SHARE FULL ARTICLE

1% Of Companies Account For 80% Of OpenAI And Anthropic Spending - And It Just Stalled

1%的企业占据了 OpenAI 和 Anthropic 全部支出的 80%——而这些企业的增长现已停滞。



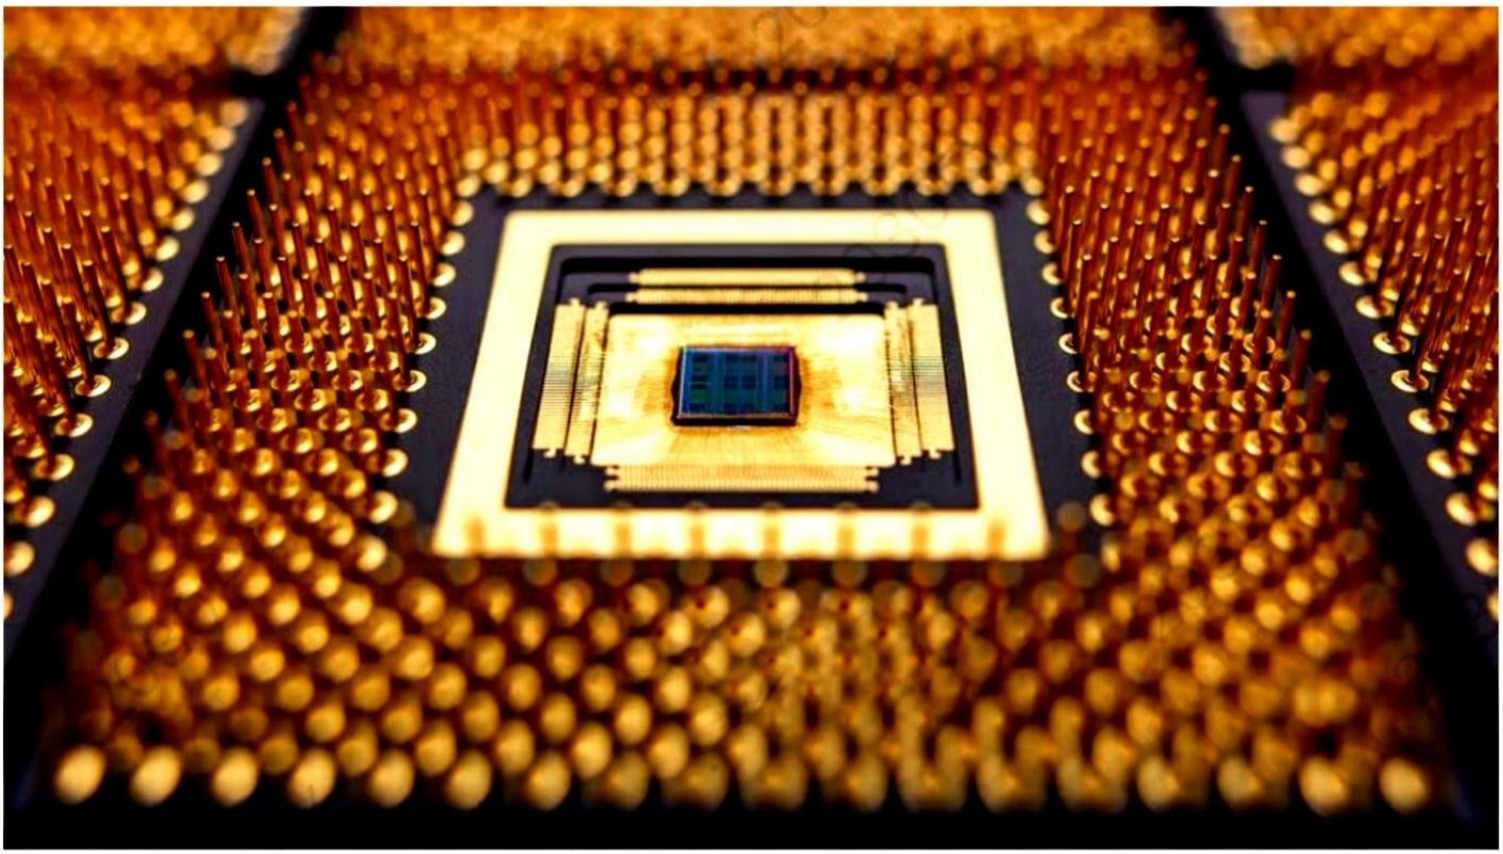
BY TYLER DURDEN

THURSDAY, OCT 22, 2026 - 08:35 AM

Just 1% of companies account for 80% of all business spending on OpenAI and Anthropic, according to corporate card giant Ramp, which tracks payments across more than 70,000 U.S. businesses. Ramp lead economist Ara Kharazian, who flagged the figure in early September, noted that spending on the two AI labs is more concentrated than in any other software category the firm tracks.

根据负责追踪美国 7 万多家企业支出情况的公司卡管理公司 Ramp 的统计，仅有 1% 的企业就占据了在 OpenAI 和

Ramp 的统计，仅有 1%的企业就占据了在 OpenAI 和 Anthropic 上的所有支出的 80%。Ramp 的首席经济学家阿拉·哈拉齐安在 9 月初提到了这一数据，他指出，对这两家人工智能公司的投资集中度，远远高于该公司所追踪的其它任何软件领域的投资集中度。



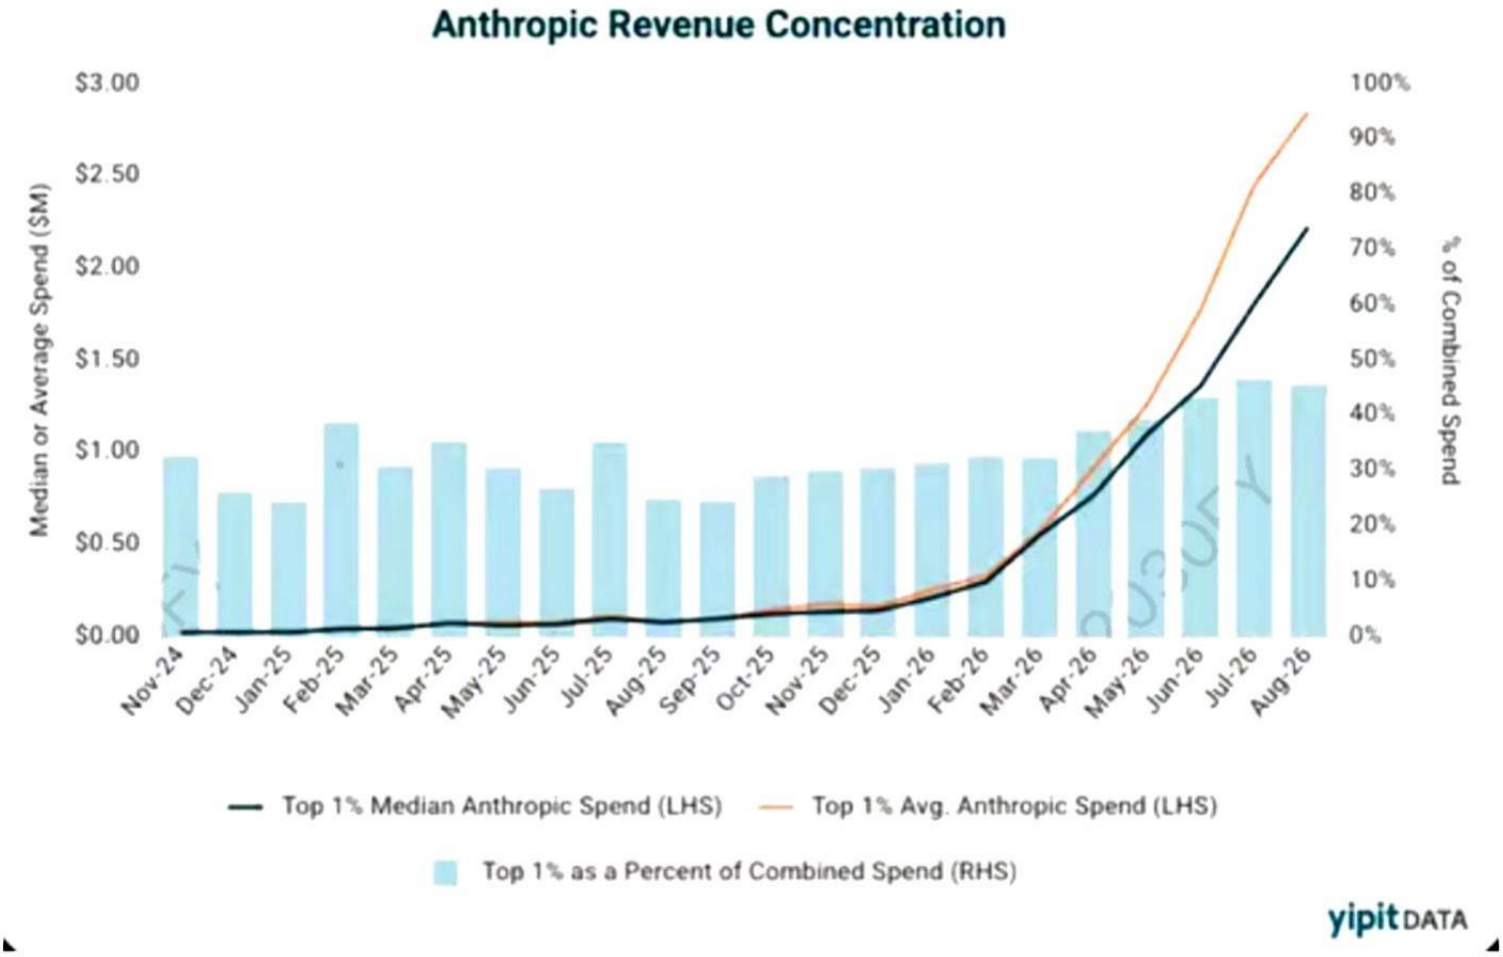
Per capita, the biggest spenders are overwhelmingly venture-backed. According to Ramp's summer report, the median VC-backed startup spent $76 per employee per month on AI - more than ten times the $7 spent at private equity-backed firms.

