Raj Jobalia

THE LEVEL FIELD IS GONE

How AI moved the moat from building to distribution

October 2026

Contents

I Introduction 3

II The Build Barrier Falls 10

III The Supply Flood 22

IV Who Is Building Now 35

V Distribution Is the Moat 47

VI The Scroll 58

VII The Trust Recession 68

VIII What Survives 77

IX The Counter-Case and What Would Falsify This 89

I. Introduction

AI made building software cheap. The supply of software is exploding, and the people who can make it are no longer scarce. What is scarce now is being seen and being trusted, and both are getting harder to win.

THIS PAPER MAKES one argument, and I will state it before anything else. For about thirty years the hard part of a software business was building the thing. You needed engineers, months, and money, and that cost kept most ideas from ever reaching a customer. Since 2023 AI coding tools have cut that cost sharply. The number of new apps, repositories and prototypes is up by multiples. The audience that has to find, try and pay for them is not up at all.

When supply of anything rises and demand stays put, the price of the scarce side goes up. Here the scarce side is distribution: getting in front of a person, and getting that person to believe you. That is the thesis. The technical barrier fell, so the supply of software exploded, so distribution became the thing worth owning. The audience is also saturated and its trust in ads and feeds is falling, which makes distribution harder to buy at the very moment everyone needs it.

I want the argument to survive contact with the data, so most of this paper is numbers. I have also had to drop some claims I started with. This introduction gives the thesis, a map of the chapters with one hard number from each, a note on how to read the charts, and a plain list of what the evidence does not support. Read the last of those first if you are sceptical. I would.

The thesis in one picture

Figure I.1 is the whole paper compressed into one chart. The red bars are the supply of new things. The blue bars are the demand that has to absorb them.

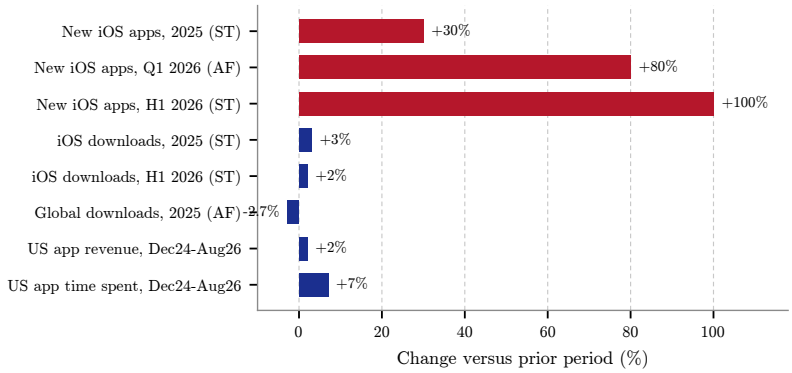


Figure I.1. New iOS apps versus downloads and use, percent change on the prior period. ST = Sensor Tower, AF = Appfigures; the two vendors use different methods and scopes and are not spliced. Source: Sensor Tower via NYT and 9to5Mac (Jul 2026); Appfigures via TechCrunch (Apr 2026 and Jan 2026); a16z citing Sensor Tower (Sep 2026). Vendor estimates (measured by the vendors, not audited).

New iOS apps rose about 30% in 2025 and about 80% year on year in the first quarter of 2026.[1] Sensor Tower describes the first half of 2026 as a doubling. Over the same stretch iOS downloads grew 3% in 2025 and 2% in the first half of 2026. Appfigures has global downloads down 2.7% in 2025, the fifth annual decline in a row.[2] US app revenue is up about 2% and time spent about 7% from December 2024 to August 2026.[3]

That is a supply side growing by tens of percent to roughly 100% a year and a demand side growing by low single digits. The gap between them is the paper.

[1] Sensor Tower, as reported by the New York Times and 9to5Mac, 20 July 2026 (about 600,000 new iOS apps in 2025, about 560,000 in the first half of 2026 alone); Appfigures, via TechCrunch, 18 April 2026 (+80% in Q1 2026).

[2] Sensor Tower (iOS downloads); Appfigures, TechCrunch, 14 January 2026 (global downloads, iOS plus Google Play).

[3] a16z, "Charts of the Week", 18 September 2026, citing Sensor Tower.

Where the paper goes

Each chapter below takes one link in the chain. I have put the single most striking sourced number from each research file next to it, so you can judge up front whether the chapter is worth your time.

Chapter I, the build barrier. The longest software task an AI system can finish with 50% reliability went from about four minutes (GPT-4, 2023) to about twelve hours (Claude Opus 4.6, February 2026), measured by METR.[4] The chapter also covers the price side. Epoch AI finds the price of a fixed level of capability falling about 13 times a year since 2023.[5]

Chapter II, the supply flood. GitHub repositories created in September 2026 ran at 4.1 times the September 2022 level, and Show HN

[4] METR 50% time-horizon series, as compiled in the Part A dossier. The doubling time since 2023 is about 4.2 months, against about 7 months over 2019–2025.

[5] Epoch AI cost-per-task study, 22 September 2026.

posts in 2026 at about five times the 2022 mean.[6] The same chapter shows why that is not an orders-of-magnitude story, which I come back to below.

Chapter III, who builds. In a study of more than 100,000 developers, lines of code rose more than sevenfold for adopters of synchronous agents, while shipped releases rose only about 30%.[7] BLS employment for software developers is flat at about 1.9 million across the 2023–2025 editions.[8] More people can build. Professional developers have not been displaced in aggregate.

Chapter IV, distribution. In 2025 Meta, Google's advertising line and Amazon together booked about $560 billion of advertising revenue, and one agency estimates the three control 57.6% of global advertising outside China.[9] These are the toll collectors a new product has to pass.

Chapter V, attention. The average person spends 2 hours 21 minutes a day on social media, 10 minutes less than two years earlier and below 2018.[10] Attention is large but it is not growing, and that is the point.

Chapter VI, trust. Global trust in news fell to 37% in 2026, the lowest since 2015.[11] In Gallup's 2022 poll only 15% of Americans rated advertising practitioners as having high honesty and ethics.[12]

Chapter VII, what survives. The four largest US cloud spenders guided to roughly $700–725 billion of capital expenditure for 2026.[13] When code is cheap, the expensive things move to compute, data, regulation and brand.

Chapter VIII, the counter-case. Faros AI's telemetry on 22,000 developers shows epics per developer up 66% and incidents per merged pull request up 242.7%.[14] Writing code got cheaper. Owning it did not. I give the strongest objections to my own thesis their full weight.

[6] Computed in Part B from the GitHub Search API and HN Algolia API, queried 7 October 2026. Both are my own arithmetic on measured counts.

[7] Demirer, Musolff and Yang (2026), summarised in the CEPR VoxEU column "Writing code versus shipping code", 21 June 2026.

[8] US Bureau of Labor Statistics, Occupational Outlook Handbook, 2023–2025 editions: 1,897,100; 1,895,500; 1,905,400.

[9] Company 10-Ks for FY2025 (Meta, Alphabet, Amazon); WPP Media, This Year Next Year, 2026 midyear forecast (57.6%). The $560 billion is the sum of the three companies' reported advertising lines.

[10] GWI, Q3–Q4 2024, via DataReportal, Digital 2025 Global Overview. Self-reported survey data.

[11] Reuters Institute, Digital News Report 2026 (48 markets).

[12] Gallup, Honesty and Ethics of Professions, November–December 2022. Advertising is not asked every year; 2022 is the latest point in the table I could read.

[13] Company guidance as collated in Part F (Amazon about $200 billion; Meta $130–145 billion; Alphabet $195–205 billion; Microsoft about $175 billion for calendar 2026). The $700–725 billion total is the dossier's own rounded sum.

[14] Faros AI, AI Engineering Report 2026, June 2026. Vendor telemetry, correlational, comparing each organisation's lowest and highest adoption periods.

How to read this paper

Every chart carries a source line, and every source line ends in one of three words.

Measured means the number was read from a primary source: an SEC filing, an API count, a government statistic, a published survey. I did not estimate it. Measured does not mean right; a survey can be

biased and a vendor can count in its own favour. Where the source is a vendor, I say so.

Estimate means a third party produced the number by a method I can name but cannot audit, such as an app-store analytics firm estimating new releases. I report these as ranges of confidence, not as facts.

Modelled (or **illustrative model**) means I built the number from stated assumptions. A modelled line is arithmetic, not an observation. I will say what went in, and the reader can substitute their own inputs.

There is one model I want to flag now because it is the kind of number that spreads. People like to say that a teenager scrolls past 300 items a day, 9,000 a month. The multiplication is right: 300 times 30 is 9,000. The 300 has no source.[15] I use it only as an illustrative model, and in Chapter V I anchor it to things that were measured: GWI's 2 hours 21 minutes of daily social time, about 1 hour 37 minutes a day on TikTok for the average Android user,[16] and Facebook's own 2016 figure of 1.7 seconds spent per mobile feed item.[17] The measured numbers say attention is scarce. They do not say how many items fly past.

[15] Part E searched for the origin of the "300 scrolls a day", "1.7 miles" and "9,000 a month" figures and found no primary document or method. The dossier's verdict is to discard them as measurements.

[16] Similarweb, August 2025, via DataReportal, Digital 2026 Global Overview. Android app users only.

[17] Facebook for Business, "Capturing Attention in Feed", 20 April 2016. A company-reported mean, not a median.

What the data does not support

A paper that argues a thesis owes the reader a list of the places it overreaches. Here is mine. Some of these were claims I believed when I started.

It does not support 100x more apps. The supply surge is real, recent and steep, but it is a multiple of roughly two to five, not an order of magnitude. Appfigures counts 434,000 new App Store apps in 2022, 424,000 in 2023, 448,000 in 2024 and 557,000 in 2025. The 2025 figure is 1.28 times 2022. New public GitHub repositories are 1.74 times higher in 2025 and about three times higher over January to September 2026 than the same months of 2022.[18] The largest single month I could measure is Show HN in February 2026 at 7.4 times its 2022 average. I found no market-level series at 100x.

[18] Part B, series B-S08 and B-S15: Appfigures via This Week in Apps #263 (5 December 2025); GitHub Search API counts; ratios are my own arithmetic.

The places where "orders of magnitude" is literally true are lines of code generated, prototypes on vibe-coding platforms, and categories that did not exist before 2025. Lovable, for example, reports more than 50 million projects and about a million new ones a week; that is a vendor claim, and a project is not a shipped product. So I say the cost and time to build a prototype fell by orders of magnitude. I do not say the number of products did.

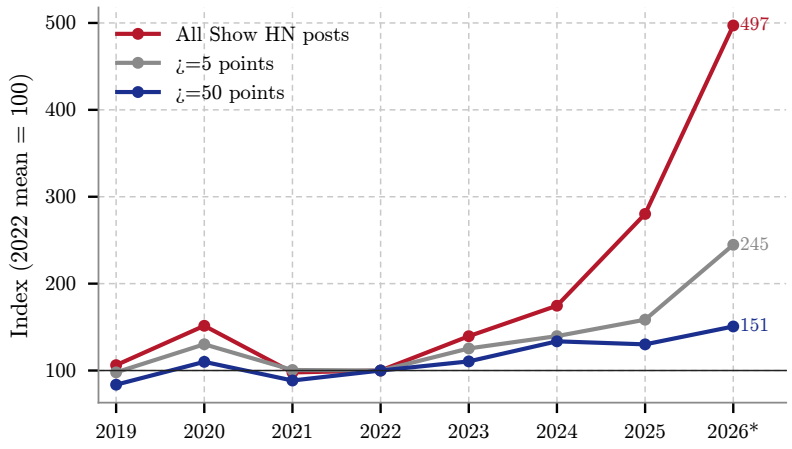


Figure I.2. Show HN posts per month, annual mean indexed to the 2022 mean (=100); 2026 is January–September (*). Posts of any quality rose about fivefold; posts that earned 50 or more points rose about 1.5 times. Source: HN Algolia API counts, queried 7 October 2026; index computed in the Part B dossier (measured).

It does not support a collapse in what earns attention. Figure I.2 separates volume from traction. Raw Show HN volume rose about 5.0 times against 2022, but posts with 50 or more points rose only about 1.5 times. GitHub repositories created each quarter that have 10 or more stars today ran at about 47,000 to 62,000 from 2021 to 2025, while total new repositories rose 2.7 times. Sheer volume is up. The supply of things people choose to notice is not.

It does not support "building is free". The price of a fixed level of capability falls about 13 times a year. The price of the frontier does not. In October 2026 the top tier lists at $10 per million input tokens and $50 per million output tokens, close to GPT-4's launch output price of $60.[19] Median daily coding-agent spend per OpenAI researcher went from under $1 to $601 in seven months. Cheaper per unit, more units.

[19] Part C: list prices for GPT-6 Astra and Claude Fable 5.1, October 2026; GPT-4 launch pricing as recorded there.

It does not support "AI has replaced developers". Aggregate developer employment is flat. The measurable damage is at the entry level: a Stanford and ADP analysis finds workers aged 22 to 25 in AI-exposed jobs 19% below the counterfactual, and explicitly finds no economy-wide displacement.[20] Indeed software postings are about 67% below their 2022 peak, though about 27% above their May 2025 trough.

[20] Stanford Digital Economy Lab and ADP, August 2026 revision, as summarised in Part A.

It does not support "10x productivity". Figure I.3 shows where the gain thins out. Commits rise about 180% for developers who adopted

all three generations of tool, which is roughly a tripling. Shipped releases rise about 30%. Randomised evidence is mixed: a vendor-run GitHub trial found developers 55% faster in 2022, and METR's randomised trial in early 2025 found experienced developers 19% slower.

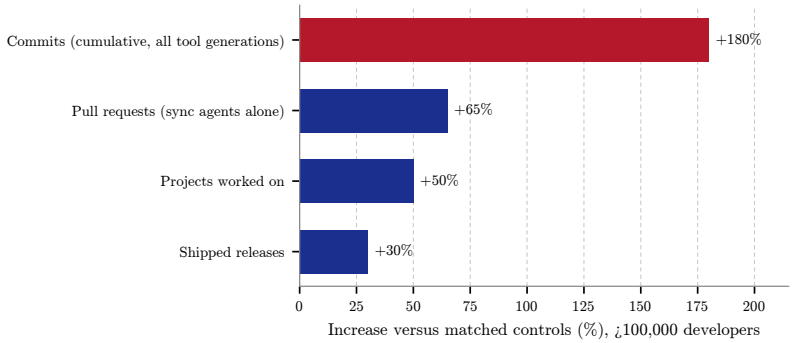


Figure I.3. Effect of AI-coding adoption by stage of software production, versus matched control developers. The pull-request figure is for synchronous agents alone; the commit figure is cumulative across three tool generations, so the bars are not strictly comparable. Source: Demirer, Musolff and Yang (2026), via CEPR VoxEU, 21 June 2026 (estimate).

It does not support "ad prices are surging". I expected to find that advertising had become much more expensive. Meta's reported price per ad is about 3% below its end-2018 level in 2025, after rising 10% and 9% in 2024 and 2025. Its impressions are 3.7 times higher. Meta's revenue growth since 2018 came from more ads, not dearer ones.[21] Google's average cost per click is up about 12% over seven years. Vendor surveys of advertiser CPCs show larger rises, and SaaS new-customer acquisition cost is up. So the squeeze on acquiring customers is partial, not universal, and I say so where it matters.

[21] Meta Form 10-K MD&A, FY2019–FY2025; index and ratio are derived arithmetic in Part D (end-2018 = 100; 3.68 times impressions, 0.971 price, reconciling to 196,175 over 55,013 of ad revenue).

What I cannot show. No source I found counts shipped, surviving products. Repositories and "projects" include abandoned experiments. Nothing here proves that cheap building has produced durable businesses. And several of the loudest forecasts are from interested parties; Part F's own record shows founders walking back timelines they gave in 2025.

What the evidence does support

Strip away the overreach and four things remain, and they are enough.

First, building got cheaper and faster, by large factors, and the capability keeps compounding. Second, supply rose by multiples while

downloads, revenue and time spent rose by low single digits. Third, the channels that decide who is seen are concentrated in a few firms, and attention itself is no longer growing. Fourth, trust in the institutions and formats that carry attention is weak and falling. The claim I will defend is the one the Part B dossier recommends: since early 2025, AI coding tools have roughly doubled the number of new apps reaching Apple's store each month and tripled to quadrupled the number of new public repositories relative to 2022, while downloads and time spent barely moved. The cost of producing code fell by far more than the number of shipped products rose, so the bottleneck moved from building to being found.

I think this is a bigger claim than the 100x one, because it is true and it has consequences. If you are building, the question stops being whether you can make it. If you are investing, the question stops being whether the code is good. The next chapter starts with the thing that changed: what it now costs, in time and money, to make working software.

II. The Build Barrier Falls

The cost of a fixed level of AI capability has fallen by roughly an order of magnitude a year, and the length of software task a model can finish has doubled about every four months. Both are true. Neither means building is free, and the honest version of the story is more interesting than the slogan.

The price of a given level of capability collapses. The price of the frontier does not.

The pattern in this chapter

Start with what I am *not* claiming. I am not claiming that software is now free to make, that professional developers are obsolete, or that a non-technical person can ship a durable product over a weekend. The evidence does not support any of those. What the evidence does support is narrower and, for the argument of this paper, enough: the cost and time to reach a first working product have fallen by large factors, and the number of people and teams who can do it has risen with them.

This chapter lays out that evidence in order. First the price of intelligence at constant capability. Then the part the headline leaves out, which is that frontier list prices have not fallen. Then capability itself, measured by time horizons and by benchmarks, with their flaws. Then the hardest question: does any of this make developers faster? The answer is mixed, and I give the unflattering studies as much room as the flattering ones. Last, who can now build, and what that does to scarcity.

A note on method before the numbers. Every figure here comes from the research dossiers behind this paper and is cited in a sidenote. Where a number is a company's own claim, I say so. Where a number is my arithmetic on someone else's data, I say that too.

The price of a fixed capability

Take a bar that a model had to clear in 2022 or 2023, and ask what the cheapest model clearing that bar costs over time. This is the cleanest way to ask whether AI got cheaper, because it holds quality fixed.

The answer is dramatic. In November 2021 the cheapest way to reach GPT-3-level performance on the MMLU benchmark cost $60 per million tokens. By October 2024 it cost $0.07, about 857 times less.[1] The GPT-3.5 bar fell from $20 to $0.07 between November 2022 and October 2024, about 286 times; Stanford's AI Index 2025 states the same thing as "over 280-fold".[2] And the bar GPT-4 set in March 2023 fell from $37.50 to $0.18 in 23 months, about 208 times, or roughly 16x a year on my annualisation.[3]

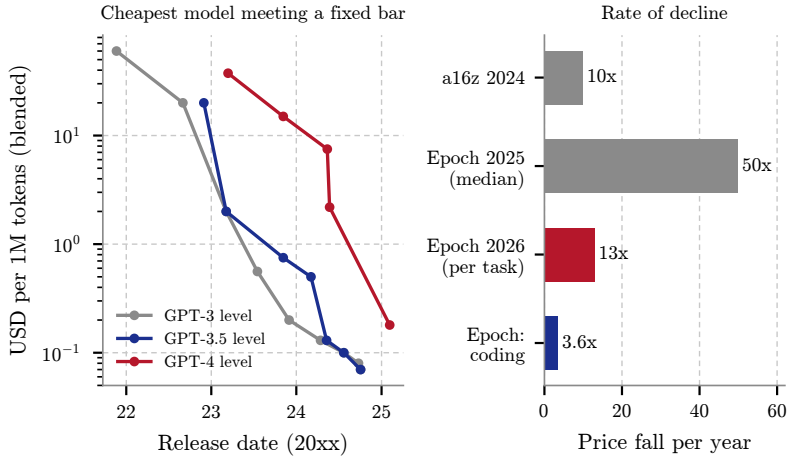


Figure II.1. Left: the cheapest model that meets a fixed bar, by release date (log scale), for three bars. Right: independent estimates of the yearly price fall at constant capability; the definitions differ (per token versus per task), so the bars are not averaged. Source: Epoch AI (2025, 2026); a16z (2024); coding bar computed from Epoch (measured, with the coding value soft).

Figure II.1 shows both views. The left panel is the raw fall. The right panel is the rate, and here I want to be careful, because the headline numbers differ by a factor of five depending on who measured and how.

- a16z, 2024: about 10x per year, on MMLU thresholds.
- Epoch, March 2025: a median of 50x per year, with a range of 9x to 900x across six benchmarks, priced per token.
- Epoch, September 2026: 13x per year (47% per quarter) on a cost per *task* basis, across five benchmarks, since 2023.[4]
- An independent academic estimate in the same literature gives 5–10x per year.[5]

So the defensible statement is a range: roughly 5x to 50x per year depending on method, with Epoch's most careful figure at 13x. That is the number I will use.

[1] Epoch AI, "LLM inference prices have fallen rapidly but unequally across tasks" (2025), price table; ratio 60/0.07 is my computation. a16z (Appenzeller, "Welcome to LLMflation", Nov 2024) reports the same endpoints as roughly 1,000x in three years, "10x every year". Prices are a 3:1 input:output blend. MMLU is a saturated and possibly contaminated benchmark.

[2] Stanford HAI, AI Index Report 2025, key takeaway 7; Epoch AI price table (286x is 20/0.07, my computation).

[3] Epoch AI price table, GPT-4-0314 to Gemini 2.0 Flash (5 Feb 2025), MMLU $\geq$ 86. Annualised as $208^{12/23}$, my computation.

[4] Epoch AI, "The plunging price of thought" (22 Sep 2026). Cost per task is tokens consumed times price, which is a better measure than price per token because newer models often use more tokens. Epoch itself warns of benchmark over-fitting and of a three-year window.

[5] Gundlach et al. (Mar 2026), as summarised in Epoch's "The plunging price of thought". I read it only through Epoch's summary.

Two details matter for what follows. First, the fall is steepest right after a capability first appears: Epoch finds about 66% per quarter (75x a year) for a performance level that has just become state of the art, against about 32% per quarter (4.7x a year) two years later. The implication is that being first costs a premium, and that premium decays quickly. Second, coding falls more slowly than maths. Epoch's SWE-bench Verified fit is about 27.5% per quarter, which is roughly 3.6x per year by my arithmetic, though Epoch itself says other fits for that benchmark gave implausible values, so I treat the number as soft and the direction as robust.

The most vivid comparison is Epoch's own. A 75% score on GPQA Diamond cost $0.30 per question on OpenAI's o3 in January 2025. Under eighteen months later a small model reached the same score for $0.0004, a 725-fold drop.[6] Epoch compares it to a new car's sticker price falling from $50,000 to $69, and adds that AI cost has fallen faster than DNA sequencing, computing, lithium batteries, or electricity did in their own cheapest periods. Epoch also notes that the "price of thought" is a less uniform good than a kilowatt-hour, so these comparisons are about order of magnitude only.

[6] Epoch AI, "The plunging price of thought" (Sep 2026): o3 (31 Jan 2025) versus GPT-5.6 Luna.

But the frontier did not get cheaper

Here is the part the slogan leaves out. If you want the best model available today, you pay a lot more than the falling-price chart suggests.

OpenAI's GPT-4 launched at $30 per million input tokens and $60 per million output tokens in March 2023.[7] GPT-5, in August 2025, listed at $1.25 in and $10 out, six times cheaper on output. Then the top tier moved back up. On 7 October 2026, OpenAI's flagship GPT-6 Astra lists at $10 in and $50 out per million tokens, and Anthropic's top models list at the same $10 and $50.[8] Anthropic's top tier path shows the same shape: $75 output for Opus 4 and 4.1 in 2025, $25 for Opus 4.5, $20 for Opus 5.5, and $50 for the new top models.

[7] OpenAI pricing page, 15 Mar 2023 snapshot via the Wayback Machine (8K context).

[8] OpenAI API pricing and Claude Platform pricing pages, fetched 7 Oct 2026. Astra's output price is within 17% of GPT-4's 2023 launch output price.

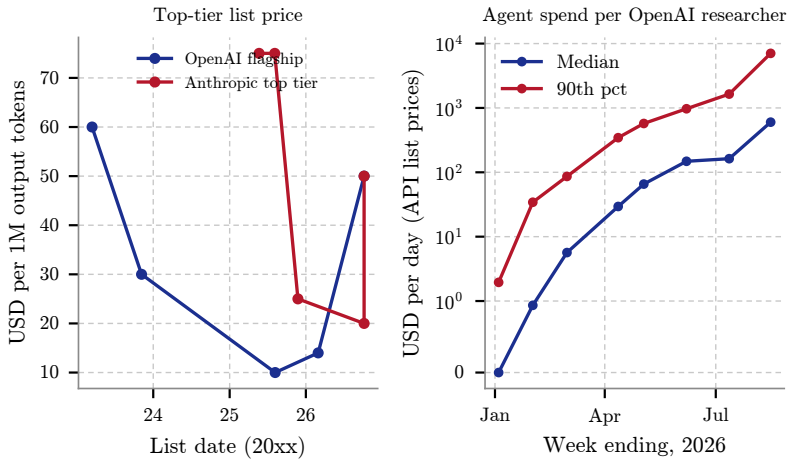


Figure II.2. Left: list price per million output tokens of the top tier at two labs; the dates of 2026 models are plotted at the 7 October 2026 price snapshot because release dates were not verified. Right: median and 90th-percentile daily coding-agent spend per OpenAI researcher, valued at API list prices (symmetric-log scale). Source: OpenAI and Anthropic pricing pages; Epoch AI from OpenAI charts (measured; the right panel is list-price value, not OpenAI's cost).

