We see three stages:

- First, AI answers questions and generates content;
- Second, AI helps users compare options and make decisions;
- Third, AI executes tasks across apps, including booking, ordering, paying, refunding, rescheduling and after-sales service.

This shift should favor platforms that own both user intent and execution rights.

- Existing in-app assistants are best positioned for closed-loop execution because they already control product catalogs, merchant relationships, payment, logistics, customer service and attribution.
- OS-level agents may become the cross-app orchestration layer, but they need app permissions and commercial agreements.
- AI-native apps may remain strong top-of-funnel intent aggregators, but they need supply integration or referral economics to capture monetization.

Stages	User behavior	Winning metric	Likely winner
1. Ask/Generate	Search, Q&A, writing, image generation	DAU, sessions, retention	AI-native apps
2. Decide	Compare products, plan travel, choose restaurants, health navigation	Recommendation acceptance, CTR (click-through rate), conversion uplift	In-app AI + AI-native apps
3. Execute	Order, book, pay, refund, reschedule, after-sales	Completed task rate, AI-attributed GMV, take rate	Existing in-app AI first; OS-level agents later

China vs. US

Entry point	China pattern	US Pattern	Key difference

AI-native apps	**Distribution- and use-case-led.** Leading apps combine free basic access with ecosystem traffic, large-scale consumer campaigns and localized everyday use cases. Product design tends to emphasize broad consumer adoption, including entertainment, companionship, search and increasingly transaction-related tasks.	**More product- and capability-led.** Free basic access is also common, while premium tiers are more established around advanced models, higher limits and productivity features. Adoption has historically been more closely associated with information, creation and productivity use cases.	**Free access itself is not the key differentiator.** China's advantage is the combination of free access with stronger ecosystem distribution, consumer campaigns and a broader set of mass-market use cases.
In-app AI	**More tightly integrated with existing consumer workflows and transaction loops.** Super apps and vertical platforms can combine AI with proprietary user data, merchant supply, payments, fulfillment and after-sales service, allowing AI to move from recommendation to transaction within the same ecosystem.	**AI is increasingly embedded across US platforms as well**, but the consumer journey is more fragmented across separate search, social, commerce, payment and local-service ecosystems, making closed-loop execution more difficult.	**China's advantage is closed-loop execution.** In-app AI may therefore have a more direct path to conversion, advertising ROI and commission monetization in China.
OS-level AI	**A multi-OEM ecosystem with growing emphasis on app calling and cross-device orchestration**. Local OEMs can combine device context with integrations into domestic apps, although fragmentation across handset brands may limit a single OS-level winner.	**More concentrated at the OS layer,** with a stronger focus on privacy, on-device processing, personal context and system-level integration. Cross-app execution still depends on third-party app permissions and APIs.	**Both markets see the OS as a potential long-term orchestration layer.** China may have deeper integration with local service ecosystems, while the US has a more concentrated OS distribution structure.

Case Study - META Muse

Meta's Sep 9 launch of Muse makes the consumer-agent debate more immediate. Muse is a free standalone consumer AI agent (formerly codenamed Hatch) designed to automate online shopping, travel, ticket purchasing, appointment/calendar management, email and other tasks. Sensitive actions such as completing purchases or payments require user verification. Importantly, at launch Muse is separate from Instagram, Facebook and Meta's other apps, and for now does not share data with Meta's advertising system (see more in Meta Platforms Inc: Muse Has Arrived; What Matters Next? (9 Sep 2026)).

This makes Muse a particularly useful test case for our framework: The product has begun as an AI-native/standalone entry point, while its potential structural advantage comes from ecosystem assets that could be connected later.

Scaled distribution and large proprietary consumer datasets are critical to personalized agentic offerings and adoption.

- For Meta, the swing factor is whether consumer adoption is followed by integration of Facebook, Instagram, Messenger and WhatsApp context, alongside third-party services and personalized datasets such as Gmail. If this happens, Muse could evolve from a standalone front end into an orchestration layer over Meta's distribution and external services. In other words, the strategic moat would not be

the standalone app itself; it would be the data, distribution and execution graph behind it.

- China's apps' also have similar advantages – WeChat can connect social context, Mini Programs and payments; Taobao/Qwen can connect product, merchant and fulfillment data; Meituan can execute local-service workflows.

Muse's most important KPIs are repeat task completion, share of high-intent shopping/travel/service queries, verified purchase/payment conversion, AI-attributed GMV or lead generation, merchant/service integration breadth, inference cost per completed task, and the extent of Meta app/data integration. Brian Nowak likewise argues that investors need evidence of adoption and monetizable behavior before assigning value to Muse. **This is consistent with our broader view that consumer-AI competition should move from traffic aggregation toward completed tasks, conversion and monetization.**

Dimensions	META Muse at launch	China AI-native apps/in-app AI assistants
Entry point	AI-native app	AI-native apps and in-app AI assistants
Data	No current sharing with Meta ad system; future Meta app context is a key swing factor	Proprietary-commerce and service data is shared, but not yet for WeChat's chat history
Execution	Shopping, travel, tickets, calendar/appointments, email; user verification for purchases/payments	Mini Programs, merchants, payment, fulfillment and service workflows.
Monetization	Not yet proven	Ads, GMV and commissions are yet to see meaningful uplift
Moat	Meta distribution + proprietary data + service integrations	Distribution + vertical data + payment/fulfillment rights

Monetization: Freemium at the User Layer, Intent Monetization at the Commercial Layer

The core monetization challenge for consumer AI is that scale creates costs before it creates revenue. Traditional internet platforms benefit from strong operating leverage as user growth spreads fixed costs over a larger base, but AI products face a structurally different cost curve: every additional interaction consumes inference compute. This weakens the traditional internet playbook of subsidizing user growth first and monetizing later.

In China, **subscription** is unlikely to become a mass-market revenue engine and may function more as a cost-sharing mechanism for heavy users; **advertising** presents a larger opportunity given Chinese users' higher tolerance for ads, but chatbot ad load is structurally limited by trust and user experience; and **transaction commissions** may offer the largest long-term upside if AI agents can move from answering questions to completing purchases.

Consumer AI is likely to accelerate existing mobile-internet monetization rather than create a separate revenue pool. We therefore start with existing consumer-internet revenue and GMV pools, and then layer in AI adoption, commercial-intent, closed-loop attribution/conversion and pricing/take-rate assumptions. On that basis, we estimate ~Rmb294bn (~Rmb300bn) of China 2C AI monetizable revenue TAM by 2030E, rising to ~Rmb1.6tn by 2040E as AI penetrates more consumer workflows.

The core challenge: Consumer AI products create costs before creating revenue

Traditional internet products are usually characterized by high fixed costs and low marginal costs. Once the product, content infrastructure, recommendation system and traffic engine are built, the cost of serving one additional user or one additional session is relatively low. Historically, this has made scale highly attractive for internet platforms: as MAU, DAU, time spent and ad inventory grow, fixed costs are diluted across a larger base and margins can expand.

Consumer AI products have a different cost curve, and do not benefit from the same degree of automatic operating leverage. Training and infrastructur
have fixed-cost characteristics, and inference efficiency can improve
core user experience is usage-based. Every additional AI query, every
every PPT creation, every data-analysis task, and every agentic work
incremental compute. As usage scales, inference cost scales with it.
therefore both a monetization opportunity and a cost driver.



This is why AI monetization should be framed as a unit economics question. AI user scale is economically valuable only if platforms can effectively monetize usage intensity, whether through subscriptions for heavy users, commercial monetization embedded within the user experience, or transaction completion and associated commissions.

Sizing China 2C AI Monetizable TAM - Definitions and Methodology

We define Consumer AI monetization revenue TAM as the platform-level revenue that can be captured from AI-assisted consumer activity across AI-native apps, in-app AI assistants, AI search, OS-level assistants and agents embedded in e-commerce, local services, travel, payment and super-app ecosystems.

We separate the opportunity into three revenue pools to avoid double counting:

- **Subscription / pay-per-use:** user-paid revenue from advanced models, higher quotas, productivity tools, faster response, image/video generation, PPT generation or data analysis.
- **Advertising:** merchant-paid sponsored exposure, paid ranking, lead generation, CPC/CPM/CPA (i.e., cost per click, thousand, or acquisition) placements or sponsored recommendations triggered by AI search, chatbot, assistant or agent surfaces.
- **Transaction commission:** take rate, CPS (cost per sale), affiliate fee or referral commission from completed AI-attributed transactions such as shopping orders, hotel bookings, travel reservations or local services.

The addressable user base also differs by monetization model. Subscription revenue is anchored to a narrower pool of users who could plausibly pay directly for AI, whereas advertising can reach a broader audience across AI-native, in-app, and AI search interfaces. We therefore use a narrower subscription-addressable AI MAU and a broader advertising-addressable user base. Users do not need to pay for AI to generate monetizable activity through sponsored recommendations, search ads, lead generation, or other merchant-funded intent.

Although advertising and commission revenue will likely become increasingly intertwined at the product level, we distinguish between the two based on the monetization trigger. We classify exposure-, ranking-, click-, or lead-driven revenue as advertising, while revenue generated only after a completed order, booking, or payment falls under transaction commissions. Together, advertising and commissions form AI intent monetization, which contributes ~99% of our 2030E base-case monetizable revenue.

Rmb bn	2030E Base	2040E Base	2040E Bear	2040E Bull
Subscription	3	33	9	84
AI Advertising	10	319	25	1,534
AI transaction commissions	281	1,241	735	1,905

Total 2C AI monetizable revenue	294	1,593	769	3,523

Exhibit 22: Definitions and Methodology

Definitions and Methodology	
Model scope	Broader AI-enabled consumer internet monetization pool, including AI-native apps, in-app AI assistants, OS-level assistants, AI search, AI shopping agents, travel agents, local-services agents, and super-app AI surfaces. It is not a standalone chatbot-only TAM.
Advertising definition	Revenue from sponsored exposure, paid ranking, AI search ads, sponsored recommendations, lead generation, or performance ads where the payment is triggered by exposure, click, lead or attributed commercial recommendation.
Transaction commission definition	Revenue from take rate, CPS, affiliate/referral fee, or platform commission when an AI workflow participates in a closed-loop transaction such as completed shopping, food delivery, hotel, ticket or local-services booking.
Boundary rule	If monetization happens because the AI surface influenced attention or intent, it is advertising. If monetization happens because the AI workflow completed or attributed a transaction, it is commission.
Double-counting rule	Ads and commissions are additive only when they are separate merchant payments. Sponsored placement / marketing spend is counted as advertising; take-rate / CPS on completed GMV is counted as transaction commission. The same payment should not be counted twice.
Advertising formula	Ad-exposed DAU x queries/day x commercial query mix x monetized commercial query rate x ad units per monetized commercial query x 365 x eCPM / 1,000.
Transaction formula	Addressable online GMV x AI-assisted GMV share x closed-loop attribution/conversion rate x effective AI take rate.
AI intent monetization	AI advertising TAM + AI transaction commission TAM. This captures the shift from monetizing attention to monetizing commercial intent.

Source: Morgan Stanley Research

Model 1: Subscription – Selective monetization for heavy users

US players popularized freemium: basic access is free and power users pay for advanced models, higher limits, faster responses or specialized generation. China follows a similar free-base / paid-premium structure, but distribution through large ecosystems, consumer campaigns and daily-life use cases has made free access more effective as a mass-adoption tool.

Subscription is the most straightforward monetization model for AI products, but it is unlikely to become the dominant mass-market revenue engine in China in the near term.

- **US** AI players have largely adopted a freemium model: basic features are free, while users pay for more advanced models, larger context windows, deeper reasoning, faster response, image or video generation, and higher usage limits. This model works best when a meaningful share of users perceive AI as a productivity tool and are willing to pay for higher-quality output.
- **China** is different. Consumer AI apps have been free from the outset, and users have been conditioned to expect free access to basic chatbot functions.

The strongest payment levers remain higher accuracy, faster response and advanced features - all of which are associated with higher inference costs. We therefore continue to model subscription/pay-per-use as selective monetization for heavy users and prosumers rather than the mass-market consumer revenue engine.

