

Three Strategic Bets on AI's Future

Maxime C. Cohen,¹ Eddy Hage-Youssef,² Daniel M. McCarthy,³ D. Daniel Sokol⁴

Idea in Brief (sidebar):

- Every major LLM launch we tracked coincided with total market growth and showed little evidence of cannibalization between models. When OpenAI releases a new model, Claude's user base grows. When Anthropic releases a new model, ChatGPT's user base grows. The "AI war" zero-sum framing appears to be incorrect.
- Three radically different monetization strategies are succeeding at once. ChatGPT pursues scale: around 79% market share with moderate revenue per user. Google Gemini pursues ecosystem: 12% share but almost no direct revenue, effectively subsidizing AI through its other businesses. Claude pursues premium: under 1% share but 40x higher revenue per user than Gemini. Each strategy faces different risks, but all three seem viable, at least for now.
- As market growth slows or capital markets become less forgiving, these strategies face different vulnerabilities. Scale leaders must convert free users. Ecosystem players must justify subsidies. Premium players must stay differentiated.
- Executives building, buying, or investing in AI should resist single-outcome bets and instead match strategy to structural advantage.

The Puzzle

Recent clashes over AI deployment boundaries, including Anthropic's standoff with the Pentagon, have made visible something executives are only beginning to recognize: the leading AI companies are not converging on a single formula for success. They are differentiating not only on capability, but also on distribution, monetization, and the types of business they are willing to pursue.

Based on worldwide mobile-app data covering seven major AI assistants from May 2023 through December 2025, we find little evidence of a zero-sum AI war. Instead, major launches have tended to expand the category, allowing three distinct strategies to gain traction at once: ChatGPT's scale play, Gemini's ecosystem play, and Claude's premium niche. For executives building, buying, or investing in AI, the key question may no longer be which model wins outright, but which strategic position is most durable as the market matures.

Based on our data, Claude, the AI assistant made by Anthropic, has less than 0.5% of the mobile AI assistant market in December 2025. In comparison, Google Gemini has around

¹ Desautels Faculty of Management, McGill University

² School of Computer Science, McGill University

³ Robert H. Smith School of Business, University of Maryland College Park

⁴ USC Gould School of Law and Marshall School of Business

12.5%. Yet Claude generates over 40 times more revenue per user than Gemini does. The monetization gap is so stark that despite having approximately 25 times fewer users, Claude actually generates 61% more total mobile revenue than Gemini.

This gap is one of the clearest findings in our analysis of consumer AI usage. It forces leaders to rethink what “success” in AI actually means. For some companies, success means converting users and activity into revenue. For others, it means creating ecosystem value by keeping users engaged in a broader platform where revenue is generated elsewhere.

Our analysis uses worldwide iOS and Android app-store data for seven major AI assistants from May 2023 through December 2025, including downloads, daily and monthly active users, and consumer revenue.⁵ Because these products distribute through the same app stores under similar payment structures, mobile data offers a clean basis for comparing consumer adoption and willingness to pay across competitors. Mobile subscriptions also unlock cross-device use, making app-store revenue a useful proxy for consumer willingness to pay. This lens does not capture enterprise contracts or API usage, but for the consumer-strategy question we ask here, it provides a clear and comparable basis for analysis.

Industry experts frequently debate whether AI will consolidate around a single dominant platform. In other words, one platform will ultimately dominate and the others will fade. Our data tells a different story. We see three distinct approaches that have all been working: ChatGPT has been succeeding by compounding first-mover advantage, scale, and continuous innovation; Google’s Gemini acts as an ecosystem defense tool, prioritizing distribution over direct revenue; and Anthropic’s Claude thrives as a premium niche, monetizing a small but high-value user base at far higher rates than its larger rivals. The remaining four platforms (Copilot, DeepSeek, Perplexity, and Grok) collectively hold about 9% market share; we focus on the three leaders, which together represent over 91% of the market and whose strategies are most instructive for understanding competitive positioning.

The question arises naturally: How can a scale leader, an ecosystem player, and a premium niche all thrive simultaneously? Part of the answer lies in a second finding: the AI assistant market is expanding rapidly, and competitive launches appear to grow the whole pie, not just redistribute it. In a market expanding this fast, three fundamentally different strategies can succeed at once. A rising tide, at least for now, is lifting all the boats.

But this raises harder questions. Which strategy will prove durable as growth slows or capital markets tighten? Can standalone AI companies survive against competitors who can subsidize losses through other business lines? And what does this divergence tell us about where value gets captured in AI?