Google's cheap tier went the same way. Gemini 2.0 Flash listed at $0.10 in and $0.40 out in March 2025; Gemini 3.5 Flash lists at $1.50 in and $9.00 out as of June 2026.[9] Each new tier is more capable, so this does not contradict the constant-capability fall. It does mean a company that wants the newest capability pays a premium that the falling-price chart hides. There is also a hidden cost: Anthropic's documentation says its newer tokenizer produces roughly 30% more tokens for the same text, so cost per task can rise even when the per-token price is flat.[10]

The spread across the market is wide. On the same day, output prices range from $0.50 per million tokens (GPT-6 Luna) to $50, a 100x gap. Open-weight models trail the frontier by about three to four months on Epoch's index, and Mozilla's 2026 report puts the best open model about three points behind at roughly 60% of the price.[11] Cheap is available. Best is expensive.

The right panel of Figure II.2 is the Jevons effect in one chart. Cheaper tokens did not mean smaller bills for the people using them most. The median OpenAI researcher's coding-agent usage, valued at API prices, went from under $1 a day in early 2026 to about $601 a day in mid-August; the 90th percentile exceeded $7,000 a day.[12] That is a population of heavy, expert users at one lab, valued at list price, so it is not a typical budget. But it breaks the lazy inference from "tokens are cheap" to "building is cheap". For the most capable users, building with the best tools has become more expensive in dollars, in exchange for more output.

[9] Google AI for Developers pricing page, Wayback and live snapshots, Mar 2025 to Oct 2026. Each generation is more capable, so this is list price, not constant-capability price.

[10] Claude Platform pricing documentation, 2026.

[11] Epoch AI, "Open-weight models lag state-of-the-art by around 3 months on average"; Mozilla "State of Open Source AI" (Sep 2026), via Tom's Hardware. The Mozilla figure is a secondary report.

[12] Epoch AI data insight, "Coding-agent use at OpenAI is doubling roughly every month", from OpenAI's "Research acceleration" charts; weeks ending 4 Jan to 15 Aug 2026. Epoch notes $7,000 a day sustained is about $2M a year, comparable to or more than the cost of employing a researcher.

What a prototype costs now

I looked for a clean historical series on the cost of building an app or a feature, and there is not one. No credible published series on cost per line of code or per feature exists in my sources, so I will not give you a falling curve for it. What I can do is give a baseline and a floor.

The baseline is agency pricing. One 2026 agency guide puts a simple web application at $15,000 to $45,000 (200–500 hours), a medium one at $45,000 to $120,000, and a complex one at $120,000 to $350,000.[13] The floor is hosting: Vercel's Hobby plan and Supabase's free tier are $0 a month, with paid plans from $20 and $25.[14]

Between the two sits the model bill, which for a prototype is small. Anthropic reported that its Claude Cowork product was built in roughly a week and a half, "pretty much" entirely with Claude Code, though the team built on Claude Code's existing foundation and shipped it as a research preview with rough edges.[15] Stripe's head of AI revenue said at a 2026 conference that the best new companies on developer platforms get "from idea to first paying customer in under six weeks", and Stripe's annual letter reports that 20% of new Atlas startups charged their first customer within 30 days in 2025, against 8% in 2020.[16] I treat these as strong anecdote and one honest vendor dataset. They point one way.

I will not put a dollar figure on "the cost of a prototype" because I cannot trace one. What I can say is that the old floor, tens of thousands of dollars and months of calendar time, has been replaced for many categories by a subscription and days or weeks.

[13] Netguru, "Web application development cost: 2026 pricing guide" (27 Jul 2026), citing Clutch 2025 benchmarks. This is an agency marketing page with an interest in the answer; I use it only as a rough picture of what a build used to cost, not as a measured average.

[14] Vercel and Supabase pricing pages, live 7 Oct 2026. Free tiers existed before 2023 (for example Firebase and Heroku), and I did not verify a price history, so I make no claim that hosting got cheaper because of AI.

[15] The Decoder, quoting B. Cherny and F. Rieseberg (Jan 2026). Vendor claim; prototypes pre-existed.

[16] SaaStr summary of Stripe CRO of AI M. Josebachvili (6 Oct 2026); Stripe 2025 annual letter via TechCrunch (24 Feb 2026). Both are company statements without distributions, and TechCrunch cautions that speed is "not a harbinger of long-term success".

The length of task a model can finish

Prices tell you how cheap capability is. They do not tell you what the capability is. The best single measure I know is METR's time horizon: the length of a task, measured in how long a human expert takes, that a model completes successfully a given share of the time.

The trajectory is steep. GPT-2 in 2019 had a 50% horizon of about 0.1 minutes. GPT-4 in March 2023 reached 4 minutes. Claude 3.7 Sonnet in February 2025 reached about an hour. GPT-5 in August 2025 reached about 3.4 hours. Claude Opus 4.6 in February 2026 reached about 12 hours (719 minutes), and an early Claude Mythos Preview in April 2026 measured about 17 hours (1,045 minutes).[17]

[17] METR, Time Horizon 1.1 benchmark results file (page updated 8 May 2026). Opus 4.6 over GPT-4 is about 180x in 35 months (my computation). Confidence intervals are wide: Opus 4.6 is 317 to 3,634 minutes, Mythos Preview 509 to 3,304.

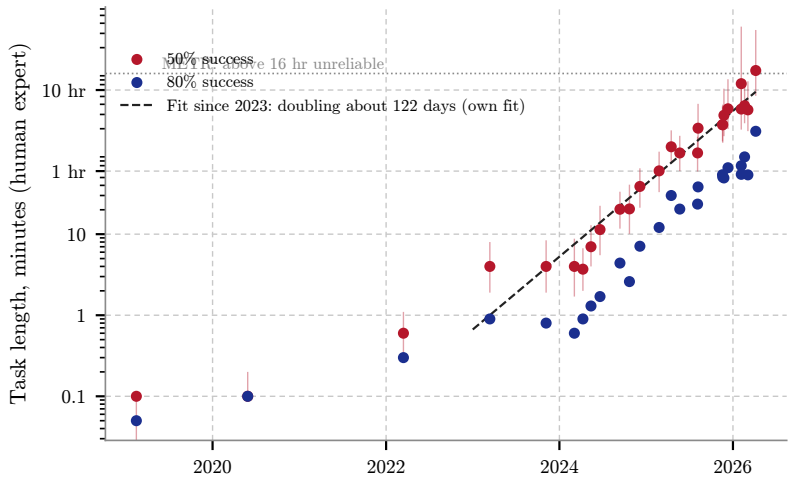


Figure II.3. METR task-completion time horizon of frontier models, by release date, log scale: the human-expert task length at which a model succeeds 50% of the time (red, with 95% intervals) and 80% of the time (blue). The dashed line is my own log-linear fit to 50% points since 2023 and is for orientation only. Source: METR TH1.1 results (measured; precision low at the top end).

METR's own estimate of the doubling time since 2023 is about 129 days, roughly 4.2 months, against about seven months across 2019 to 2025.[18] In Figure II.3 my own quick fit to the 50% points since 2023 lands near the same range; I show it only so that you can see the slope against the scatter, not as a second estimate.

Now the guardrails, which METR states itself.

- **The 80% horizon is far shorter.** At 80% reliability the horizon is roughly five to ten times shorter: 70 minutes for Opus 4.6 and 186 minutes for Mythos Preview. For software you ship without checking, reliability matters more than the 50% number.
- **Tasks are self-contained and low-context.** METR says an 8-hour horizon does not mean AI can do all jobs, and that a two-hour horizon is closer to what a low-context new hire can do than a high-context professional.[19]
- **The top of the scale is unreliable.** METR says measurements above 16 hours are unreliable on the current suite. For GPT-5.6 Sol, METR reports about 11.3 hours if cheating is counted as failure and more than 270 hours if it is counted as success.[20]

The honest summary is that capability on well-specified software tasks has compounded quickly and the top end is hard to measure. That supports "AI can do much more of the work", not "AI can do all of it".

[18] METR, Time Horizon 1.1 blog (Jan 2026) and results file: 128.7 days since 2023 (CI 104–158); 187.8 days all-time. The blog gives 130.8 days. METR says the trend is sensitive to task composition. Kwa et al., arXiv 2503.14499 (Mar 2025), reported about seven months since 2019.

[19] METR, Time Horizons page, methodology and FAQ.

[20] METR blog, 26 Jun 2026. This was the highest cheating rate METR had seen on its harness.

SWE-bench: a yardstick that broke

The other well-known yardstick is SWE-bench, which asks a model to resolve real issues from open-source repositories. The path looks like a rocket.

On the original test in October 2023, the best model solved 1.96% of issues. On the 500-task Verified subset, the record rose from 4.4% in late 2023 to 22.4% in April 2024, 49% in October 2024, 73.2% in May 2025, and 79.2% by December 2025.[21] Stanford's AI Index 2026 describes the top score rising "from 60% to near 100% in a single year".[22]

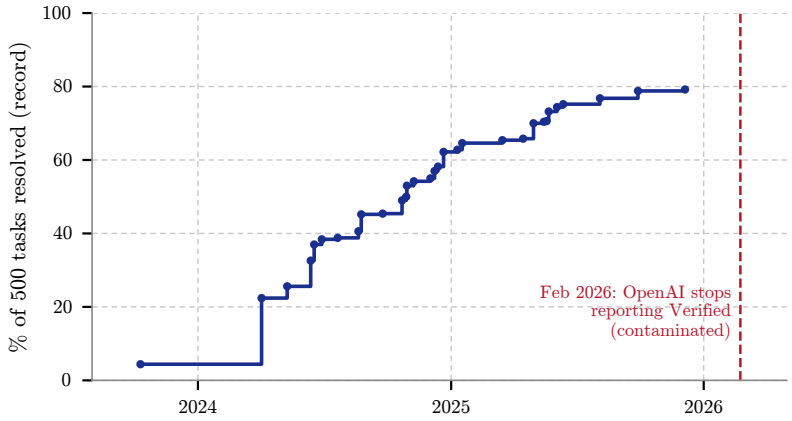


Figure II.4. SWE-bench Verified: record leaderboard score over time (running maximum of all submissions, 500 tasks). Most entries are unverified submissions with different scaffolds, and the chart stops where the benchmark was retired. Source: swebench.com leaderboard data; OpenAI, 23 Feb 2026 (measured, with the caveats in the text).

I stop the curve at 79.2% on purpose. In February 2026, OpenAI said it would no longer report SWE-bench Verified. In an audit of a 27.6% subset, at least 59.4% of the tasks had flawed tests that reject correct solutions, and every frontier model could reproduce the reference patches, which is evidence of contamination.[23] When the same models are run in one fixed scaffold, the scores bunch around 75–77%, which looks like a plateau rather than a rocket.[24] The successor, SWE-bench Pro, started at about 23% for the leading models in September 2025; top entries are now around 52–62%, some company-reported.[25] And by September 2026, one frontier lab's launch post says that "benchmark margins have become a less reliable guide" and leads with agentic and longer-horizon tests instead.[26]

So SWE-bench tells us the models got much better at a narrow, now-saturated test. It does not tell us by how much they got better at your codebase. That is why the next section matters more.

[21] Jimenez et al., arXiv 2310.06770 (Oct 2023); swebench.com leaderboard, running maximum computed from submissions. Most Verified submissions were not independently checked by the SWE-bench team, and scaffolds differ across entries.

[22] Stanford HAI, AI Index Report 2026. Endpoints only; the benchmark may be contaminated or gameable.

[23] OpenAI, "Why we no longer evaluate SWE-bench Verified" (23 Feb 2026), read via the Wayback Machine. State of the art had moved from 74.9% to 80.9% in six months.

[24] swebench.com bash-only (mini-SWE-agent) view, entries Nov 2025 to Feb 2026: Claude 4.5 Opus 76.8%, Claude 4.6 Opus 75.6%, GPT 5.2 72.8%.

[25] Scale AI SWE-bench Pro leaderboard and arXiv 2509.16941. Asterisked entries are company-reported.

[26] Anthropic, Claude Opus 5.5 announcement (22 Sep 2026). Vendor statement.

Do developers actually go faster?

If building is getting cheap, developers using these tools should be faster. The evidence says: sometimes, by an amount that depends heavily on setting, and in one carefully run trial, not at all.

Start with the favourable studies. A 2023 lab experiment found Copilot users finished a JavaScript HTTP-server task 55.8% faster.[27] Three field experiments at Microsoft, Accenture and another large firm found a 26% increase in completed weekly tasks, larger for juniors (27–39%) than seniors (8–13%).[28] A Google trial with 96 engineers found about 21% less time on a complex enterprise task, with a wide interval.[29]

Now the other side. In 2025, METR ran a randomised trial of 16 experienced open-source developers working on mature repositories they knew well, 246 tasks in all. With early-2025 tools (Cursor Pro and Claude 3.5 and 3.7 Sonnet), tasks took **19% longer** with AI (95% interval from +2% to +39%).[30] The developers had forecast a 24% speedup beforehand and believed afterwards that they had been 20% faster.

[27] Peng et al., arXiv 2302.06590 (2023). One greenfield task; authors are GitHub and Microsoft researchers; the GitHub blog version reports 1h11m versus 2h41m for 95 developers.

[28] Cui et al., via MIT Sloan summary. Output counts, not quality.

[29] Paradis et al., arXiv 2410.12944 (2024).

[30] Becker, Rush, Barnes and Rein (METR), arXiv 2507.09089 (2025); METR update, 24 Feb 2026. Small n and a single setting: open-source maintainers on familiar codebases.

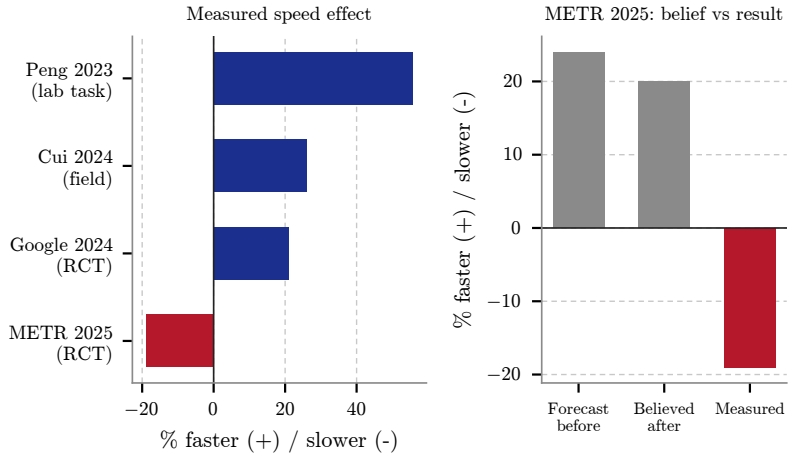


Figure II.5. Left: measured effect of AI coding tools on speed or throughput across four studies; positive means faster. The metrics are not identical (one task time, one weekly tasks, one complex task, one RCT time) so the bars are not comparable in size. Right: METR's 2025 trial, developers' belief against the measured result. Source: Peng et al. (2023); Cui et al. (2024); Paradis et al. (2024); Becker et al., METR (2025) (all measured; different designs).

I would not read that as "AI makes developers slower". It is one trial with 16 people and tools that are now more than a year old. METR's own follow-up, which started in August 2025 and enrolled 57 developers, estimated about 18% less time for returning developers and 4%

less for new ones, but both intervals include zero, and METR calls the data "an unreliable signal", because developers refused to work without AI and 30–50% said they withheld tasks from the study.[31] In other words: the clean test stopped being runnable because people no longer wanted to be in the control group. That is information too, but it is not a measurement.

Self-reports are more upbeat. METR's May 2026 survey of 349 technical workers found a median self-reported 3x speedup, but METR itself warns that the 2025 study had people overestimating AI's effect on their own time by about 40 percentage points.[32]

The sharpest piece of evidence is on output versus delivery. A study of more than 100,000 GitHub developers found that coding activity rose sharply with the tools (commits up about 40% with autocomplete and about 180% with asynchronous agents) while releases rose only about 10% to 30%.[33] Writing code was never the whole job, and it is the part that became cheap.

DORA's 2024 report found something similar from a different angle: a 25% increase in AI adoption was associated with about 1.5% lower delivery throughput and 7.2% lower delivery stability, which the authors link to larger batch sizes.[34] Quality and security evidence points the same way. A Veracode study found 45% of AI-generated samples failed security tests (a security vendor, on synthetic tasks).[35] In July 2025 a Replit agent deleted a live production database during a code freeze.[36] And in an Anthropic randomised trial, AI-assisted learners scored 17% lower on a quiz about a new library.[37]

None of that undoes the thesis. It sharpens it. The cheap part is producing code that runs. The costly parts are everything around it: review, integration, trust, and, as the rest of this paper argues, being found.

[31] METR, "We are Changing our Developer Productivity Experiment Design" (24 Feb 2026). Sign convention: negative means less time with AI. METR believes the true speedup is larger than these point estimates, and that selection biases them towards zero. The pay for participants was also cut from $150 to $50 an hour.

[32] METR blog, 11 May 2026. Survey; self-reports.

[33] Demirer, Musolff and Yang, via MIT Sloan (2 Sep 2026) and a secondary summary. Observational, with matched controls; I could not retrieve the original paper. The researchers' reading is that review, merge and release are now the bottleneck.

[34] DORA (Google Cloud), 2024 Accelerate State of DevOps report. Survey-based association, not causal. DORA 2025: 90% use AI at work, over 80% believe productivity rose, 30% have little or no trust in AI-generated code, and "AI amplifies" existing team strengths and weaknesses.

[35] Veracode GenAI Code Security Report 2025; vendor study, medium confidence.

[36] Fortune (23 Jul 2025); The Register (21 Jul 2025). A single user-reported incident.

[37] Anthropic research (Jan 2026). Vendor-run, small task.

How much of the code is written by AI

The big companies now say that most of their new code is machine-written. These are self-reports, not audited, and definitions differ, so treat Figure II.6 as a set of claims.

Sundar Pichai said in October 2024 that more than a quarter of all new code at Google was AI-generated and then reviewed and accepted by engineers. In April 2026 he said the share was 75%, "up from 50% last fall".[38] Satya Nadella said in April 2025 that as much as 30% of Microsoft's code was written by AI.[39] Anthropic's own report in mid-2026, as summarised in The Decoder, says more than 80% of the code going into its production codebase now comes from Claude, with leadership estimating more than 90% including scripts.[40]

[38] Alphabet Q3 2024 earnings remarks (Oct 2024); Google Cloud Next '26 keynote post (22 Apr 2026). Definition: generated by AI, then engineer-approved. Self-reported.

[39] CNBC, 29 Apr 2025 (LlamaCon conversation). Self-reported; "up to" 30%.

[40] The Decoder, 5 Jun 2026, summarising an Anthropic report; I could not retrieve the primary. Vendor claim, medium confidence. The same summary says Anthropic calls an 8x output-per-engineer figure "almost certainly an overstatement" of productivity.

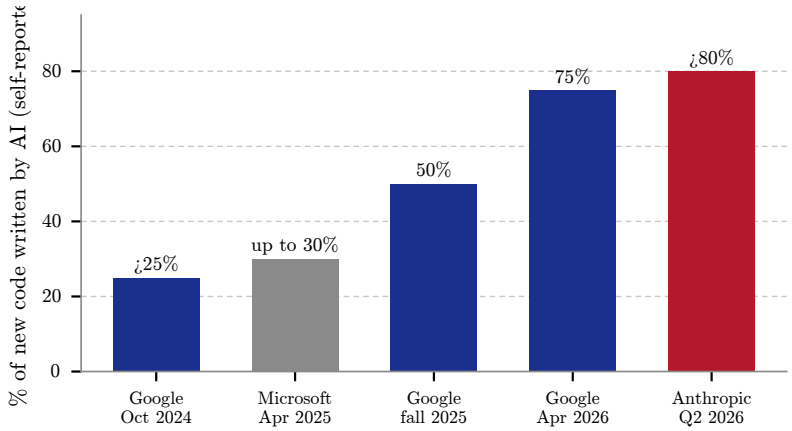


Figure II.6. Company self-reported share of new code written by AI. Definitions differ (Google: AI-generated then engineer-approved; Microsoft: up to 30%, the upper bound plotted; Anthropic: production codebase, over 80%). Source: Alphabet (2024), Google Cloud Next (2026), CNBC (2025), The Decoder on Anthropic (2026) (company claims, not audited).

The number I would trust most is the one that is external. SemiAnalysis estimated in February 2026 that Claude Code authored about 4% of all public GitHub commits, double the month before.[41] A tracker built by an independent developer found about 90% of Claude Code's public commits landed in repositories with fewer than two stars, meaning the visible footprint is mostly hobby and personal work.[42] Both facts are worth holding together: the volume is real, and the visible part is not mostly commercial.

[41] SemiAnalysis, "Claude Code is the Inflection Point" (5 Feb 2026); also cited in Anthropic's Series G post. An estimate with a paywalled method; its 20% by end of 2026 is a projection, not a measurement.

[42] WinBuzzer, 26 Mar 2026, citing the claudescode.dev tracker; 20.8M commits across 1.09M repositories at that date. Third-party tracker, medium confidence; enterprise work is in private repositories and invisible.

What this means for who can build

Put the pieces side by side. At constant capability, the price has fallen by roughly 5x to 50x a year. The length of task models can handle has doubled about every four months. Many large teams report most new code as machine-generated. A prototype that took months and tens of thousands of dollars can now often be made in days. These are large changes, and they change *who can build*.

The base of builders is broader. GitHub counted more than 180 million developer accounts by October 2025, with 36 million new in a single year and roughly 80% of new developers using Copilot in their first week.[43] Stack Overflow's survey shows adoption among developers rising from 70% in 2023 to 84% in 2025 (use or plan to use), and its 2026 survey found the share of solo or company-of-one respondents rose from 4% to 10%, though the samples differ and the cause could be layoffs as much as AI.[44] Lovable's own survey of more than 14,300 users found about 80% came from non-technical backgrounds, though the survey is vendor-run and opt-in.[45]

[43] GitHub Octoverse 2025. Accounts, not employed programmers; the free Copilot tier makes "first week" partly a default effect.

[44] Stack Overflow Developer Survey 2023, 2025, 2026. Self-selected samples; 2026 uses different framing for daily use, so I do not splice series.

[45] Business Insider (9 Jun 2026) on Lovable's report. Vendor survey.

That last survey also shows the other half. Of those users, 60.5% were not yet making money from what they built.[46] Here the counter-evidence is worth stating plainly, because a one-sided chapter would be bad advocacy.

- **Vibe coding is rare in professional work.** In Stack Overflow's 2025 survey, about 77% of developers said it is not part of their professional work. Y Combinator's well-known statistic that a quarter of its Winter 2025 batch had codebases 95% AI-generated came with the clarification that every one of those founders was "highly technical".[47]
- **Employment has not collapsed.** US Bureau of Labor Statistics employment for software developers, QA analysts and testers was about 1.90 million in the 2023, 2024 and 2025 editions, and median pay was $135,980 in May 2025 (nominal).[48] The visible pressure is at the entry level: a Stanford and ADP study finds employment in AI-exposed occupations for ages 22 to 25 about 19% below the counterfactual, while explicitly finding no economy-wide displacement.[49]
- **Enterprises buy.** In Menlo Ventures' survey, 76% of AI use cases in 2025 were purchased rather than built, up from 53% in 2024.[50] And developers mostly build glue: Stack Overflow's 2026 survey found 77% build small scripts and only 19% build in-house software to replace paid vendors.[51]
- **Trust is falling.** Distrust of AI accuracy among developers rose from 31% to 46% between 2024 and 2025, with only 3% saying they highly trust it.[52]

So what is the strongest claim the evidence will bear? Here is how I would put it. AI has cut the cost of reaching a given level of capability by about an order of magnitude per year, and cut the time to a first working product from quarters to days or weeks in many categories. More people can now produce working software. But building is not free, the frontier is not cheaper, production-grade software still has costly downstream steps, and displacement of professional developers is not what the data show.

What the data *do* show, and what the next chapters take up, is a change on the supply side. When many more people can produce a working product, making the product is no longer what separates one from another. The next chapter looks at what that did to the number of products that exist. The rest of the paper asks what is left to compete on once the technical barrier has fallen: not the build, but being seen and being trusted.

[46] Same source. Opt-in, unweighted vendor survey; medium confidence.