从人均支出来看，支出最高的企业大多是由风险投资支持的。根据 Ramp 的夏季报告，那些由风险投资支持的初创企业，每名员工每月在人工智能方面的支出为 76 美元，这一数字是那些由私募股权支持的企业的 7 美元的十多倍。

That top-heavy dynamic mirrors what YipitData uncovered at Anthropic specifically: the lab's top 1% of customers

generated 46% of its revenue in August, nearly doubling from 25% a year earlier. **How many of those whale customers, we asked at the time, are funded by the very same VCs backing Anthropic?**

这种由少数大客户主导的局面，与 YipitData 在 Anthropic 那里发现的状况如出一辙：8 月份，该公司排名前 1%的客户创造了其 46%的营收，这一比例较上一年的 25%几乎翻了一番。我们当时还询问过，这些大客户中，有多少是得到了与 Anthropic 相同的风险投资机构的资助?



Anthropic's IPO prospectus - the updated version of which we dissected yesterday - confirms the concentration at the very top: just two unnamed customers generated nearly a quarter of its 2025 revenue, according to Reuters. Their identities aren't hard to guess. Last year, two AI coding tools - Cursor and Microsoft's GitHub Copilot - drove roughly $1.2 billion of Anthropic's $5 billion annualized revenue run-rate, as

VentureBeat reported.

Anthropic 的 IPO 招股书——我们昨天已经对其进行了分析——显示了收入来源的高度集中：据路透社报道，2025 年时，仅有两家未具名的客户就为该公司创造了近四分之一的营收。其实这些客户的身份并不难猜。据 VentureBeat 的报道，去年，两款人工智能编码工具——Cursor 以及微软的 GitHub Copilot——就为 Anthropic 带来了约 12 亿美元的收入，而该公司当年的年度总营收则为 50 亿美元。

Follow the money. Cursor's parent company, Anysphere, raised $2.3 billion at a $29.3 billion valuation last November in a round co-led by Accel and Coatue, alongside Thrive Capital, Andreessen Horowitz, Nvidia, and Google. Three months later, Coatue and Singapore's GIC led Anthropic's massive $30 billion Series G, with Nvidia again participating. **Coatue's cash flows into Cursor, Cursor spends it on Claude tokens**, those tokens book as Anthropic revenue, and that revenue underwrites the markups on Coatue's Anthropic stake.

跟着钱走。Cursor 的母公司 Anysphere 去年 11 月以 293 亿美元估值融资 23 亿美元，由 Accel 和 Coatue 联合领投，Thrive Capital、Andreessen Horowitz、英伟达和谷歌参投。三个月后，Coatue 和新加坡的 GIC 领投了 Anthropic 高达 300 亿美元的 G 轮融资，英伟达再次参与。Coatue 的资金流入 Cursor，Cursor 将其用于购买 Claude 代币，这些代币计入 Anthropic 的收入，而这笔收入又支撑了 Coatue 所持 Anthropic 股份的估值溢价。

Even the scorekeeper is inside the loop. Ramp itself raised $1.85 billion at a $60 billion valuation last week in a round led

by Dragoneer, Thrive, Founders Fund, Sands Capital, and Avenir Growth. Its cap table also features Coatue, D1, ICONIQ, GIC, Fidelity, T. Rowe Price, General Catalyst, and Lightspeed - every single one of which, alongside Dragoneer and Founders Fund, backed Anthropic's Series G.