Our model therefore treats subscription and pay-per-use primarily as a compute-cost sharing mechanism. The formula is:

AI MAU x paying user ratio x monthly ARPPU x 12

By 2040E, we assume 1,073mn AI MAU, a 3.2% paying user ratio and Rmb81 monthly ARPPU, implying Rmb33bn subscription/pay-per-use SAM. This represents only 2.1% of total 2C AI SAM.

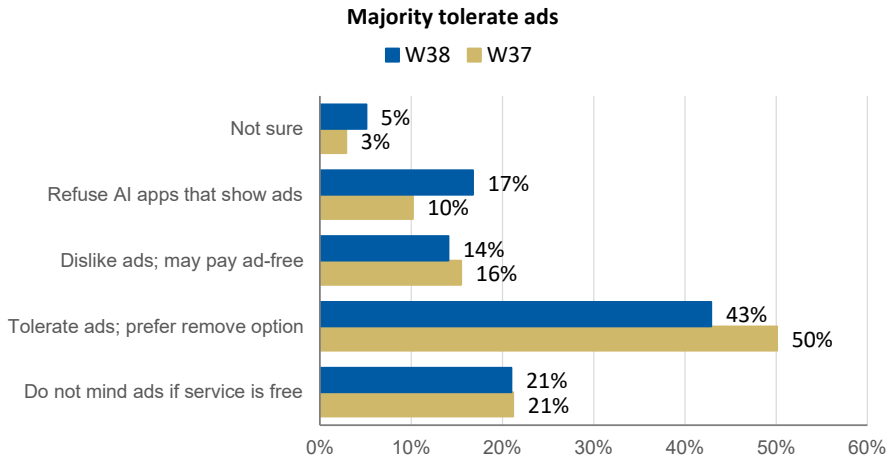
Model 2: Advertising – Supported by high ad tolerance

AI's impact on advertising is more about "reallocation + efficiency gains"

The total size of the traditional internet advertising market is fundamentally constrained by corporate marketing budgets, total consumer GMV, advertising ROI, macro consumption trends, and merchants' profit margins. The emergence of AI, at least in the short term, is unlikely to cause companies to materially increase their overall advertising budgets. Rather, it means that over time, a meaningful share of existing advertising revenue may shift from "traditional keyword or feed-based ads" to "AI/intent-driven ads."

Advertising appears more naturally suited to China's consumer internet market, given users' relatively higher tolerance for ads and the maturity of China's advertising ecosystem. ~64% of our survey respondents either do not mind ads in a free AI service or are willing to tolerate a certain level of advertising if an ad-free option exists; only ~17% indicated they would avoid AI apps that display ads. Ad tolerance has softened slightly in the past three months, but remains high in absolute terms.

Exhibit 23: Attitudes toward ads or sponsored results in AI tools - Majority tolerate ads



Source: AlphaWise, Morgan Stanley Research. W38 is conducted in June 2026, W37 is conducted in March 2026

However, The total size of the internet advertising market is ultimately constrained by corporate marketing budgets, consumer GMV, advertising ROI, macro consumption trends and merchants' profit margins. AI does not automatically cause advertisers to spend more. **In the near term, AI is more likely to reallocate advertising budgets and improve efficiency than create a completely new advertising TAM.**

We see three mechanisms.

- **AI can improve advertising ROI on existing platforms.** Better targeting, ranking, creative generation and conversion optimization can allow platforms to increase ad prices and capture budget share from peers. This is especially relevant for platforms with large existing ad businesses such as ByteDance, Tencent, Kuaishou,

Baidu and Alibaba.

- **AI can turn previously hard-to-monetize user intent into commercial queries.** Traditional search and feed ads monetize explicit keywords or inferred interest. AI conversations can surface longer, more complex and more context-rich intent. A user asking an AI assistant to plan a family trip, compare skincare products, choose a laptop, or identify the best hotel near a theme park is revealing high-quality purchase intent. This represents not only additional advertising inventory, but also higher-quality intent capture.
- **AI can shorten the path from intent to transaction.** As AI moves from answering questions to recommending products, comparing options and completing purchases, monetization shifts from CPM and CPC toward CPA, CPS, affiliate revenue and transaction commissions. **This is where the distinction between advertising and commissions begins to blur.**

The more attractive opportunity lies in intent-based advertising. Under this model, the AI platform does not need to maximize ad impressions. It needs to identify high-value commercial intent and match users with relevant merchants, products or services. A single AI-assisted travel-planning query, for example, may be more valuable than multiple low-intent feed impressions if it leads to hotel a booking, flight purchase or local-service transaction.

We therefore model AI advertising using commercial intent rather than raw time spent:

ad-exposed DAU x AI interactions/day x commercial query mix x monetized commercial query rate x ad units per monetized query x eCPM x 365

By 2040E, our base case assumes 320mn ad-exposed DAU, 13 queries per day, 36.5% commercial query mix, 63% monetized commercial query rate, 4.8 ad units per monetized query and Rmb190 eCPM. This yields Rmb318bn AI advertising SAM, or 7.2% of China online advertising TAM.

This is deliberately more conservative than assuming AI becomes a another feed. The value of AI advertising comes from higher-quality intent signals rather than very high ad load. A travel-planning, product-comparison or food-delivery query may be worth significantly more than a passive feed impression, but the number of sponsored units per answer has a natural ceiling.

For advertising, AI is more likely to improve the monetization of existing traffic than create a new inventory pool. Better intent capture and conversion can support higher pricing and faster share gains.

Model 3: Transaction commissions – The most attractive long-term monetization path

Transaction commissions may offer the most attractive long-term monetization path for consumer AI, because they connect revenue directly to user intent and completed transactions. Unlike subscriptions, commissions do not require users to pay directly for AI. Unlike display advertising, commissions do not depend on high ad load. Instead, the platform monetizes when AI helps users make or complete a purchase, and this is particularly valuable in categories with high information friction, frequent comparison, or fragmented supply.

Our AlphaWise survey already shows evidence of this behavior. 50% of respondents always or often use AI recommendations for purchases, including 9% who always use them and 41% who often use them. Usage is higher among prime-age and higher-income consumers: 53% of consumers aged 30-39 and 61% of consumers earning above Rmb20k/month always or often use AI recommendations for purchases.

Sizing AI transaction commission SAM: The key distinction is between AI-assisted GMV and AI-attributed closed-loop GMV. Many users may consult AI before purchasing, but not every purchase can be monetized by the AI platform. Commission revenue should apply only when the AI flow can be attributed to a completed order, booking or payment. Our formula is addressable online GMV **x** AI-assisted GMV share **x** closed-loop attribution/capture rate **x** effective AI take rate.

By 2040E, our base case includes ~Rmb1.26tn of AI-addressable commission SAM, led by e-commerce, local services and OTA/travel. Much of this represents existing platform economics shifting into AI-assisted journeys; the incremental upside comes from higher conversion, transaction frequency and take rates than under a no-AI baseline.

China's leading AI apps are already moving in this direction.

- **Doubao** embeds Douyin Mall product cards in replies and can potentially earn commissions from transactions completed through Douyin Mall.
- **Qwen** is integrated with Alibaba's ecosystem, including Taobao, Taobao Instashopping, Alipay, Fliggy, Freshippo, Amap and Damai, giving it a stronger path from question to action.
- **Yuanbao**'s standalone app remains less differentiated, but a future **WeChat AI agent** could connect AI intent with Tencent's social graph, mini-programs, payments and merchant ecosystem (currently Trip.com, Tongcheng, Meituan, JD.com, KFC, etc. have announced partnerships with WeChat AI agent).

However, transaction monetization is harder than advertising because it requires more than intent recognition. It requires product catalogs, real-time pricing, inventory, merchant relationships, payment, fulfillment, customer service, attribution and after-sales support. **Standalone chatbots usually lack these capabilities.** They can recommend products, but they cannot easily guarantee price accuracy, complete checkout, resolve disputes or manage the transaction lifecycle.

This is why the reverse path may be more practical: Instead of standalone AI apps trying to become e-commerce platforms – e-commerce, OTA, local service, payment and super-app platforms can **embed AI deeply into their native transaction flows.** In this model, AI

becomes a smart shopping guide inside an existing commercial ecosystem. The platform already owns user identity, transaction data, merchant supply, payment, logistics and after-sales service, making the monetization loop much more natural.

The boundary between advertising and commission will likely become increasingly blurred. A sponsored product recommendation inside an AI answer may begin as an ad, but if the user completes the purchase through an embedded checkout flow, the monetization may become CPA, CPS, or commission. Strategically, both models represent the same shift: from monetizing attention to monetizing intent.

Stock Implications

Tencent — WeChat's moat should be enhanced, with AI adding monetization upside

Tencent is one of the clearest positive read-throughs in our consumer-AI framework. We do not expect Doubao or other AI-native apps to displace WeChat's social operating layer: the social graph, messaging, Official Accounts, Video Accounts, Mini Programs, payments and merchant services create high switching costs and a proprietary execution environment. Yuanbao remains behind peers in standalone scale and differentiation, but its strategic value increasingly lies in generating consumer feedback and reusable capabilities for Hy and Tencent's broader product stack. The more important opportunity is Xiaowei inside WeChat, which can connect user intent to WeChat content, native functions and Mini Program services without requiring users to build a new standalone habit.

AI can also translate into revenue upside through the existing WeChat ecosystem. As LLM-powered Weixin Search expands query coverage and identifies more commercial intent, Tencent can monetize a larger pool of explicit-intent traffic through sponsored search, product recommendations and qualified leads. Longer term, if Xiaowei converts recommendations into attributable Mini Programs, Mini Shops or Weixin Pay transactions, Tencent could participate in referral or commission-sharing, payment and merchant-service economics. We would not assume a marketplace-style take rate across all Mini Program GMV; net capture will depend on whether Tencent controls the transaction handoff, attribution and partner economics. The investment case is therefore a combination of moat protection, advertising-yield uplift and longer-dated transaction optionality, rather than standalone Yuanbao share alone.

The caveat is ROI discipline. Tencent's new AI products and compute procurement are significant investments, and the company separately disclosed lower operating profit when including New AI Products. Encouragingly, management expects Xiaowei's cost envelope to remain below the investment in Yuanbao over the past year and described Xiaowei costs as manageable, and expects returns to outweigh investment as the product experience improves. This is exactly the allocation principle we think makes sense: spend more where AI deepens an existing moat and monetization loop, and maintain a higher bar for standalone chatbot share.

Vertical leaders – Domain data and execution rights are the real moat

Our thesis is constructive for vertical leaders because a general model does not automatically own the workflow.

- Meituan knows real-time restaurant/local-service supply and can execute delivery;
- Taobao owns product and merchant data and a large advertising/payment/fulfillment stack;
- OTA platforms own inventory and booking workflows;
- recruitment marketplaces, such as BOSS Zhipin, own job/candidate supply, matching context and employer workflows.

In each case, AI can be an interface upgrade on top of a proprietary data and execution moat.

The ROI can be measured through conversion uplift, lower search friction, better matching, higher ad yield, merchant productivity or service automation. These are economically more attractive than paying to acquire a user into a new assistant, then paying inference costs, then trying to recreate the vertical supply and checkout layer.

Alibaba – Positive on commerce-embedded AI, more selective on standalone Qwen apps spend

Alibaba has arguably the strongest consumer transaction-loop advantage among the AI-native app owners. We view Qwen Shopping Assistant inside Taobao and merchant agents positively. These investments improve an existing commercial funnel.

However, the economics of heavy standalone Qwen-app investment warrant continued discipline. We see weak post-subsidy retention/engagement: Qwen's DAU peaked at 74mn on Feb 7, 2026, and has now stabilized at 29mn (March 2026), which is still well-above pre-campaign levels. However, Qwen's user engagement metrics significantly lag peer levels – average daily time spent per user was 4min (vs. peers at 10-17min), monthly sessions was 17 (vs. Doubao at 69, Deepseek at 61, Yuanbao at 29), and DAU to MAU was 17% (vs. Deepseek at 25% and Doubao at 42%).

Baidu – AI search threatens the incumbent search distribution layer

Our more cautious view on AI from a consumer perspective is about search disruption: answer engines and AI assistants can satisfy informational intent without a traditional results page, potentially shifting query share and reducing the inventory available to legacy search advertising. Baidu's online marketing revenue fell 19% YoY in 2Q26 while management prioritized AI-search experience over monetization, illustrating the near-term transition pressure.