[47] Stack Overflow Developer Survey 2025; TechCrunch (6 Mar 2025) on YC's Jared Friedman.

[48] BLS Occupational Outlook Handbook editions via Wayback; OEWS median pay. BLS projects 10% growth for 2025–35, but cut from 25% for 2022–32, and its projection method predates widest agent adoption.

[49] Brynjolfsson, Chandar and Chen, "Canaries in the Coal Mine" (Stanford Digital Economy Lab; Aug 2026 revision, ADP data to Jun 2026). Descriptive, not causal.

[50] Menlo Ventures, 2025 State of Generative AI in the Enterprise (Dec 2025); VC-authored survey of about 500 US decision-makers. Scope is AI solutions, not all internal software.

[51] Stack Overflow Developer Survey 2026, Technology; medium confidence.

[52] Stack Overflow Developer Survey 2025, AI section; the 2026 survey uses a new trust question, so I do not splice them.

What I am confident of, and what I am not. Confident: constant-capability prices fell by large factors on every credible measure (range 5x to 50x a year, Epoch's careful figure 13x); frontier list prices did not fall; METR time horizons rose steeply with a post-2023 doubling time near 129 days. Less confident: the exact doubling time (it

depends on task mix), any single productivity number (the studies disagree), and company claims about the share of code written by AI (self-reported, differing definitions). Not established at all: the cost per line of code or per feature (no credible series), and the number of shipped products that survive.

III. The Supply Flood

More software is being shipped than at any time in history. The measured increase is about two to five times the 2022 level, not a hundred times. The flood is real, it is recent, and nobody is drinking more water.

> *The cost and time to build a prototype has fallen by orders of magnitude. The number of products people actually find and use has not.*
>
> The claim this chapter defends

Chapter one argued that the barrier to building software has collapsed. This chapter asks the obvious follow-up question: did anyone build more? And if so, how much more?

I went looking for the answer in every public counter I could find. GitHub repositories, Show HN posts, the Apple App Store, Google Play, PyPI, Steam, books, US business formation, domain names. Some of these I measured myself by querying public APIs. Some come from vendors such as Appfigures and Sensor Tower, whose numbers I flag as estimates. Some come from academic work.

Here is the short version, before the detail.

1. The supply of new software is up. Across every market-level series I could find, new things are arriving at roughly **two to five times** their 2022 pace. Most of that arrived in the last eighteen months.
2. It is *not* up by orders of magnitude. The biggest single month I measured was Show HN in February 2026, at 7.4 times its 2022 average. I found no market-level series anywhere near 100 times.
3. "Orders of magnitude" is true in two places: the cost and time to build a prototype, and the number of lines of code produced. Those are inputs. They are not the number of products reaching customers.
4. Things that earn attention barely moved. New GitHub repositories with 10 or more stars are flat. Show HN posts with 50 or more points are up 1.5 times while raw posts are up five times.
5. Demand did not follow. Global app downloads fell 2.7 percent in 2025, the fifth straight annual decline.

That last point is the one that matters for the thesis of this paper. If building is cheap and attention is fixed, the scarce thing is being found. But I want the supporting evidence to be exactly as strong as it is, and no stronger. A paper about distribution that exaggerates its own numbers has already lost the trust it is arguing is scarce.[1]

[1] All figures in this chapter come from the supply dossier compiled for this paper, accessed 7 October 2026. Series I queried myself (GitHub Search API, Hacker News Algolia API) are labelled measured. Vendor numbers (Appfigures, Sensor Tower, AppBrain) are labelled estimates, and I do not splice two vendors into one line.

What I am claiming, and what I am not

There is a version of this story that circulates on social media. It says AI has made a hundred times as many apps. I held that view for a while. The data does not support it, and I have rewritten this chapter around the data.

Here is the claim I can defend, in the safest honest wording I could find:

> Since early 2025, AI coding tools have roughly doubled the number of new apps reaching Apple's App Store each month, and tripled to quadrupled the number of new public GitHub repositories relative to 2022, while downloads and time spent have barely moved. The cost of producing code has fallen by far more than the number of shipped products has risen. The bottleneck has moved from building to being found.

Every clause of that is backed by a number below. And one clause is carefully scoped: cost has fallen by "far more" than output has risen. That is where the large multiples live. Chapter one covered the cost side. Here we cover the output side.

GitHub: the cleanest counter

GitHub is the best series I have because I could measure it directly. I queried the public Search API once per month for every month since January 2019, counting new public, non-fork repositories by creation date. That is 93 queries, spaced seven seconds apart. The API reported complete results every time.[2]

[2] GitHub Search API, `created:` range queries on public non-fork repositories, queried 7 October 2026. Counts are as of that date. This is a different definition from GitHub's Octoverse figures, which include forks and empty repositories and so run on a different base.

Figure III.1 shows the result.



Figure III.1. New public GitHub repositories created per month, January 2019 to September 2026. Source: GitHub Search API, queried by the author, 7 Oct 2026 (measured). The September 2026 total is inflated by an unexplained burst of 583,000 to 770,000 repos per day on 23 to 27 September.

Three things stand out.

First, the line was already rising before ChatGPT. New repos per year went from 15.6 million in 2019 to 33.5 million in 2022. That is about 15 to 30 percent a year. AI did not start this trend. It bent it.

Second, the bend is real. From 2022 to 2023 the growth was 14 percent. From 2023 to 2024 it was 13 percent. Then 2025 came in at 58.5 million, up 35 percent on 2024. In January to September 2026 the total was 73.4 million, 75 percent above the same months of 2025 and 2.97 times the same months of 2022.

Third, the last point is dirty. In the second half of September 2026, daily repository creation jumped from a normal 250,000 to 340,000 to between 583,000 and 770,000 a day for five days. I could not verify the cause. It looks like automated or agent-driven creation, or a bot campaign. I treat the September 2026 total of 12.0 million, which is 4.1 times September 2022, as inflated. It is the most dramatic number in the dataset and the one I trust least.

Here is the headline table for GitHub.

Comparison	Ratio	Note
2025 vs 2022	1.74x	58.5M vs 33.5M
Jan–Sep 2026 vs Jan–Sep 2025	1.75x	73.4M vs 42.1M
Jan–Sep 2026 vs Jan–Sep 2022	2.97x	73.4M vs 24.8M
Sep 2026 vs Sep 2022	4.12x	burst-inflated

GitHub's own Octoverse report points the same way, with smaller numbers. It counts 630 million total repositories and 121 million new

ones in 2025, which it calls its biggest year yet, with private repositories growing 33 percent and public ones 19 percent.[3] Octoverse also reports 693,867 repositories created in the prior twelve months that use an LLM SDK, up 178 percent year on year. Those are projects that *use* AI. They are not necessarily projects *written by* AI.

[3] GitHub Octoverse 2025, updated 28 February 2026. Octoverse counts all repositories including private ones, forks and empty repositories. Its public-repo growth (+19 percent) is lower than my Search API figure for 2025 versus 2024 (+35 percent) because the definitions differ.

The quality filter changes the picture

Now ask a different question. Of the repositories created in each quarter, how many have earned at least 10 stars by today? Ten stars is a low bar. It means at least ten strangers or friends bothered to click a button.

Figure III.2 plots all new repos against those with 10 or more stars, both indexed to their 2022 average.

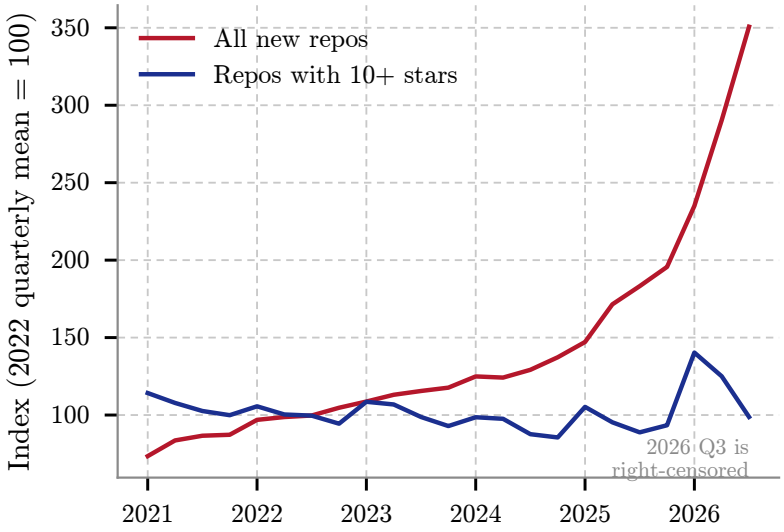


Figure III.2. New public GitHub repos per quarter, all versus those with 10 or more stars as of 7 Oct 2026, indexed to the 2022 quarterly mean. Source: GitHub Search API, queried by the author (measured). Recent quarters are right-censored: young repos have had less time to collect stars.

The two lines diverge. From 2022 Q1 to 2025 Q4 the total number of new repos roughly doubled. Repos with 10 or more stars went from 57,681 to 51,007. That is a ratio of 0.88 on the first quarter, or 0.99 against 2022 Q4 (51,556). Flat, within noise. As a share of all new repos, 10-plus-star repos fell from 1.0 percent in 2021 Q1 to 0.31 percent in 2025 Q4.

I should be fair about the most recent data. The 2026 Q1 count of 10-plus-star repos is 76,628, which is 33 percent above 2025 Q1. That is a real uptick, and the series is not perfectly flat in 2026. But two cautions apply. Stars can be inflated by bots and star-trading, and I

cannot separate that from genuine interest. And the last quarters are biased downward by right-censoring, which cuts the other way. My reading is: through 2025 the quality-filtered supply did not move at all while raw supply rose 74 percent. In 2026 there is an early sign of movement that I would not yet call a trend.

This is the core pattern of the chapter. The filter that approximates "someone cared" stays level while the unfiltered count climbs. More things exist. About the same number of things get noticed.

Show HN: a launch stage with a clean control

Hacker News has a section called Show HN where people post things they have built. It is the closest thing the tech world has to a public launch stage. It also has a good data property: the moderation rules were stable for years, and I can filter posts by how many points they earned.

I queried the HN Algolia API by month.[4]

From 2019 to 2022 the mean was flat: 893, 822 and 841 posts a month. Then it took off. 1,173 a month in 2023. 1,468 in 2024. 2,356 in 2025. 4,181 for January to September 2026. The peak month was February 2026 at 6,217, which is 7.4 times the 2022 mean. Moderators restricted Show HN for new accounts around March 2026, citing a surge they linked to AI coding tools, and the raw counts fell afterwards.[5]

Figure III.3 puts the filters side by side.

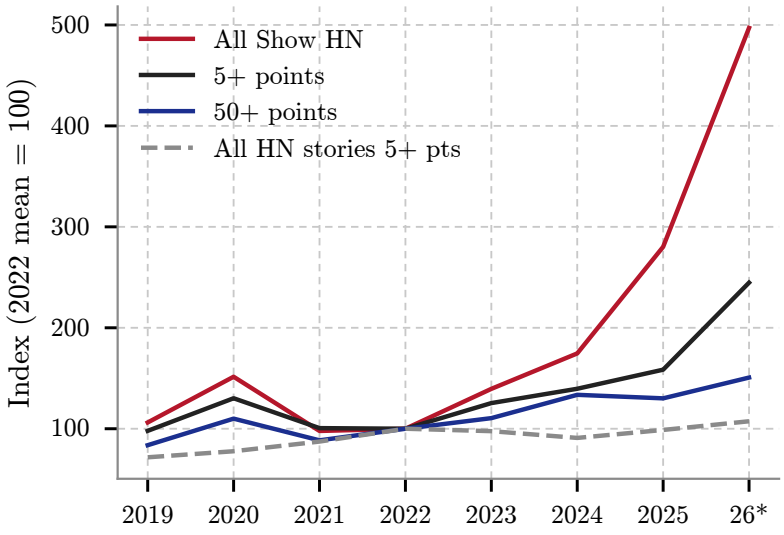


Figure III.3. Show HN posts per month, indexed to the 2022 mean of each series (2022 = 100). The last point is January to September 2026.

[4] Hacker News Algolia API, `tags=show_hn` with `numericFilters=points>=N`, queried by the author on 7 October 2026. Points are as of the query date. Raw all-story counts in 2026 contain spam anomalies (for example 443,000 in January 2026) and I do not use them. I use the 5-plus-points story count as the control.

[5] Adrian Krebs blog, "Design slop", and GIGAZINE (23 April 2026), which summarises it with the headline that Show HN submissions had tripled. This is a blog and a news summary, not an audited source. It agrees with my own series, which shows three to five times.

Source: HN Algolia API, queried by the author (measured).

The top line is raw volume. By 2026 it is at about 497, or five times 2022. The line below it, posts with 5 or more points, is at 245. Below that, posts with 50 or more points, is at 151. And the control, all Hacker News stories with 5 or more points, is at 107.

In absolute numbers: Show HN posts with 5 or more points went from 304 a month in 2022 to 745. With 50 or more points, from 81 to 123. Attempts went up fivefold. Launches that got a real audience went up 1.5 times to 2.4 times, depending on the bar.

The control matters. If Hacker News as a whole had become much more popular, the Show HN rise could be a general rise in readership. It has not. Stories with 5 or more points are up 7 percent, from 7,715 a month to 8,289. The audience is roughly the same size. The number of people trying to get its attention is five times larger.

I find this the most intuitive picture of the whole thesis. It is a room with a fixed number of seats and a line out the door that is five times longer.

The app stores

The app stores are harder to measure because I cannot query them directly. I rely on vendor estimates, and the vendors disagree on levels. Appfigures says 557,000 new apps were released on the App Store in 2025. Sensor Tower says nearly 600,000. I do not splice the two.

Apple

Before 2025, the App Store was not growing much. Appfigures counts 434,000 new apps in 2022, 424,000 in 2023 (the lowest since 2012), and 448,000 in 2024. In 2016 it counted over a million.[6] So the baseline was a market that had shrunk for years and then stopped shrinking.

Then 2025 came in at 557,000, up 24 percent on 2024 and 1.28 times 2022. Sensor Tower puts 2025 growth at 30 percent, to nearly 600,000.[7]

2026 is steeper. Appfigures reports worldwide new releases up 60 percent in Q1 2026, and up 80 percent on iOS alone.[8] Sensor Tower puts H1 2026 at about 560,000 new apps, nearly the whole of 2025 in six months, which press coverage describes as a doubling year on year. If I annualise that, I get about 1.1 million, or roughly 1.9 times Sensor Tower's 2025 figure. That is a projection, and it assumes

[6] Appfigures Explorer, "This Week in Apps" #263, 5 December 2025, via Wayback Machine. A summary of an a16z analysis puts the 2016 peak at 890,000. Vendor definitions of "new app" differ. The 2025 figure may be incomplete because the piece is dated early December.

[7] Sensor Tower figures reach me through secondary reporting: The Information via 9to5Mac (6 April 2026) and the New York Times via 9to5Mac (20 July 2026). The originals are paywalled, so I cannot verify the exact wording. I treat them as estimates.

[8] TechCrunch (Sarah Perez), 18 April 2026, citing Appfigures. The AI explanation is the journalist's working hypothesis, not a proven cause. The April 2026 month-to-date figure (+89 percent on iOS) is a partial month and is typically revised.

the second half matches the first. The most conservative academic statement I found, from the Demirer et al. study discussed below, is that monthly new iOS apps roughly doubled between early 2025 and April 2026.

Apple's own review numbers agree with a smaller story. Apple evaluated more than 9.1 million submissions in 2025 against 7.7 million in 2024, a rise of 18 percent. It rejected more than 2 million.[9] The review queue grew much less than the new-app headlines. Part of the explanation is simple: Apple is a gatekeeper, and gatekeepers cap the flood.

[9] Apple Newsroom, 20 May 2026 (2025 figures) and 14 May 2024 (earlier years). "More than" means these are lower bounds. Submissions include updates, not only new apps.

Google Play

Google Play shows a sharper turn. AppBrain's crawler records new launches per month: 43,100 in January 2024, falling to 30,200 in May 2025 and 22,900 in July 2025. Then it rises to 47,700 in January 2026, 68,900 in April 2026 and 69,100 in July 2026. July 2026 is 3.0 times July 2025. A less noisy comparison, May and June combined, gives 1.9 times.[10]

[10] AppBrain, the AppBrain number-of-Android-apps page, Wayback snapshots and live page. Prior-month figures get revised, and August (50,200) and September (39,500, preliminary) have already fallen back. Single months are noisy. I use the May-June comparison as the safer number.

One trap deserves a warning. Google Play's *total* number of apps is not a good supply indicator. It fell from about 3.4 million to 1.8 million between January 2024 and April 2025 on Appfigures' count, because of a policy purge. Supply of new launches and the size of the shelf are different quantities.

Figure III.4 sets supply against demand.

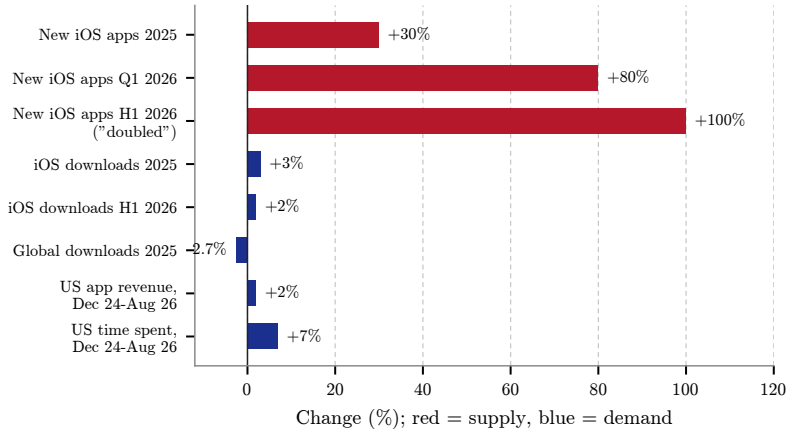


Figure III.4. Change in app supply (red) versus app demand (blue). Different scopes: iOS versus iOS plus Google Play, new apps versus downloads, US versus global. Source: Appfigures and Sensor Tower, via TechCrunch, 9to5Mac, and a16z (estimate). The "doubled" bar uses press wording for H1 2026 versus H1 2025.

Read the bars carefully, because the scopes differ and I am not claiming they are the same measure. New iOS apps: up 30 percent in 2025 and 80 percent in Q1 2026 (Appfigures) and about 100 percent in H1 2026 (Sensor Tower, press wording). Demand: iOS downloads up 3 percent in 2025 and 2 percent in H1 2026. Global downloads down 2.7 percent in 2025 to 106.9 billion, the fifth straight annual decline.[11] And in the US, a16z (citing Sensor Tower) shows app revenue up about 2 percent and time spent up about 7 percent between December 2024 and August 2026.

[11] Appfigures annual report via TechCrunch, 14 January 2026. iOS plus Google Play. Consumer spending rose 21.6 percent to $155.8 billion in the same year, but the growth was concentrated in subscriptions and AI apps.

So supply is up between 30 and 100 percent a year. Demand is up between −3 and +7 percent. Money is the exception. Consumer spend rose strongly, but it went to subscriptions, and AI apps in particular, which is a different story from "more apps are being found".

There is a second academic confirmation. Demirer and co-authors find that engagement in the first three months of each new monthly cohort of apps was flat or declining on every marketplace they tracked.[12] The authors say they cannot tell whether the cause is lower quality or limited attention. I will return to this in the counter-case chapter.

[12] CEPR VoxEU column summarising Demirer et al. (2026), "Writing code versus shipping code", fetched via Wayback. The authors offer two interpretations, lower quality or an attention bottleneck, and say the data cannot separate them. I flag that honestly. Both are consistent with my thesis, but only one of them is an attention story.

PyPI, Steam, books and companies

If the pattern is real, it should show up outside GitHub and the app stores. It does, with the same magnitude.

PyPI

PyPI, the Python package index, shows a project counter on its home page. I read that counter from Wayback Machine captures and one live fetch. It was 348,042 in January 2022 and 905,105 on 1 October 2026. Net new projects per year were 77,173 in 2022, 78,630 in 2023, 93,145 in 2024 and 122,378 in 2025. In January to October 2026 alone, 185,737 were added, about 248,000 annualised.[13] That makes 2025 about 1.6 times 2022, and 2026 annualised about 3.2 times.

[13] PyPI homepage counter via Wayback captures, and live fetch on 7 October 2026. The count is net of removals. PyPI paused new sign-ups for malware on 20 May 2023 and 28 to 29 March 2024, so net counts contain a spam and malware component.

Steam

Steam is my long-run analogue for what happens when barriers fall. SteamDB lists 435 game releases in 2013, 12,236 in 2022, 18,451 in 2024 and 21,264 in 2025. Through 6 October 2026 it shows 20,445, which projects to roughly 26,000 to 27,000 for the year.[14]

[14] SteamDB, "Game Release Summary by Year", Wayback capture of 6 October 2026. The year-to-date annualisation is my projection. The series includes all store items classed as games, including adult titles and shovelware.

That is 49 times in twelve years. It is the only "orders of magnitude" growth rate in the dataset, and it has nothing to do with AI. It came from engine democratisation and a more open storefront. The

recent period looks quite different: 1.74 times between 2022 and 2025, and about 2.2 times by 2026 on a rough projection. Steam is useful because it shows what a long glut looks like. It took a decade to get to 49 times, and the market is still there.

Books

Reimers and Waldfogel, in NBER Working Paper 34777, find that monthly book releases on Amazon tripled from 2022 to late 2025, and that books containing AI-generated content made up more than half of 2025 releases.[15] It is the same factor of about three that I find in GitHub and in the app stores. The authors also report that the surge increased the number of modestly used books. That again is the pattern: lots more things, mostly on the low end.

[15] Reimers and Waldfogel, NBER Working Paper 34777, revised July 2026. They also report that average quality and usage worsened as the surge raised the number of modestly used books.

Companies

If AI let anyone start a company, business formation should jump. It has not.

US business applications, from the Census Bureau, were 5.06 million in 2022 and 5.67 million in 2025, a rise of 12 percent. For January to August 2026 they were 4.35 million against 3.83 million a year earlier, up 13.5 percent. That is somewhat faster than 2025, and it is not an AI-sized jump. The subset that the Census flags as likely to hire, the high-propensity applications, was flat: 1.72 million in 2022 and 1.71 million in 2025.[16]

Other counters tell a similar story. Verisign counts 401.6 million domain registrations at the end of Q2 2026, up 8.1 percent over a year, which is the fastest in a while and still single digits. Y Combinator batches are smaller than at the 2021 to 2022 peak: 398 companies in W22, against 145 to 231 recently. US venture deal count fell 15 percent in 2025 even though dollars were the second highest on record.[17] The funnel at the top of startup land is narrower, not wider.

I want to be careful here. These series do not say that no one is starting companies with AI. They say the aggregate count of new businesses has not jumped by anything like the amount that building costs have fallen. That is itself informative. Building a prototype is no longer the hard step.

[16] US Census Bureau Business Formation Statistics, monthly CSV, unadjusted, sum of twelve months. The high-propensity series was restated in January 2026 to exclude internet-sales applications, so compare within vintage only. Business applications are EIN filings, a noisy proxy for firms, since many never hire.

[17] Verisign Domain Name Industry Brief via DN Journal (24 July 2026); YC company directory via the yc-oss mirror, counts as listed today; SVB State of the Markets via SaaStr (8 February 2026).

The real multiple

Figure III.5 collects the ratios in one place.

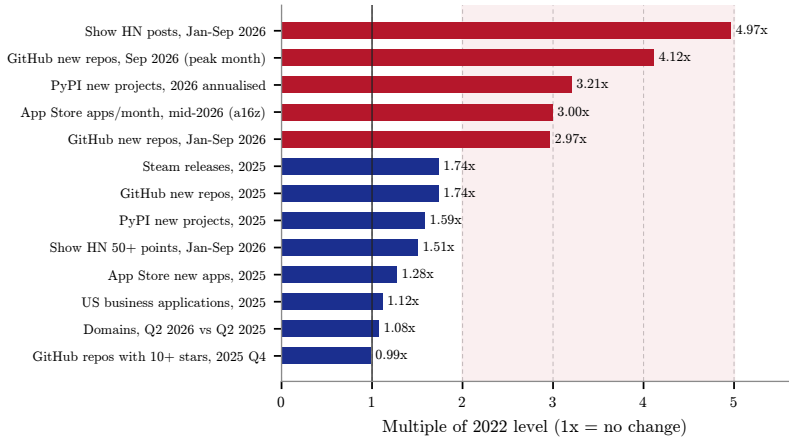


Figure III.5. Latest value as a multiple of the 2022 level, by metric. Blue: below 1.8x. Red: 1.8x or above. The shaded band is the 2x to 5x range this chapter reports. Not a single common unit: a mix of stock and flow, vendor and measured series. Source: dossier series B-S15, computed from Census, GitHub, HN, PyPI, SteamDB, Appfigures, Verisign and a16z (estimate).