⁵ Our data comes from the market intelligence firm, [Sensor Tower](#)

Three Strategic Positions

The gap between market share and monetization in the AI assistant market reveals three fundamentally different bets about how value gets captured and, crucially, who actually needs to capture it. Figure 1 plots the major platforms on two dimensions: share of DAUs (horizontal axis) and revenue per user per day (vertical axis) as of December 2025. The three leaders occupy entirely different quadrants, and understanding why sheds light on the strategic logic of each approach.

AI Assistant Market: Three Divergent Strategies

Mobile App Data, December 2025

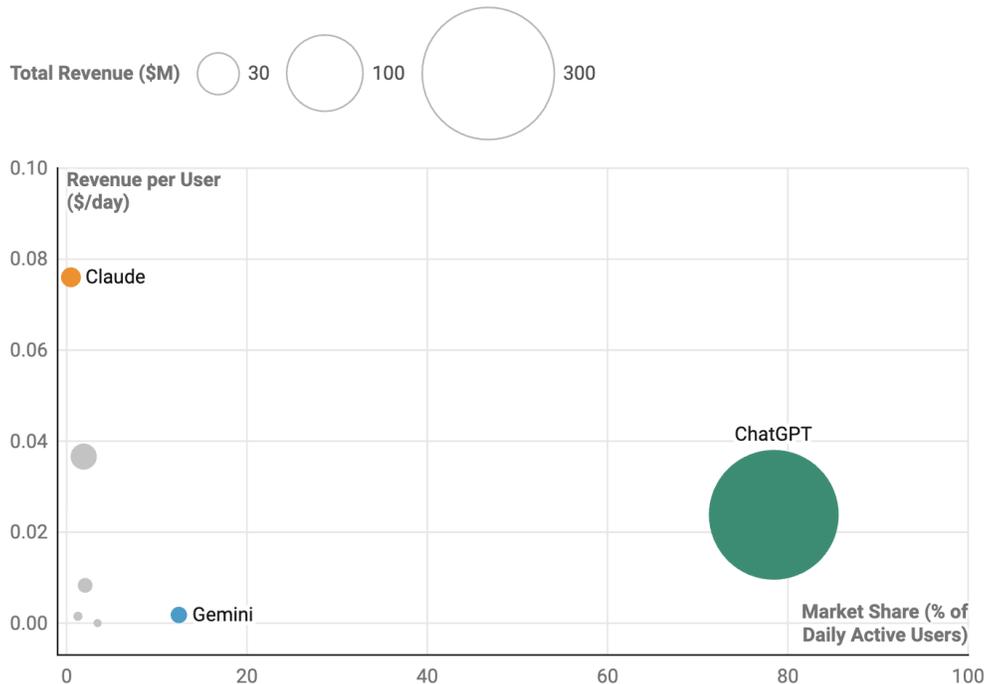


Figure 1. Market share vs. revenue per user in AI assistants (bubble size represents total daily revenue; data: iOS and Android, December 2025).

As of December 2025, Claude’s users are 37.5% iOS versus 62.5% Android, the highest iOS concentration among major platforms, while Gemini skews dramatically toward Android at 87.4%. Geography points in the same direction: Claude skews heavily toward high-income markets (especially the United States), while Gemini skews toward large Android-first markets such as India. That mix likely contributes to monetization differences, because subscription pricing and conversion rates vary across countries. These patterns aren’t coincidental; they reveal each platform naturally attracting the demographics that align with its monetization strategy. Claude concentrates where willingness to pay is highest. Gemini expands where user acquisition is cheapest. ChatGPT balances both combined with a first-mover advantage, maintaining a meaningful presence across developed and emerging markets alike.

The Scale Leader (ChatGPT)

It is no surprise that ChatGPT leads the AI assistant market by almost any measure. With more than 390 million daily active users as of December 2025, OpenAI's flagship product commands around 78.5% of the mobile market. In fact, ChatGPT experienced explosive growth in Q1 2025, with daily user additions averaging 1.2 million throughout the quarter and peaking at 5 million DAU per day in late March 2025. In addition, ChatGPT reached [one million users in just five days](#) after its launch on November 30, 2022, [garnered around 57 million users](#) by the end of its first month, and [reached 100 million monthly active users](#) in two months. This meteoric adoption marked a record for the fastest user acquisition in digital history. The next largest competitor, Google Gemini, has only 16% as many users.