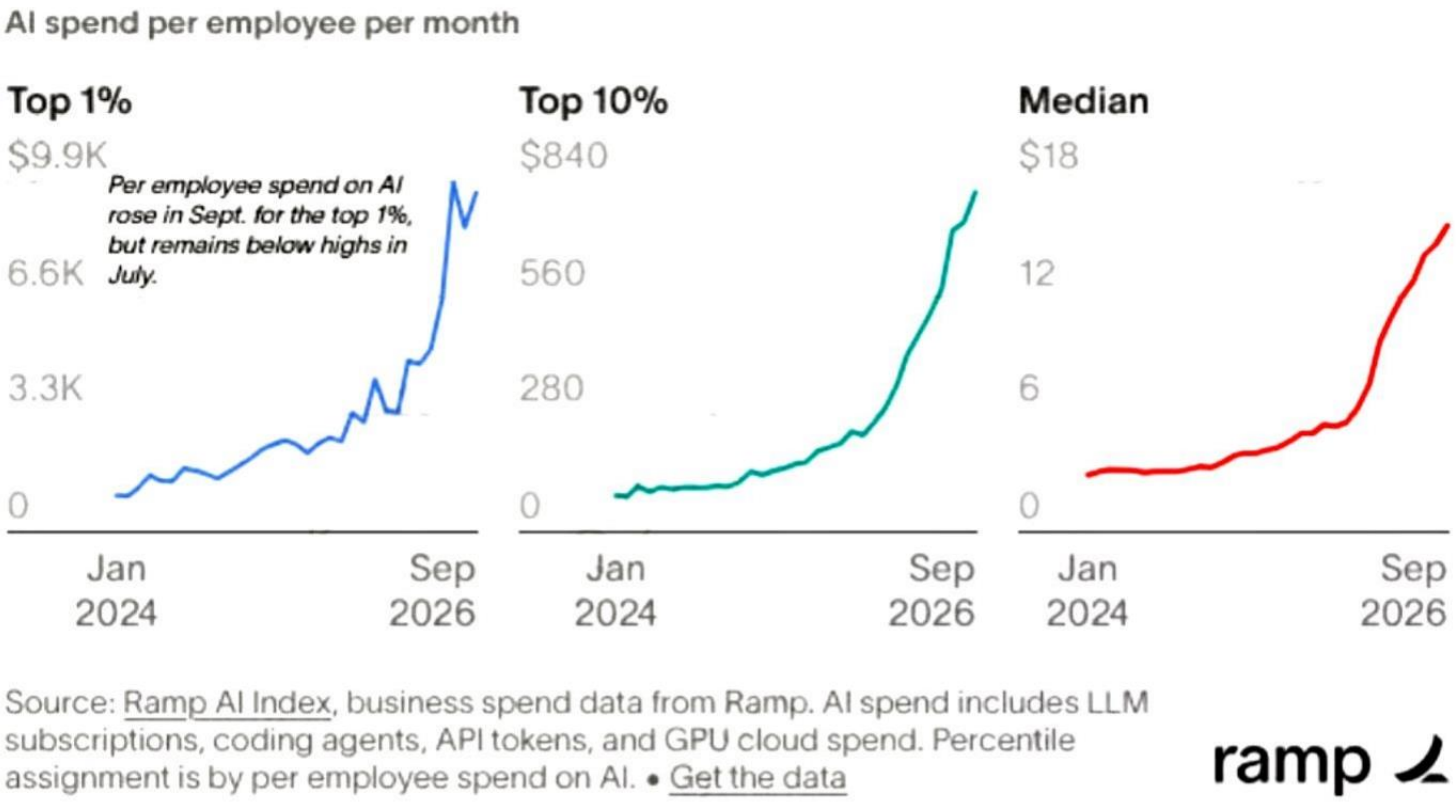
就连记分员也在这个圈子里。上周，在 Dragoneer、Thrive、Founders Fund、Sands Capital 以及 Avenir Growth 的领投下，Ramp 成功筹集了 18.5 亿美元的资金，其企业估值则为 600 亿美元。其股东名单中还包括 Coatue、D1、ICONIQ、GIC、Fidelity、T. Rowe Price、General Catalyst 和 Lightspeed 等投资机构——这些机构都与 Dragoneer 和 Founders Fund 一起，曾投资过 Anthropic 的 G 轮融资。

Volumes Up, Dollars Flat 销量上升，美元汇率持平

Yet it is Ramp's own economist who is documenting the cracks in the facade. In his October AI Index - pointedly titled "Cracks in the AI Thesis Part 3" - Kharazian reported that the top 1% of AI-spending firms shelled out $9,568 per employee in September. While that was up 11.7% from August, it remained 3.1% below July's peak of $9,873.

然而，正是兰普自己的经济学家揭示了这一表面下的问题。在 10 月发布的《AI 指数》中——该报告的标题明确为“AI 理论中的漏洞 第三部分”——卡哈齐安指出，9 月份，那些在 AI 领域投入最多的企业，每名员工所对应的支出为 9,568 美元。这一数字虽比 8 月增长了 11.7%，但仍比 7 月的最高水平 9,873 美元低 3.1%。

AI spend stutters for the Top 1% of companies

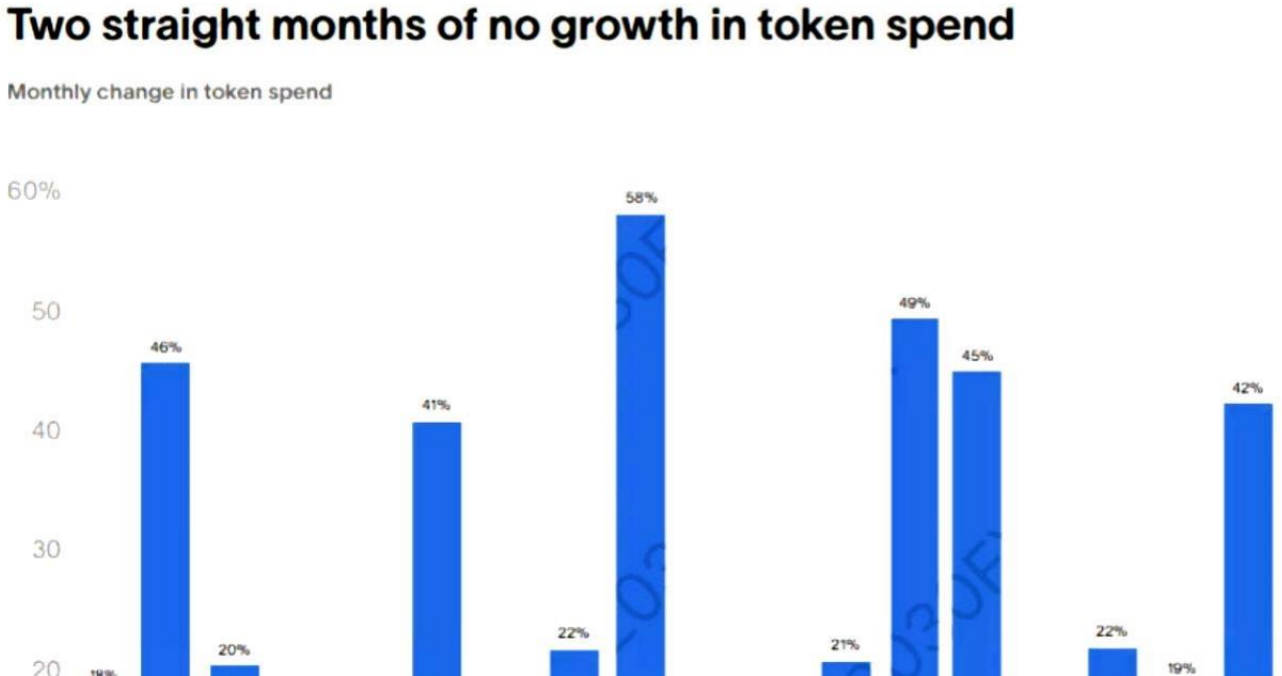


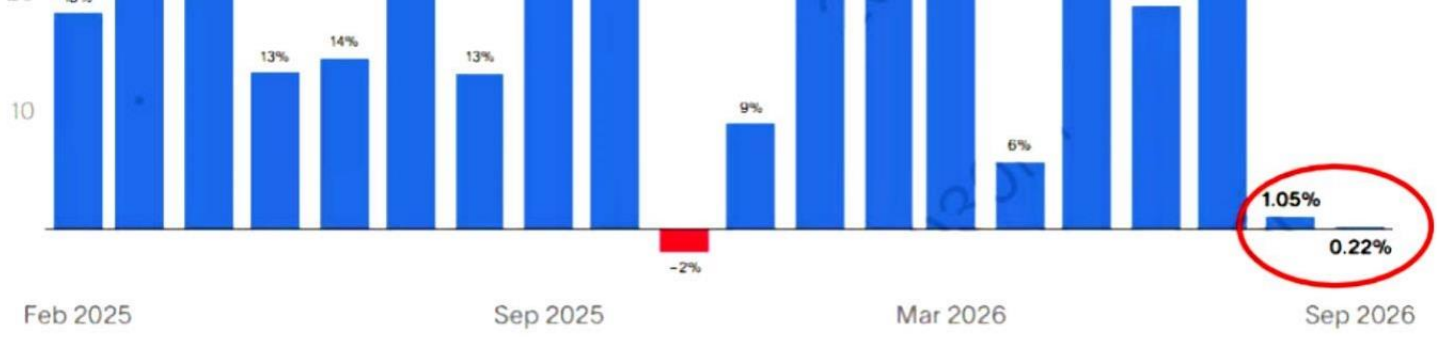
Source: Ramp AI Index, business spend data from Ramp. AI spend includes LLM subscriptions, coding agents, API tokens, and GPU cloud spend. Percentile assignment is by per employee spend on AI. • Get the data

ramp

Across the broader cohort, the stall is even starker. **In August and September, Ramp recorded "two straight months of virtually no growth in token spend,"** even as raw token volumes continued to climb.