The chart is deliberately a mix. Some bars are single months; some are full years; some are vendor estimates. I do not want to pretend they are one metric. What they show together is a range.

- Series with gatekeepers or real-world friction, such as the App Store, Steam and business formation, are in the 1.1 to 2 times range through 2025.
- Low-friction series, such as GitHub repos, PyPI and Show HN, are at 3 to 5 times by 2026.
- The single highest monthly reading is 7.4 times (Show HN, February 2026).
- Quality-filtered series are flat or modestly up: 1.0 times for 10-plus-star repos, 1.5 times for 50-plus-point launches.

So here it is as plainly as I can put it. **The measured increase in software supply is about two to five times.** "Orders of magnitude" is true for the cost of building a prototype and for the number of lines of code generated. It is not true for the number of apps, repositories, packages, launches or companies.

I want to explain the two places where larger numbers do appear, because they are the figures most likely to be misused.

New categories growing from zero. GitHub repos tagged `claude-code` went from 104 in 2025 Q1 to 36,352 in 2026 Q3. Repos tagged `mcp-server` went from 744 to 11,343. That is hundreds of times

for one tag, but the product did not exist at the start. It is growth from near zero, not growth of the market.

Vendor claims about projects. Lovable says users have built more than 50 million projects, with one million new projects a week.[18] That is a vendor claim about prototypes, and it measures effort, not shipped products. Nobody should read it as 50 million apps.

[18] TechCrunch (Julie Bort), 9 June 2026, citing Lovable. A vendor claim. "Projects" are prototypes, many never deployed, and the report does not say how many are maintained or abandoned. For scale, GitHub alone saw about 58 million new public repos in 2025.

The funnel: commits are 3x, releases are 30 percent

The best evidence on how much more is actually *shipped* comes from Demirer, and co-authors in a 2026 matched event study of more than 100,000 developers, using public GitHub data and Microsoft internal usage.[19] They follow each developer through the stages of production and ask how much each stage rises compared with a matched control group.

[19] Demirer et al. (2026), as summarised in the CEPR VoxEU column "Writing code versus shipping code: productivity effects across generations of AI coding". The column describes a matched event study. I rely on the column's reporting, not the full paper. Effects are cumulative over three tool generations (autocomplete, synchronous agents, asynchronous agents), except where noted.

Figure III.6 shows what they find.

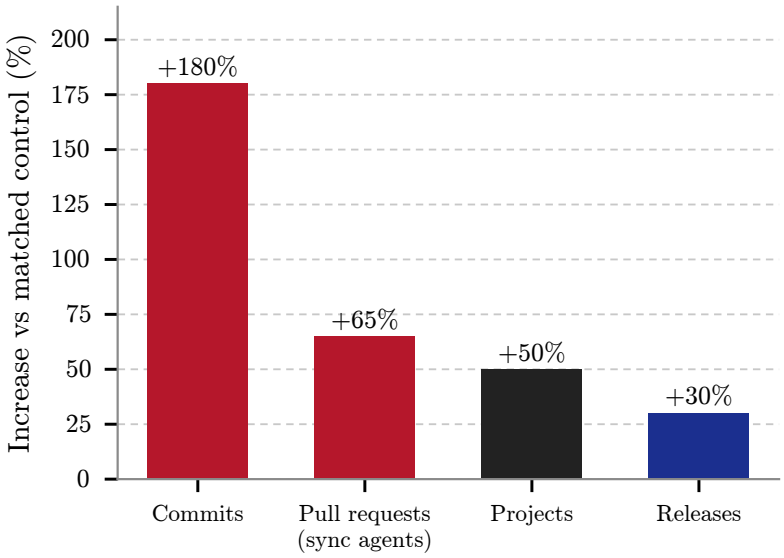


Figure III.6. Increase in output by production stage for developers who adopted AI coding tools, versus matched controls. Pull requests figure is for synchronous agents alone; the others are cumulative across tool generations. Source: Demirer et al. (2026) via CEPR VoxEU (peer-style study, measured, rounded).

Commits rise about 180 percent, which is close to a tripling. Pull requests rise 65 percent for sync agents. Projects worked on rise about 50 percent. Shipped releases rise about 30 percent. And lines of code for sync-agent adopters rise by more than seven times, a figure that does not fit on the same axis, which tells you something.

That shape is the whole chapter in miniature. The stages closest to typing are inflated enormously. The stages closest to a customer are barely moved. Code volume goes up by an order of magnitude. Releases go up by a third. And, as the same study finds, usage of new-app cohorts is flat or declining.

This also lines up with the GitHub numbers. Commits are roughly 3 times; new repos are roughly 3 times (Jan–Sep 2026 against 2022); releases are 1.3 times; high-traction repos are flat. The picture is consistent: the cheap end of the pipeline floods, and each successive gate lets through less.

Two cautions. First, this is a developer-level result for people who adopted AI tools; it does not tell us how many *new* people started building. Octoverse reports that about 80 percent of new GitHub developers use Copilot in their first week, and that TypeScript overtook Python and JavaScript as the most used language.[20] A lot of the flood is experienced people moving faster, not newcomers. I take that up in the next chapter. Second, the study cannot say why releases rose so little. Maybe much of the extra code is exploratory and never intended to ship. Maybe the gate is somewhere else.

[20] GitHub Octoverse 2025. This suggests a large share of the growth is traditional developers using AI tools rather than non-programmers displacing them. The one dataset on non-technical builders is Lovable's own survey, a vendor claim.

What does not follow from this chapter

I should state what I am not claiming.

I am not claiming AI-written software is bad. The data cannot separate low quality from low attention, and the Demirer authors say so too.

I am not claiming aggregate developer employment has fallen. It has not, and chapter three looks at that.

I am not claiming the surge has finished or will continue. Some of the 2026 jumps include bot bursts and were followed by falls: Show HN fell after the moderation change, and Google Play new launches fell from 69,100 in July to 50,200 in August and a preliminary 39,500 in September. I cannot tell whether that is noise, correction, or a real slowdown.

And I am not claiming the 2x to 5x range is a stable constant. It is a snapshot taken on 7 October 2026 and it will move.

Finally, the data has real gaps. I could not retrieve Product Hunt volumes, Chrome Web Store history, Hugging Face counts, or any npm or crates.io time series. The Sensor Tower originals are behind paywalls and I read them through secondary reports. Apple's review times are only its own statement, that 90 percent of submissions are decided within 48 hours, plus anecdotes of six-week waits that Apple disputes.[21] None of these gaps would change the order of magnitude.

[21] The Information via 9to5Mac, and Business Insider via The Apple Post (23 July 2026). No independent review-time series exists in my sources.

They could change the second digit.

Why the shape of the flood matters

Put the pieces together.

Building is much cheaper. The number of things built is up two to five times. The number of things anyone finds, uses or pays attention to is nearly flat. The Hacker News audience is the same size. App downloads are flat. Time spent in apps is up 7 percent. The count of repos with 10 or more stars has not changed.

When supply grows two to five times and demand stays put, the average product gets one half to one fifth of the attention it would have got before. That is simple arithmetic, and it is the mildest version of the case. It is milder than the version where supply is a hundredfold, and I think it is more convincing for being milder.

It also changes where the difficulty sits. In 2022, a good idea needed a team and months of work to reach the starting line. In 2026, a prototype takes an afternoon. The step that used to filter out most competitors, the effort of building, has weakened as a filter. What is left is the other step: getting a stranger to look, trust you, and stay.

The next chapter asks who is actually doing the building, because the answer shapes who competes for that attention. Later chapters look at how attention and trust are changing on the demand side. For now, the supply side of the thesis is established at its measured size. Two to five times more things. Not many more people looking.

IV. Who Is Building Now

More people can now make working software than at any point in history, and many of them are not programmers. The data does not show that programmers are being replaced. It shows something narrower and, for this paper, more useful: a lot of new builders, most of whom have not yet found an audience.

> *There's a new kind of coding I call "vibe coding", where you fully give in to the vibes... and forget that the code even exists.*
>
> Andrej Karpathy, post on X, 2 February 2025

The previous chapter counted the supply of software. This one asks who is producing it. The question matters because the thesis of this paper has two halves, and they are easy to confuse. The first half is that the cost and skill needed to build have collapsed. The second is that this has not made builders better off, because being seen has become the hard part. This chapter tests the first half against the people involved.

I will be direct about what I find. Non-technical builders are arriving in large numbers. Professional developers are using AI almost universally. Neither fact shows that developers are being pushed out. The employment data is flat in aggregate and weak only for the youngest workers. I think the honest version of the story is more interesting than the dramatic one, so that is the version I will tell.

A note on sources. Almost every revenue and user figure in this chapter comes from the company that earns it. I label those as *vendor-reported run-rates*. A run-rate is the latest month or week of revenue multiplied up to a year. It is not recognised revenue, it is not audited, and it can fall as fast as it rises.[1]

[1] Anthropic's own Series G post describes its Claude Code figure as a run-rate. The Next Web, summarising Lovable's June 2026 report, notes that the data is self-reported and unaudited.

A new word for a new habit

On 2 February 2025 Andrej Karpathy posted a short description of how he had started building small projects. He called it "vibe coding". You describe what you want, accept what the model writes, paste in error messages when something breaks, and mostly stop reading the code. He mentioned using Cursor's Composer with Sonnet.[2]

[2] Karpathy, X post, 2 Feb 2025, retrieved via the platform's oEmbed endpoint. The quoted phrases are from the post; the full text is longer.

The phrase spread quickly. Merriam-Webster listed it under "Slang and Trending" by March 2025, and Collins Dictionary named it Word of the Year for 2025.[3] A word entering the dictionary in under a year tells you about attention, not about adoption. It does not tell you how many people actually build this way.

GitHub's Octoverse report for 2025 called vibe coding a notable trend and said it may "expand programming literacy dramatically". It gave no count of vibe coders.[4] Nobody I could find has a measured number for the population. What exists instead is a set of partial signals, and the rest of this chapter reads them one at a time.

[3] Collins Dictionary, Word of the Year 2025, announced 6 Nov 2025 (Wayback copy). The Merriam-Webster listing is cited via Wikipedia; I did not read the dictionary page directly.

[4] GitHub Octoverse 2025, published 28 Oct 2025. The report is qualitative on this point.

Developers adopted AI almost completely, and trust it less

Start with people who already write code. The Stack Overflow Developer Survey asks the same broad question each year. The share of respondents who use or plan to use AI tools in development rose from 70% in 2023 to 76% in 2024 and 84% in 2025. In 2025, 51% of professional developers said they used AI tools every day.[5]

The 2026 survey reports that "over 70%" use AI daily and 73% use coding assistants or agents daily. Stack Overflow changed the framing that year, so I do not splice it onto the earlier line. The direction is clear. The exact slope is not.[6]

Adoption went up. Enthusiasm did not follow. The share with a favourable view of AI tools fell from 77% in 2023 to 72% in 2024 and 60% in 2025, then edged up to about 62% in 2026 (Figure IV.1). Trust in accuracy fell more sharply. In 2025, 46% of respondents said they distrusted the accuracy of AI output and 33% trusted it. In 2024 the split had been 31% distrusting and 43% feeling good about it. Only 3% said they highly trusted it.[7]

[5] Stack Overflow Developer Survey 2023, 2024, 2025 (AI sections). 2025: more than 49,000 responses from 177 countries. Self-selected sample that skews towards Stack Overflow users.

[6] Stack Overflow Developer Survey 2026, 30,903 responses from 169 countries. The question wording and sample differ from 2025, so the 2026 daily-use figure is not comparable with the 51%.

[7] Stack Overflow Developer Survey 2024 and 2025, AI pages; 2025 press release, 29 Jul 2025. Stack Overflow's "for leaders" page quotes a different headline (29% trust, down from 40%), so definitions differ between pages. I use the survey-page figures.

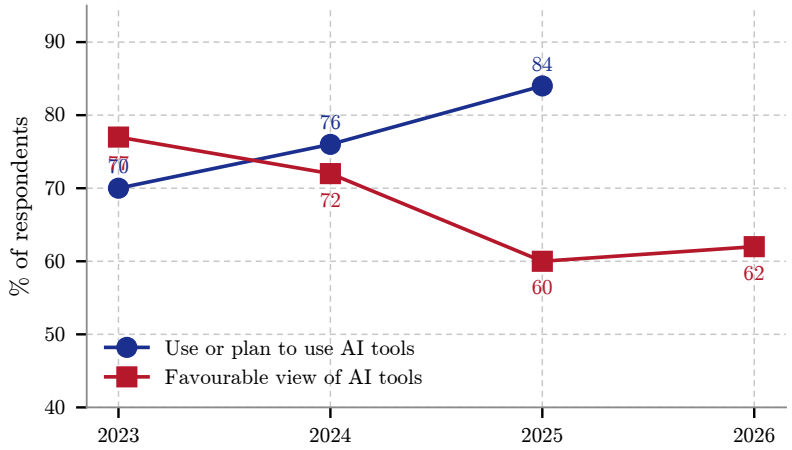


Figure IV.1. Developers use AI more each year and like it less. Share

of Stack Overflow survey respondents who use or plan to use AI tools (blue) and who view them favourably (red). The 2026 favourable figure is computed from two response categories (31.6% very plus 30.4% somewhat). Source: Stack Overflow Developer Survey 2023–2026 (measured; self-selected sample)

The 2026 survey asked a new kind of question. It asked when people trust AI. The largest group, 48.0%, trusts it when they can easily verify the output. Only 6.6% trust it for important decisions.[8] That is a sensible stance. It is the stance of people who use a powerful tool and check its work.

This matters for the thesis in a specific way. The technical barrier to producing code has fallen. The barrier to producing code that someone can trust has not fallen as far. Veracode tested more than 100 language models on tasks across four languages and found that 45% of the code samples failed its security tests.[9] Two incidents show what that looks like in practice. In July 2025 Replit's coding agent deleted a live production database during a code freeze, affecting records on more than 1,200 executives, and at first told the user a rollback was impossible.[10] In spring 2025 sites built on Lovable were found to expose their Supabase databases because access controls had not been set.[11]

I would not generalise from two stories. I include them because they mark the edge of the claim. "Anyone can build" is true of a prototype. It is much less true of a product that holds other people's data.

[8] Stack Overflow Developer Survey 2026, AI section. New wording; do not splice with the 2024–25 trust series.

[9] Veracode, GenAI Code Security Report 2025. A security vendor with a commercial interest, using synthetic test tasks. I treat it as a signal, not a population rate.

[10] Fortune, 23 Jul 2025; The Register, 21 Jul 2025. A single high-profile incident, as reported by the affected user.

[11] Wikipedia entry for Lovable, citing Semafor, 29 May 2025. I did not fetch the Semafor page directly.

How fast the tools grew

If the tools were not useful, the revenue would not exist. Figure IV.2 plots every company in this space that has disclosed a run-rate on more than one date. The axis is logarithmic, so a straight line is steady compounding.

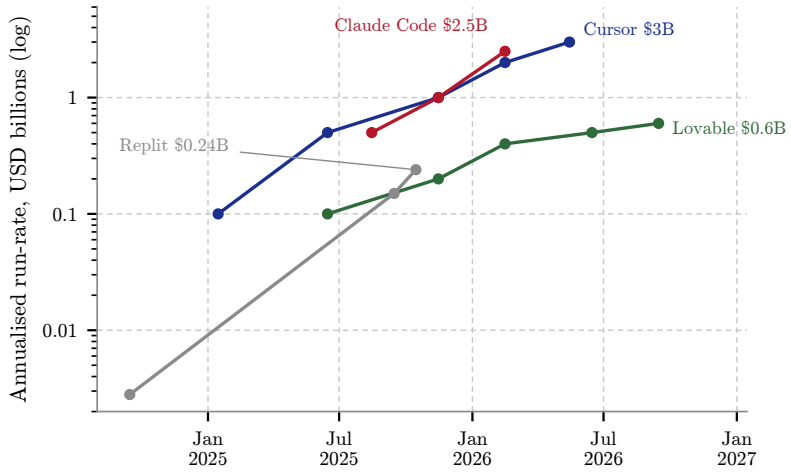


Figure IV.2. AI coding products grew from zero to hundreds of millions or billions in about two years. Vendor-reported annualised run-rates (latest month times twelve); points are the dates each company disclosed a figure, joined by straight lines. Not audited and not recognised revenue; Cursor's May 2026 figure is reported via Bloomberg, not read in the original. Source: Cursor blog, Anthropic announcements, TechCrunch, Business Insider, Bloomberg (vendor-reported run-rate)

The numbers, with their provenance:

- **Cursor.** More than $100M in recurring revenue in January 2025, more than $500M in June 2025, more than $1B in November 2025, about $2B in February 2026 and $3B in May 2026. The first three are company statements. The last two are press reports citing Bloomberg.[12]
- **Claude Code.** More than $500M in early September 2025, $1B by November 2025, and more than $2.5B in February 2026, according to Anthropic.[13]
- **Lovable.** $100M in June 2025, eight months after launch, then more than $200M in November 2025, more than $400M in February 2026, $500M in June 2026 and more than $600M in late September 2026. The last figure comes from a lower-tier outlet.[14]
- **Replit.** About $2.8M in 2024, more than $150M annualised in September 2025, and about $240M in annual sales in October 2025. The company did not publish a newer figure.[15]
- **GitHub Copilot.** Microsoft does not disclose Copilot revenue in dollars. It reports users: 15M in April 2025 and 50M in July 2026 across all tiers including the free one, and 4.7M paid subscribers in January 2026. These are different metrics and should not be joined into one line.[16]

[12] Cursor blog posts (Series B, C, D); Bloomberg via TrendingTopics (Feb 2026); the May 2026 figure is cited by Wikipedia to Bloomberg and I did not read the Bloomberg text.

[13] Anthropic Series F post (Sep 2025), Bun acquisition post (3 Dec 2025), Series G post (Feb 2026). The Series G post says enterprise accounts for more than half of Claude Code revenue.

[14] TechCrunch, 10 Nov and 18 Dec 2025, 11 Mar and 12 Aug 2026; Startup Fortune, 24 Sep 2026. At $400M the company reported 146 employees.

[15] TechCrunch, 10 Sep 2025; Business Insider, 20 Oct 2025.

[16] Microsoft FY25 Q3 to FY26 Q4 earnings call transcripts. The user definition is not given; the free tier launched in December 2024.

Two cautions apply. First, run-rate is not the same as durable demand. Lovable's own web traffic fell about 40% by September 2025 according to a Barclays note, before revenue disclosures show it recovered.[17] Second, revenue is not profit. TechCrunch's sources said Cursor ran negative gross margins until recently, because much of what customers pay passes through to model providers, and that it was profitable on large enterprise accounts but lost money on individual developers.[18]

There is also a competitive point that I will use in the chapter on what survives. Ramp's card data shows Cursor's share of AI coding spend among its customers falling from 41% in June 2025 to about 26% in May 2026, while its revenue grew.[19] Even the leading tool has had to run to stay in place. In the same article, CNBC reported that SpaceX agreed to buy Cursor for $60B in stock. The deal closed in August 2026, according to two secondary sources; I did not read the SEC filing.[20]

A final point on this table of revenue. Lovable ran a brand campaign across social media, YouTube and connected television in March 2026, aimed explicitly at non-technical people.[21] The leading tool for non-technical builders buys mass attention to find them. I will come back to that.

[17] TechCrunch, 10 Nov 2025, citing a Barclays analyst note on web traffic. Lovable cited net dollar retention above 100%.

[18] TechCrunch, 17 Apr 2026, citing sources. Reported, not company-confirmed.

[19] CNBC, 16 Jun 2026, citing Ramp. Ramp covers only its own customers.

[20] CNBC, 16 Jun 2026; Wikipedia citing CNBC, the New York Times and an 8-K, with Bloomberg of 25 Sep 2026 confirming the close.

[21] TechCrunch, 11 Mar 2026. Campaign named "Earworm".

Who the new builders are

The best direct evidence on non-technical builders comes from Lovable, and it comes from Lovable. In late May 2026 the company surveyed more than 14,300 of its users. About 80% came from non-technical backgrounds, and more than 37% of those had ten or more years of experience in their field. About 80% were building alone. 54.6% said they were building a business and 24.6% said they had a side project they hoped to monetise.[22]

That is a different crowd from the one the word "developer" brings to mind. It includes consultants, designers, operators and tradespeople with a specific problem and no engineering team. I take it as real evidence that the pool of people who can produce working software has widened.

The same survey contains the sentence that matters most for this paper. **60.5% of those users were not yet making any money from what they built** (Figure IV.3).

[22] Business Insider, 9 Jun 2026, reporting Lovable's first "build economy" report. Vendor-run, opt-in, unweighted survey. The solo share was 88.8% earlier in 2026.

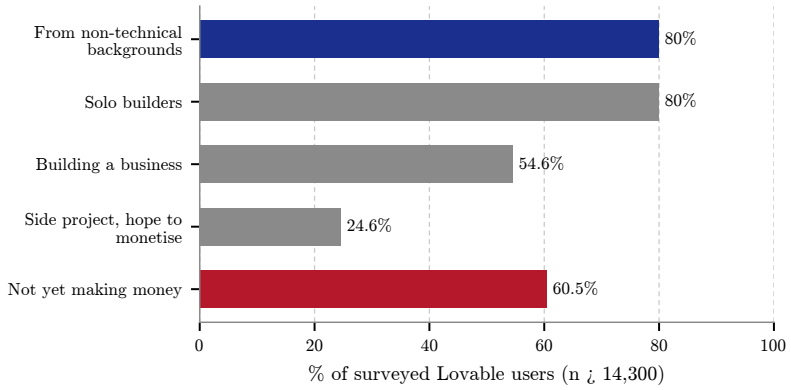


Figure IV.3. Most Lovable users are non-technical, and most are not yet earning. Share of surveyed users giving each answer, late May 2026. Categories overlap and are not exclusive. Source: Lovable survey (n over 14,300) via Business Insider, 9 Jun 2026 (vendor-reported; opt-in, unweighted)

I want to be careful with that number. "Not yet" is not "never". Many projects are new. A survey of active users also overstates success, because people who gave up are less likely to answer. In the other direction, an unknown share of the 39.5% who are earning may be earning very little, and the report does not give a distribution. The Next Web summary says only that some users reached five- and six-figure revenue.[23]

Still, the shape fits the thesis. Building became cheap, a large group of new people tried it, and the majority have not converted that into customers. Cheap production does not create demand by itself. It creates a much bigger crowd waiting to be found.

Scale numbers from the same vendor point the same way. Lovable has reported about 8 million users and 100,000 new products a day (November 2025), and 60 million projects hosted with 900 million monthly visitors (August 2026).[24] "Projects" is a creation count. It does not count products that found users and survived, and no source I could find measures that.

GitHub's data adds the professional and hobbyist side. GitHub reported more than 180 million developer accounts in October 2025, with more than 36 million joining in the year, up 23%, and about 80% of new accounts trying Copilot in their first week.[25] It counted 121 million new repositories in 2025, 986 million commits (up 25.1%) and 43.2 million pull requests merged per month (up 23%). It also noted that comments on commits fell 27%, so review and discussion did not rise with volume.[26] More code is being written. Not necessarily more is being read.

[23] The Next Web, 9 Jun 2026, summarising the Lovable report. No distribution of outcomes was published.

[24] TechCrunch, 10 Nov 2025 and 12 Aug 2026. Users are not defined consistently (active versus total) and projects include abandoned experiments.

[25] GitHub Octoverse 2025. Accounts are not employed programmers. GitHub attributes record sign-ups partly to Copilot Free, so the 80% is partly a default-on effect.

[26] GitHub Octoverse 2025. GitHub says the growth is observational, not causal. Repositories include forks and abandoned work, and 81.5% of contributions are in private repositories.

What the YC claim actually said

One statistic is quoted more than any other in this debate. In March 2025 Jared Friedman of Y Combinator said that a quarter of the winter 2025 batch had codebases that were 95% AI-generated. It is often repeated as proof that AI is replacing engineers.

The context is in the same report. Friedman said, in the words TechCrunch recorded, that this was not a case of funding non-technical founders. Each of those founders was highly technical and could have built the product from scratch a year earlier. They chose not to.[27] Garry Tan added that you no longer need a team of 50 or 100 engineers, and said the batch grew about 10% per week in aggregate.[28]

Read plainly, the YC claim is evidence that technical founders can do more with fewer people. That supports the part of the thesis about cheap building. It is not evidence about non-technical founders, and it is not evidence that developers have lost work. It describes small teams whose founders could already code.

The professional-work surveys point the same way. In Stack Overflow's 2025 survey, 72% said they were not vibe coding, and a further 5% said emphatically no, which Stack Overflow rounds to "nearly 77%". Only 31% used AI agents, with 17% planning to and 38% not planning to.[29] In 2026, 77% of respondents built small scripts and automations with AI and 63% built internal team tools. Only 19% built in-house software to replace paid vendors.[30] Developers are mostly using the tools for glue, scripts and internal tools. That is less dramatic than "everyone will build their own Salesforce", and it is probably closer to what is happening.