OpenAI's position rests on having built the market. ChatGPT was the product that introduced most consumers to AI assistants, and that first-mover advantage created a brand that competitors have struggled to displace. The company offers tiers ranging from free to \$20 per month for regular users, all the way up to \$200 per month for power users, yet the blended revenue per user is just \$0.024 per day, or roughly \$9 per user annually. This moderate average reflects the reality that most users never pay. Comparing total DAU to implied subscribers from revenue metrics, we estimate that the vast majority—over 90%—of users never pay. At this scale, even moderate monetization generates substantial revenue. Our data suggests that ChatGPT produces over \$9 million per day in mobile app revenue alone, before accounting for desktop, API, and enterprise revenue streams. In addition, OpenAI recently introduced an [\\$8/month “Go” plan](#) and previewed an advertising approach for free users, suggesting a push to monetize at scale without fully paywalling usage.

Defending the lead naturally requires constant innovation. ChatGPT released four major models during our study period (May 2023 to December 2025), more than any other platform. Each release generated media attention and user growth, though with diminishing returns: o1 Preview (released on September 12, 2024) drove 26.3% DAU growth, while GPT-5 (released on August 7, 2025) produced just 5.0%. The challenge of exciting users when you already have ~400 million of them is considerable. At scale, innovation becomes necessary but insufficient.

The risks are significant. Unlike Google, OpenAI must make money from AI directly - it has no adjacent business to subsidize potential short-term losses. The company must continue outpacing rivals on capability while also converting enough users to paid tiers (or to explore alternative business models, such as advertisements) to sustain its massive research investment. That pressure shows up in monetization experimentation (new tiers, and now advertising initiatives) as well as fundraising. OpenAI has raised billions in venture capital, including a [\\$40 billion funding round](#) led by SoftBank in March 2025 at a \$300 billion valuation, and investors will eventually expect returns commensurate with that scale. In fact, OpenAI is preparing for an IPO at a [valuation of up to \\$1 trillion](#). If competitors narrow the quality gap, ChatGPT's moderate monetization may prove insufficient to support frontier AI development, especially if growth slows or capital markets become less forgiving. Scale is only an advantage if monetization grows faster than costs.

The Ecosystem Play (Gemini)

Google's Gemini occupies a puzzling position: the second-largest user base in the market combined with almost no direct revenue from those users. With over 60 million DAUs and about 12.5% market share, Gemini has built a substantial presence. And mobile likely understates Google's true reach – Gemini is also embedded across Google surfaces like Search, Chrome, and Workspace – many of which monetize indirectly instead of through an explicit assistant subscription. The gap between Google's AI footprint and its AI monetization is probably even larger than our data suggests. Gemini monetizes at just \$0.0018 per user per day, or about 60 cents per user annually. That is more than 40 times less than Claude and 13 times less than ChatGPT.

This may look like strategic malpractice until you consider Google's broader business. Google does not need Gemini to generate direct revenue, at least not in the short run. Google earns [over \\$260 billion annually from advertising](#), the vast majority of its total 2024 revenue of \$350 billion, and its AI assistant serves primarily to keep users within the Google ecosystem. If Gemini prevents even a small fraction of searches from migrating to ChatGPT or Claude, it protects far more advertising revenue than any subscription fee could generate.

This is the ecosystem play: treat the AI assistant as a loss leader that supports the core (lucrative) business. Google can afford to give Gemini away because the alternative of ceding the AI assistant market to competitors threatens Google's core search advertising business. When users ask ChatGPT 'What's the best restaurant near me?' instead of searching on Google, it represents a lost ad impression and a loss of user intent data. Gemini doesn't need to make money directly; it serves to keep users asking questions inside Google's ecosystem, where those interactions can monetize on Youtube, Maps, and through ads. Low direct monetization is the strategy and not a failure of execution. This is a luxury only Google can afford. OpenAI and Anthropic must monetize AI directly; Google doesn't.

Our data shows Gemini growing aggressively during the study period. All three Gemini launches we tracked (2.0 Pro, Nano Banana, and 3.0 Pro) produced market share gains totaling nearly three percentage points. The Nano Banana launch in August 2025 drove a remarkable 373% increase in downloads, the largest single-launch acquisition spike for an established platform in our dataset. In short, Google is playing defense, using distribution to protect its advertising franchise. The fact that competitors are now openly exploring advertising-supported tiers for AI assistants may further blur the line between “assistant revenue” and “ecosystem revenue.”

The issue is that this strategy only works for companies with profitable adjacent businesses to protect. Standalone AI companies cannot subsidize free products indefinitely. Google's approach may prove sustainable for them but unreplicable for others. It also creates competitive dynamics that raise antitrust questions as a strategic constraint: when a company with a strong position in one market uses that position to subsidize competition in another, where is the line between aggressive competition and anticompetitive conduct? The legal and economic analysis of such arrangements in platform markets may be complex, but the strategic implications are clear. Standalone AI companies cannot match this subsidy indefinitely.