在更广泛的群体中，这种停滞现象更为明显。8 月和 9 月期间，Ramp 的代币使用量连续两个月几乎没有增长，而代币的总体发行量却仍在持续上升。



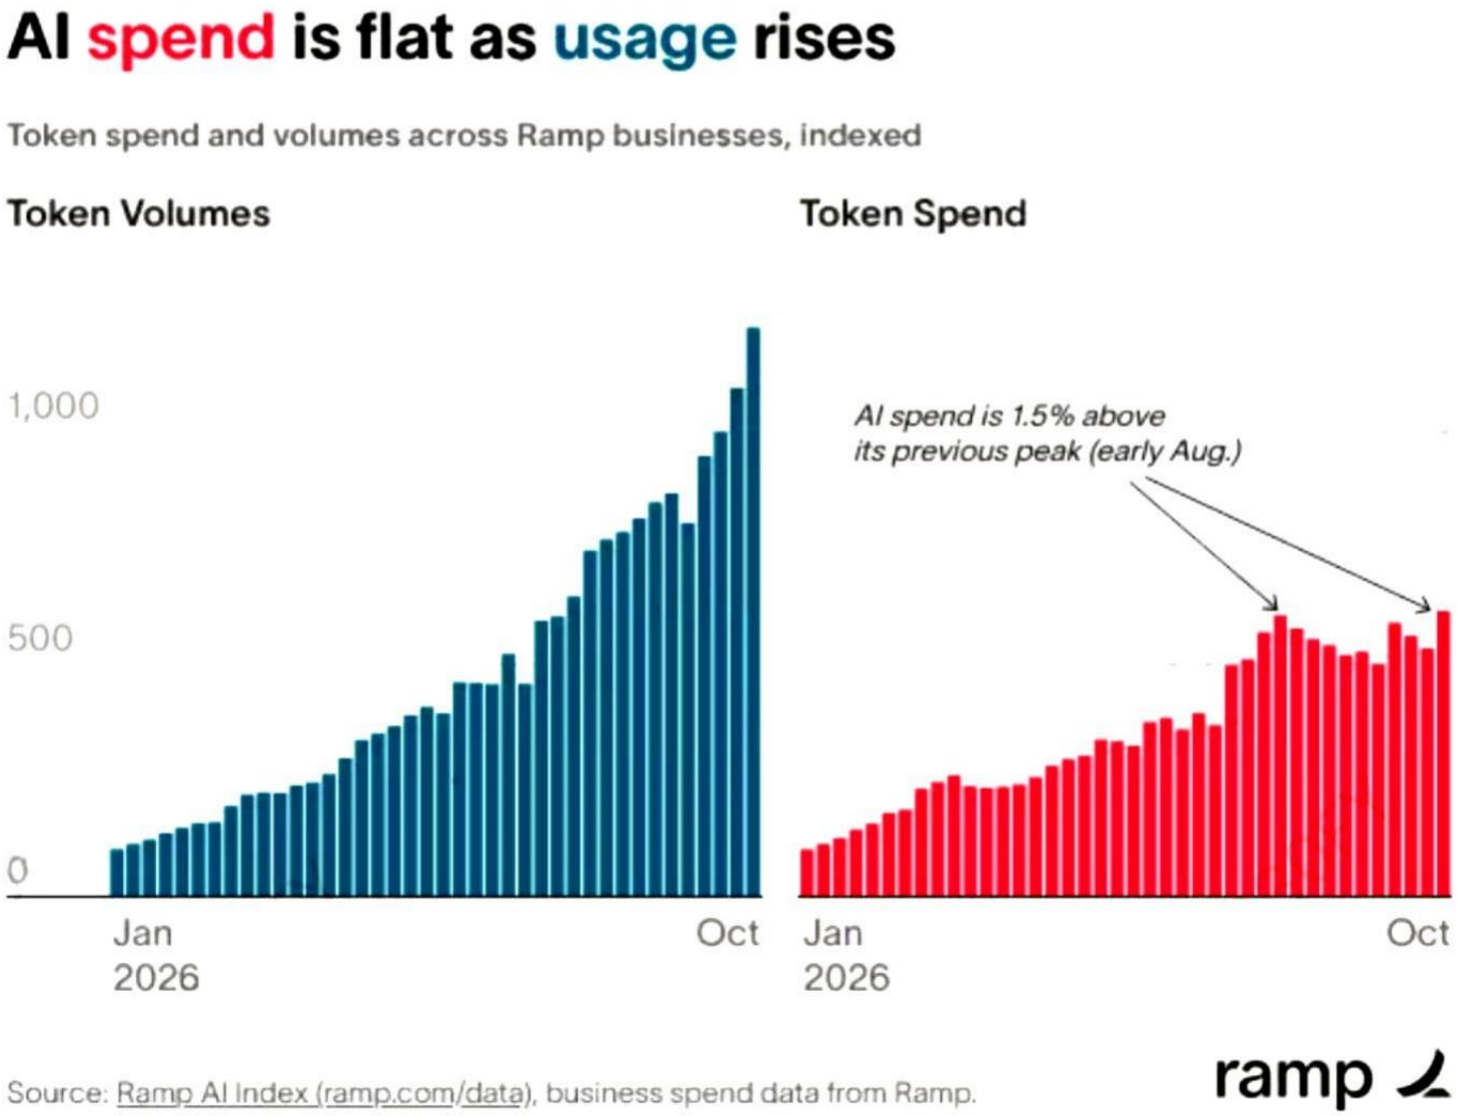


Source: Ramp AI Index (ramp.com/data), business spend data from Ramp. • Get the data

ramp

Total dollar spend did eke out a new high in early October, but cleared its previous August peak by a paltry 1.5%.