TME – AI remix music damages the moat built by copyright

TME can use AI for playlisting, discovery, engagement, and efficiency, which should improve user experience and stickiness. The disruption risk, however, lies on the supply side: AI music tools lower the cost and skill threshold for music remixing and creation. With weak IP protection (or regulatory change lagging far behind industry developments), emerging AI-generated music (especially remixes of existing songs) could hit the market and take share from licensed content. Put another way, AI could become a powerful tool for music piracy in China. This would hurt the moat music platforms have built through decades of heavy investment in copyrights, so market leader TME could suffer significantly, as evidenced by churn at the lower end of the music market.

iQIYI – AI video creation lowers production barriers and intensifies competition for attention

iQIYI can benefit materially from lower production costs and faster workflows. But the same tools are available to competitors and creators. AI-generated video sharply lowers the cost and time of content creation, expands supply across short-form and long-form formats, and increases competition for consumer time. This can fragment viewing share and weaken the scarcity value of traditional production capabilities. We therefore see a real cost-efficiency benefit but a more negative competitive read-through as AI broadens

video supply and pressures market share.

Key Insights from AlphaWise Survey

AI usage & behavior: Adoption broadens beyond early adopters. 80% of respondents report at least weekly personal AI use, and 77% of employed respondents use AI weekly at work, up a respective ~6ppt and ~3ppt from Wave 37; about one-third now use AI daily. China remains well ahead of the US in June 2026, by ~29ppt for personal weekly usage and ~25ppt for work.

Cohort gaps are narrowing. Weekly personal usage among ages 40–49 and 50–59 rose to a respective 81% and 72% (from 65% and 56%); work usage for the same age groups rose to 80% and 74%. Higher-income users remain the most engaged.

Multi-homing deepens as AI becomes a routine layer. 64% used five or more AI tools in the past month versus 32% in Wave 37. Different features/task strengths remain the main reason for multi-tool usage, but experimentation and output-comparison reasons have declined, suggesting behavior is becoming more purposeful.

Doubao widens its consumer lead. Doubao reached 77% past-month usage; 49% named it their most frequently used AI tool, and 68% placed it in their top three. Qwen, Yuanbao and DeepSeek follow at 40%, 35% and 28% past-month usage.

Use cases are maturing from novelty toward outcomes. Efficiency (55%), better/creative results (50%) and better decisions (49%) are the top reasons for AI use. Privacy/security remains the top barrier among infrequent/non-users (64%), while lack of know-how is less prominent than in Wave 37.

AI-assisted shopping: In-app AI is showing measurable commerce utility. 64% of consumers used AI to research/compare products, and ~45% completed a purchase with AI support. 75% find AI search/chatbot features inside shopping apps useful, while 57% say these features would make them purchase more often.

Perceived usefulness converts into higher purchase frequency. Among consumers who find shopping-app AI useful, ~72% say they would purchase more often. High-income users show the strongest response: 91% find the features useful, and 80% say they would buy more often.

Consumer subscription monetization remains challenging. Paid penetration softened to 35%, from 41% in Wave 37, while average maximum willingness-to-pay stayed around Rmb41/month. Only 21% prefer subscription; one-time purchase and pay-per-use are similarly important.

Ads and high-intent transactions may be the more scalable monetization path. 64% still do not mind or can tolerate ads in AI tools, despite softer ad tolerance versus Wave 37. Taken together with the shopping results, the survey supports merchant-funded monetization around sponsored recommendations, advertising/lead generation and transaction-linked economics rather than broad consumer subscriptions.

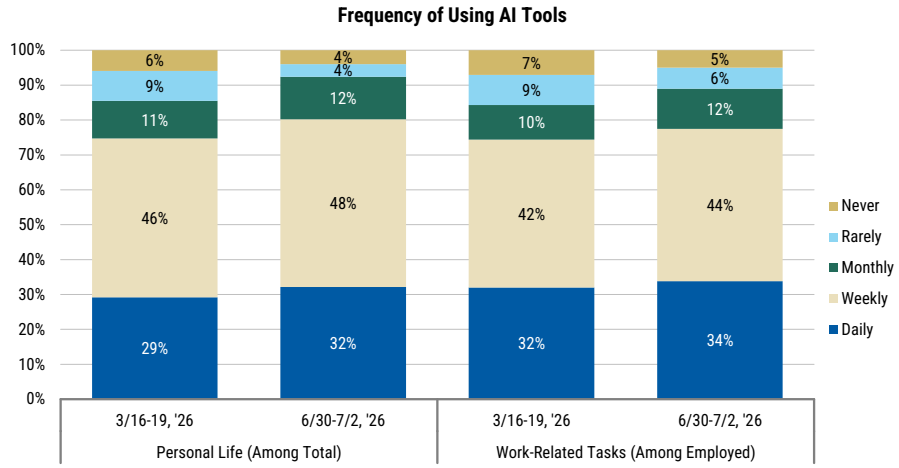
Consumers remain constructive on the AI outlook. About half expect their personal and work AI usage to increase further. 60% believe future generations will be better off because of AI, versus only ~5% who expect them to be worse off.

1. AI Adoption & Engagement

1.1 AI usage continues to deepen, with broader cohort participation

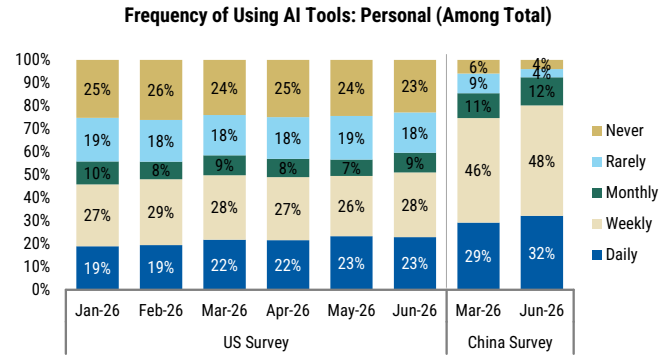
- **Weekly use moves higher.** At-least-weekly personal usage rose to 80% from 75% in W37; work-related weekly usage rose to 77% from 74%. Daily usage is now ~32% in personal life and ~34% at work.
- **Growth is increasingly broad-based.** The largest gains came from ages 40–59 and blue-collar users rather than the already-high 30–39 cohort.

Exhibit 24: ~80% use AI at least weekly in personal life and ~77% at work; daily usage is at about one-third



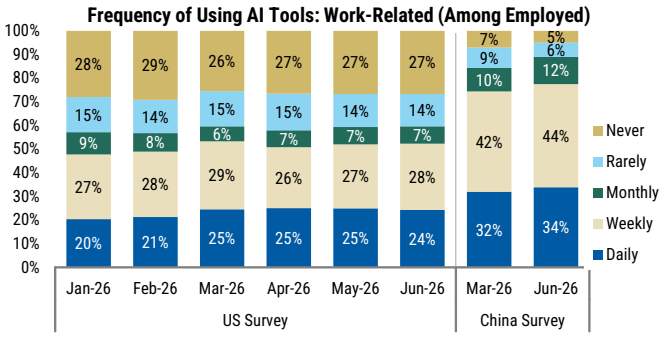
Source: AlphaWise, Morgan Stanley Research.

Exhibit 25: Personal AI use: China vs. US



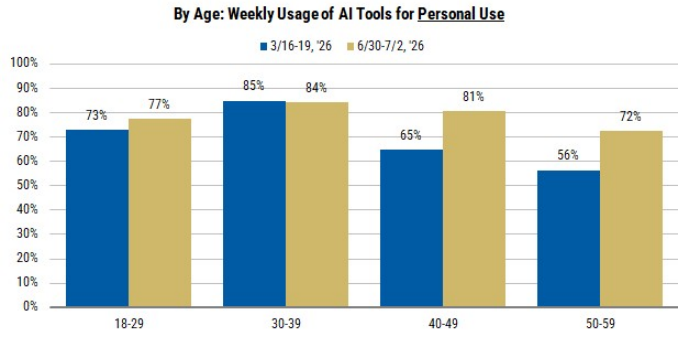
Source: AlphaWise, Morgan Stanley Research.

Exhibit 26: Work-related AI use: China vs. US



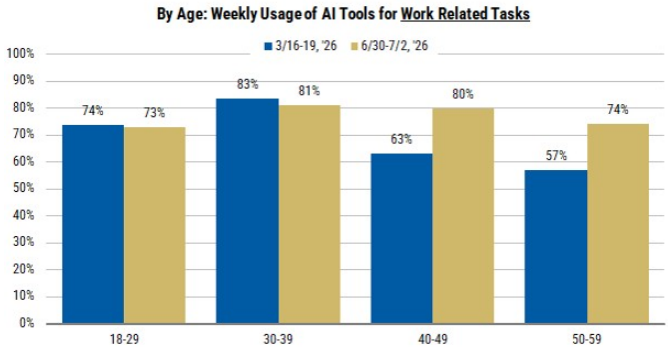
Source: AlphaWise, Morgan Stanley Research.

Exhibit 27: By age: weekly personal AI use



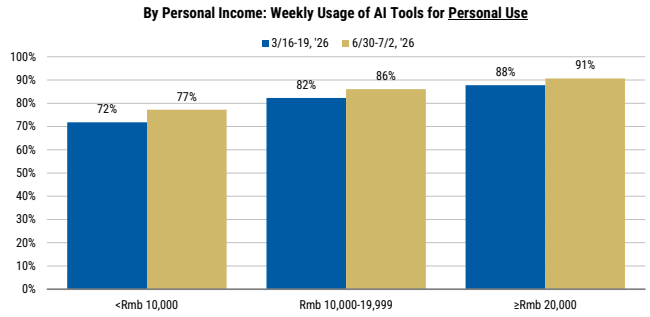
Source: AlphaWise, Morgan Stanley Research.

Exhibit 28: By age: weekly work-related AI use



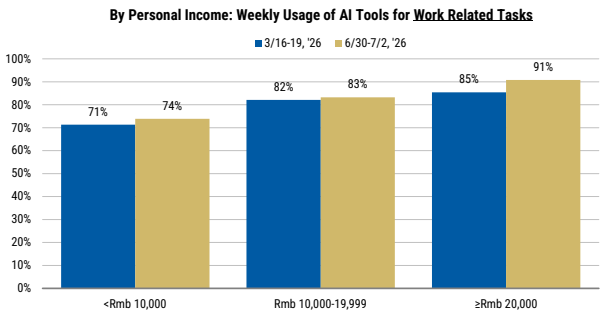
Source: AlphaWise, Morgan Stanley Research.

Exhibit 29: By income: weekly personal AI use



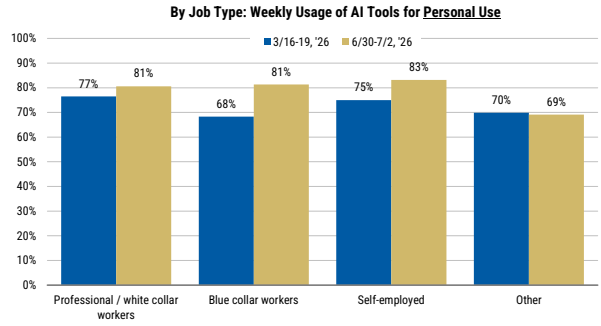
Source: AlphaWise, Morgan Stanley Research.

Exhibit 30: By income: weekly work-related AI use



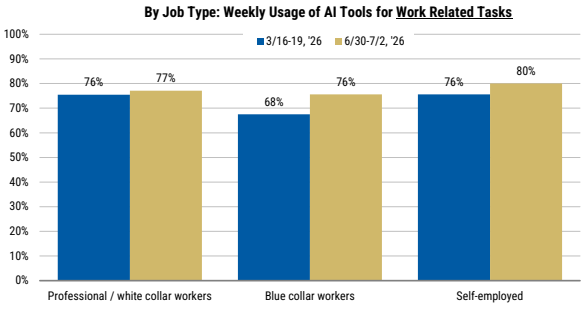
Source: AlphaWise, Morgan Stanley Research.