The Premium Niche (Claude)

Claude represents the opposite extreme: a small user base paying premium prices. With roughly 2.5 million DAUs and around 0.5% market share, Anthropic's assistant is barely visible in market share statistics. But Claude's users generate \$0.076 per user per day, or about \$27.75 per user annually. That is more than three times ChatGPT's rate and more than 40 times Gemini's.

How does Anthropic achieve this? The company appears to have found a specific niche audience: professional users who value Claude's particular strengths—especially in coding-adjacent workflows—and who convert at unusually high rates. Anthropic's subscription pricing is comparable to ChatGPT's, but a much larger fraction of Claude's users convert to paid tiers. Where ChatGPT has hundreds of millions of free users diluting its average revenue, Claude's smaller, more committed user base drives premium economics.

This premium strategy requires clear differentiation. Anthropic has made deliberate product choices that favor depth over breadth. Unlike ChatGPT, which spans image generation, voice, and video, Claude stays more focused on text-heavy tasks such as coding, document analysis, and complex reasoning. That focus appears to attract users who treat AI as a work tool worth paying for rather than a general-purpose novelty to sample for free. This is a classic niche strategy: cede scale to competitors but capture disproportionate value from a devoted high-value segment. Anthropic's differentiation may also extend beyond product features to deployment boundaries. Its willingness to forgo some uses in order to preserve tighter constraints reinforces the brand it is building with professional users. Whether that stance ultimately attracts or repels particular segments remains uncertain, but it further distinguishes Claude from competitors pursuing broader market access. Like OpenAI, Anthropic still must monetize AI directly. The difference is that it is pursuing depth of monetization before breadth of adoption.

Our data shows Claude launches generating strong engagement gains. Claude 3.5 Sonnet in June 2024 produced a 102% increase in DAUs, the largest percentage gain of any launch we tracked in our analyses for established platforms (excluding new market entrants). More notably, revenue grew in proportion to users, suggesting that new users converted to paid subscribers at rates comparable to the existing base. This is the premium flywheel: quality attracts discerning users who are willing to pay for quality.

The risks are also clear. Premium positioning requires continuous differentiation. If ChatGPT or Gemini match Claude's quality while offering free or cheaper alternatives, the premium segment may collapse. Anthropic also faces a scale challenge: AI development is expensive, and without an ecosystem to subsidize losses, a 2.5 million user base may not generate sufficient revenue to fund frontier research and future developments. Most recently, [the company raised \\$13 billion in September 2025](#) at a \$183 billion valuation, betting that its premium users can eventually support a sustainable business. Whether that bet pays off depends on how much differentiation persists as all models keep improving. In summary, premium strategies are often fragile in markets with fast-moving quality convergence.

Market share and revenue tell only part of the story. Stickiness (which can be defined as the ratio of daily to monthly users) reveals how deeply these products embed into users' routines. ChatGPT's dominance extends beyond sheer volume. With 43.5% stickiness as of December 2025, a significant portion of ChatGPT's monthly users return on any given day. Gemini's low stickiness (17%) confirms what the revenue data suggested: despite a massive user base, engagement remains shallow. Google's distribution advantage gets people to install Gemini, but it hasn't yet convinced them to use it regularly.

Claude presents a more nuanced picture. Its stickiness stands at 19.5% in December 2025 even as revenue per user remains more than 40 times higher than Gemini's. This confirms that Claude is valued for specific professional tasks (for example, coding and other professional workflow) rather than everyday utility. Users engage less frequently but pay significantly more, consistent with a product serving specialized, high-value use cases.

Why All Three Strategies Work (For Now)

Our data reveals the mechanism behind this coexistence.

Every major launch we tracked—fifteen in total across ChatGPT, Claude, DeepSeek, Gemini, Grok, and Perplexity—coincided with total market growth. For instance, the market expanded by 8.2% during Anthropic's Claude 3.5 Sonnet's launch and by 19.1% during OpenAI's o3 release (see Figure 2). On average, total DAUs increased by approximately 13.7% in the 30-day window following each launch.⁶

This surprised us. The "AI wars" framing suggests zero-sum competition: when one platform gains, others typically lose. That's not what the data shows in this market. When Anthropic released Claude 3.5 Sonnet in June 2024, ChatGPT's user base grew. When OpenAI released o3 in April 2025, Claude's user base grew too. Competitive launches appear to expand the category, not just shuffle users between apps. During these launches, revenue also grew: all launches showed positive total market gains, averaging 12%. Downloads tell a more mixed story. Only half the launches showed positive total market download growth, with wide variance.