10 月初，美元支出总额再次创下新高，但仅比 8 月的最高值高出 1.5%，涨幅极为微小。



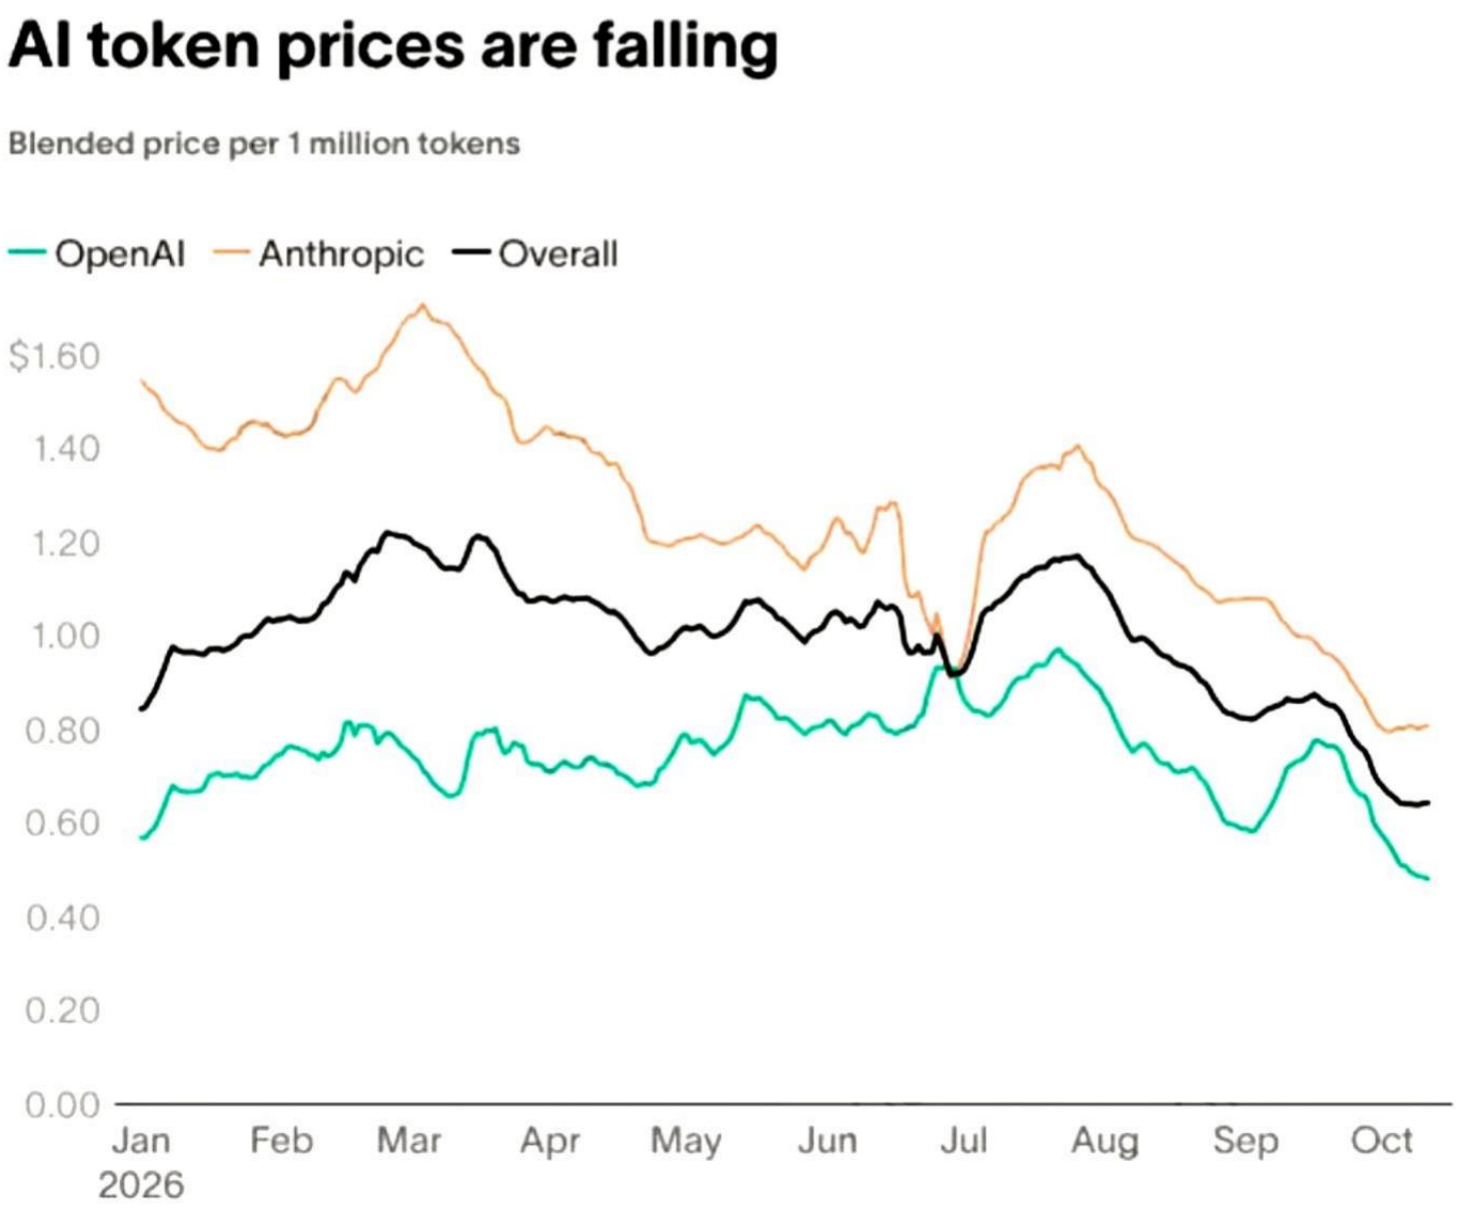
Source: Ramp AI Index (ramp.com/data), business spend data from Ramp.

ramp

Source: Ramp Rate **来源：上升速率**

The culprit, according to Kharazian, is deflation: "the main driver is price declines at the frontier, in addition to efficiency and quality improvements for models across intelligence levels, that are driving the blended price per million tokens down."

据卡拉齐安所说，罪魁祸首是通缩现象：“其主要原因在于各层级智能模型在效率与质量方面的提升，再加上前端产品价格的下跌，这些因素共同导致了每百万个标记的综合价格下降。”