Exhibit 31: By job type: weekly personal AI use



Source: AlphaWise, Morgan Stanley Research.

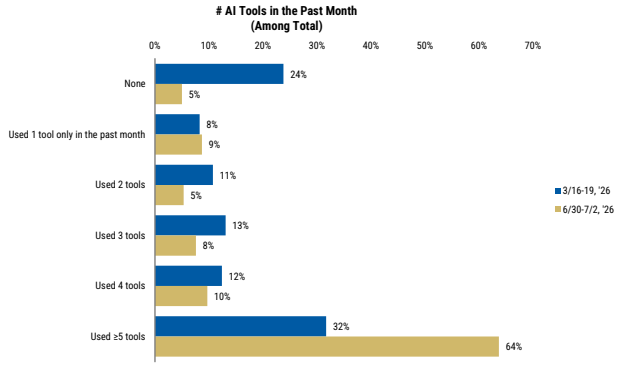
Exhibit 32: By job type: weekly work-related AI use



Source: AlphaWise, Morgan Stanley Research.

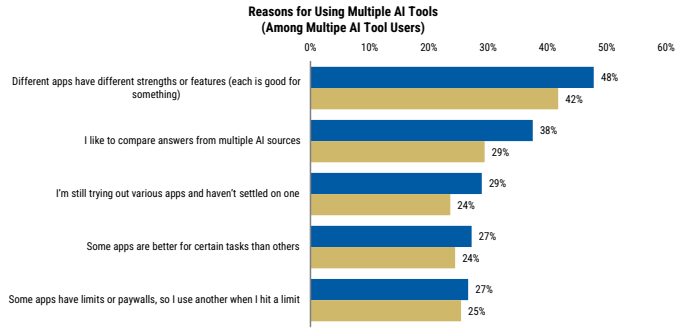
1.2 Multi-homing deepens; behavior is becoming more outcome-driven

Exhibit 33: 64% used five or more AI tools in the past month



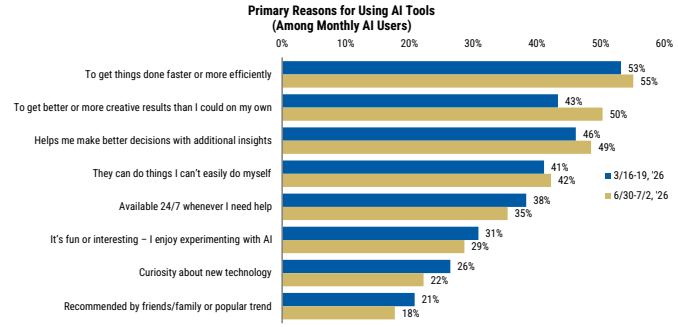
Source: AlphaWise, Morgan Stanley Research.

Exhibit 34: Multi-tool usage remains driven by differentiated features and task fit



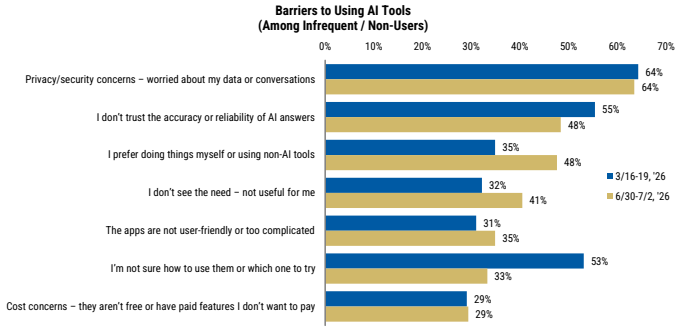
Source: AlphaWise, Morgan Stanley Research.

Exhibit 35: Efficiency, creative quality and decision support are the top usage drivers



Source: AlphaWise, Morgan Stanley Research.

Exhibit 36: Privacy remains the top barrier; "don't need it" and preference for non-AI tools rise



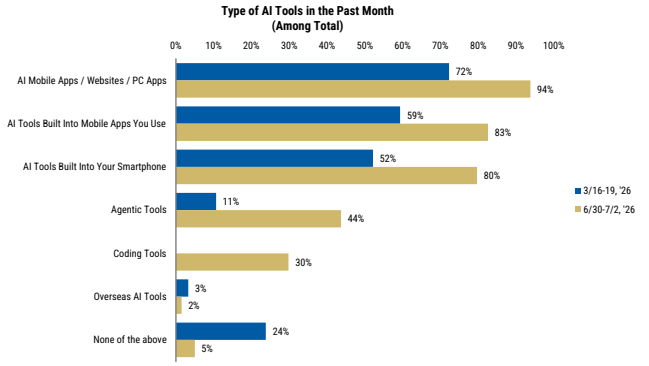
Source: AlphaWise, Morgan Stanley Research.

2. Competitive Landscape and Use Cases

2.1 Doubao extends its lead, while AI becomes embedded across more entry points

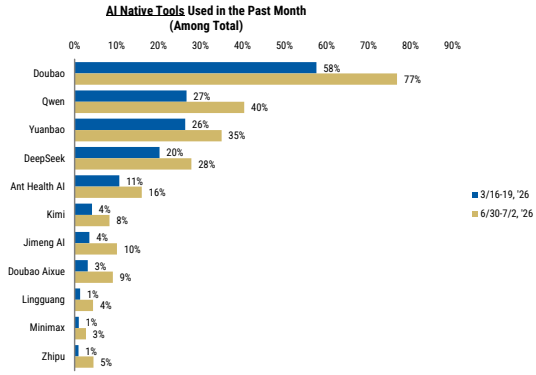
W38 shows broad gains across AI-native apps, in-app AI and on-device assistants. Some of the aggregate increase reflects a broader W38 tool list, but the direction is consistent: AI is becoming a distributed feature layer rather than a single-app market.

Exhibit 37: Usage broadened across AI-native, in-app and on-device categories



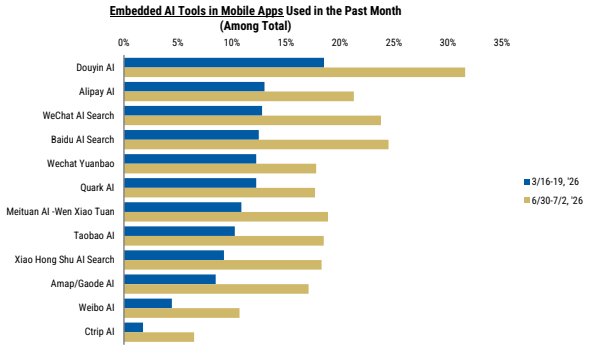
Source: AlphaWise, Morgan Stanley Research. .

Exhibit 38: Doubao widens the lead among AI-native tools



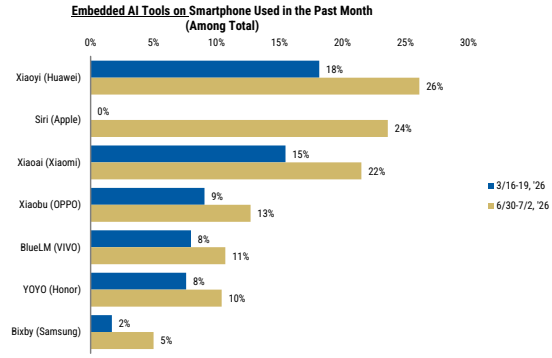
Source: AlphaWise, Morgan Stanley Research.

Exhibit 39: In-app AI assistant usage rises across major mobile apps



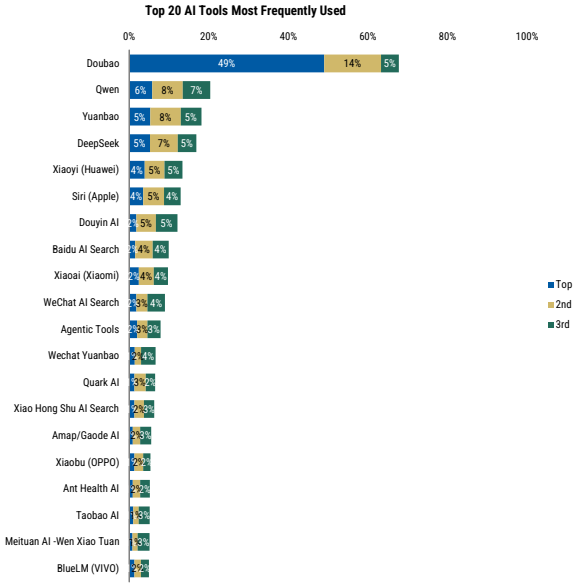
Source: AlphaWise, Morgan Stanley Research.

Exhibit 40: On-device AI usage broadens across smartphone ecosystems



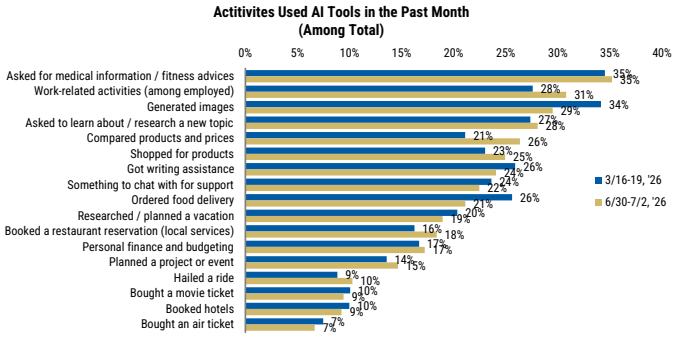
Source: AlphaWise, Morgan Stanley Research.

Exhibit 41: Doubao dominates top-three most frequently used AI tools



Source: AlphaWise, Morgan Stanley Research.

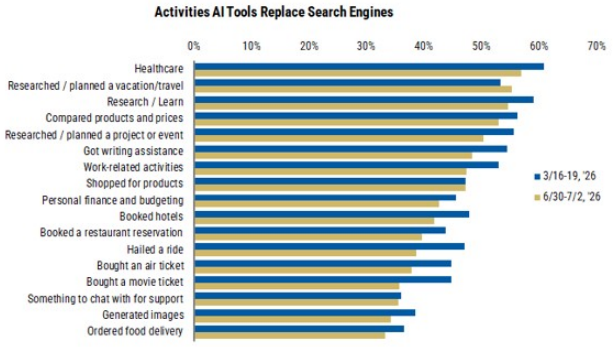
Exhibit 42: Health, work and image generation remain leading past-month use cases



Source: AlphaWise, Morgan Stanley Research.

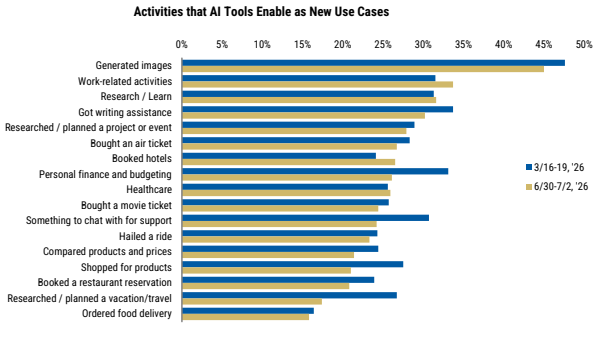
2.2 AI still displaces conventional search most strongly in information discovery, while creation and work remain the most important incremental use cases

Exhibit 43: Search substitution remains strongest in healthcare, travel and research



Source: AlphaWise, Morgan Stanley Research.

Exhibit 44: Image generation and work are the leading "new-use-case" categories



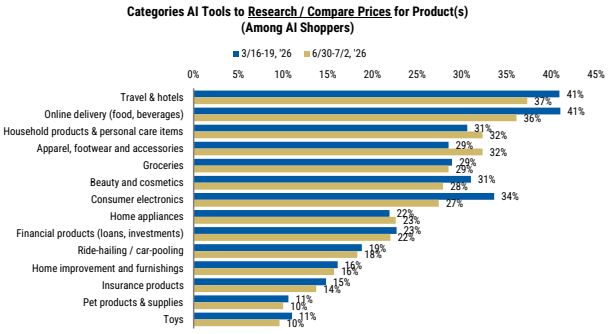
Source: AlphaWise, Morgan Stanley Research.