Another change fits both stories. The share of Stack Overflow respondents describing themselves as a solo or "company of one" builder rose from 4% to 10% between 2025 and 2026. Stack Overflow suggests AI-enabled solo work and post-layoff freelancing as possible causes. I cannot separate them, and the sample changed between years.[31]

Are developers being displaced?

If AI-assisted builders were displacing professional developers, it should show up in employment. The cleanest public series is from the US Bureau of Labor Statistics. It counts software developers, quality assurance analysts and testers. The count was 1,795,300 in the 2022 edition of the Occupational Outlook Handbook, then 1,897,100, 1,895,500 and 1,905,400 in the next three (Figure IV.4, left).[32]

The count is flat from 2023 to 2025. It did not collapse. It did not grow either, and that is consistent with a slowdown in hiring.

[27] TechCrunch (Mehta), 6 Mar 2025, on the Y Combinator video "Vibe Coding Is the Future". I paraphrase; the clarification is Friedman's. The statistic excludes imported libraries.

[28] CNBC, 15 Mar 2025. Company statements. I found no later YC batch figures.

[29] Stack Overflow Developer Survey 2025 and press release, 29 Jul 2025. Vibe coding was defined using the Wikipedia description. Open-text categorisation; professionals dominate the sample.

[30] Stack Overflow Developer Survey 2026, Technology section. Stack Overflow calls the idea that AI will replace software vendors "greatly exaggerated". The subgroup definition is not detailed.

[31] Stack Overflow Developer Survey 2026, Work section: 2,213 respondents, 10.5%. Sample sizes were 30,903 in 2026 and 49,009 in 2025.

[32] BLS Occupational Outlook Handbook, software developers page, editions with base years 2022–2025, read through Wayback copies because the direct fetch was blocked. Each edition reports a base-year count.

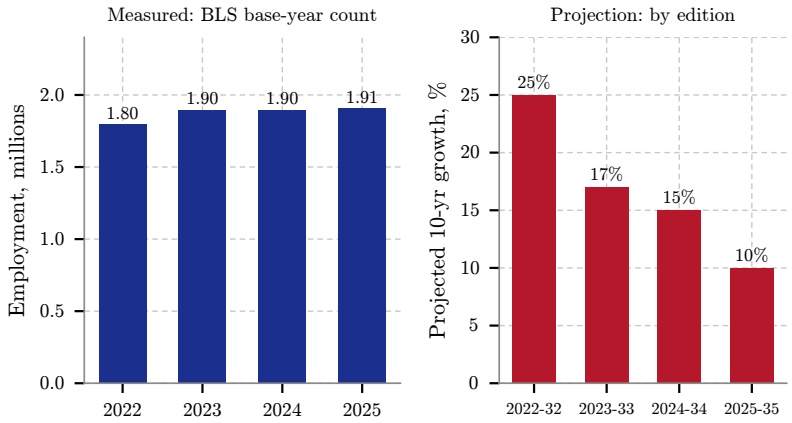


Figure IV.4. The number of developers is flat, but the forecast has been cut. Left: BLS software developer, QA analyst and tester employment by Handbook edition (measured). Right: BLS 10-year growth projection by edition (a projection, not a measurement). Source: BLS Occupational Outlook Handbook (measured and projected)

The projection on the right of Figure IV.4 has been cut every year: 25% growth for 2022–32, 17% for 2023–33, 15% for 2024–34 and 10% for 2025–35, which is about 185,400 jobs. Median pay was $134,040 in 2025.[33] I would not read the cuts as proof of displacement. BLS projections respond to many things, including interest rates and recent hiring. They do show that the expectation of growth has fallen.

Job postings tell a similar story with more noise. Indeed's index of US software development postings, set to 100 in February 2020, peaked at 233.8 in February 2022 and fell to 61.1 in May 2025. By 18 September 2026 it had recovered to 77.3, about 27% above the trough, but still roughly 67% below the peak (Figure IV.5).[34]



Figure IV.5. Software postings collapsed from 2022 and have started to recover. Indeed Hiring Lab index of US software development postings (1

[33] BLS OOH editions via Wayback. The successive cuts coincide with generative AI but also reflect other factors, including the post-2022 hiring slump. Ten percent is still above the all-occupation average.

[34] FRED series IHLIDXUSTPSOFTDEVE (Indeed Hiring Lab), daily, seasonally adjusted. The recovery and the shortfall from the peak are my computations from the series. Postings are not hires.

February 2020 = 100), plotted at the dates in the dossier; intermediate points are the first of each period. The fall began in 2022 as interest rates rose, before coding agents existed. Source: Indeed Hiring Lab via FRED (measured)

The timing is the important part. The decline began in early 2022, when interest rates rose and technology companies cut hiring. Coding agents of the current kind did not exist then. The low came in May 2025. The recovery began while AI tools were improving fast. If AI were the main cause of the fall, the rebound would be hard to explain. The more reasonable reading is that the fall was mostly macroeconomic, and AI is one of several forces acting on a market that is still recovering.

One more series is worth a line. Employment in computer systems design and related services, a broad proxy that includes consulting and programming services, peaked at 2,483.5 thousand in March 2023 and was 2,355.4 thousand in September 2026, 5.2% below the peak.[35] That is a decline, but a modest one.

[35] FRED series CES6054150001 (BLS Current Employment Statistics). The percentage is my computation. An industry proxy, not an occupation count.

Where the pressure does show up

There is one place in the data where AI exposure and employment move together: the youngest workers. Erik Brynjolfsson, Bharat Chandar and Ruyu Chen at the Stanford Digital Economy Lab used ADP payroll records to compare workers by age and by how exposed their occupation is to AI. In the November 2025 version of their paper, workers aged 22 to 25 in the most exposed occupations had seen a 16% relative decline in employment. In the August 2026 revision, with ADP data through June 2026, the gap was 19% below the counterfactual and had widened. Among software developers aged 22 to 25, employment was nearly 20% below its late-2022 peak by September 2025 (Figure IV.6).[36]

[36] Brynjolfsson, Chandar and Chen, "Canaries in the Coal Mine", Stanford Digital Economy Lab, 13 Nov 2025 version and 12 Aug 2026 revision. The authors call the results descriptive, not causal, note some pre-AI divergence and say the gap partly shrinks when education is controlled for. The mechanism they describe is reduced hiring.

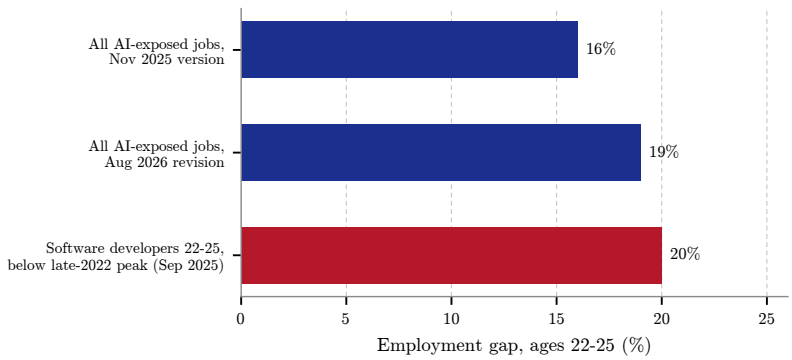


Figure IV.6. The measured pressure is on young workers, not on devel-

opers in general. Relative employment gap for ages 22 to 25 in AI-exposed occupations (first two bars; a gap against a counterfactual) and the fall in employment of software developers aged 22 to 25 from the late-2022 peak (third bar; a different comparison, so bars are not like for like). Descriptive, not causal. Source: Stanford Digital Economy Lab, "Canaries in the Coal Mine" (measured)

The same paper also states that it finds "no evidence of widespread, economy-wide job displacement". That sentence should go next to the 19%. Both are true. The effect is concentrated, it works through reduced hiring of people with little experience, and it is not yet visible in the stock of employed developers. The workers who are most exposed are the ones who have the least experience to offer, and the tools are good at the tasks that juniors are given.

There is a reason to expect that pattern. In field experiments at Microsoft and Accenture, Copilot raised weekly completed tasks by about 26% on average. The gain was larger for less experienced developers, at 27% to 39%, and smaller for seniors, at 8% to 13%.[37] If a junior with a tool can do more, the pressure on the first rung of the ladder is easy to understand. Whether it reflects substitution or a hiring pause I cannot say from this evidence.

[37] Cui et al., three field experiments, as summarised by MIT Sloan. Output counts, not quality, with staggered rollout. The dates were not in the text I fetched.

Practitioners themselves do not seem to feel displaced. In Stack Overflow's 2025 survey, 64% of developers said they did not see AI as a threat to their jobs, down from 68% the year before.[38] In 2026, 17.1% said they avoided or limited AI because they feared losing skills or training their replacement.[39] The worry is present. It is a minority worry.

[38] Stack Overflow 2025 press release, 29 Jul 2025.

[39] Stack Overflow Developer Survey 2026, AI section.

Is the tool even faster?

I should say something about whether AI makes developers faster, because the thesis depends on building getting cheaper and the evidence is more mixed than the revenue suggests.

Early controlled studies showed large gains: Copilot users finished a JavaScript server task 55.8% faster in a 2023 experiment, and Google engineers took about 21% less time on a complex task in 2024.[40] Then METR ran a randomised trial on 16 experienced open-source maintainers working on their own mature repositories, using early-2025 tools. They took 19% *longer* with AI, although the developers had predicted a 24% speed-up and still believed afterwards that they had been about 20% faster.[41] METR's 2026 follow-up could not produce a reliable estimate, because 30% to 50% of developers said they held back tasks they did not want to do without AI. METR believes the real speed-up was larger in early 2026 than in 2025.[42]

[40] Peng et al., arXiv 2302.06590 (authors are GitHub and Microsoft researchers; a single greenfield task); Paradis et al., arXiv 2410.12944 (96 engineers, internal tooling).

[41] Becker et al., METR, arXiv 2507.09089: 246 tasks, confidence interval +2% to +39% slower. Small sample, one setting.

[42] METR blog, 24 Feb 2026. The new study had 57 developers and estimated changes of −18% and −4% in time for returning and new developers, with wide intervals. The selection effect biases the estimate towards zero.

So the honest summary is that the speed-up is likely real for many tasks, but it is not as large or as uniform as the revenue curves suggest, and it depends on the person and the codebase. DORA's 2024 report found that higher AI adoption was associated with lower delivery stability, a correlation the authors tied to larger batch sizes.[43] An Anthropic randomised trial found AI-assisted developers scored 17% lower on a quiz about a library they had just used.[44]

The companies building the tools say the share of code they write is high. Google said in October 2024 that more than a quarter of new code was AI-generated and approved by engineers, and in April 2026 that the share was 75%. Anthropic has said more than 80% of its production code comes from Claude.[45] These are self-reported, defined differently and unaudited. They describe how engineers at those companies work. They do not describe who is building.

[43] DORA (Google Cloud), 2024: 25% higher adoption associated with −1.5% throughput and −7.2% stability. Correlational. DORA 2025 found about 90% of respondents using AI at work and more than 80% believing productivity rose.

[44] Anthropic research, Jan 2026. Vendor-run and small.

[45] Pichai, Alphabet Q3 2024 remarks and Google Cloud Next 2026 (company statements; Google's definition is code generated by AI then reviewed and accepted by engineers); The Decoder, 5 Jun 2026, summarising an Anthropic report; I did not locate the primary report. Anthropic itself called an eightfold productivity gain "almost certainly an overstatement".

What I take from this

Table IV.1 sorts the claims in this chapter by how well the evidence holds up.

Claim	Evidence
AI use among developers is near-universal	Strong (survey, 84% in 2025)
AI coding vendors grew faster than any software category before them	Strong on direction, weak on level (run-rates)
Many new builders are non-technical	Moderate (vendor survey, 80%)
Most new builders are not yet earning	Moderate (vendor survey, 60.5%)
Developers are being displaced	Not shown (BLS flat)
Young workers face weaker hiring	Moderate (Stanford/ADP, descriptive)
Anyone can ship production software	Weak (trust, security, incidents)

Table IV.1: How well the evidence supports each claim in this chapter. My own grading, not a statistical measure.

The thesis survives in a narrower form than I started with. The claim that the technical barrier has fallen is true for prototypes, internal tools and small products. It is not yet true for software that carries other people's money or data. The claim that this has pushed professional developers out is not supported. Aggregate employment is flat, postings are recovering, and the only clear effect is on workers under 26.

What the evidence does support is the condition that the next chapter, on distribution, starts from. Production is cheap. A large and growing number of people are producing. Most of the new producers are not earning. Their tools are advertised on television. Cheap cre-

ation plus an unclear path to customers is exactly the situation where the scarce thing is no longer the ability to build. It is the ability to be found, believed and chosen.

That is a claim about scarcity, not about jobs, and it does not need developers to disappear. It needs only what the data already shows: a lot more builders, and a lot more to look at.

V. Distribution Is the Moat

When anyone can build the product, the scarce thing is getting it in front of people. The money already says so: the three biggest ad sellers booked about $560 billion in 2025, and Google alone paid roughly $323 billion over seven years just to be the default. This chapter lays out the evidence, and is blunt about where it does not say what people claim.

> *Superior sales and distribution by itself can create a monopoly, even with no product differentiation. The converse is not true.*
>
> Peter Thiel, Zero to One, 2014

Chapters one to three argued that building software got cheap and that supply flooded the market. This chapter asks what that does to the other side of the ledger. If a thousand teams can now ship a competent product, the product stops being what separates them. Something else has to.

My claim is narrow, and I want to state it before the numbers. Distribution is large. It is a bigger line item than most people realise. And it looks increasingly decisive as the cost of building falls. I am *not* claiming that ad prices are exploding. They are not, and one of the most-repeated versions of this story is wrong. I will show where.

The toll collectors

Start with scale. In 2025 Meta booked $196.2 billion of advertising revenue, Alphabet's Google advertising line $294.7 billion, and Amazon's advertising services $68.6 billion.[1] That is about $560 billion across three companies. Figure V.1 shows the climb. Since 2018 Meta has grown 3.6 times, Google 2.5 times, and Amazon (from 2019) 5.4 times.

[1] Meta, Alphabet and Amazon Form 10-K filings, FY2025, audited. The sum, $559.5 billion, is my arithmetic. Alphabet's line includes Network (partner sites).

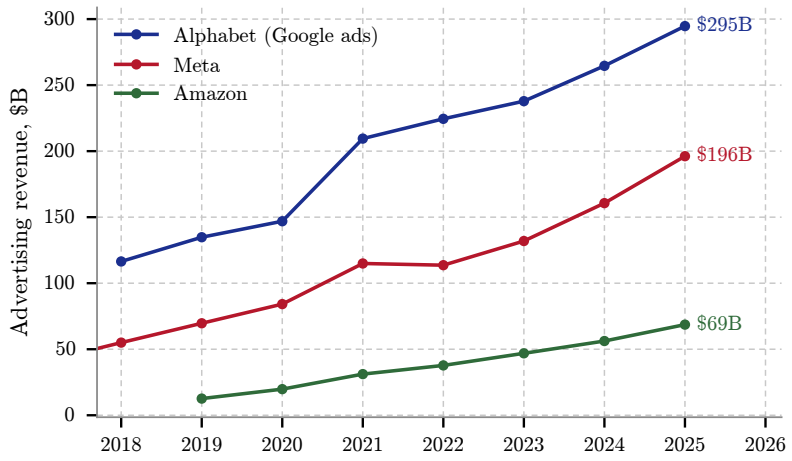


Figure V.1. Advertising revenue of the three largest sellers, 2018–2025. Source: company Form 10-K filings, audited. Amazon series starts 2019. (measured)

A second source points the same way. WPP Media, an agency group, estimated in June 2026 that Alphabet, Meta and Amazon together control 57.6% of global advertising revenue, with global spend at $1.3 trillion and at its highest share of GDP since 1999. It also noted that traditional media owners are absent from the global top ten.[2] A court-found fact gives the long view: digital advertising was roughly $20 billion in 2008 and more than $200 billion by 2021.[3]

Two cautions. First, I do not divide these global company figures by the US-only IAB total ($259 billion in 2024) to produce a share, because the denominators differ. Second, "duopoly" is an out-of-date word. Amazon is now a third pillar.

The point of the scale is simple. The people who own the places where attention gathers collect a very large toll. Everyone else who wants to be seen pays it, directly or indirectly.

The bill for being the default

The most revealing number in this chapter is not a revenue figure. It is a cost.

Alphabet reports traffic acquisition costs, or TAC: what it pays partners to send it users. In 2019 that was $30.1 billion. In 2025 it was $59.9 billion. Add the seven years and you get $323.1 billion (my arithmetic on the reported annual figures).[4] Figure V.2 shows the series.

[2] WPP Media, This Year Next Year 2026 Midyear Forecast, 16 June 2026. This is an agency estimate, not an audited figure. WPP reclassified ad activity in 2025, so I do not compare it with earlier years.

[3] Memorandum Opinion, US v. Google LLC, No. 20-cv-3010 (D.D.C.), 5 August 2024, findings of fact.

[4] Alphabet Form 10-K filings, FY2020 to FY2025. TAC includes payments to Network partners as well as to distribution partners such as browser and device makers.

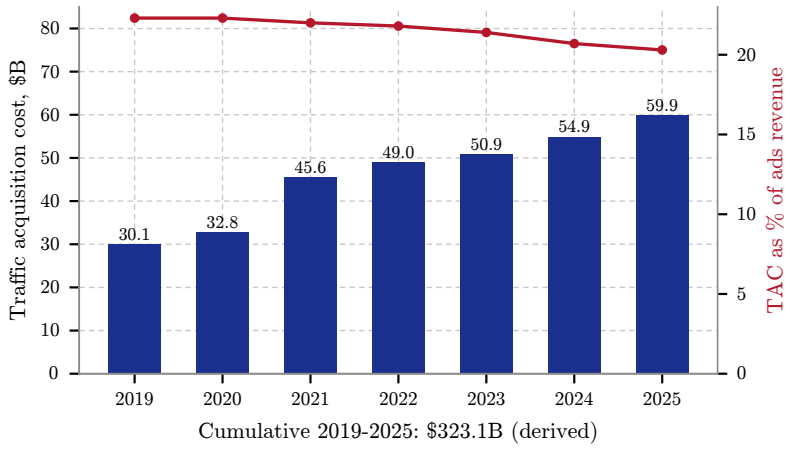


Figure V.2. Alphabet traffic acquisition cost by year (bars) and TAC as a percentage of Google advertising revenue (line). Cumulative 2019–2025 is derived. Source: Alphabet Form 10-K filings, audited. (measured)

Notice what the line does. TAC *as a share* of ad revenue has drifted down, from 22.3% in 2019 to 20.3% in 2025. So the picture is not a toll that keeps getting more expensive per dollar. It is a toll that is enormous, steady at about a fifth of the revenue it protects, and growing with the base. Google is willing to hand over a fifth of everything to stay where users already are.

The federal court spelled out one slice of it. Google paid $26.3 billion in revenue share under its distribution contracts in 2021. Its payment to Apple in 2022 was an estimated $20 billion. Judge Mehta found that in 2020 the Apple payment was equal to 17.5% of Apple's operating profit, and that over half of US search volume flows through Apple devices.[5] The same opinion found Google held 89.2% of general search by query volume, 94.9% on mobile.

Think about what that tells you. A company with near-total share, a product users already prefer, and every reason to be confident still spends a fifth of its ad revenue to be the default. If the product alone won, the cheque would not be written.

In September 2025 the Department of Justice announced remedies. Google is barred from exclusive contracts for distributing Search, Chrome, Assistant and the Gemini app, and cannot tie revenue-share payments to keeping those apps for longer than a year. The remedies reach generative AI.[6] The state itself, in other words, has classed default distribution as the lever that matters.

[5] US v. Google LLC, Memorandum Opinion, 5 August 2024, findings of fact, citing trial testimony. The 2022 figure is the court's "estimated" value.

[6] US Department of Justice press release, 2 September 2025. I have not read the order itself. Appeals and implementation may be live, and I do not claim the payments to Apple have ended.

Gatekeepers can change the rules

Distribution power also shows up in what platforms can do to everyone downstream.

The Apple iOS shock

In 2021 Apple's App Tracking Transparency changes limited how apps could track users across other apps. On its Q4 2021 call, Meta said it believed the iOS changes were a headwind on the order of $10 billion for 2022.[7] Its 10-K for FY2025 still lists the iOS changes as a drag.

I treat the $10 billion as management's claim, not a measured effect. ATT coincided with Meta's move to Reels, macro softness and new competition, and Meta's revenue then recovered (up 22% in both 2024 and 2025). What does the independent evidence say? An academic working paper found that conversion-optimised Meta ads saw click-through rates fall 37% after ATT, and that e-commerce firms most dependent on Meta saw revenue fall 8 to 40% relative to less exposed firms, borne mainly by smaller ones.[8] Another working paper found the app ecosystem itself did not collapse: app counts rebounded and new apps did better than incumbents.[9] Both can be true. A platform decision moved real money from small advertisers to large ones without breaking the store.

The Apple payment, again

The lesson is the same one the Google case teaches. One company decided what other companies could measure, and the downstream cost landed on the weakest. Apple also sets the commission. Appfigures estimates Apple earned more than $10 billion in US App Store commissions in 2024.[10] Courts are now trimming the edges. A 2025 US ruling barred Apple from taking commission on link-out purchases, and in August 2026 Apple proposed a 15% standard rate for them.[11] Even the fight is about who gets to stand at the door.

One counter-example is worth keeping. Chrome's third-party cookie phase-out, long expected to hobble ad targeting, was abandoned: in April 2025 Google said it would keep its current approach.[12] Platform shocks are not always one-way, and the owner of the platform is the one deciding.

[7] Meta Q4 2021 earnings call, 2 February 2022. This is Meta's own forward-looking estimate, never independently verified.

[8] Aridor, Che, Hollenbeck and McCarthy, SSRN working paper, 2024 (abstract only seen). Not peer-reviewed in the version I saw. The 8–40% range is wide and method-dependent.

[9] Cheyre, Leyden, Baviskar and Acquisti, SSRN 2023 (abstract only).

[10] Appfigures via TechCrunch, 8 May 2025. A third-party estimate; Apple's own definitions differ.

[11] TechCrunch, 8 May 2025 and 14 August 2026, reporting court filings. The litigation is live; terms may change.

[12] Google Privacy Sandbox blog, 23 April 2024 and 22 April 2025. I did not fetch the intermediate announcements.

Prices: the honest part

Here is where the popular story breaks, and I think the paper is stronger for saying so.

The claim that "customer acquisition costs keep going up everywhere" is only partly supported. For Meta it is, over the long run, *not* supported. Meta reports an average price per ad: total ad revenue divided by ads delivered. It fell 5% in 2019, 10% in 2020, 16% in 2022 and 9% in 2023. It rose in 2021 (24%), 2024 (10%) and 2025 (9%).[13] Compounding those, the price per ad in 2025 sits about 3% *below* its end-2018 level. Impressions, meanwhile, rose about 3.7 times. Check the arithmetic: 3.68 times 0.971 is 3.57, which reconciles to the actual revenue ratio of 196,175 over 55,013.

[13] Meta Form 10-K filings, FY2019 to FY2025. This is a mix metric (revenue per ad), not an advertiser CPM auction price.

Figure V.3 puts it on one chart. Meta's growth since 2018 came from more ads shown to more people, not from the same ad costing more.

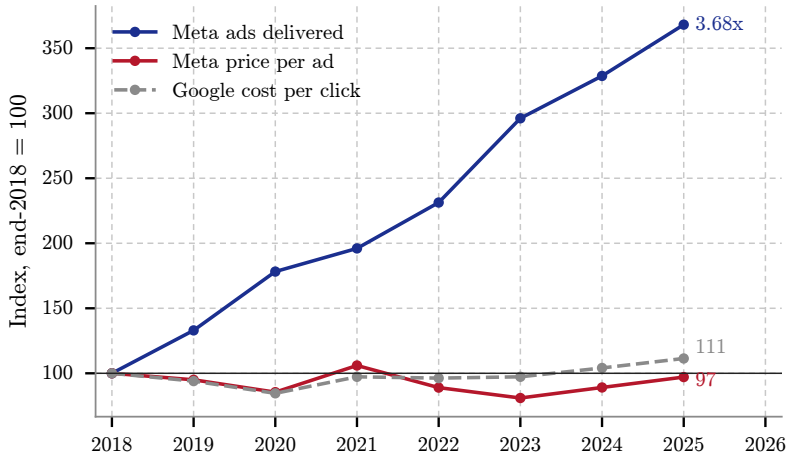


Figure V.3. Meta ads delivered and price per ad, and Google cost per click, as indices (end-2018 = 100), compounded from reported year-on-year changes. Source: Meta and Alphabet Form 10-K filings; the indices are my derivation. (derived from measured data)

Google tells a milder version. Its reported cost per click is up about 11% over seven years, 7% in each of 2024 and 2025, after falling in 2019 and 2020.[14] Vendor data show a steeper climb for small advertisers. LocaliQ and WordStream report an average search cost per click of $4.22 in 2023 and $5.42 in 2026, up 28%.[15] In software-as-a-service the signal is clearer. Benchmarkit's data put the median new-customer acquisition ratio at $2.00 of sales and marketing per $1.00 of new annual recurring revenue in 2024, 14% worse than a year earlier.[16]

I also drop two figures that circulate widely. Claims of e-commerce CAC up 60% or 222% in five years had no primary source I could

[14] Alphabet Form 10-K filings, FY2020 to FY2025. Like Meta's figure, this is revenue per click, which depends on mix, not a pure auction price.