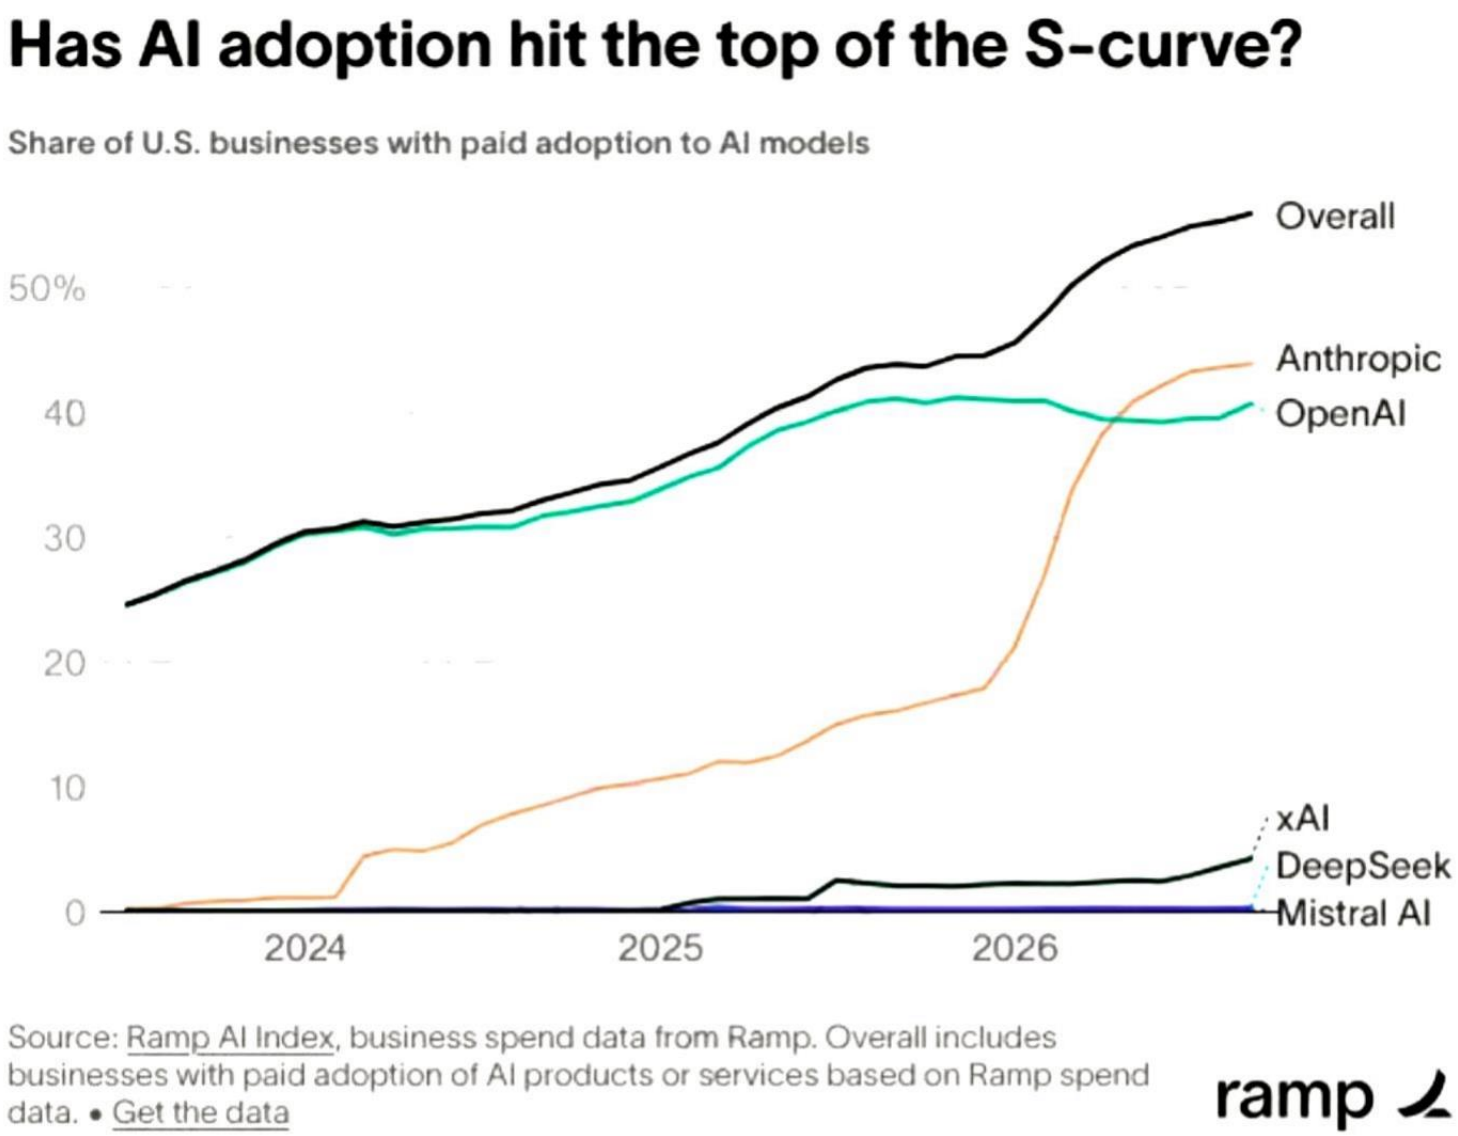
Source: Ramp Rate 来源：上升速率

Nor have the AI labs been able to lock in their biggest customers. Asked whether OpenAI and Anthropic are using enterprise contracts to guarantee revenue, Kharazian noted that buyers hold all the leverage. Enterprise deals are rarely signed on the assumption of higher-than-normal spend, and when they are, "the labs are willing to provide concessions in the form of discounts on models from the sticker price." His takeaway: "These contracts could contribute to price cuts without driving additional volume."

这些人工智能实验室也未能留住自己的核心客户。当被问及 OpenAI 和 Anthropic 是否通过企业级合同来确保收入时，Kharazian 指出，实际掌握主动权的是客户方。企业级合同很少是基于“客户会增加支出”这一前提而签订的；即便有这样的约定，这些实验室也愿意以降低模型价格的形式给予优惠。他的结论是：“这类合同虽然可能导致价格下降，但并不会带来更多的业务量。”

The customer base isn't widening much, either. Paid AI adoption across Ramp's corporate network nudged up just 0.62 percentage points to 55.9% in September. "It feels like we're at the top of the S-curve," Kharazian wrote, "where the growth in the market will have to come from more intense adoption rather than more customers."

客户群体也在没有显著扩大。9 月时，Ramp 的企业网络中采用付费 AI 服务的比例仅上升了 0.62 个百分点，达到 55.9%。Kharazian 写道：“看来我们已经处于 S 型曲线的顶点，此阶段的增长只能依靠更深入的采用率提升，而非新增客户。”



Source: Ramp Rate 来源：上升速率

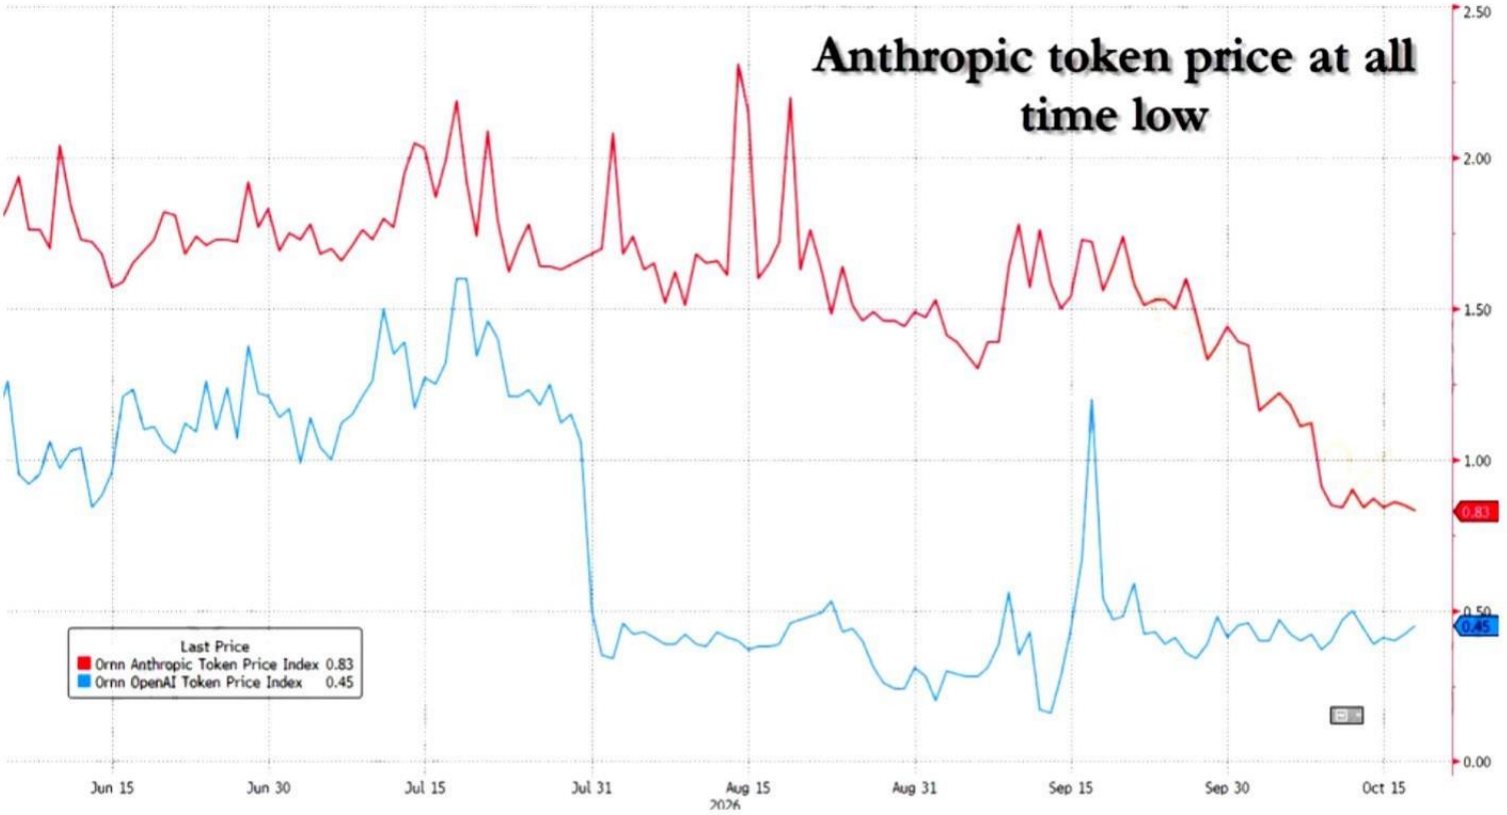
Anthropic Joins The Price War

Anthropic 加入价格战

Meanwhile, Anthropic's realized token price has been cut in half in just a month. The Ornn Anthropic Token Price Index - which measures the actual dollars customers pay per token consumed - has cratered to $0.83 per million tokens from over $1.70 in mid-September, marking a record low since the benchmark launched in June.