3. Monetization: AI-Assisted Shopping

3.1 AI is increasingly part of the consumer decision funnel, but transactions remain concentrated in high-frequency categories

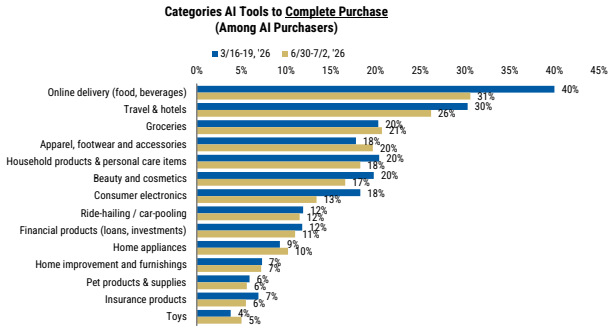
The share of consumers using AI to research/compare products rose to ~64% from ~58%, while the share completing a purchase with AI remained ~45%. Food delivery and travel remain the clearest transaction categories.

Exhibit 45: AI shopping research is broadest in travel, delivery and retail categories



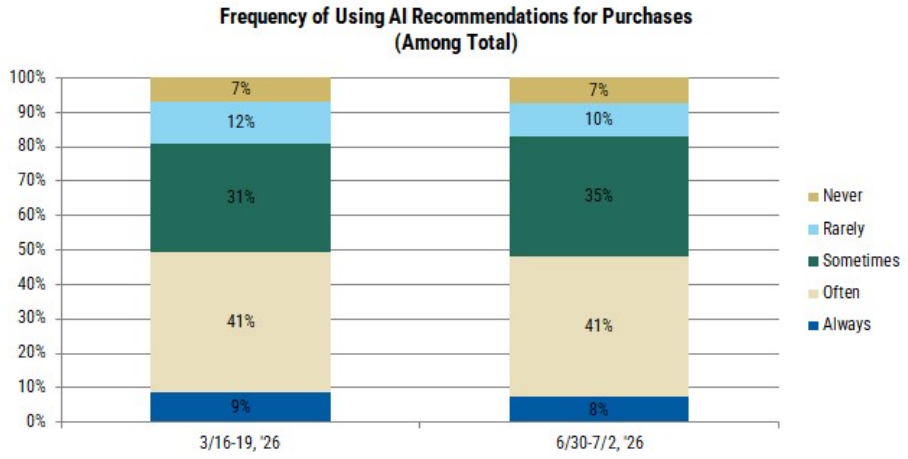
Source: AlphaWise, Morgan Stanley Research.

Exhibit 46: AI-supported transactions remain concentrated in delivery and travel



Source: AlphaWise, Morgan Stanley Research.

Exhibit 47: 48% always/often use AI recommendations for purchases

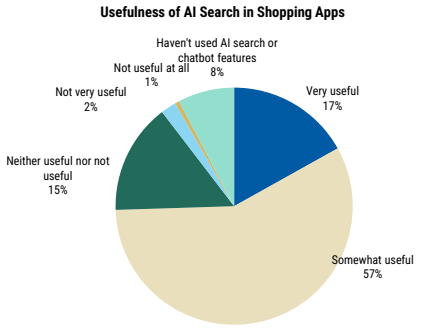


Source: AlphaWise, Morgan Stanley Research.

3.2 In-App AI Shows Clear Commerce Utility

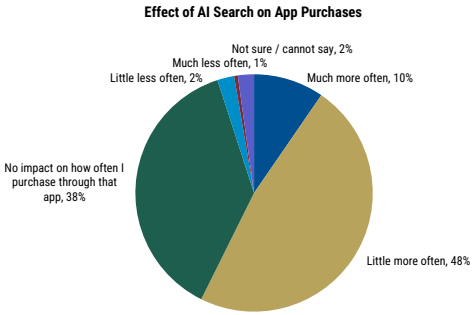
Our AlphaWise survey results strengthen the case for in-app AI: usefulness translates into higher purchase frequency. High-income consumers show the strongest response: ~91% find shopping-app AI useful, and ~57% say it would increase purchase frequency.

Exhibit 48: 75% find AI search/chatbot features inside shopping apps useful



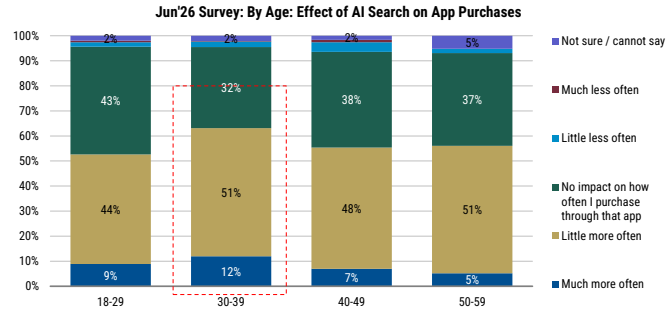
Source: AlphaWise, Morgan Stanley Research.

Exhibit 49: 57% say in-app AI would make them purchase more often



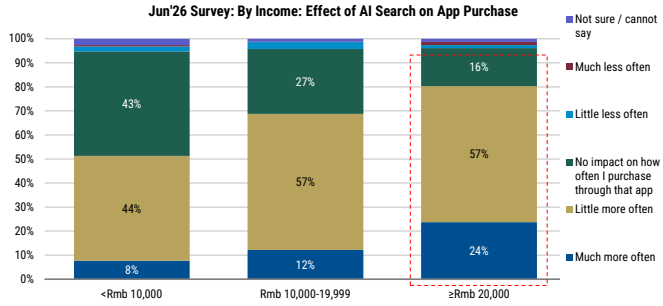
Source: AlphaWise, Morgan Stanley Research.

Exhibit 50: By age: Effect of AI search on purchases



Source: AlphaWise, Morgan Stanley Research.

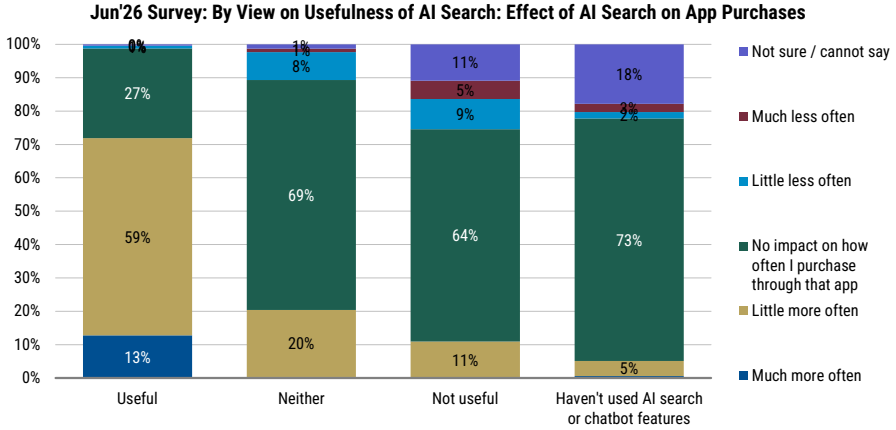
Exhibit 51: By income: Effect of AI search on purchases



Source: AlphaWise, Morgan Stanley Research.

Exhibit 52:

Among users who find shopping AI useful, ~72% say they would purchase more often

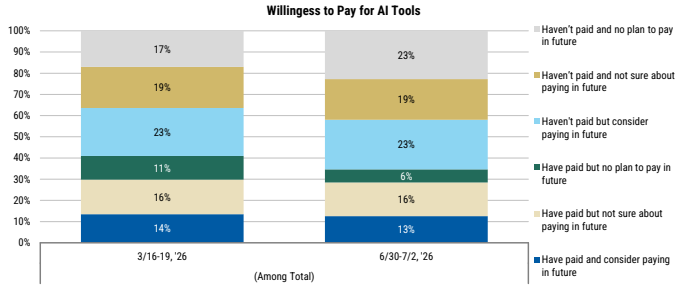


Source: AlphaWise, Morgan Stanley Research.

3.3 Broad consumer subscription remains difficult to scale

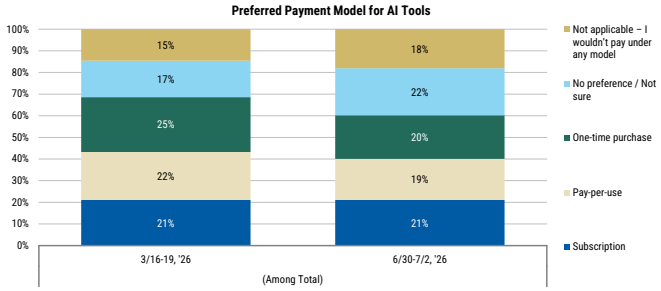
Paid penetration fell to ~35% from 41% in W37, while the share considering payment remains ~36%. Average maximum willingness to pay (WTP) is broadly stable at ~Rmb41/ month. Subscription is only one of several payment models, rather than the clear default.

Exhibit 53: Paid penetration softened to 35%; another 36% would consider paying



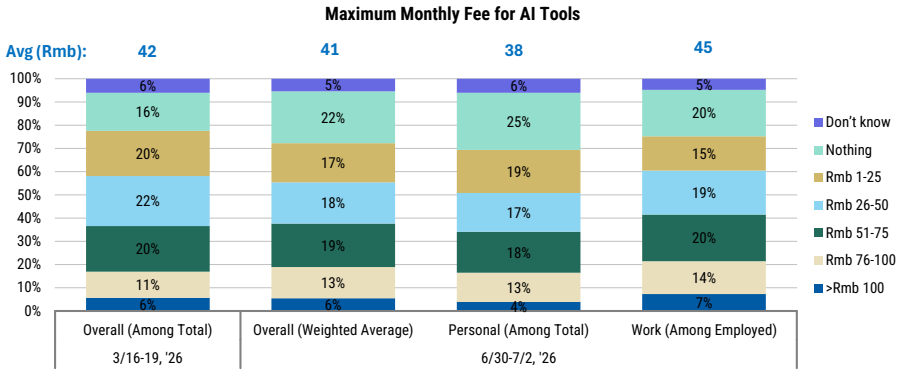
Source: AlphaWise, Morgan Stanley Research.

Exhibit 54: Subscription is only one of several preferred payment models



Source: AlphaWise, Morgan Stanley Research.

Exhibit 55: Average maximum willingness-to-pay remains ~Rmb41/month



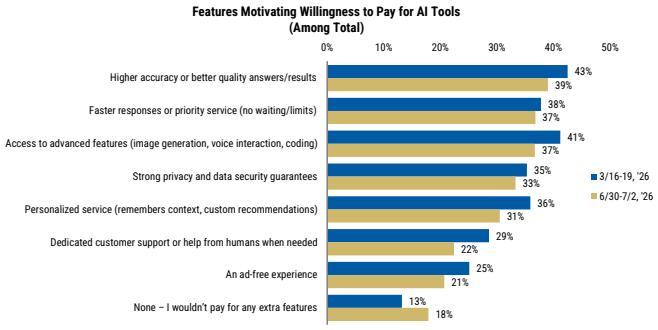
Source: AlphaWise, Morgan Stanley Research.

3.4 Ads and Intent May Scale Better Than Subscription

The survey still points to a sizeable ad-compatible audience, although tolerance softened in W38. Taken together with the shopping results, this supports merchant-funded monetization around sponsored recommendations, advertising/lead generation and transaction-linked economics.

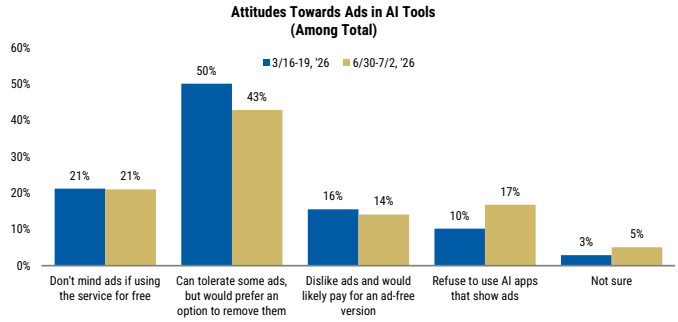
Our read-through: For consumer AI, mass-market subscription looks hard to scale. A more natural model is to keep the core service free, monetize high-intent moments through ads/ sponsored recommendations and commissions, and reserve paid tiers for power users or compute-heavy features.