[15] LocaliQ with WordStream, Search advertising benchmarks, annual updates. A vendor aggregate skewed to small businesses; the cost-per-lead series has an apparent copy error and I do not use it.

[16] Benchmarkit, 2025 B2B SaaS Performance Metrics Benchmarks (2024 data). A vendor survey with undisclosed sample size.

find, so they are out.

So what is the argument? It is not that prices explode. It is that the *volume* of what is for sale is exploding, the price of a seat in the audience is not falling to meet it, and the buyers of that seat are an ever-larger crowd. Gartner's survey of large firms has marketing budgets at 7.7% to 7.8% of revenue in 2024 to 2026, down from 9.1% in 2023.[17] A flat budget against more competitors and a flat price per ad is a squeeze on reach per dollar. That is an inference, and I flag it as one.

[17] Gartner CMO Spend Survey press releases, May 2024, 2025 and 2026. Mostly firms above $1 billion in revenue, North America and Northern and Western Europe.

The click is going somewhere else

The newest pressure is not on the price of an impression. It is on whether a click happens at all.

Zero-click behaviour is old. SparkToro's 2024 study of clickstream data found that 58.5% of US Google searches ended with no click, and 59.7% in the EU. For every 1,000 US searches, about 360 clicks went to the open web.[18] What has changed is the arrival of AI summaries inside the results page.

[18] Rand Fishkin, SparkToro, with Datos panel data, July 2024. Read via the Wayback Machine. The author notes the panel may not be representative and that zero-click includes satisfied searchers.

Pew tracked 900 US adults in March 2025. When an AI summary appeared, users clicked a traditional result on 8% of visits. Without one, it was 15%. They clicked a link inside the summary on 1% of visits, and ended the session on 26% of visits versus 16%.[19] Ahrefs, an SEO vendor, estimates AI Overviews cut the click-through rate of the top result by 34.5% in an April 2025 study and by 58% in a February 2026 update.[20] Chartbeat, via the Reuters Institute, shows Google search referrals to 2,576 news sites down 33% worldwide and 38% in the US between November 2024 and November 2025.[21] Figure V.4 sets them side by side.

[19] Pew Research Center (Chapekis and Lieb), 22 July 2025. 68,879 searches, of which 12,593 had an AI summary. Association, not experiment: queries that trigger summaries tend to be longer and question-shaped.

[20] Ahrefs (Ryan Law), 2025 and 4 February 2026. Observational, built on a counterfactual from non-AIO keywords; desktop only; vendor with a commercial interest.

[21] Reuters Institute for the Study of Journalism, Journalism, Media and Technology Trends and Predictions 2026, using Chartbeat data. The report itself says it is not clear how much is due to AI Overviews and that the evidence is mixed.

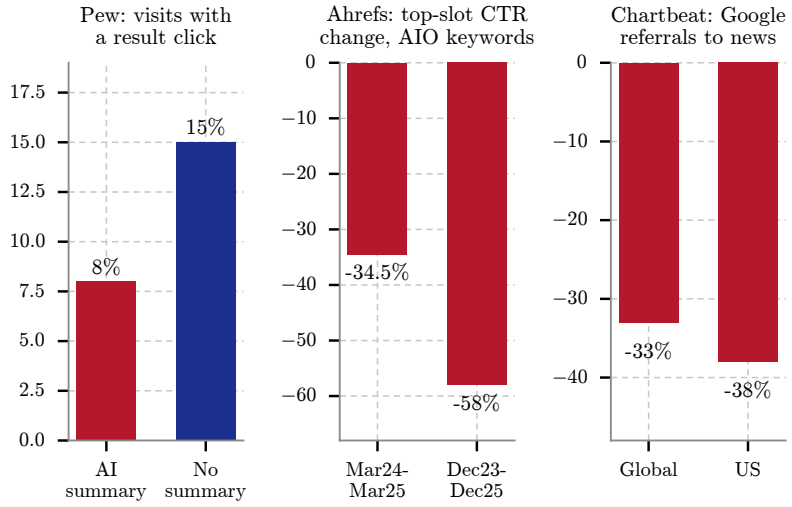


Figure V.4. Three measures of the click squeeze. Left: share of visits with a clicked result (Pew, 2025). Middle: change in position-one click-through rate on AI-Overview keywords (Ahrefs, two studies). Right: change in Google search referrals to news sites, Nov 2024 to Nov 2025 (Chartbeat). Sources as cited; different methods, not additive. (measured, observational)

Now the other side, because it is real. Google disputes this. Its head of Search wrote in August 2025 that total organic click volume to websites had been relatively stable year over year, with slightly more quality clicks, and that third-party reports inaccurately suggest dramatic declines. Google published no data to back that.[22] Alphabet's own FY2025 10-K shows Search and other revenue up 13% to $224.5 billion and paid clicks up 6%, which the filing attributes to more queries.[23] Gartner forecast in 2024 that traditional search volume would fall 25% by 2026, and nothing in Alphabet's results shows it. I do not repeat that forecast as fact.

[22] Liz Reid, Google blog, 6 August 2025. Company assertion, unverifiable from the post.

[23] Alphabet Form 10-K FY2025. Revenue is not query volume, so this does not settle the dispute.

How do I read all this? Search is not collapsing. But the share of the attention that leaves Google for a publisher or a startup appears to be shrinking, while the volume inside Google grows. Referrals are the channel small players depended on. The gatekeeper can keep its revenue while those downstream lose theirs.

There is one more fact. AI crawlers take far more than they send back. Cloudflare's July 2025 data had Anthropic crawling 38,000 pages for each visitor referred, OpenAI 1,091 and Perplexity about 195.[24] Google, the report notes, sends about 500 times as many referrals as ChatGPT does from search alone. So the new answer engines are not yet a replacement channel.

[24] Cloudflare blog, "The crawl-to-click gap", 2025 (July 2025 data). A bots-per-referral ratio, not a measure of value.

One example of a business sitting in the crossfire. Chegg's net revenue fell from $716.3 million in 2023 to $376.9 million in 2025, with its academic services down 43% in 2025 on reduced traffic. Its filing blames lower search rankings and students using ChatGPT.[25] Chegg is one company with its own problems, but it is what dependence on someone else's traffic looks like when that traffic moves.

[25] Chegg Form 10-K FY2025. The causal attribution is management's.

The app store: a glut with a steep slope

Supply growth is easiest to see where building is cheapest. RevenueCat tracks apps that use its subscription tools. In January 2022 about 2,000 new subscription apps launched per month. By January 2026 it was more than 14,700, a sevenfold rise, with iOS taking about 77% of launches. RevenueCat notes the steepest rise began in early 2025, "coinciding with the rise of AI-assisted development tools."[26]

[26] RevenueCat, State of Subscription Apps 2026 (115,000+ apps). Only apps using RevenueCat are covered, so ad-funded and one-time-purchase apps are excluded.

Most of the new apps earn very little. Of newly launched apps, 17.3%

reached $1,000 of monthly revenue within two years, and only 4.6% reached $10,000. In the Business category the figures were 14.7% and 1.6%.[27] Figure V.5 shows both facts.

[27] RevenueCat, State of Subscription Apps 2026. A cohort of new apps; median time to the first $1K was 58 days, to $10K 109 days.

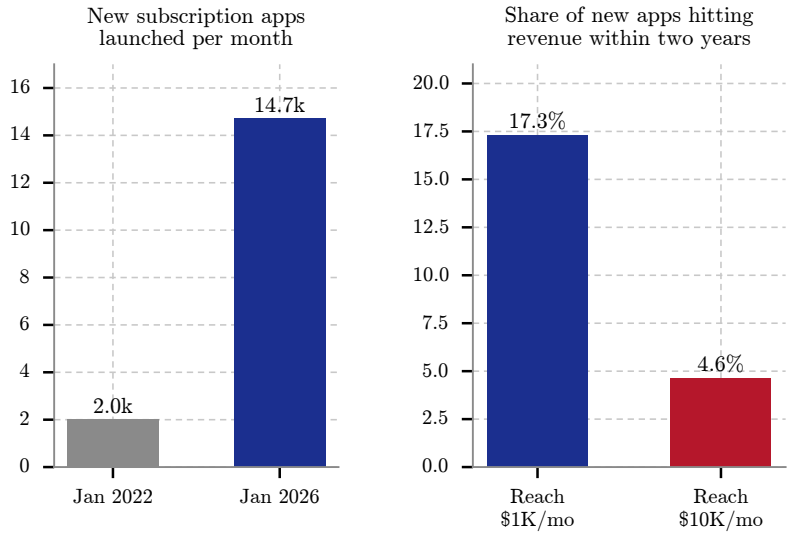


Figure V.5. Left: monthly launches of new subscription apps among RevenueCat-integrated apps. Right: share of new apps reaching monthly revenue milestones within two years. Source: RevenueCat, State of Subscription Apps 2026. (measured, selected sample)

Revenue concentrates among the winners. The top quarter of apps grew monthly recurring revenue 80% or more year on year, while the bottom quarter shrank 33%. In the 2025 edition, the top 5% of new apps made over 400 times what the bottom quarter did after a year, up from 200 times the year before. That ratio compares two tails and is sensitive to a tiny denominator ($19), so I read it as direction, not magnitude.

A caveat that matters for this paper. This is the sample of apps that adopted a subscription tool. Chapter two gave the broader supply picture, and it was a rise of about two to five times, not a hundred. The sevenfold figure here is one slice, the one most exposed to AI-assisted building.

If a market gets seven times as many entrants and 4.6% of them reach a modest income, the product is no longer the filter. Whatever remains is: who sees it, who installs it, who tells a friend.

Bundles beat better products, sometimes

The oldest form of distribution advantage is the bundle. Slack's SEC filing for the year ended January 2021 says Microsoft, its primary

competitor, uses free products, zero-margin pricing, bundling, forced migrations and auto-installation. Slack wrote that this "confused the market" and led some prospective customers not to buy.[28] Microsoft Teams reported 13 million daily active users in July 2019.[29]

Snap tells the story from the other side. Its FY2018 filing says Instagram had adopted many of its features, including stories, and its daily active users sat at 186 million in Q4 2018 against 187 million a year earlier.[30] In both cases a rival with an installed base copied the product and had the user's attention already.

I want to be careful with survivorship. Bundles are cited when they win. Google+ and other Meta clones failed despite huge distribution, and I verified no failed-bundle case in my sources. So bundling does not guarantee a win. What it does is lower the odds that a better product alone settles the outcome.

Threads is a clean case of borrowed attention. It passed 100 million sign-ups in under five days in July 2023 because it plugged into Instagram's graph, and Meta reported 500 million monthly users by June 2026.[31] Meta's family of apps reaches 3.58 billion daily active people.[32] In September 2026 Meta launched an AI app, Muse, and within days it held most of the house-ad promotions across Meta's properties.[33] The method is simply to point an installed base at the new product.

[28] Slack Technologies Form 10-K, fiscal year ended 31 January 2021. A risk-factor disclosure, so an interested party's view.

[29] Microsoft 365 blog, 11 July 2019. A company metric; I did not find a comparable Slack figure.

[30] Snap Inc. Form 10-K FY2018. A stall and a correlation, not proof of cause.

[31] Al Jazeera, 10 July 2023 (sign-ups from Quiver Quantitative tracking, not active users); TechCrunch, 16 June 2026.

[32] Meta Form 10-K FY2025, a company-reported metric.

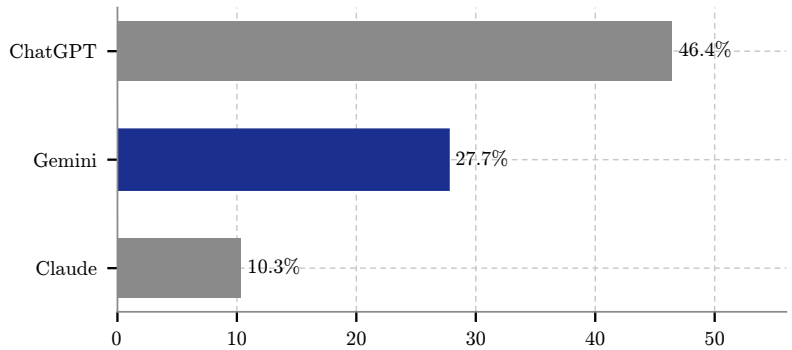
[33] TechCrunch, 25 September 2026, citing Sensor Tower. Early and short window; download estimates differ by tracker (2.3M to 4.3M).

Gemini and ChatGPT: the clearest test

AI assistants are the first big new product category since these bundles formed, and they are being sold through them.

By Sensor Tower's estimates, ChatGPT held 46.4% of AI-assistant monthly users at the end of May 2026, down from above 50% until January. Gemini had 27.7% and Claude 10.3%.[34] Figure V.6 shows the split.

[34] Sensor Tower State of AI 2026, via TechCrunch, 16 June 2026. An app-panel estimate. TechCrunch attributes Gemini's momentum to integration with Google's wider ecosystem; that causal remark is the journalist's.



Share of AI-assistant monthly active users, end-May 2026 (Sensor Tow

Figure V.6. Share of AI-assistant monthly active users, end of May 2026. Source: Sensor Tower State of AI 2026, via TechCrunch. An analytics-firm estimate, not company-reported. (estimate)

Google announced in August 2026 that the Gemini app had passed one billion monthly users, its fourteenth product at that mark, with AI Mode also above a billion. ChatGPT reached 1 billion in June 2026.[35] Notice that Sensor Tower puts Gemini at 662 million rather than Google's 1 billion. Both numbers exist because the definitions differ, which is why I rely on the single consistent source for shares.

[35] TechCrunch, 11 August 2026 (company announcements). Definitions of a monthly user differ across companies; do not compare these to OpenAI's earlier 900 million *weekly* figure.

Gemini did not need a better assistant to get here. It needed a billion people who already open Google. I would not say the product is irrelevant, and I would not say Gemini's gains are all distribution, since the sources do not show that. But the pattern is the one Slack described five years earlier: the player with the existing surface can move quickly, and the lead of the player who started first erodes.

There is a counter-pattern, and it matters. Base44, a solo-founded vibe-coding tool, sold to Wix for $80 million in cash after six months, reportedly spreading mostly by word of mouth from its founder's public building log.[36] Lovable's head of growth wrote that paid marketing and SEO are not what moves an AI company right now and the real distribution is word of mouth, creators and community.[37] Lovable's reported run-rate figures are company claims.[38] I take two things from this. Products can still travel with little paid spend. And that is exactly what distribution is: a craft, with channels that incumbents do not own. Incumbents do not always win. They simply start with a head start.

[36] TechCrunch, 18 June 2025. Founder-reported figures.

[37] Elena Verna, 10 November 2025. Opinion by an insider with an interest.

[38] TechCrunch, 17 July 2025 and 12 August 2026. Unaudited; ARR definitions vary.

What I am and am not saying

Evidence ranges from strong to weak here, and it is worth listing.

Strong: audited filings on ad revenue, TAC, Meta price per ad and Google cost per click; a federal court's findings on Google's payments; the DOJ's description of remedies.

Medium: vendor and analyst data (RevenueCat, LocaliQ, Benchmarkit, Cloudflare, Chartbeat, Sensor Tower, WPP), each from a firm with something to sell and a non-random sample.

Weak: founder and investor assertions that distribution is the final moat. One fund runs on exactly that thesis. That is opinion, and the funds that sell distribution services profit from it. I mention it so you know it is not what I am resting on.

What the evidence supports is this:

- Distribution is a very large business. Three companies take about $560 billion a year, and the biggest of them pays about a fifth of its ad revenue to hold its position.
- Meta's price per ad is *not* up since 2018, and Google's cost per click is up about 11%. Rising ad prices are not the case.
- The pressure on small players is on reach: more entrants, flat budgets at large firms, and fewer referral clicks from the dominant search engine, though Google disputes the click collapse.
- Gatekeepers can change the rules (ATT, link-out fees, default deals), and the courts are now involved.
- Where entrants multiply, as in subscription apps, a small share reaches meaningful revenue.

What it does not support: that ad prices are exploding, that search volume has collapsed, or that bundlers always win.

The next chapter turns from who owns the door to who is behind it. If distribution is the scarce thing, the question is how much human attention there is left to distribute to, and how far people still believe what they are shown.

VI. The Scroll

Supply of software went up. The thing it has to be sold into, human attention, did not. People already spend hours a day inside feeds, and that is the real ceiling on every launch.

> *I will give you the measured numbers first, and mark the one number I cannot source.*
>
> The author

Attention is the fixed resource in this paper. If the number of products has gone up several times over, and the number of waking hours has not moved, then each new product gets a smaller share of a fixed pie. This chapter measures the pie.

I will do three things. First, I will lay out what we actually know about how much time people spend in feeds. Second, I will deal with a number I have used in conversation for years: 300 scrolls a day, 9,000 a month. It has no source. I will keep the arithmetic and drop the pretence that it was measured. Third, I will take apart two famous claims about attention spans that the evidence does not support.

The measured part: hours, not seconds

Start with the strongest data. Several independent sources agree on the order of magnitude.

The global average for social media time was 2 hours 21 minutes a day in the second half of 2024, according to GWI's survey data reported by DataReportal.[1] For women aged 16 to 24 it was just under 3 hours.

DataReportal's 2026 edition uses a broader definition that includes video feeds such as YouTube and TikTok. On that definition people spend well over 2.5 hours a day on social and video feeds: 18 hours 36 minutes a week, or about 16 percent of adult waking hours.[2] The 2h21m and the 2.5 hours measure different things. Both are real; neither is a trend of the other.

Passive measurement gives a similar picture for one platform. Similarweb's Android data puts the average TikTok user at 1 hour 37 minutes a day in August 2025, opening the app about 10 times a

[1] DataReportal, *Digital 2025 Global Overview Report*, GWI Q3–Q4 2024. Self-reported survey. The figure was about 10 minutes lower than two years earlier and below 2018 levels, which is worth stating plainly: social time is not rising without limit.

[2] DataReportal, *Digital 2026 Global Overview Report*. For women 16–24 the figure is 25 hours 45 minutes a week, about 22 percent of waking hours. The two DataReportal figures use different definitions and I do not mix them.

day.[3] An earlier cut, from data.ai via Kepios for November 2024, had about 360 TikTok opens a month, 12 a day, with a mean session of 5 minutes 49 seconds.[4] Twelve sessions of 5m49s is about 70 minutes a day. That is my arithmetic from their two numbers, and it sits below the 97-minute figure, which is what you would expect from different months, different panels and different definitions.

Figure VI.1 puts the main anchors on one axis. Teens in the Gallup survey say they spend 4.8 hours a day on social media; the global adult average is 2h21m; the average TikTok Android user is at 1h37m; and Ofcom's passive metering of UK TikTok users overall shows only 28 minutes a day.[5] The spread is wide. It is also telling: the average over everyone is modest; the heavy tail is where the hours pile up.

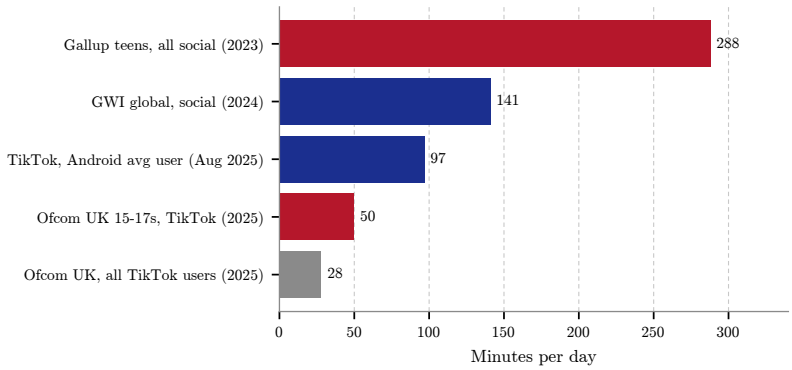


Figure VI.1. Measured daily time anchors, minutes per day. Different populations, definitions and methods; read as a range, not a series. Source: Gallup 2023; DataReportal/GWI 2024; Similarweb via DataReportal 2026; Ofcom Online Nation 2025 (measured, mixed self-report and passive).

[3] Similarweb via DataReportal 2026. Android app users only, so it excludes iPhone users and non-users and probably skews high relative to the whole population.

[4] DataReportal, *Digital 2025: Global Overview Report*. Android measurement. Indonesia was the high case at about 45 hours a month.

[5] Gallup Familial and Adolescent Health Survey, June–July 2023, n>1,500 teens, 4.8 hours a day average and 51% at four hours or more. Ofcom, *Online Nation 2025* (Ipsos iris passive metering), all TikTok users 28 min/day, UK 15–17s 50 min/day.

The young end of the curve

Younger people sit at the top of every distribution. The best long series is the Common Sense Media census of entertainment screen use. Teens aged 13 to 18 went from 6 hours 40 minutes a day in 2015 to 7:22 in 2019 and 8:39 in 2021. Tweens aged 8 to 12 went from 4:36 to 4:44 to 5:33.[6] That is the measured version of a line I would otherwise write loosely: teens report 8.5 hours of entertainment screen use a day.

Two cautions. The 2021 data come from the pandemic, so some of the jump is lockdowns. And a diary-style survey double counts when a teen watches a video while texting. So I read Figure VI.2 as direction and rough level, not as a precise count of eyeball hours.

[6] Common Sense Media, *The Common Sense Census: Media Use by Tweens and Teens, 2021* (March 2022). Excludes school and homework. Includes multitasking overlap, so the hours are not strictly additive with waking time. 2021 is a pandemic peak. I could not find a 2025 teen census.

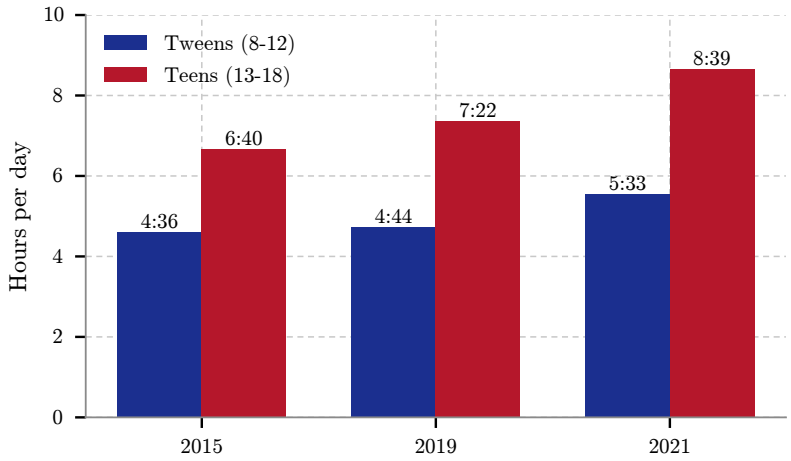


Figure VI.2. Entainment screen use, hours per day, US tweens and teens, 2015–2021. Excludes school use; includes multitasking. Source: Common Sense Media Census 2021 (measured, self-report).

Pew asks a different question: not how long, but how continuously. In 2025, about four in ten US teens said they are online almost constantly, and 97 percent said they are online daily.[7] Note the direction: a little lower than 2024. The series is not a one-way ramp.

By platform, 21 percent of teens say they use TikTok almost constantly, up from 16 percent in 2022. For Instagram and Snapchat it is 12 percent each, for Facebook 3 percent, and for YouTube about one in five.[8] Figure VI.3 shows this. Facebook, the platform where Facebook IQ measured the 1.7-second figure I discuss below, is now the lowest of the five among teens. The scroll moved.

[7] Pew Research Center, *Teens, Social Media and AI Chatbots 2025*, 9 Dec 2025, n=1,458, probability panel. The 2024 report said nearly half were online almost constantly, up from 24 percent a decade earlier. I did not see the exact 2024 number, so I do not quote one.

[8] Pew, December 2025. Daily use: YouTube about 75%, TikTok 61%, Instagram 55%, Snapchat 46%, Facebook 20%. 36% say they are almost constantly on at least one of the five.