与此同时，**Anthropic** 的代币实际价格在短短一个月内就下

降了一半。用于衡量用户每购买一个代币所需支付的实际金额的 Ornn Anthropic 代币价格指数，从 9 月中旬的 1.70 美元/百万代币，跌至 0.83 美元/百万代币，这一数值为该指标自 6 月推出以来的最低水平。



Even after that plunge, Anthropic remains pricier than OpenAI, whose index sits at $0.45. OpenAI's realized pricing collapsed in late July after it slashed rates on its GPT-5.6 Luna and Terra models by up to 80% - just three days after Chinese lab Moonshot released the open weights for Kimi K3, the model blamed for sending chip stocks into a bear market earlier that month. OpenAI has since reportedly leaned toward delaying its IPO until 2027.

即便经历了股价暴跌之后，Anthropic 的估值依然高于 OpenAI——后者的股价仅为 0.45 美元。7 月下旬，由于 OpenAI 将其 GPT-5.6、Luna 和 Terra 等模型的价格降低了多达 80%，其股价也随之大幅下跌。而就在三天前，中国的 Moonshot 实验室发布了 Kimi K3 模型的开源代码，这一模型正是导致那个月芯片相关股票市场陷入低迷的原

因。据报道，OpenAI 现在考虑将 IPO 时间推迟到 2027 年。

Jevons Isn't Riding To The Rescue

杰文斯不会去救援。

Jevon is an offset and total demand is rising... until it hits physical constraints of hundreds of delayed or canceled data centers. https://t.co/BqONarz80n

杰冯效应正在发挥作用，总需求也在不断上升……直到遇到数百座数据中心延期建设或被迫关闭所带来的实际限制。https://t.co/BqONarz80n

— zerohedge (@zerohedge) October 14, 2026

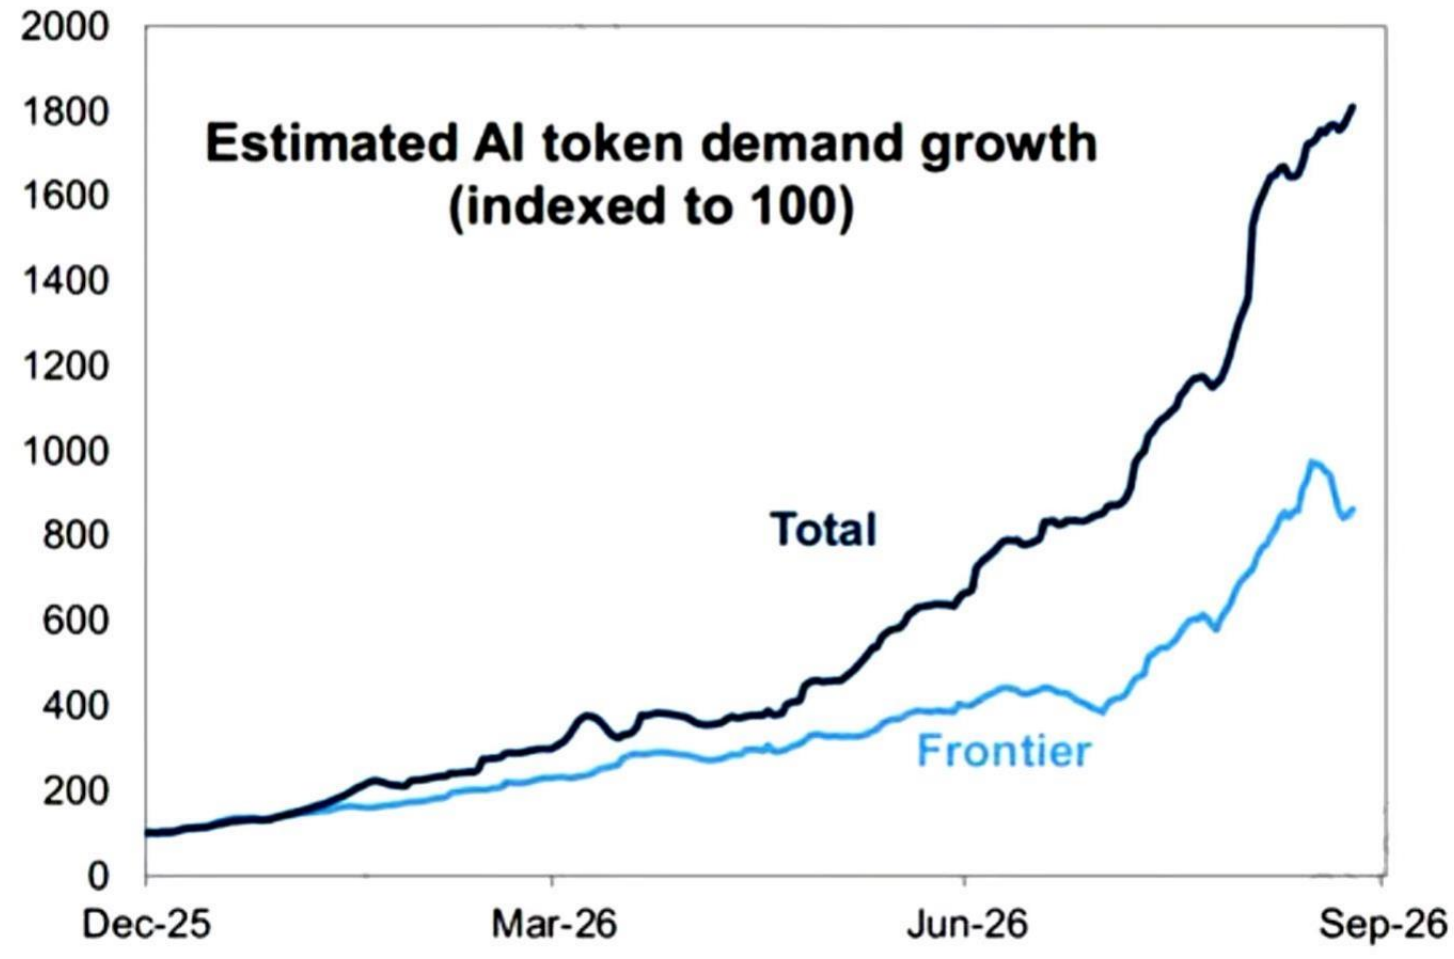
— zerohedge (@zerohedge) 2026 年 10 月 14 日

The bulls' reflex answer to relentless price cuts is Jevons paradox: make tokens cheap enough, and usage will explode so fast that total revenue rises anyway. The data isn't cooperating. In a Sept. 25 note, Goldman Sachs warned that "token demand growth will need to outpace declining token prices to support continued growth in investment spending," noting that while "frontier models are currently a key source of demand for hyperscaler compute," cheap open-source rivals are rapidly gaining ground. By Goldman's estimate Anera and OpenRouter data - total token demand roughly 18-fold since December, but demand for h frontier models has grown only half as fast.