Some caveats are in order. Correlation is not causation. Model releases come with press coverage, social media buzz, and renewed public attention to AI capabilities. Any of these factors could partially drive the growth we observe. What we can say is that in our study period, competitive launches and market expansion went hand in hand. For executives worried about a new competitor's launch cannibalizing their user base, this should be reassuring.

⁶ Methodological note: Our launch analysis measures 30-day pre-launch and 30-day post-launch windows. Important limitations include: (1) temporally clustered launches create overlapping measurement windows, potentially attributing the same market growth to multiple events; (2) time-series autocorrelation means that current market size naturally predicts future market size independent of launches; (3) we cannot isolate individual launch effects from confounding factors like seasonal trends or media coverage. Our analysis demonstrates correlation between launches and market growth but cannot establish causation. The consistent pattern across all 15 launches suggests robustness, but more sophisticated econometric techniques would be needed to rigorously isolate causal effects.

The expansion dynamic explains why three strategies can work simultaneously. When the pie grows fast enough, you don't need to take share to build a successful business. ChatGPT's share fell from 89% to around 79% during our study period, but its absolute user count more than tripled from 122 million to 392 million DAUs during the same period as the market tripled in size.

The structural advantages we described earlier—ChatGPT's first-mover brand, Google's ecosystem subsidy, Claude's differentiation—provide explanations on *which* strategy each company can pursue. And crucially, whether they even need to make money from AI at all (at least directly and in the short run).

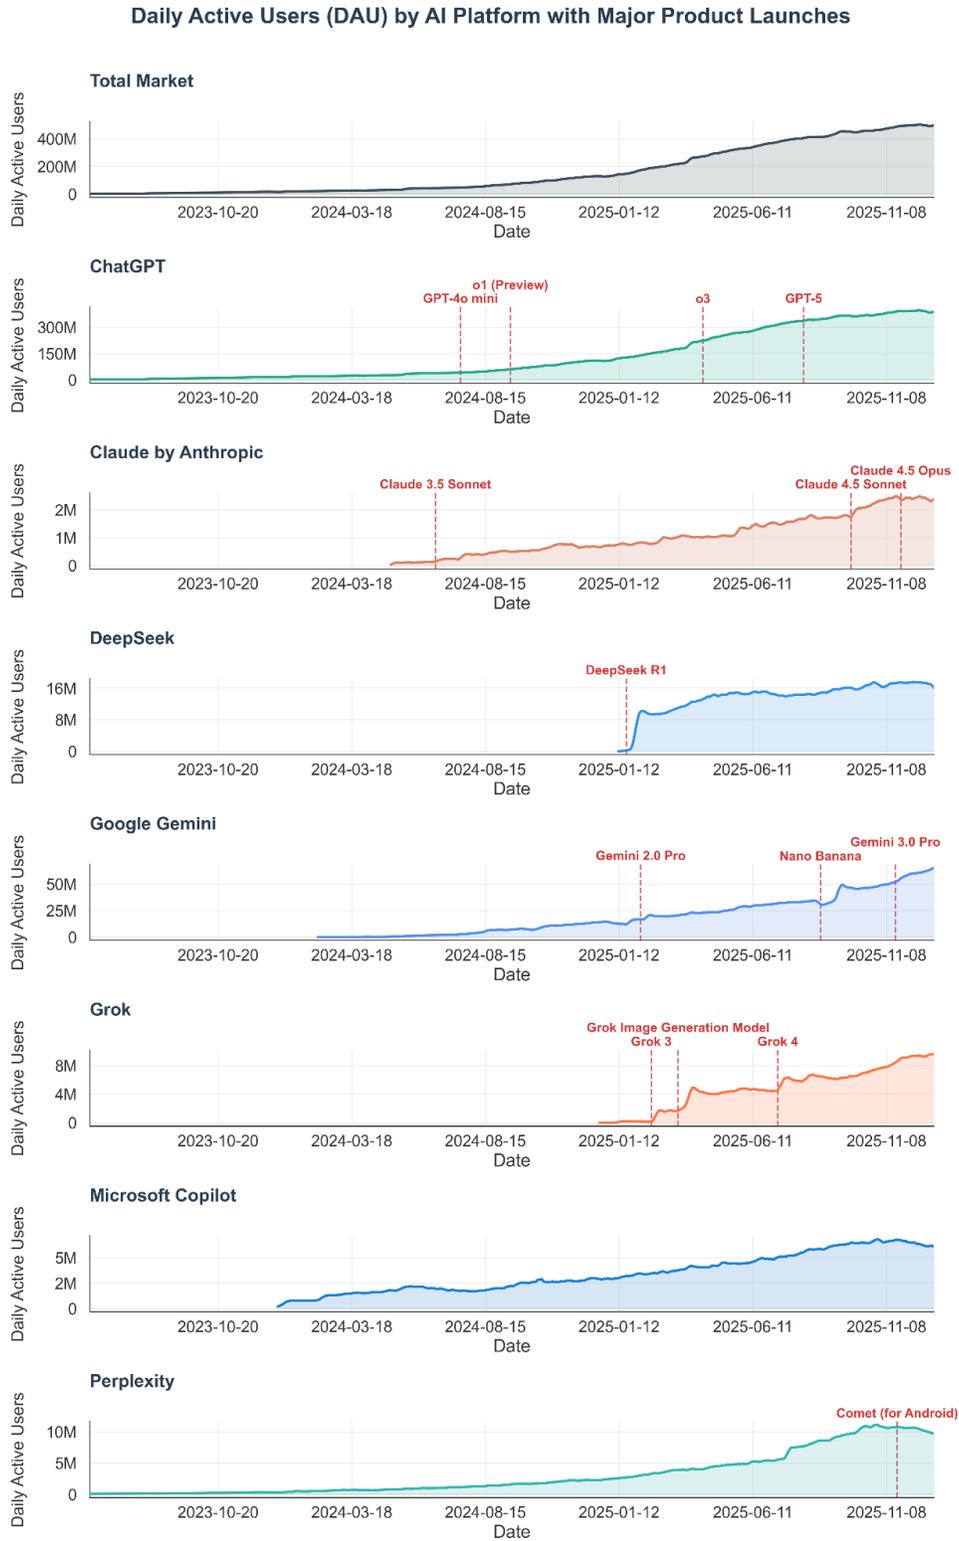
The question is what happens when growth slows or capital markets pull back.