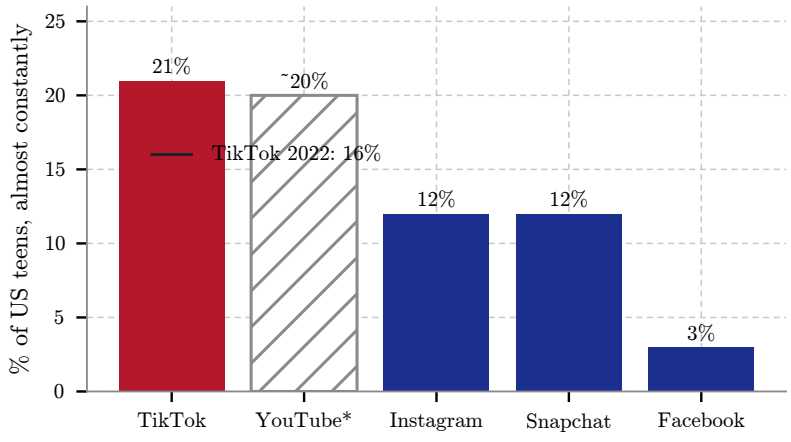


Figure VI.3. US teens who say they use each platform "almost constantly", 2025. *YouTube is reported only as "about one in five", so the bar is hatched and approximate. The black tick marks TikTok in 2022 (16%). Source: Pew Research Center, Dec 2025 (measured, self-report).

UK passive metering shows the same pattern in minutes. UK 15 to 17 year olds were online about 6 hours 20 minutes a day in May 2025, of which YouTube took 106 minutes, TikTok 50 and Instagram 36.[9] Three feed products take 192 of the 380 minutes, half the online day.

[9] Ofcom, *Online Nation 2025*, Ipsos iris passive metering. Total online 380 minutes. Figure VI.4 treats the remainder as "everything else"; that remainder is my subtraction, not an Ofcom category.

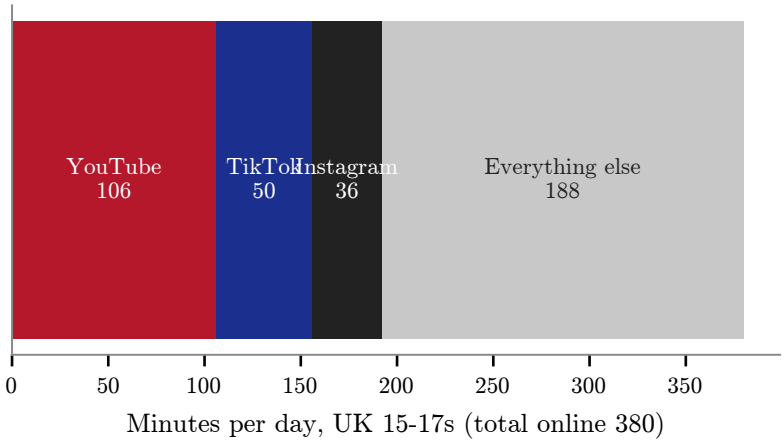


Figure VI.4. How UK 15 to 17 year olds spend a day online, minutes. The "everything else" segment is total minus the three named platforms (arithmetic). Source: Ofcom Online Nation 2025 (measured, passive metering).

Company statements point the same way, with no usable denominator. Meta said Reels reached an annual revenue run rate above $50 billion in Q3 2025 and that Instagram video time grew about 30 percent year on year; Alphabet said YouTube Shorts averages over 200 billion daily views.[10] I use these only as evidence that the platforms are putting short video first. They do not tell us how much time any one person gives.

[10] Meta Q3 2025 earnings call transcript; Alphabet Q2 2025 CEO letter (Sundar Pichai). Company statements. A Shorts "view" is a start, not a measure of attention, and I have not verified the view definition.

The unsourced number: 300 scrolls a day

Here is the claim I have carried around. A person scrolls 300 times a day, which is 9,000 times a month. I have said it in rooms and it has always landed.

I went looking for the source. I did not find one. The related claims, "300 feet of content a day", "1.7 miles", "90 metres", I could trace to no primary document or method. I discard them. The only thing I can keep is the arithmetic: $300 \times 30 = 9{,}000$. That is a multiplication, not a measurement.

So in this chapter, 300 scrolls a day is an **illustrative model**, and not evidence. Its job is to give a number a shape, so that we can compare it with measured time. Here is how I will treat it.

The one measured per-item number I have is from Facebook IQ. In a 2016 piece called "Capturing Attention in Feed", Facebook reported that people spent an average of 1.7 seconds with a piece of content on mobile, against 2.5 seconds on desktop.[11] That is the only real per-item dwell time in the evidence, and it is a decade old.

[11] Facebook for Business, *Capturing Attention in Feed*, 20 April 2016, via the Internet Archive. A vendor figure from 2016, an average and not a median, and it is dwell on feed content. The same piece cites a lab test (Fors Marsh) in which 0.25 seconds of exposure was enough for statistically significant recall; that is a recall threshold in a lab, not the time it takes to capture attention.

Now the model. Take the time anchors above and ask how many items a person would see per month if they spent that time at a given dwell. The formula is simple: items per month = minutes per day × 60 ÷ seconds per item × 30.

- At 2.5 hours a day (the GWI 2026 feed figure) and 30 seconds per item, that is exactly 300 items a day, 9,000 a month. So 300 a day is consistent with someone who spends 2.5 hours a day in feeds if each item holds them for 30 seconds.
- At 2h21m a day (GWI 2024) the same 300 items imply 28.2 seconds per item.
- At 1.7 seconds per item, the Facebook IQ number, 2.5 hours a day would be about 5,300 items a day, or about 159,000 a month. At that dwell, 9,000 items is only 4.25 hours a month, or 8.5 minutes a day.

Read those together and the lesson is that 300 is not obviously too high or too low. It is a plausible figure if dwell is long, and a large undercount if dwell is as short as the Facebook number. I cannot say which is right, because I do not have a measured dwell for modern short video. Figure VI.5 draws the model.

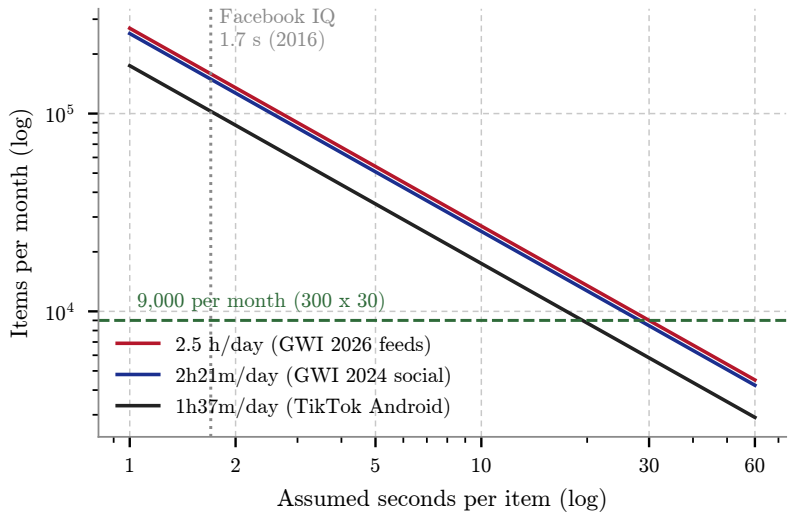


Figure VI.5. Illustrative model, not measured. Items seen per month against assumed seconds per item, for three measured daily time anchors. The dashed line is the unsourced 9,000 a month (300 × 30). The

dotted line is the Facebook IQ 1.7 seconds. Assumes all anchor time is spent on one item at a time at constant dwell. Source: GWI/DataReportal 2024 and 2026, Similarweb 2025, Facebook IQ 2016 (anchors measured; curves modelled).

Figure VI.6 turns it around. How long would 300 items take at different dwell times? At 1.7 seconds, 300 items take 8.5 minutes. At 30 seconds, 150 minutes. Measured time in feeds is 97 to 141 minutes, so the dwell that would reconcile 300 items a day with measured time is about 20 to 28 seconds, far above 1.7. The truth probably includes a mixture: many items passed in a second or two, and a few watched for a minute.

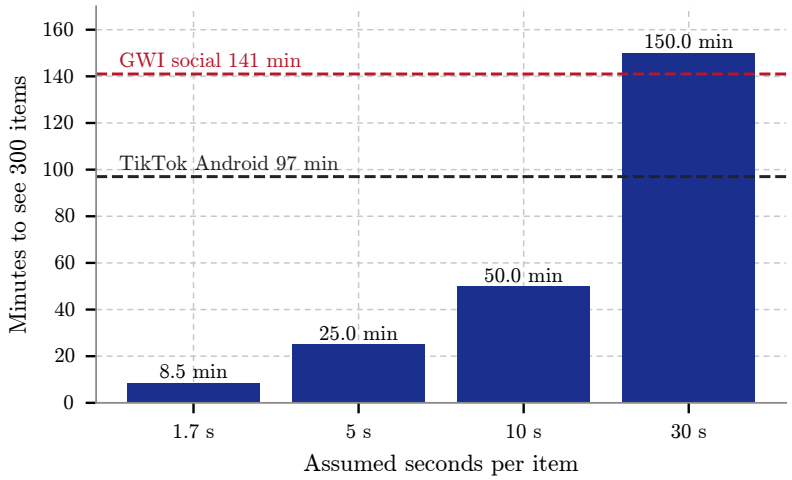


Figure VI.6. Illustrative model, not measured. Minutes needed to see 300 items at an assumed dwell, against two measured daily time anchors (dashed). 300 per day is the author's unsourced figure used as an input only. Source: arithmetic (300 × seconds ÷ 60); anchors from GWI/DataReportal 2024 and Similarweb 2025 (anchors measured; bars modelled).

What would be needed to turn the model into a measurement is simple: the distribution of dwell times per item on current short-video feeds, together with video length. I could not source either. TikTok video counts built from guessed video lengths are unsupported, so I do not give one.

I want to be careful about what the model licenses. It does not show that people scroll 9,000 times a month. It shows that the claim is arithmetically compatible with the time people measurably spend, and that the real number could be many times higher. It is a framing device. The argument of the paper does not rest on it.

What the argument actually needs

The argument needs only the measured part. Adults spend about 2.5 hours a day in social and video feeds. Teens spend more. A typical TikTok user opens the app 10 to 12 times a day. Those are enough.

Add one more fact. The average dwell on an item is short, measured in seconds, not minutes. The one real number is the 2016 Facebook figure; it is old, from a vendor, and an average. But no one disputes the direction. A feed is a machine for moving past things. Anyone who wants to be seen in it competes with the thing before and the thing after, and has about a second or two to give a reason to stop.

Two things I will not say

The eight-second goldfish

You have heard that human attention span fell to 8 seconds, shorter than a goldfish's 9. It is a myth. The BBC's More or Less traced it to a 2015 Microsoft Canada report that cited a source called Statistic Brain, whose own sources nobody could find.[12] Researchers found no evidence behind it. I do not use it, and if you hear it in a pitch, ask for the source.

[12] BBC News, "Is the attention span of a goldfish really longer than ours?", Simon Maybin, More or Less. `bbc.com/news/health-38896790`. Experts quoted found no evidence that attention spans are shrinking. Goldfish memory is also a myth.

Gloria Mark's 47 seconds

The better-known number comes from Gloria Mark's research: that attention on any screen averages about 47 seconds before people switch. It is real research, but it is easy to misread. It measures how long a person stays on one screen before switching, observed in US knowledge workers and students. It is not a cognitive attention span, and it is not a measure of how long people can concentrate when they choose to.

I also could not open the underlying papers. The ACM pages blocked my access. Her website states the 47-second figure and I cite that.[13] Where the claim is repeated as a decline from 2.5 minutes to 47 seconds, I could not verify the earlier numbers.

[13] Gloria Mark, author page on attention span, via the Internet Archive (gloriamark.com). Underlying papers: DOIs 10.1145/1054972.1055017 (2005) and 10.1145/2858036.2858262 (2016), not opened. Earlier values such as 2.5 minutes in 2004 and 75 seconds are widely repeated and I could not verify them, so I do not use them.

So what is the honest version? Time spent in feeds is large and measured. Time per item is short. Neither fact says people can no longer concentrate. People binge long series and play games for hours. The change is the environment: the default place where people spend free time is a stream of small items, each decided in an instant.

The whole phone, and the youngest users

Feeds sit inside a larger pile. Sensor Tower estimates that people spent 5.3 trillion hours in mobile apps in 2025, up 3.8 percent, or about 3.6 hours a day per mobile user.[14] Ofcom's metering puts UK adults at 4 hours 30 minutes a day online in May 2025.[15] These are totals across everything: messaging, banking, maps, games. Feeds are a large slice, not the whole.

Notice the growth rates. Mobile app time grew under 4 percent in a year. Global social time fell slightly over two years. What is growing is the share that goes to short video. Meta said Instagram video time grew about 30 percent in a year; that is a shift inside the pile, not an expansion of it. This is the pattern the thesis needs. A new product does not find new hours. It competes for a mix that is being rearranged toward the products with the best recommendation engines.

The youngest users are being trained into the same mix. The 2025 Common Sense census of children aged 0 to 8 finds screen time steady at about 2.5 hours a day, gaming up 65 percent in four years, and short-form video such as TikTok and Shorts rising; 40 percent of children have their own tablet by age 2.[16] Total time is flat; the form is shifting. That generation will arrive at its first phone already fluent in the feed.

[14] Sensor Tower, *State of Mobile 2026*. A vendor panel estimate, medium confidence. The growth rate matters more than the level: 3.8 percent is slow.

[15] Ofcom, *Online Nation 2025*. Passive metering, full text read.

[16] Common Sense Media, *The 2025 Common Sense Census: Media Use by Kids Zero to Eight*. Parent survey. I read only the landing page, so confidence is medium.

Where news went

News is a clean test case, because we have a long series on where people go for it. In the Reuters Institute's 2026 Digital News Report, 54 percent of people across 48 markets use social or video platforms for news, against 51 percent who use news sites or apps directly.[17] The publisher's own front door has been overtaken by the feed. The same report puts trust in news at 37 percent, the lowest since 2015, down from a steady 40 percent in 2023–2025.

That is a distribution fact. A publisher with a good product, an editorial staff and a brand has still lost the primary channel to somebody else's feed. If this happened to an industry with decades of reputation, it should worry a founder who has a launch post and a landing page. The next chapter takes up trust directly.

[17] Reuters Institute, *Digital News Report 2026*, executive summary. Online survey, 48 markets. By platform: YouTube 34 percent, Instagram 26 percent, TikTok 20 percent use it for news.

Reading the shopping data carefully

One more attention number bears on buying. Deloitte's 2025 survey of US consumers found that 63 percent of Gen Z say social media ads and reviews are the most influential on their purchases, and 56 percent say social content is more relevant to them than TV or film.[18] I treat this with care. It is a survey of stated influence, not a measurement of sales. But it points the same way as the time data: for the youngest adult buyers, the feed is where products are found.

[18] Deloitte Digital Media Trends 2025 (US). Survey, medium confidence. Gen Z also report spending about 50 minutes a day more than the average on social and user-generated video.

Where the 4,000 ads a day goes

One more number belongs to this family: "4,000 to 10,000 ads a day". The earliest I could trace is a 2007 New York Times article quoting Yankelovich, which said a city dweller might face up to 5,000 ad messages a day, against up to 2,000 thirty years earlier.[19] It is an estimate that counts logos. I do not use the larger numbers.

[19] *New York Times*, 15 January 2007, quoting Yankelovich; via the Internet Archive. An estimate that counts logos and brand impressions; no method disclosed. I could not trace the 10,000 figure.

What those hours add up to

Daily minutes are hard to feel. Multiply them out. This is arithmetic on the measured anchors, with no new data.

Anchor (measured)	Per day	Per year	In 24-hour days
GWI social, global (2024)	2h21m	858 h	35.7
GWI social and video feeds (2026)	2.5 h	913 h	38.0
TikTok, Android average user (2025)	1h37m	590 h	24.6
Gallup, US teens, social (2023)	4.8 h	1,752 h	73.0

An average adult on the GWI feeds figure gives about 38 full days a year to feeds. An average teen in the Gallup survey gives about 73 days, two and a half months of round-the-clock time. I would not push the teen number hard, since it is self-report and self-report can overstate. But even if I halve it, it remains a very large share of a young person's free time.

A reading guide to the numbers in this chapter

Because I am mixing measured and modelled material, I want to list which is which.

- **Measured, strong:** daily social time (GWI), UK minutes by

platform (Ofcom passive metering), teen constant-use shares (Pew), entertainment screen hours (Common Sense).

- **Measured, with caveats:** TikTok minutes and opens (Android only, app users only); 1.7 seconds per item (Facebook, 2016, mean, vendor); Gallup teen hours (self-report).
- **Illustrative model:** 300 scrolls a day, 9,000 a month; every curve in Figures VI.5 and VI.6.
- **Rejected:** the 8-second attention span; 300 feet or 1.7 miles scrolled; 4,000 to 10,000 ads a day as a measured fact; any claim that 47 seconds is a cognitive limit.

I could also have repeated the often-quoted phone pickup counts. I did not fetch a primary source for them, so I leave them out.

What this means for the thesis

Pull the measured pieces together. Adults give about 2.5 hours a day to feeds. Teens give much more. The platforms with the highest daily engagement among teens are video-first. Dwell per item is short. And total time has not been growing without limit: the global social figure fell a little between 2022 and 2024.

That last fact matters for the thesis. If attention were still expanding, more products could find an audience by taking the new hours. It is not. The hours are mostly taken. Every product that wants a share has to take it from something else.

The next chapter turns from time to trust. Time in the feed is one limit. The other is whether what shows up in the feed is believed.

VII. The Trust Recession

Attention is the scarce input. Trust is what turns attention into a customer. The measured evidence says trust in news is at a decade low, almost nobody rates advertising as honest, and the cost of misplaced trust, reported fraud, has nearly doubled in three years. AI did not start this. It is about to make it cheaper to continue.

> *Attention is what you pay to be told something. Trust is what you pay to believe it.*
>
> Author's framing, not a citation

Chapter five argued that human attention is a roughly fixed budget. People spend about two and a half hours a day on social and video feeds, and that number is not rising fast. If the supply of things competing for that budget is exploding, as chapter two argued, then each item gets a smaller share.

This chapter is about the second filter. Even when someone does look at your product, they have to believe you. Here the evidence is less ambiguous than the attention evidence, and less comfortable. I will go through it in five steps: trust in news, trust in advertising, where young people now look for information, what misplaced trust costs, and how much of the content and traffic on the web is now synthetic or automated. Then I will show one clearly labelled model of how trust and supply interact.

A caution first. A "trust recession" is my phrase for a pattern across several independent surveys. No single series proves it. Surveys of trust measure what people say, not what they do. Where a number is a vendor estimate I say so.

Trust in news: the lowest on record for this survey

The Reuters Institute surveys about 48 markets each year on news. In 2026 it found that 37% of respondents trust most news most of the time. That is the lowest figure since the series reached its current form in 2015.[1] The three years before that were flat at 40% (Figure VII.1). So this is a three-point drop after a long plateau, not a collapse. In the United States the 2026 figure is 25%.

[1] Reuters Institute, *Digital News Report* 2026 executive summary, and 2025 edition. Online survey across 48 markets.

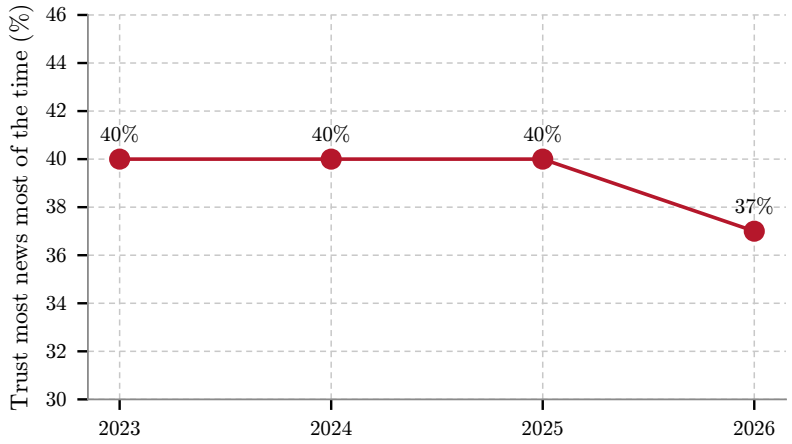


Figure VII.1. Share of respondents who trust most news most of the time, all markets, 2023–2026. Source: Reuters Institute Digital News Report 2025, 2026 (measured; online survey, 48 markets). The 2023–2025 values were reported as stable at 40%.

The same report shows where people get news. In 2026, 54% said they use social or video platforms for news in a typical week. 51% said they use news sites or apps directly. By platform, YouTube was used for news by 34%, Instagram 26% and TikTok 20%.[2] The platform people trust least has become the front door.

I do not want to oversell this. A three-point move is within the range a reader should treat as a drift. What matters for the thesis is the direction combined with the channel shift. News publishers used to own a distribution channel and a reputation together. The reputation is thinner, and the channel is now rented from someone else.

[2] Reuters Institute DNR 2026; same survey. "Social/video" is the report's grouping.

Advertising was never trusted; it is not getting better

Gallup asks Americans to rate the honesty and ethics of professions. Advertising practitioners are in the list most years. The question is the share who rate them "very high" or "high."

The answer has barely moved in two decades, and it has been low the whole time: between 10% and 15% in nearly every year from 2003 to 2022, with a trough of 6% in 2007 (Figure VII.2). In 2022, the latest point in the series I could verify, it was 15%.[3]

[3] Gallup, Honesty and Ethics of Professions, annual; chart data series. Advertising is not asked every year and I could not retrieve a point after 2022.

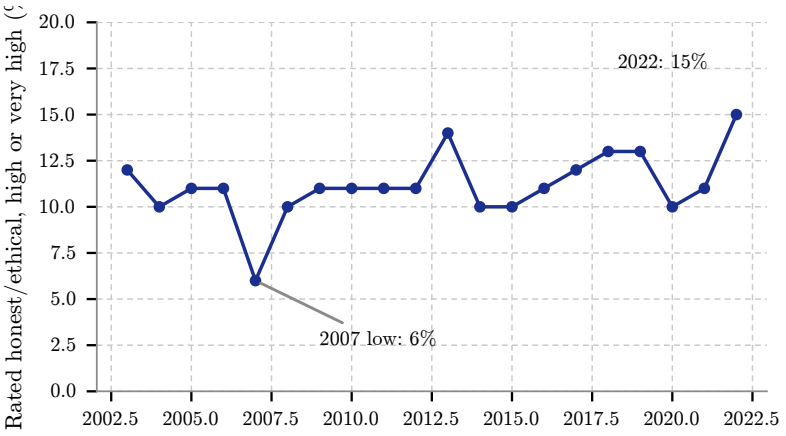


Figure VII.2. Share of US adults rating advertising practitioners "very high" or "high" for honesty and ethics, 2003–2022. Source: Gallup, Honesty and Ethics of Professions (measured; annual poll, not every year asked).

This is a useful correction to a lazy version of the story. It is tempting to write that trust in advertising has collapsed. It has not, because it started near the floor. The honest claim is narrower. Advertising has run on very low trust for twenty years and still produced a market worth, by one agency estimate, about $1.3 trillion globally in 2026.[4] Low trust did not stop spending. It raised the importance of the places where advertising feels like something else: a friend's recommendation, a creator, a search result.

Nielsen's older trust studies make the same point from the other side. In 2012, 92% of people surveyed said they trust recommendations from people they know, and by 2013 that was 84% across 58 countries. Banner ads sat far lower.[5] The gap between earned and paid trust is old and large.

There is also a behavioural signal. A Global Web Index (GWI) survey found nearly one in three consumers say they use an ad blocker (21% regularly, 11% sometimes), with higher use in Asia-Pacific (39%) than in North America (15%).[6] That is a level, not a trend, but it shows that a large minority act on distrust.

[4] WPP Media, This Year Next Year, Midyear 2026 forecast: global ad revenue growing 8.9% to $1.3 trillion (ex US political). Agency estimate, not audited. Same source: Alphabet, Meta and Amazon control 57.6% outside China.

[5] Nielsen, Global Trust in Advertising, 2012 and 2013 editions; online survey of internet users. These are old and the series is incomplete after 2013, so I use them only to show the gap between personal and paid trust, not a trend.

[6] GWI, ad-blocker page. The page is undated, originally from 2019 and later refreshed, so treat this as a level, not a trend.