面对无情的降价，看涨者的本能回应是杰文斯悖论：只要把代币价格压得足够低，使用量就会爆炸式增长，总收入无论如何都会上升。但数据并不配合。在 9 月 25 日的一份报告中，高盛警告称，“代币需求的增长将需要超过代币价格的下降，才能支撑投资支出的持续增长”，并指出，尽管“前沿模型目前是超大规模计算需求的关键来源”，但廉价的开源竞争对手正在迅速追赶。根据高盛的估算——基于 Anera 和 OpenRouter 的数据——自去年 12 月以来，代币总需求激增了约 18 倍，但高利润率前沿模型的需求增速只有其一半。

Exhibit 6: Estimated token demand growth
lines represent equal-weighted averages of growth implied by Anera and OpenRouter data



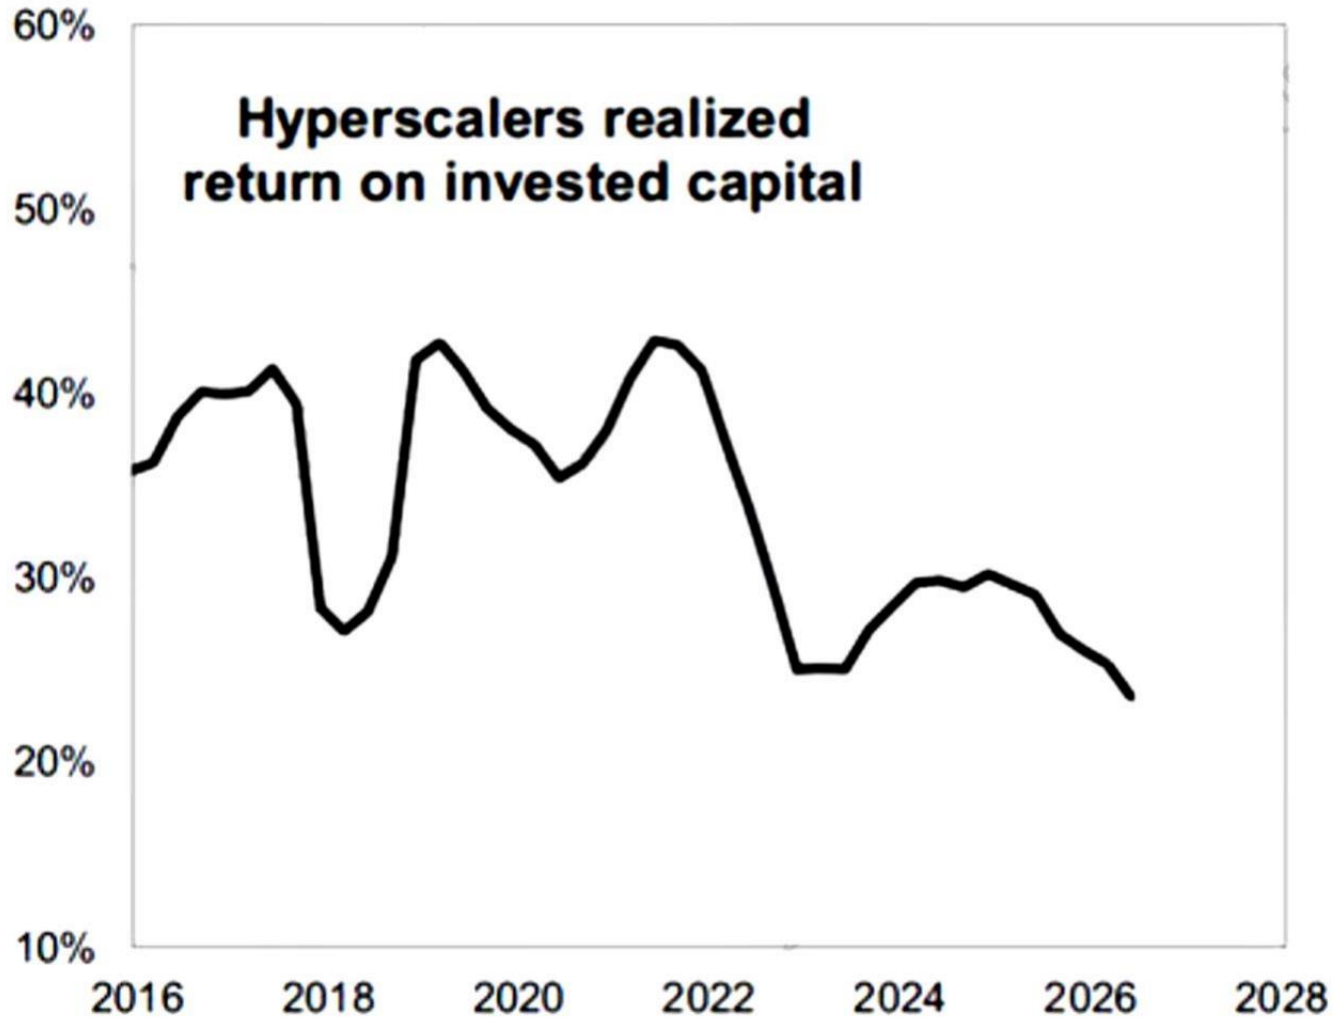
Source: Anera, OpenRouter, Goldman Sachs Global Investment Research

That margin squeeze is now flowing directly upstream to the companies building the data centers. Goldman calculates that realized return on invested capital (ROIC) across Amazon, Alphabet, Meta, Microsoft, and Oracle has fallen from 40% to

25% - the lowest reading in its dataset going back to 2016.

这种利润压力正直接传导给那些负责建设数据中心的公司。高盛测算显示，亚马逊、谷歌、Meta、微软和甲骨文这五家公司的实际资本回报率已从 40%下降至 25%，这一数值是自 2016 年以来记录中的最低水平。

Exhibit 7: Realized hyperscaler return on invested capital has declined from 40% to 25%



Hyperscalers include AMZN, GOOGL, META, MSFT, ORCL. ROIC = (EBIT less taxes) / (equity plus debt less cash).

Source: Compustat, Goldman Sachs Global Investment Research

On Tuesday - the same day Citadel's Ken Griffin told

Bloomberg that markets are "almost certainly" in an AI bubble - Apollo chief economist Torsten Slok diagnosed the underlying demand problem. "Silicon Valley operates on the premise that compute demand is unlimited, yet the recent slowdown in token spending growth suggests that this appetite may exist only when compute is effectively free," Slok wrote, citing Ramp's figures. Instead of inelastic hunger, compute faces "a flat, highly price-sensitive demand curve, where even modest changes in token prices can drive buyers to throttle usage or substitute cheap open-source alternatives."

周二，就在 Citadel 的肯·格里芬向彭博社表示市场“几乎可以肯定”正处于人工智能泡沫之中时，Apollo 公司的首席经济学家托尔斯滕·斯洛克指出了背后的需求问题。他引用 Ramp 的数据写道：“硅谷一直认为对计算能力的需求是无限的，但近期代币支出增长的放缓表明，这种需求其实只存在于计算成本近乎为零的情况下。”因此，计算能力的需求曲线并非缺乏弹性，而是一条平坦且对价格极为敏感的曲线——哪怕代币价格有轻微变化，也会导致用户减少使用量或转而选择廉价的开源替代方案。

"Ultimately, the number-one variable governing the viability of data center construction models is not physical capacity or power; it is the actual price users are willing to pay for the compute being built," Slok concluded. Corporate buyers will only pay a premium if they can translate tokens into "higher top-line growth, bigger profit margins or substantial cost savings, little of which is showing up at scale yet. Maybe that will come in the future. Maybe not."

斯洛克总结道：“归根结底，决定数据中心建设模式能否成功的最关键因素并非其物理容量或能耗，而是用户愿意为所

模显现。也许未来会有所改变，也未必。

Back in June, Ramp product director Adam Sommer boasted to a New York panel that "our largest AI spenders at Ramp are seeing their spend double month-over-month for the last several months." Three months later, the top 1% were spending less per employee than they had in July.

早在 6 月时，Ramp 的产品总监亚当·索默就在纽约的一个研讨会上宣称：“在 Ramp，那些最大量使用人工智能技术的客户，其支出在过去几个月里每月都在翻倍。”三个月后，那些排名前 1%的客户的每名员工的支出却比 7 月时有所下降。

27,742 150