The Risks Ahead

The coexistence we've documented depends on rapid market expansion. As growth slows or capital markets tighten, the alternative strategies face naturally different vulnerabilities. These vulnerabilities mean that executives need to make strategic organizational choices.

The math gets harder. We're already seeing early signs of this with ChatGPT. The platform continues to add users in absolute terms, but the percentage growth rates are compressing. This is the natural arithmetic of market maturation: when you've already acquired 400 million DAUs, growing by 10% is harder than when you've already acquired 40 million. The remaining non-adopters are, almost by definition, less interested in AI assistants than those who came before them. They'll require more convincing, more use cases and higher precision, enhanced security protocols and reassurance, and more reasons to change their behavior. This is true for every platform, but it matters most for the scale strategy. ChatGPT's position depends on continued dominance at scale. If growth stalls while competitors keep expanding, the share erosion we've already observed could accelerate.

The “who needs to monetize” question gets harder. Google can afford to give Gemini away because it makes money elsewhere. The AI assistant doesn't need to be profitable (at least in the short term) if it keeps users in Google's ecosystem, generates training data, or simply prevents competitors from gaining ground. Anthropic and OpenAI don't have that luxury. They're standalone AI companies that must generate revenue from AI directly (and eventually, profits). That means the pressure to monetize can rise even if user growth stays strong, simply because capital markets and cash burn force action. The income statements of these companies are not pretty right now. OpenAI reportedly [lost an estimated \\$5 billion in 2024](#) and projected losses of roughly [\\$14 billion in 2026](#), with cumulative cash burn expected to reach [\\$115 billion through 2029](#). The company doesn't expect to turn profitable until [2029 or 2030](#). Anthropic raised [\\$13 billion in September 2025](#) at a \$183 billion valuation, and by December 2025 had signed a [term sheet for another \\$10 billion round](#) at a reported \$350 billion valuation. The company projects breaking even by 2028, but the trajectory depends on revenue growth that would take it from about [\\$9 billion in 2025 to \\$26 billion or more in 2026](#).



Source: Sensor Tower Data

Figure 2. Total market growth coincides with competitive launches (DAUs increased an average of 13.7% in the 30-day window following major model releases across platforms).