Exhibit 56: Accuracy, speed and advanced features are the main levers to unlock payment



Source: AlphaWise, Morgan Stanley Research.

Exhibit 57: 64% still accept/tolerate ads, though ad tolerance softened



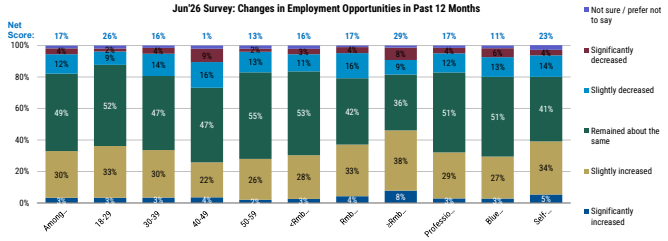
Source: AlphaWise, Morgan Stanley Research.

4. Employment and AI Outlook

4.1 Remain constructive on AI's labor-market and long-term impact

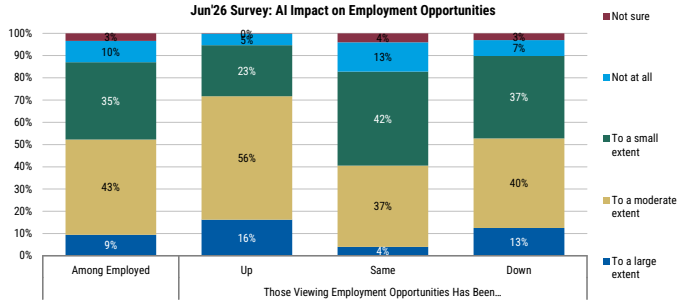
Net sentiment is positive: 33% saw job opportunities rise vs. 16% decline, and 60% expect future generations to be better off compared with ~5% worse off.

Exhibit 58: Employment opportunities are perceived as net improving over the past year



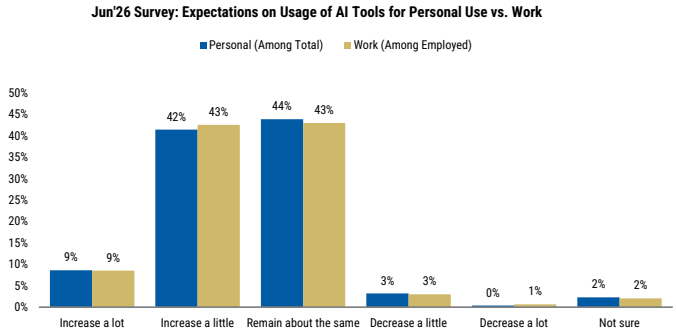
Source: AlphaWise, Morgan Stanley Research.

Exhibit 59: AI is seen as a meaningful contributor to changing employment opportunities



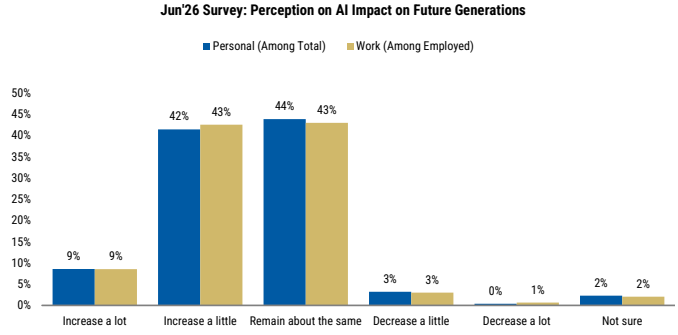
Source: AlphaWise, Morgan Stanley Research.

Exhibit 60: ~50% expect to increase AI use further in personal life and at work



Source: AlphaWise, Morgan Stanley Research.

Exhibit 61: 60% think future generations will be better off because of AI



Source: AlphaWise, Morgan Stanley Research.

Appendix

History of AI-native Apps

Doubao / ByteDance

- Aug 2023: Doubao app launched with text Q&A, summarization, translation and polishing; subsequent releases moved quickly into voice, image and customizable agents.
- May 2024: ByteDance unified the consumer-app and underlying-model brand around Doubao and sharply reduced API pricing, while marketing spend accelerated.
- Early 2025: ByteDance strengthened foundational research and separated model research from application/growth responsibilities, supporting faster product iteration.
- 2H25: Organizational changes put more emphasis on DAU growth, product commercialization and deeper integration with ByteDance products; Doubao also connected to Douyin Mall for shopping-related recommendations.
- Late 2025–2026: Doubao continued to scale through Douyin/Toutiao distribution, multimodal product iteration and broad mass-market positioning.
- 2026 CNY: Volcano Engine was the CCTV Gala AI cloud partner; Doubao continued to grow after the holiday rather than simply reverting to the pre-campaign base.
- Jun 24, 2026: Doubao Professional formally launched, based on Doubao 2.1, with Rmb68/Rmb200/Rmb500 monthly tiers and a continuing free version.
- Jul 2026 QuestMobile: ~399mn MAU, ~168mn DAU and ~42% DAU/MAU, maintaining a clear scale lead.

Qwen / Alibaba

- 2023–1H25: Qwen developed first as a foundation-model family and enterprise/cloud asset, while Tongyi/Quark were explored as consumer AI surfaces.
- Nov 2025: Alibaba reorganized the consumer AI effort and renamed Tongyi app as Qwen, concentrating resources on an AI-native entry point.
- Jan 2026: Qwen integrated Taobao Instant Commerce, Alipay, Freshippo, Amap, Damai and other services, creating a differentiated question-to-action loop.
- Feb 2026 CNY: Large consumer subsidies drove trial and transaction-oriented usage, but post-subsidy engagement remained the key issue to watch.
- May 2026: Qwen Shopping Assistant launched inside Taobao, spanning inspiration, search/recommendation and after-sales; June-quarter management described rapid adoption.
- Aug 3, 2026: Qwen3.8-Max launched as a 2.4T-parameter flagship focused on coding, office productivity and long-horizon tasks.
- Aug 2026: Qwen app introduced paid office-assistant tiers at Rmb19/Rmb49/Rmb128 per month; basic access remained free.
- 2Q26: AI Labs and Applications adjusted EBITA loss was Rmb13.9bn, primarily reflecting AI investment and Qwen app inference costs; management said Qwen

marketing expense had come down sequentially.

Tencent

- May 2024: Yuanbao launched with a relatively muted acquisition strategy.
- Feb 2025: Tencent rapidly integrated DeepSeek-R1 across products, which became the major inflection point for Yuanbao awareness and usage.
- Apr 2025: Yuanbao became a Weixin friend, allowing content/messages to be forwarded to the assistant and reducing incremental distribution cost.
- 2025–1H26: Tencent reorganized Hunyuan research and AI infrastructure, while Yuanbao scaled through DeepSeek integration, Weixin distribution and CNY marketing; standalone engagement remained below Doubao and DeepSeek.
- 2Q26: Tencent began small-scale prototype testing of Xiaowei inside Weixin. Xiaowei uses WeLM, a customized model focused on privacy, Weixin-specific use cases and inference efficiency.
- 2Q26 earnings call: Management described Mini Program/merchant agents interacting with user agents over time, creating an architecture for agent-to-agent transactions.
- Aug 28, 2026: Tencent released an open-source Hy4 preview, a 770B-total / 49B-active model with >1M context focused on coding, office and scientific productivity; it launched across WorkBuddy, CodeBuddy, Yuanbao and ima.

Meituan

- Apr 2026: Xiaotuan entry moved to the Meituan homepage, improving distribution without requiring a separate app.
- Jun 2026: Meituan released LongCat 2.0, a trillion-parameter model trained and served on domestic compute, with emphasis on coding and agentic tool use.
- 2Q26: Xiaotuan was upgraded from answer/search to task execution including ordering, ride-hailing and reservations; Meituan also introduced merchant CatPaw agents and a rider safety AI assistant.

DeepSeek

- 2023–2024: DeepSeek was built as a model-first lab backed by High-Flyer, with unusually high engineering focus and little consumer marketing.
- Jan 2025: DeepSeek R1 triggered the breakout moment for Chinese reasoning models and drove rapid app adoption through technical reliability and word-of-mouth.
- Dec 2024–Apr 2026: DeepSeek released V3, R1, V3.2 and then V4, combining MoE architecture, reasoning and increasingly long context while maintaining an efficiency-first engineering philosophy.
- DeepSeek remains the model-first outlier: Consumer scale came primarily from technical prowess and word-of-mouth rather than paid user acquisition.
- Aug 13, 2026: DeepSeek V4-Pro reached general availability (GA) on app, web and API with materially stronger agent capabilities, three reasoning-effort levels and a new peak/off-peak API pricing framework.

This report references U.S. Executive Order 14032 and/or entities or securities that are designated thereunder. U.S. persons may be prohibited from buying certain securities of entities named in this report. Readers are solely responsible for ensuring that their investment activities are carried out in compliance with applicable laws.

This report references export controls and/or entities that may be subject to export control restrictions. Readers are solely responsible for ensuring that their investment or trade activities are carried out in compliance with applicable laws.

This report references U.S. Executive Order 14105 and/or entities that may be in scope of such order. U.S. persons may be prohibited from engaging in certain transactions or otherwise require certain other transactions be notified to the U.S. Department of Treasury. Readers are solely responsible for ensuring that their investment or trade activities are carried out in compliance with applicable laws.

Disclosure Section

The information and opinions in Morgan Stanley Research were prepared or are disseminated by Morgan Stanley Asia Limited (which accepts the responsibility for its contents) and/or Morgan Stanley Asia (Singapore) Pte. (Registration number 199206298Z) and/or Morgan Stanley Asia (Singapore) Securities Pte Ltd (Registration number 200008434H), regulated by the Monetary Authority of Singapore (which accepts legal responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research), and/or Morgan Stanley Taiwan Limited and/or Morgan Stanley & Co International plc, Seoul Branch, and/or Morgan Stanley Australia Limited (A.B.N. 67 003 734 576, holder of Australian financial services license No. 233742, which accepts responsibility for its contents), and/or Morgan Stanley Wealth Management Australia Pty Ltd (A.B.N. 19 009 145 555, holder of Australian financial services license No. 240813, which accepts responsibility for its contents), and/or Morgan Stanley India Company Private Limited having Corporate Identification No (CIN) U22990MH1998PTC115305, regulated by the Securities and Exchange Board of India ("SEBI") and holder of licenses as a Research Analyst (SEBI Registration No. INH000001105); Stock Broker (SEBI Stock Broker Registration No. INZ000244438), Merchant Banker (SEBI Registration No. INM000011203), and depository participant with National Securities Depository Limited (SEBI Registration No. IN-DP-NSDL-567-2021) having registered office at Altimus, Level 39 & 40, Pandurang Budhkar Marg, Worli, Mumbai 400018, India; Telephone no. +91-22-61181000; Compliance Officer Details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: tejarshi.hardas@morganstanley.com; Grievance officer details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: msic-compliance@morganstanley.com which accepts the responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research, and their affiliates (collectively, "Morgan Stanley"). Morgan Stanley India Company Private Limited (MSICPL) may use AI tools in providing research services. All recommendations contained herein are made by the duly qualified research analysts.

For important disclosures, stock price charts and equity rating histories regarding companies that are the subject of this report, please see the Morgan Stanley Research Disclosure Website at www.morganstanley.com/eqr/disclosures/webapp/generalresearch, or contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY, 10036 USA.

For valuation methodology and risks associated with any recommendation, rating or price target referenced in this research report, please contact the Client Support Team as follows: US/Canada +1 800 303-2495; Hong Kong +852 2848-5999; Latin America +1 718 754-5444 (U.S.); London +44 (0)20-7425-8169; Singapore +65 6834-6860; Sydney +61 (0)2-9770-1505; Tokyo +81 (0)3-6836-9000. Alternatively you may contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY 10036 USA.

Analyst Certification

The following analysts hereby certify that their views about the companies and their securities discussed in this report are accurately expressed and that they have not received and will not receive direct or indirect compensation in exchange for expressing specific recommendations or views in this report: Lydia Lin; Yang Liu; Andy Meng, CFA; Tom Tang; Eddy Wang, CFA; Rebecca Xu; Gary Yu.