When investor patience wears thin, the pressure to monetize will intensify. Claude's premium strategy looks prudent in this light: better to have fewer users who pay real money than millions who don't. But as mentioned, the premium play only works if the product stays differentiated enough to justify the price. And in such a fast-moving market, that's far from guaranteed. That said, companies are exploring various alternative revenue streams beyond subscriptions: OpenAI recently began testing [advertising in ChatGPT](#), while enterprise contracts for specialized models in healthcare, finance, and legal sectors are growing rapidly for all three firms. In fact, a large portion of Claude's revenue seems to be coming from enterprise deals. However, these monetization experiments remain in early stages, and none have yet demonstrated the scale needed to offset infrastructure costs.

Commoditization looms. The strategic question at the heart of AI competition is whether these AI assistants will become interchangeable utilities or remain differentiated products. Commoditization doesn't mean models stop improving – it means they all improve together and the gaps narrow over time. If that happens, the scale and ecosystem strategies become more viable, while premium positioning gets harder to defend. Companies like Google and Microsoft can reach users at zero acquisition cost through existing products, and absorb operating losses through profitable adjacent businesses. This lets them treat AI as a feature within a broader portfolio, and not as a standalone business. Standalone AI companies like Anthropic and OpenAI would face the same fate as standalone word processors or spreadsheet programs: absorption into larger platforms. But commoditization isn't inevitable. If AI assistants stay differentiated, namely, if users genuinely prefer one model's reasoning or writing or personality over another, then premium positioning remains viable. Deployment scope and usage boundaries may prove to be another durable axis of differentiation, one that is harder to converge on than model quality because it reflects distinct business models and customer bases. The early evidence is mixed. [Benchmark performance gaps seem to be narrowing](#). Yet user behavior suggests meaningful differences: Claude's 40x average revenue per user (ARPU) advantage implies that users find something worth paying for.

Two factors could slow commoditization. First, switching costs may build over time and lock in users. As AI assistants accumulate user preferences, saved memories, searchable conversation histories, and custom instructions, the friction of starting over with a new provider increases. The more users invest in a specific assistant, the more valuable it becomes to retrieve and build on that history. A professional who has spent months teaching Claude their writing style or ChatGPT their coding preferences faces real costs in switching, even if a competitor offers marginally better performance. Second, user inertia is well-documented in technology adoption: once users settle into a workflow, they resist change even when better alternatives exist. ChatGPT's 43.5% daily stickiness suggests habit formation is already occurring at scale. These dynamics could create moats that preserve differentiation even as underlying capabilities converge.

The investment cycle may not cooperate. Hundreds of billions of dollars have flowed into AI infrastructure: data centers, chips, and power plants. Much of this investment is predicated on growth assumptions that may not materialize, or may materialize more slowly than expected. If the AI market disappoints, the capital flowing to AI startups could dry up before they reach

profitability. This would force consolidation, partnership, or outright failure. The companies that will survive may be those with either sustainable unit economics or deep-pocketed sponsors willing to absorb losses until their competitors “tap out.”

None of this means any particular strategy will fail. Instead, it means the environment that made all three viable simultaneously is unlikely to persist into the indefinite future. The market will eventually mature, growth will slow and investors will become more selective, and the competitive dynamics are likely to shift. The question is not whether this happens, but when, and which strategies will adapt best when it does.

Implications for Executives

The coexistence of three viable strategies offers a practical framework for anyone navigating the AI landscape. But to be practical, we need to be more specific.

If you're building AI products

*Do you have distribution that can reach hundreds of millions of users at near-zero marginal cost? **The scale play*** requires infrastructure to acquire users at scale at a cheap rate, not just being first to market. OpenAI built this through a first-mover advantage. Google owns both Android and Search, a massive user base that can be converted to Gemini. Most companies have neither advantage. If your user acquisition strategy depends on paid marketing or viral growth, then the scale play is probably not available to you.

*Do you have a profitable business to protect and subsidize losses with? **The ecosystem play*** only works if AI engagement supports revenue you're already generating elsewhere, and you can afford to treat AI as a cost center for some time. Google can give Gemini away (and absorb short-term losses) because it keeps users in the advertising funnel. Microsoft bundles Copilot to protect Office and Azure relationships. If you don't have a cash-generating business that an AI assistant would support, this strategy isn't available for you either.

*Can you serve a segment that will pay premium prices, and can you resist the temptation to dilute it? **The premium play*** is more accessible to smaller more-specialized companies, but it demands discipline and a clear, constant differentiation. Chasing growth metrics that expand your user base into non-paying segments can undermine the economics that make the strategy viable. Anthropic appears to have found professionals who [value reasoning quality](#). What segment would value your differentiation enough to pay 3x the market rate?