.

Global Research Conflict Management Policy

Morgan Stanley Research has been published in accordance with our conflict management policy, which is available at www.morganstanley.com/institutional/research/conflictpolicies. A Portuguese version of the policy can be found at www.morganstanley.com.br

Important Regulatory Disclosures on Subject Companies

The analyst or strategist (or a household member) identified below owns the following securities (or related derivatives): Andy Meng, CFA - Alibaba Group Holding(GDR), Kuaishou Technology(common or preferred stock), Meituan(common or preferred stock), NetEase, Inc(GDR), PDD Holdings Inc(GDR), Trip.com Group Ltd(GDR), Vipshop Holdings Ltd(GDR); Kathy Zhu - Meituan(common or preferred stock), Tencent Holdings Ltd.(common or preferred stock).

As of August 31, 2026, Morgan Stanley beneficially owned 1% or more of a class of common equity securities of the following companies covered in Morgan Stanley Research: Alibaba Group Holding, Baidu Inc, Bilibili Inc, Full Truck Alliance Co. Ltd, HUYA Inc, IQIYI Inc, JD.com, Inc., JOYY Inc., Kanzhun Ltd, Kuaishou Technology, Meituan, PDD Holdings Inc, Trip.com Group Ltd.

Within the last 12 months, Morgan Stanley managed or co-managed a public offering (or 144A offering) of securities of Alibaba Group Holding, Kanzhun Ltd, Tencent Holdings Ltd..

Within the last 12 months, Morgan Stanley has received compensation for investment banking services from Alibaba Group Holding, Bilibili Inc, Meituan, Tencent Holdings Ltd..

In the next 3 months, Morgan Stanley expects to receive or intends to seek compensation for investment banking services from Alibaba Group Holding, Autohome Inc, Baidu Inc, Bilibili Inc, Full Truck Alliance Co. Ltd, IQIYI Inc, JD.com, Inc., JOYY Inc., Kanzhun Ltd, KE Holdings Inc, Kuaishou Technology, Meituan, NetEase, Inc, PDD Holdings Inc, Tencent Holdings Ltd., Tongcheng Travel Holdings, Trip.com Group Ltd, Vipshop Holdings Ltd, Weibo Corp.

Within the last 12 months, Morgan Stanley has received compensation for products and services other than investment banking services from Alibaba Group Holding, Baidu Inc, Bilibili Inc, Full Truck Alliance Co. Ltd, HUYA Inc, JOYY Inc., Kanzhun Ltd, KE Holdings Inc, Meituan, NetEase, Inc, Tencent Holdings Ltd., Tongcheng Travel Holdings, Trip.com Group Ltd, Vipshop Holdings Ltd, Weibo Corp.

Within the last 12 months, Morgan Stanley has provided or is providing investment banking services to, or has an investment banking client relationship with, the following company: Alibaba Group Holding, Autohome Inc, Baidu Inc, Bilibili Inc, Full Truck Alliance Co. Ltd, IQIYI Inc, JD.com, Inc., JOYY Inc., Kanzhun Ltd, KE Holdings Inc, Kuaishou Technology, Meituan, NetEase, Inc, PDD Holdings Inc, Tencent Holdings Ltd., Tongcheng Travel Holdings, Trip.com Group Ltd, Vipshop Holdings Ltd, Weibo Corp.

Within the last 12 months, Morgan Stanley has either provided or is providing non-investment banking, securities-related services to and/or in the past has entered into an agreement to provide services or has a client relationship with the following company: Alibaba Group Holding, Baidu Inc, Bilibili Inc, Full Truck Alliance Co. Ltd, HUYA Inc, IQIYI Inc, JOYY Inc., Kanzhun Ltd, KE Holdings Inc, Meituan, NetEase, Inc, Tencent Holdings Ltd., Tongcheng Travel Holdings, Trip.com Group Ltd, Vipshop Holdings Ltd, Weibo Corp.

Morgan Stanley & Co. LLC makes a market in the securities of Autohome Inc, IQIYI Inc, JOYY Inc., NetEase, Inc, Weibo Corp.

The equity research analysts or strategists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality of research, investor client feedback, stock picking, competitive factors, firm revenues and overall investment banking revenues. Equity Research analysts' or strategists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

Morgan Stanley and its affiliates do business that relates to companies/instruments covered in Morgan Stanley Research, including market making, providing liquidity, fund management, commercial banking, extension of credit, investment services and investment banking. Morgan Stanley sells to and buys from customers the securities/instruments of companies covered in Morgan Stanley Research on a principal basis. Morgan Stanley may have a position in the debt of the Company or instruments discussed in this report. Morgan Stanley trades or may trade as principal in the debt securities (or in related derivatives) that are the subject of the debt research report.

Certain disclosures listed above are also for compliance with applicable regulations in non-US jurisdictions.

STOCK RATINGS

Morgan Stanley uses a relative rating system using terms such as Overweight, Equal-weight, Not-Rated or Underweight (see definitions below). Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold and sell. Investors should carefully read the definitions of all ratings used in Morgan Stanley Research. In addition, since Morgan Stanley Research contains more complete information concerning the analyst's views, investors should carefully read Morgan Stanley Research, in its entirety, and not infer the contents from the rating alone. In any case, ratings (or research) should not be used or relied upon as investment advice. An investor's decision to buy or sell a stock should depend on individual circumstances (such as the investor's existing holdings) and other considerations.

Global Stock Ratings Distribution

(as of August 31, 2026)

The Stock Ratings described below apply to Morgan Stanley's Fundamental Equity Research and do not apply to Debt Research produced by the Firm.

For disclosure purposes only (in accordance with FINRA requirements), we include the category headings of Buy, Hold, and Sell alongside our ratings of Overweight, Equal-weight, Not-Rated and Underweight. Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold, and sell but represent recommended relative weightings (see definitions below). To satisfy regulatory requirements, we correspond Overweight, our most positive stock rating, with a buy recommendation; we correspond Equal-weight and Not-Rated to hold and Underweight to sell recommendations, respectively.

	Coverage Universe		Investment Banking Clients (IBC)			Other Material Investment Services Clients (MISC)	
Stock Rating Category	Count	% of Total	Count	% of Total IBC	% of Rating Category	Count	% of Total Other MISC
Overweight/Buy	1555	42%	450	48%	29%	767	44%
Equal-weight/Hold	1577	43%	395	42%	25%	770	44%
Not-Rated/Hold	3	0%	1	0%	33%	1	0%
Underweight/Sell	559	15%	93	10%	17%	216	12%
Total	3,694		939			1754	

Data include common stock and ADRs currently assigned ratings. Investment Banking Clients are companies from whom Morgan Stanley received investment banking compensation in the last 12 months. Due to rounding off of decimals, the percentages provided in the "% of total" column may not add up to exactly 100 percent.

Analyst Stock Ratings

Overweight (O). The stock's total return is expected to exceed the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Equal-weight (E). The stock's total return is expected to be in line with the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Not-Rated (NR). Currently the analyst does not have adequate conviction about the stock's total return relative to the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Underweight (U). The stock's total return is expected to be below the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Unless otherwise specified, the time frame for price targets included in Morgan Stanley Research is 12 to 18 months.

Analyst Industry Views

Attractive (A): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be attractive vs. the relevant broad market benchmark, as indicated below.

In-Line (I): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be in line with the relevant broad market benchmark, as indicated below.

Cautious (C): The analyst views the performance of his or her industry coverage universe over the next 12-18 months with caution vs. the relevant broad market benchmark, as indicated below.

Benchmarks for each region are as follows: North America - S&P 500; Latin America - relevant MSCI country index or MSCI Latin America Index; Europe - MSCI Europe; Japan - TOPIX; Asia - relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Important Disclosures for Morgan Stanley Smith Barney LLC Customers

Important disclosures regarding any material conflict of interest that can reasonably be expected to have influenced Morgan Stanley Smith Barney LLC's choice of a third-party research provider or the subject company of a third-party research report, are available on the Morgan Stanley Wealth Management disclosure website at www.morganstanley.com/online/researchdisclosures.

For Morgan Stanley specific disclosures, you may refer to https://www.morganstanley.com/eqr/disclosures/webapp/generalresearch.

Each Morgan Stanley research report is reviewed and approved on behalf of Morgan Stanley Smith Barney LLC. This review and approval is conducted by the same person who reviews the research report on behalf of Morgan Stanley. This could create a conflict of interest.

Other Important Disclosures

A member of Research who had or could have had access to the research prior to completion owns securities (or related derivatives) in the Alibaba Group Holding, Baidu Inc. This person is not a research analyst or a member of research analyst's household.

Morgan Stanley Research policy is to update research reports as and when the Research Analyst and Research Management deem appropriate, based on developments with the issuer, the sector, or the market that may have a material impact on the research views or opinions stated therein. In addition, certain Research publications are intended to be updated on a regular periodic basis (weekly/monthly/quarterly/annual) and will ordinarily be updated with that frequency, unless the Research Analyst and Research Management determine that a different publication schedule is appropriate based on current conditions.

Morgan Stanley is not acting as a municipal advisor and the opinions or views contained herein are not intended to be, and do not constitute, advice within the meaning of Section 975 of the

Dodd-Frank Wall Street Reform and Consumer Protection Act.

Morgan Stanley produces an equity research product called a "Tactical Idea." Views contained in a "Tactical Idea" on a particular stock may be contrary to the recommendations or views expressed in research on the same stock. This may be the result of differing time horizons, methodologies, market events, or other factors. For all research available on a particular stock, please contact your sales representative or go to Matrix at http://www.morganstanley.com/matrix.

Morgan Stanley Research is provided to our clients through our proprietary research portal on Matrix and also distributed electronically by Morgan Stanley to clients. Certain, but not all, Morgan Stanley Research products are also made available to clients through third-party vendors or redistributed to clients through alternate electronic means as a convenience. For access to all available Morgan Stanley Research, please contact your sales representative or go to Matrix at http://www.morganstanley.com/matrix.

Any access and/or use of Morgan Stanley Research is subject to Morgan Stanley's Terms of Use (http://www.morganstanley.com/terms.html). By accessing and/or using Morgan Stanley Research, you are indicating that you have read and agree to be bound by our Terms of Use (http://www.morganstanley.com/terms.html). In addition you consent to Morgan Stanley processing your personal data and using cookies in accordance with our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html), including for the purposes of setting your preferences and to collect readership data so that we can deliver better and more personalized service and products to you. To find out more information about how Morgan Stanley processes personal data, how we use cookies and how to reject cookies see our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html). Please use the provided link to review the Terms and Conditions and Most Important Terms and Conditions for Morgan Stanley India Company Private Limited (https://www.morganstanley.com/assets/pdfs/about-us-global-offices/india/Terms_and_conditions.pdf) and the following link to review the audit report (https://ny.matrix.ms.com/eqr/research/webapp/researchdocs/MSICPL_Morgan_Stanley_Research_Audit_Report.pdf).

If you do not agree to our Terms of Use and/or if you do not wish to provide your consent to Morgan Stanley processing your personal data or using cookies please do not access our research.

Morgan Stanley Research does not provide individually tailored investment advice. Morgan Stanley Research has been prepared without regard to the circumstances and objectives of those who receive it. Morgan Stanley recommends that investors independently evaluate particular investments and strategies, and encourages investors to seek the advice of a financial adviser. The appropriateness of an investment or strategy will depend on an investor's circumstances and objectives. The securities, instruments, or strategies discussed in Morgan Stanley Research may not be suitable for all investors, and certain investors may not be eligible to purchase or participate in some or all of them. Morgan Stanley Research is not an offer to buy or sell or the solicitation of an offer to buy or sell any security/instrument or to participate in any particular trading strategy. The value of and income from your investments may vary because of changes in interest rates, foreign exchange rates, default rates, prepayment rates, securities/instruments prices, market indexes, operational or financial conditions of companies or other factors. There may be time limitations on the exercise of options or other rights in securities/instruments transactions. Past performance is not necessarily a guide to future performance. Estimates of future performance are based on assumptions that may not be realized. If provided, and unless otherwise stated, the closing price on the cover page is that of the primary exchange for the subject company's securities/instruments.

The fixed income research analysts, strategists or economists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality, accuracy and value of research, firm profitability or revenues (which include fixed income trading and capital markets profitability or revenues), client feedback and competitive factors. Fixed Income Research analysts', strategists' or economists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

The "Important Regulatory Disclosures on Subject Companies" section in Morgan Stanley Research lists all companies mentioned where Morgan Stanley owns 1% or more of a class of common equity securities of the companies. For all other companies mentioned in Morgan Stanley Research, Morgan Stanley may have an investment of less than 1% in securities/instruments or derivatives of securities/instruments of companies and may trade them in ways different from those discussed in Morgan Stanley Research. Employees of Morgan Stanley not involved in the preparation of Morgan Stanley Research may have investments in securities/instruments or derivatives of securities/instruments of companies mentioned and may trade them in ways different from those discussed in Morgan Stanley Research. Derivatives may be issued by Morgan Stanley or associated persons.

With the exception of information regarding Morgan Stanley, Morgan Stanley Research is based on public information. Morgan Stanley makes every effort to use reliable, comprehensive information, but we make no representation that it is accurate or complete. We have no obligation to tell you when opinions or information in Morgan Stanley Research change apart from when we intend to discontinue equity research coverage of a subject company. Facts and views presented in Morgan Stanley Research have not been reviewed by, and may not reflect information known to, professionals in other Morgan Stanley business areas, including investment banking personnel.

Morgan Stanley Research personnel may participate in company events such as site visits and are generally prohibited from accepting payment by the company of associated expenses unless pre-approved by authorized members of Research management.

Morgan Stanley may make investment decisions that are inconsistent with the recommendations or views in this report.

To our readers based in Taiwan or trading in Taiwan securities/instruments: Information on securities/instruments that trade in Taiwan is distributed by Morgan Stanley Taiwan Limited ("MSTL"). Such information is for your reference only. The reader should independently evaluate the investment risks and is solely responsible for their investment decisions. Morgan Stanley Research may not be distributed to the public media or quoted or used by the public media without the express written consent of Morgan Stanley. Any non-customer reader within the scope of Article 7-1 of the Taiwan Stock Exchange Recommendation Regulations accessing and/or receiving Morgan Stanley Research is not permitted to provide Morgan Stanley Research to any third party (including but not limited to related parties, affiliated companies and any other third parties) or engage in any activities regarding Morgan Stanley Research which may create or give the appearance of creating a conflict of interest. Information on securities/instruments that do not trade in Taiwan is for informational purposes only and is not to be construed as a recommendation or a solicitation to trade in such securities/instruments. MSTL may not execute transactions for clients in these securities/instruments.

Certain information in Morgan Stanley Research was sourced by employees of the Shanghai Representative Office of Morgan Stanley Asia Limited for the use of Morgan Stanley Asia Limited. Morgan Stanley is not incorporated under PRC law and the research in relation to this report is conducted outside the PRC. Morgan Stanley Research does not constitute an offer to sell or the solicitation of an offer to buy any securities in the PRC. PRC investors shall have the relevant qualifications to invest in such securities and shall be responsible for obtaining all relevant approvals, licenses, verifications and/or registrations from the relevant governmental authorities themselves. Neither this report nor any part of it is intended as, or shall constitute, provision of any consultancy or advisory service of securities investment as defined under PRC law. Such information is provided for your reference only.

Morgan Stanley Research is disseminated in Brazil by Morgan Stanley C.T.V.M. S.A. located at Av. Brigadeiro Faria Lima, 3600, 6th floor, São Paulo - SP, Brazil; and is regulated by the Comissão de Valores Mobiliários; in Mexico by Morgan Stanley México, Casa de Bolsa, S.A. de C.V which is regulated by Comision Nacional Bancaria y de Valores. Paseo de los Tamarindos 90, Torre 1, Col. Bosques de las Lomas Floor 29, 05120 Mexico City; in Japan by Morgan Stanley MUFG Securities Co., Ltd. and, for Commodities related research reports only, Morgan Stanley Capital Group Japan Co., Ltd; in Hong Kong by Morgan Stanley Asia Limited (which accepts responsibility for its contents) and by Morgan Stanley Bank Asia Limited; in Singapore by Morgan Stanley Asia (Singapore) Pte. (Registration number 199206298Z) and/or Morgan Stanley Asia (Singapore) Securities Pte Ltd (Registration number 200008434H), regulated by the Monetary Authority of Singapore (which accepts legal responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research) and by Morgan Stanley Bank Asia Limited, Singapore Branch (Registration number T14FC0118J); in Australia to "wholesale clients" within the meaning of the Australian Corporations Act by Morgan Stanley Australia Limited A.B.N. 67 003 734 576, holder of Australian financial services license No. 233742, which accepts responsibility for its contents; in Australia to "wholesale clients" and "retail

clients" within the meaning of the Australian Corporations Act by Morgan Stanley Wealth Management Australia Pty Ltd (A.B.N. 19 009 145 555, holder of Australian financial services license No. 240813, which accepts responsibility for its contents; in Korea by Morgan Stanley & Co International plc, Seoul Branch; in India by Morgan Stanley India Company Private Limited having Corporate Identification No (CIN) U22990MH1998PTC115305, regulated by the Securities and Exchange Board of India ("SEBI") and holder of licenses as a Research Analyst (SEBI Registration No. INH000001105); Stock Broker (SEBI Stock Broker Registration No. INZ000244438), Merchant Banker (SEBI Registration No. INM000011203), and depository participant with National Securities Depository Limited (SEBI Registration No. IN-DP-NSDL-567-2021) having registered office at Altimus, Level 39 & 40, Pandurang Budhkar Marg, Worli, Mumbai 400018, India; Telephone no. +91-22-61181000; Compliance Officer Details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: tejarshi.hardas@morganstanley.com; Grievance officer details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: msic-compliance@morganstanley.com. Morgan Stanley India Company Private Limited (MSICPL) may use AI tools in providing research services. All recommendations contained herein are made by the duly qualified research analysts; in Germany and the European Economic Area where required by Morgan Stanley Europe S.E., authorised and regulated by Bundesanstalt fuer Finanzdienstleistungsaufsicht (BaFin) under the reference number 149169; in the US by Morgan Stanley & Co. LLC, which accepts responsibility for its contents. Morgan Stanley & Co. International plc, authorized by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, disseminates in the UK research that it has prepared, and research which has been prepared by any of its affiliates, only to persons who (i) are investment professionals falling within Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (as amended, the "Order"); (ii) are persons who are high net worth entities falling within Article 49(2)(a) to (d) of the Order; or (iii) are persons to whom an invitation or inducement to engage in investment activity (within the meaning of section 21 of the Financial Services and Markets Act 2000, as amended) may otherwise lawfully be communicated or caused to be communicated. RMB Morgan Stanley Proprietary Limited is a member of the JSE Limited and A2X (Pty) Ltd. RMB Morgan Stanley Proprietary Limited is a joint venture owned equally by Morgan Stanley International Holdings Inc. and RMB Investment Advisory (Proprietary) Limited, which is wholly owned by FirstRand Limited. Morgan Stanley Research is disseminated in Canada by Morgan Stanley Canada Limited and Morgan Stanley Wealth Management Canada Inc., members of the Canadian Investment Regulatory Organization (CIRO). The information in Morgan Stanley Research is being disseminated by Morgan Stanley Saudi Arabia, regulated by the Capital Market Authority in the Kingdom of Saudi Arabia, and is directed at Sophisticated investors only.

Morgan Stanley Hong Kong Securities Limited is the liquidity provider/market maker for securities of Alibaba Group Holding, Baidu Inc, Bilibili Inc, JD.com, Inc., Kuaishou Technology, Meituan, NetEase, Inc, Tencent Holdings Ltd., Trip.com Group Ltd listed on the Stock Exchange of Hong Kong Limited. An updated list can be found on HKEx website: http://www.hkex.com.hk.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (DIFC Branch), regulated by the Dubai Financial Services Authority (the DFSA) or by Morgan Stanley & Co. International plc (ADGM Branch), regulated by the Financial Services Regulatory Authority Abu Dhabi (the FSRA), and is directed at Professional Clients only, as defined by the DFSA or the FSRA, respectively. The financial products or financial services to which this research relates will only be made available to a customer who we are satisfied meets the regulatory criteria of a Professional Client. A distribution of the different MS Research ratings or recommendations, in percentage terms for Investments in each sector covered, is available upon request from your sales representative.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (QFC Branch), regulated by the Qatar Financial Centre Regulatory Authority (the QFCRA), and is directed at business customers and market counterparties only and is not intended for Retail Customers as defined by the QFCRA.

As required by the Capital Markets Board of Turkey, investment information, comments and recommendations stated here, are not within the scope of investment advisory activity. Investment advisory service is provided exclusively to persons based on their risk and income preferences by the authorized firms. Comments and recommendations stated here are general in nature. These opinions may not fit to your financial status, risk and return preferences. For this reason, to make an investment decision by relying solely to this information stated here may not bring about outcomes that fit your expectations.

The trademarks and service marks contained in Morgan Stanley Research are the property of their respective owners. Third-party data providers make no warranties or representations relating to the accuracy, completeness, or timeliness of the data they provide and shall not have liability for any damages relating to such data. The Global Industry Classification Standard (GICS) was developed by and is the exclusive property of MSCI and S&P.

Morgan Stanley Research, or any portion thereof may not be reprinted, sold or redistributed without the written consent of Morgan Stanley.

Indicators and trackers referenced in Morgan Stanley Research may not be used as, or treated as, a benchmark under Regulation EU 2016/1011, or any other similar framework.

The issuers and/or fixed income products recommended or discussed in certain fixed income research reports may not be continuously followed. Accordingly, investors should regard those fixed income research reports as providing stand-alone analysis and should not expect continuing analysis or additional reports relating to such issuers and/or individual fixed income products. Morgan Stanley may hold, from time to time, material financial and commercial interests regarding the company subject to the Research report.

Registration granted by SEBI and certification from the National Institute of Securities Markets (NISM) in no way guarantee performance of the intermediary or provide any assurance of returns to investors. Investment in securities market are subject to market risks. Read all the related documents carefully before investing.

INDUSTRY COVERAGE: China Internet and Other Services

COMPANY (TICKER)	RATING (AS OF)	PRICE* (09/11/2026)
Eddy Wang, CFA		
Autohome Inc (ATHM.N)	E (02/09/2023)	US$21.68
Full Truck Alliance Co. Ltd (YMM.N)	O (07/05/2023)	US$8.17
JD.com, Inc. (JD.O)	U (11/10/2025)	US$27.06
Kanzhun Ltd (BZ.O)	O (08/04/2021)	US$16.44
KE Holdings Inc (BEKE.N)	O (03/16/2022)	US$16.93
PDD Holdings Inc (PDD.O)	O (03/02/2023)	US$77.81
Vipshop Holdings Ltd (VIPS.N)	E (02/24/2022)	US$12.47
Gary Yu		
Alibaba Group Holding (BABA.N)	O (02/24/2025)	US$109.30
Baidu Inc (BIDU.O)	U (08/19/2026)	US$91.40
Meituan (3690.HK)	O (08/29/2024)	HK$75.05

Tencent Holdings Ltd. (0700.HK)	O (03/19/2020)	HK$430.60
Rebecca Xu		
HUYA Inc (HUYA.N)	E (05/16/2024)	US$2.04
IQIYI Inc (IQ.O)	E (01/19/2023)	US$0.95
JOYY Inc. (JOYY.O)	E (06/02/2022)	US$77.26
Weibo Corp (WB.O)	U (05/17/2024)	US$6.63
Yang Liu		
Bilibili Inc (BILI.O)	++	US$15.54
Kuaishou Technology (1024.HK)	E (08/20/2026)	HK$30.58
NetEase, Inc (NTES.O)	O (01/08/2025)	US$115.61
Tongcheng Travel Holdings (0780.HK)	O (01/04/2019)	HK$10.79
Trip.com Group Ltd (TCOM.O)	O (05/17/2021)	US$39.02

Stock Ratings are subject to change. Please see latest research for each company.
* Historical prices are not split adjusted.

© 2026 Morgan Stanley