For most companies, the honest answer is sobering: you're neither the incumbent who built the market, nor a giant with an ecosystem to subsidize your new product. The AI assistant market is consolidating around well-capitalized leaders with structural advantages that new entrants will struggle to match. Building on top of existing platforms may be more realistic than competing with them.

If you're buying AI services

Enterprise buyers often focus on capability benchmarks and per-seat pricing. Our data suggests that they should also stress-test strategic sustainability.

What happens if investor patience runs thin? As mentioned, OpenAI reportedly projects losses of \$14 billion in 2026. Anthropic is raising billions at valuations that assume explosive revenue growth. If growth disappoints, these companies will face pressure to raise prices, identify diversified strategies, cut development, or seek acquirers. Ask your provider: What is your path to profitability, and what happens to our pricing and roadmap if you don't hit it?

What am I trading for "free"? When your AI assistant is subsidized by an advertising or cloud computing business, your usage data becomes part of that business model. The product roadmap will be shaped by the parent company's priorities, which may not align with yours. Google's AI investments likely serve Google's strategic interests first. You should ask: What data do you retain from my usage and how do you use it? How does your parent company's strategy affect product direction?

Is this provider differentiated enough to survive? A 2.5-million user base, even at high ARPU, may not generate sufficient revenue to fund frontier research. If Claude's differentiation narrows, Anthropic's premium economics collapse. You should ask: What is your sustainable competitive advantage, and how do you plan to maintain it as competitors improve?

The safest approach may be deliberate diversification: avoid deep dependence on a single provider until the market structure stabilizes.

If you're investing

Our data reveals a disconnect between the metrics that capture headlines and the metrics that drive business value:

Growth rate or revenue quality? Gemini's 373% download spike from Nano Banana made news. But those downloads generated minimal revenue. Claude's modest user growth came with strong revenue growth, which is a healthier signal.

Market share or unit economics? The ARPU gap we documented (Claude at 40x Gemini's rate, 3x ChatGPT's) reflects real differences in business strategies. A company adding millions of non-paying users is building something very different from one adding thousands who convert immediately.

Reported metrics or sustainable metrics? Ask what happens to engagement when the marketing spend stops. Ask what the retention curve looks like at 90 days. Ask whether ARPU is rising or falling as the user base expands.

The value of a business strongly depends on future cash flow – profit, not revenue. In markets where category leaders can grow quickly by spending aggressively, it's easy to confuse scale

with sustainability. [WeWork reached a \\$47 billion valuation](#) before filing for bankruptcy in 2023 when its growth-at-any-cost strategy failed to produce sustainable unit economics. Blue Apron followed a similar trajectory in meal kits: rapid revenue growth, majority market share, then a collapse from a \$3.2 billion IPO target to a \$103 million acquisition after a 95% stock price decline. These are not isolated cases. Many companies have become market leaders while burning cash unsustainably, only to fail when capital markets lost patience. In a market where AI leaders are collectively burning tens of billions annually, these precedents warrant caution.

Conclusion

We set out to understand competitive dynamics in the AI assistant market. What we found was more surprising than we expected.

Three fundamentally different strategies have emerged. ChatGPT defends its position as scale leader. Gemini leverages Google's ecosystem to compete without needing direct short-term revenue. Claude pursues premium users who are likely to pay in exchange for differentiation. Each reflects a different hypothesis about where value gets captured in the AI market, and for now, all three appear viable. The market expanded throughout our study period, with every major model launch coinciding with overall growth. In an environment this dynamic, different approaches can coexist in the short run without destroying each other.

But coexistence is not equilibrium. The conditions that make all three strategies viable simultaneously, rapid market expansion, abundant capital, and unsettled user preferences, will clearly not last forever. As growth slows or capital becomes less available, the different strategies will face different tests. The scale leader will need to convert users to revenue. The ecosystem player will need to justify its subsidies. The premium niche will need to maintain strong differentiation against improving competitors.

We do not know which strategy will win. It's possible that all three prove durable, serving different segments of a maturing market. It's also possible that one approach will dominate as conditions change. The honest answer is that anyone claiming to know the outcome with certainty is likely to be overconfident.

What we can say is this: the "winner-takes-all" framing that pervades AI coverage appears wrong, at least for now. The market is large enough and growing fast enough to support multiple approaches. Executives building, buying, or investing in AI should resist the temptation to bet everything on a single outcome. The companies that thrive will likely be those that understand their own structural advantages, choose strategies that fit those advantages, and adapt as the market evolves.

We are watching three different hypotheses about the future of AI play out in real time. The experiment is far from over.