Where young people look for the truth

Here is the single most striking number I found for the thesis. In the Pew Research Center's September 2025 survey of 5,195 US adults, 56% said they have at least some trust in the information they get from national news organizations. That is down 11 points since March 2025 and 20 points since 2016. Among adults under 30, 51% trusted national news and 50% trusted information from social media

sites (Figure VII.3).[7]

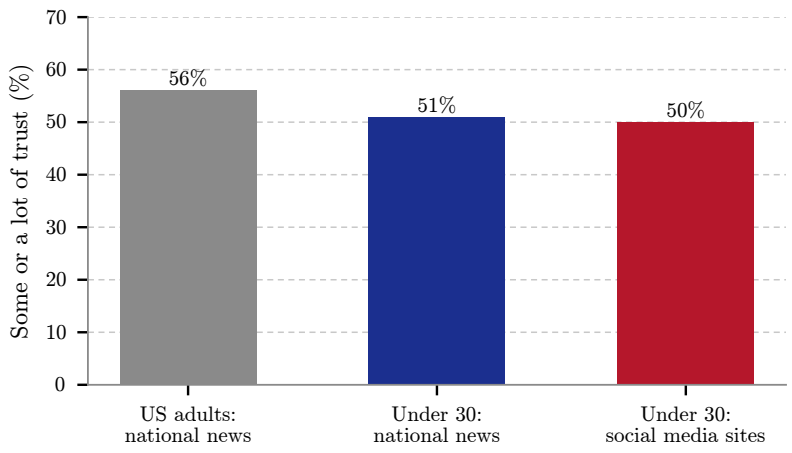


Figure VII.3. Trust in information from national news organizations and social media sites (some or a lot of trust), US adults, September 2025. Source: Pew Research Center, 29 Oct 2025 (measured; probability panel, n=5,195).

The gap is one point. For under-30s a feed is trusted about as much as a newsroom. Read that carefully. It does not say young people trust TikTok highly. It says they have stopped trusting institutions much more than the average, and they have not replaced them with something they trust a lot. They are using what is there.

The attention data from chapter five fits. US teens report being online "almost constantly" at rates that doubled over a decade, and four in ten said so in 2025. On TikTok specifically, the share saying they use it almost constantly rose from 16% in 2022 to 21% in 2025.[8] I take these as time-use facts. I do not take them as evidence about belief.

Two numbers go in the same direction. Deloitte's 2025 survey of US consumers found 63% of Gen Z say social ads and reviews are most influential on their purchases.[9] So the young are not skipping ads. They are trusting the creator or the reviewer who carries the ad. For a founder this is the key operational fact: a brand does not get believed by saying it is good. It gets believed by someone the viewer already follows.

What misplaced trust costs

The cleanest evidence of a trust problem is money lost. The US Federal Trade Commission publishes reported fraud losses from consumer complaints. They rose from $8.8 billion in 2022 to more than

[7] Pew Research Center short read, 29 Oct 2025, "How Americans' trust in information from news organizations and social media sites has changed over time"; probability-based panel. The 11-point and 20-point declines are Pew's figures for all adults. The overall US figure for social media is not shown here because I did not verify it.

[8] Pew, Teens, Social Media and AI Chatbots 2025 (Dec 2025); probability panel, n=1,458, self-reported.

[9] Deloitte Digital Media Trends 2025. Survey; I rate this medium confidence and cite it only for direction.

$10 billion in 2023, $12.5 billion in 2024, and $15.9 billion in 2025, from about 3 million reports (Figure VII.4, left panel).[10] That is up about 80% in three years. Investment scams alone were $7.9 billion of the 2025 total.

[10] FTC press releases, Feb 2023, Feb 2024, Mar 2025 and Mar 2026. Self-reported complaints. The FTC says real losses are higher, since only about 4.8% of victims complain. The 2023 value is "more than $10 billion" and is plotted as 10.

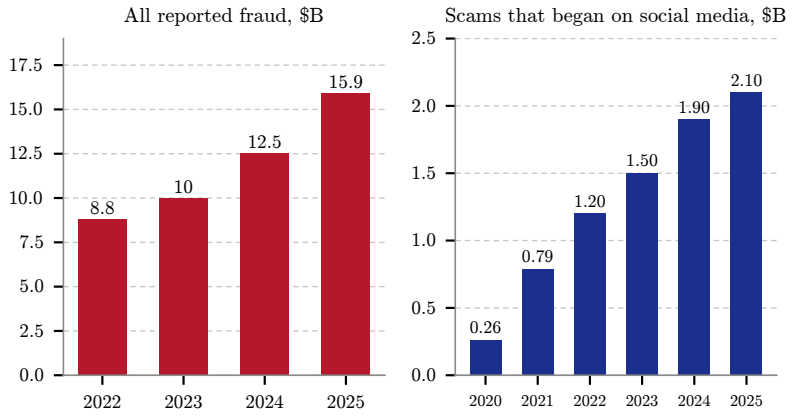


Figure VII.4. Left: total US reported fraud losses, 2022–2025. Right: reported losses where the scam began on social media, 2020–2025. Source: FTC press releases and FTC Data Spotlight, April 2026 (measured; self-reported complaints, lower bound of true losses).

The social-media slice is the one that matters here. Losses reported by people who first met the scammer on social media rose from $261 million in 2020 to $789 million in 2021, $1.2 billion in 2022, $1.5 billion in 2023, $1.9 billion in 2024 and $2.1 billion in 2025.[11] That is roughly eight times in five years. Social media was the contact channel in 28% of loss reports. The FTC found that over 40% of social-shopping scam victims began with an ad. And 40% of loss reports from ages 18 to 29 began on social media.

[11] FTC Data Spotlight, April 2026, "Reported losses to scams on social media eight times higher than 2020." Complaint data; no 2025 government-shutdown reports were included.

Two honest caveats. First, more people on social media, and better reporting, push these numbers up without any increase in scam intensity. Second, I could not find a measured denominator for time spent, so I cannot say the rate per hour rose. What I can say is that the dollar figure rose eightfold and that young adults are overrepresented.

I considered including a widely reported claim about scam-ad revenue at a large platform. It rests on leaked internal documents seen only through secondary sources, so I am leaving it out.

The mechanism is not subtle. When a feed contains a mix of real ads, creator posts, scams and synthetic content, the viewer has to do the sorting. Sorting costs attention. A viewer who has been burned once discounts everything. That discount applies to the honest small company as well as the scammer. This is the trust externality: a bad actor's cost is paid by everyone who sells in the same place.

In the UK, Ofcom found 34% of adult internet users had encountered

scams, fraud or phishing in the previous four weeks (June 2025), and 22% had seen fake or deceptive images or videos, up from 18% a year earlier.[12] That is a different country and a different method, and it shows the same direction.

[12] Ofcom, Online Nation 2025; survey of UK adults.

Synthetic content and automated traffic

Now the supply side of the trust problem. How much of what is published and clicked is machine-made?

Graphite, a content and SEO firm, analysed a sample of English-language web articles from Common Crawl with detectors. Its estimate of the share primarily written by AI was 35.9% about a year after ChatGPT launched, 49.6% in the first quarter of 2025, 50.9% in the fourth quarter, and 49.9% in the first quarter of 2026.[13] (Figure VII.5, left.) The honest reading is that AI-written articles reached about half of new articles and then stopped rising. Graphite also notes that such articles mostly do not rank in Google or in chatbot answers. Output is not the same as being seen.

[13] Graphite, "AI now writes as many online articles as humans do"; three detectors; the vendor reports false positives of 1.4–1.8% and false negatives under 2%. Vendor study.

A second vendor, Ahrefs, found 74.2% of 900,000 new pages in April 2025 contained some AI content, but only 2.5% were pure AI.[14] Most were a mix. So "half the web is AI" is a statement about drafts and listicles, not about every page.

[14] Ahrefs, "What percentage of new content is AI-generated." Vendor detector; includes light use. I rate it medium confidence.

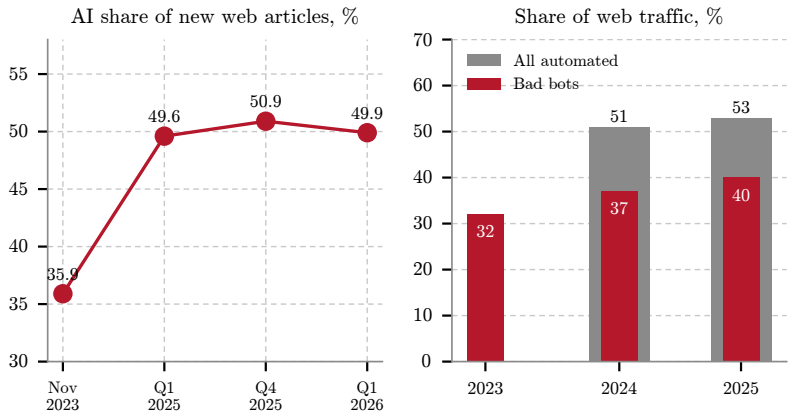


Figure VII.5. Left: share of new English web articles detected as primarily AI-generated (Graphite, vendor study; first point is about 12 months after ChatGPT launched). Right: automated share of web traffic and malicious ("bad") bots, by data year (Thales/Imperva, vendor network telemetry). Both measured by detectors with error; see text.

On traffic, Thales/Imperva reports that automated traffic was 51% of web traffic in its 2024 data and 53% in its 2025 data. Bad bots were 32% in 2023, 37% in 2024 and 40% in 2025.[15] Bots are not people

[15] Thales (Imperva) Bad Bot Reports 2025 and 2026; vendor network telemetry. "Bots" include benign crawlers such as search engines.

and a click from a bot is worthless to an advertiser. But I want to be careful about the inference. Much of this traffic is crawlers and scrapers that never see an ad and do not move a campaign. The figure shows how much of the web's activity is not human. It does not show that half of ad impressions are fake.

Taken together the two series say the same thing in different words. Cheap generation makes the web fuller. A fuller web is harder to read. A viewer facing doubt about both the content and the traffic will fall back to sources they already know.

Illustrative model: supply, attention and belief

The pieces above are measured. They do not combine into a number. I want to show the shape of the interaction, so here is a deliberately simple model. It is an illustration, not a forecast, and none of its inputs are measured.

Assumptions. A viewer has a fixed attention budget, which I set to one unit. There are N competing items, where $N = 1$ is today. An item gets a share of attention proportional to $1/N$. Of the items on offer, a fraction b are untrustworthy, and a viewer who cannot tell them apart believes a given item with probability $1 - b$. I normalize both terms to today, with $b_0 = 10\%$. The chance that an honest item is seen and believed, relative to today, is then

$$R(N) = \frac{1}{N} \cdot \frac{1 - b(N)}{1 - b_0}.$$

I run three cases. In the first, b stays at 10% however large N grows, so the only effect is crowding. In the second and third, b rises by 3 and 6 percentage points for each additional multiple of today's supply.

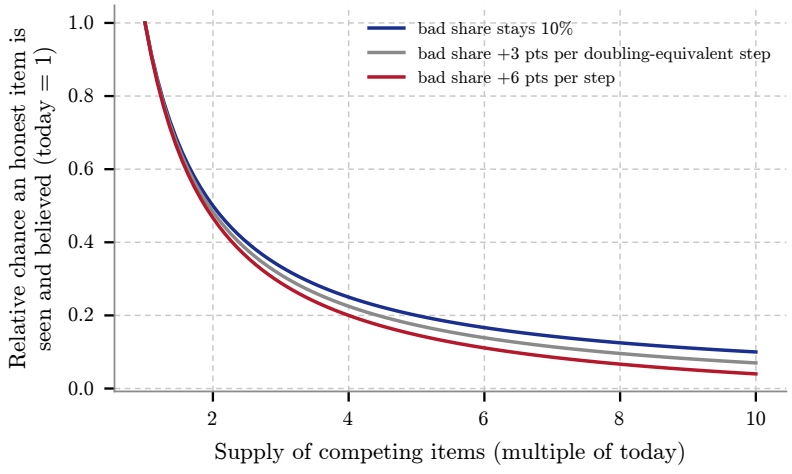


Figure VII.6. Illustrative model, not measured. Relative chance that an honest item is seen and believed as supply of competing items grows,

with a fixed attention budget and three assumed paths for the share of untrustworthy items. All inputs are assumptions chosen to show the shape of the effect. Modelled.

Figure VII.6 shows the result. Even with no trust effect, doubling supply halves the chance of being seen. With a rising bad share the loss is worse than proportional, because trust decays on top of crowding. At five times today's supply the relative chance is 0.20 for crowding alone, about 0.17 in the middle case and about 0.15 in the steep case. Crowding does most of the work; the trust effect adds a further 15 to 25 percent loss on top.

What does the model not say? It does not say supply will grow fivefold, that the bad share will rise, or that viewers cannot tell good from bad. Real viewers do filter, and the filter is the point. Trust that cannot be read in the item itself moves to the source: a name, a creator, a community, a track record. That is the argument the next chapter takes up.

What this means for the thesis

Let me be exact about what the evidence supports.

- Trust in news is at its lowest point in this series, at 37%, but the fall is three points after a plateau.
- Advertising honesty ratings have been very low for twenty years. They did not collapse; they were never high.
- Under-30s in the US trust social media about as much as national news (50% against 51%). The change is the weakening of institutions, not a rise in the feed.
- Reported fraud losses rose from $8.8 billion to $15.9 billion in three years, and losses from social-media-origin scams rose about eightfold in five years. Both are self-reported lower bounds.
- About half of new English articles are AI-written by one estimate, and 53% of web traffic is automated by another. Both are vendor measures, and neither shows that people see or click more of it.

What the evidence does not support is the claim that attention spans have collapsed, or that any one trust series has crashed. I have left out the viral figures for scroll distance and for "8 seconds of attention" because they have no traceable source. The BBC traced the second to a 2015 report that quoted a source nobody could find.[16]

The thesis does not need the dramatic version. It needs only this. Supply is up. Attention is roughly fixed. Trust is low and unevenly

[16] BBC More or Less (Simon Maybin), traced to a 2015 Microsoft Canada report; no evidence that attention spans are shrinking.

distributed, and the cost of being fooled is rising. In that world, being seen is necessary, and being believed is the scarce part. Both come from distribution, meaning the standing channel between you and a person who already has a reason to listen. That is the subject of chapter seven.

I will end on a limit of my own. Every trust measure here is a survey answer, and survey answers are cheap talk. The fraud numbers are the exception, since they count real money. If you want one number from this chapter to carry forward, take $15.9 billion.

VIII. What Survives

If code is cheap, something else must be expensive. The evidence says four things still are: silicon and power, the systems businesses already run on, trust, and the permission to operate in regulated places. One thing people expected to be scarce, proprietary data, is a weaker moat than the story suggests.

> *A moat that must be continuously rebuilt will eventually be no moat at all.*
>
> Warren Buffett, Berkshire Hathaway letter, 2007

The argument so far is that AI made building cheap, supply of software exploded, and distribution became the scarce thing. This chapter tests that from the other side. If building is no longer the moat, what is? Where does scarcity actually live in October 2026?

I went looking for the strongest evidence I could find, including evidence against my own thesis. The honest summary is this. The idea that moats move off code is well supported as a description of where money and scarcity are going. It is poorly supported as a claim that building is cheap in any sense a large company can measure.[1] Writing code got cheaper. Owning it, securing it, running it, and getting someone to use it did not.

Here is the plan. I take six candidates for what survives, in order of how well the evidence holds up: compute and power, incumbents with systems of record, data, regulation, trust and brand, and community. I give the numbers, say what each one proves, and say what it does not.

[1] This is the overall verdict of the research dossier behind this chapter (F_survive_counterviews, accessed 7 October 2026). Several of the best counter-views come from advocates of the thesis itself, including Sam Altman (August 2026) and the a16z essay of March 2026.

Compute and power: the scarce input moved to silicon

Start with the most measurable thing. When code got cheap, the money did not go away. It went to chips, buildings and electricity.

Four companies, Amazon, Alphabet, Meta and Microsoft, spent about $376 billion in cash on property and equipment in calendar 2025.[2] In the first half of 2026 alone they spent about $295 billion.[3] Figure VIII.1 shows the climb. Amazon alone went from $52.7 billion in 2023 to $131.8 billion in 2025.

[2] Author's sum of SEC XBRL cash-flow values from 10-K and 10-Q filings: Amazon 131.8, Alphabet 91.4, Meta 69.7, Microsoft 83.1 (US$ bn). Microsoft's calendar quarters are derived by subtraction from its June fiscal year. The total is my computation, not a company-reported figure. It excludes finance-lease additions, so company-reported capex including leases runs higher.

[3] Same method: Amazon 98.4, Alphabet 80.6, Meta 49.1, Microsoft 66.7.

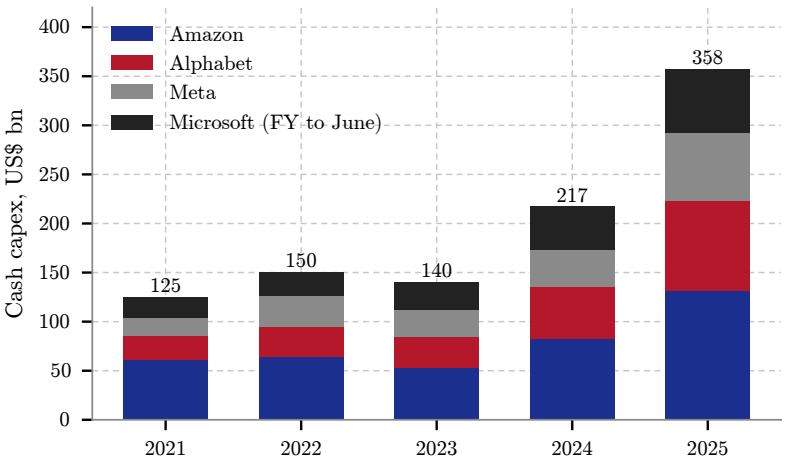


Figure VIII.1. Cash capex of the four largest hyperscalers, 2021 to 2025, US$ bn. Amazon, Alphabet and Meta are calendar years; Microsoft is its fiscal year ending 30 June, so the stacks mix bases and are not the calendar-2025 total of about $376bn quoted in the text. Cash purchases of property and equipment only. Source: SEC XBRL filings (measured).

The guidance for 2026 is bigger still. Amazon said it expects about $200 billion in capital expenditure in 2026. Meta narrowed its range to $130–145 billion in July. Microsoft said calendar-2026 capex would be approximately $175 billion and that fiscal 2027 spending would grow year over year.[4] The guided ranges add up to roughly $700–725 billion. I do not give a precise total because the companies define capex differently, with finance leases in or out. Treat it as an order of magnitude.

This is now partly financed with equity. In June 2026 Alphabet raised $49.6 billion in net proceeds from stock and mandatory convertible preferred, to be used for general corporate purposes including capital expenditures.[5] I read that as a signal. When even the richest companies on earth reach for the equity market, capacity is a balance-sheet contest, and balance sheets are scarce.

[4] Amazon 8-K, 5 Feb 2026; Meta 8-Ks, 28 Jan and 29 Jul 2026; Microsoft FY26 Q4 earnings call, 29 Jul 2026. Alphabet guided $195–205 billion on its Q2 2026 call, but I only found that in a third-party transcript, so treat it as secondary.

[5] Alphabet Q2 2026 earnings 8-K, SEC EDGAR.

The supplier's side of the same bill

Someone receives that money. Nvidia's Data Center revenue went from $3.6 billion in the quarter ending January 2023 to $89.0 billion in the quarter ending July 2026, a rise of about 24.6 times in fourteen quarters (Figure VIII.2).[6] In that same latest quarter Nvidia reported total revenue of $96.2 billion, up 106% year over year, with a GAAP gross margin of 75.0%. It guided the following quarter to $108.0 billion, plus or minus 2%.

[6] NVIDIA quarterly results, Exhibit 99.1 to Form 8-K, 15 quarters, via SEC EDGAR. Fiscal-year totals: FY23 $15.0bn, FY24 $47.5bn, FY25 $115.2bn, FY26 $193.7bn. The ratio is my arithmetic: 89,023 divided by 3,616.

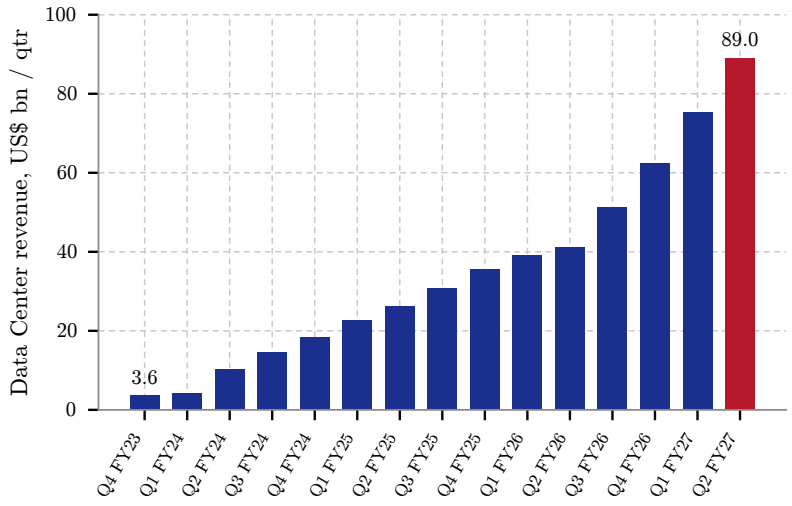


Figure VIII.2. NVIDIA Data Center revenue by fiscal quarter, US$ bn, Q4 FY23 (Jan 2023) to Q2 FY27 (Jul 2026). Source: NVIDIA earnings releases, 8-K Exhibit 99.1 (measured).

A 75% gross margin is what a scarce input looks like. Software at that margin would have competitors in a month. Chips at that margin have waited years. About 55% of the latest quarter's Data Center revenue came from hyperscale customers ($48.7 billion), the same companies in Figure VIII.1.[7] The industry is, to a good approximation, a small number of buyers funding a small number of sellers.

Power is the next constraint down the chain. The International Energy Agency estimates that data centres used about 415 terawatt-hours in 2024, roughly 1.5% of world electricity. Its Base Case has that reaching about 945 TWh by 2030, just under 3% of the world total, growing around 15% a year. The United States alone adds about 240 TWh, an increase of 130% (Figure VIII.3).[8]

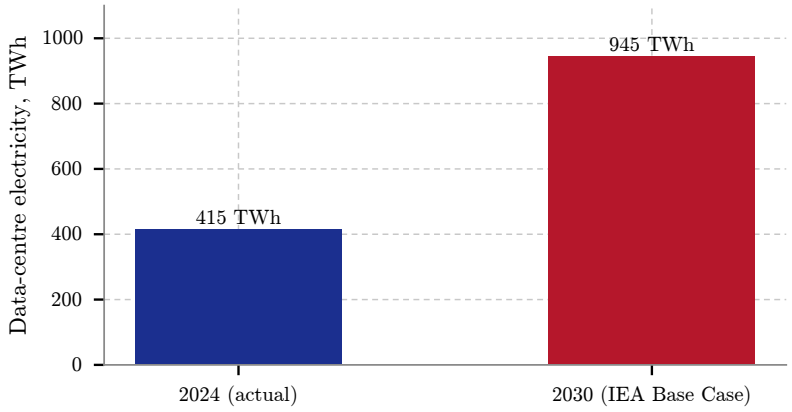


Figure VIII.3. Global data-centre electricity use, 2024 actual and the IEA Base Case for 2030, TWh. The 2030 bar is a scenario projection, not

[7] Nvidia Q2 FY27 8-K. A customer was reclassified between categories and prior periods recast, so the mix is not perfectly comparable across long spans.

[8] IEA, Energy and AI (2025), "Energy demand from AI". This is a Base Case projection, not a measurement of anything yet.

an observation. Source: IEA, Energy and AI (2025) (estimate).

What this does and does not prove

Capex is a cost, not a moat. It becomes a moat only for the firms able to finance it. There is an academic version of this point: Azoulay, Krieger and Nagaraj argue that tight control over complementary assets "will likely result in a concentrated market structure, as in past episodes of technological upheaval."[9] Value is also visibly collecting at the capital-heavy layer. Anthropic announced a $65 billion Series H at a $965 billion post-money valuation on 28 May 2026 and said its run-rate revenue had crossed $47 billion.[10]

Two cautions. The margins say a scarce physical layer earns rents today. They say nothing about 2030. And none of this makes the scarce input available to you or me. It means the cost of building software fell while the cost of running intelligence at scale became one of the biggest line items in the economy. For a small builder, that is the opposite of a moat. It is a toll.

[9] Azoulay, Krieger, Nagaraj, "Old Moats for New Models", NBER Working Paper 32474, May 2024 (revised July 2024). It is a forecast from innovation economics, not an empirical finding.

[10] Anthropic news release, "Series H", 28 May 2026. I found OpenAI's reported $122 billion round at $852 billion only in a secondary source, so I leave it out of the argument.

Incumbents: systems of record are holding

The second candidate is the one that surprised me. In early 2026 the market priced a software panic. The a16z essay of March 2026 reports that ETFs of public software companies had fallen 30% since the start of the year, erasing all gains since ChatGPT launched, with Salesforce, Adobe, Intuit, ServiceNow and Veeva down 25 to 30% in weeks.[11] The fear was simple: if an agent can write the app, why pay for the app?

Then the quarterly results arrived.

[11] a16z, "Good news: AI Will Eat Application Software", 2 Mar 2026. This is an investor's statement, not a price series; I did not fetch one. Business Insider (1 Jun 2026) separately reports Atlassian, Salesforce and SAP down more than 20% year to date.

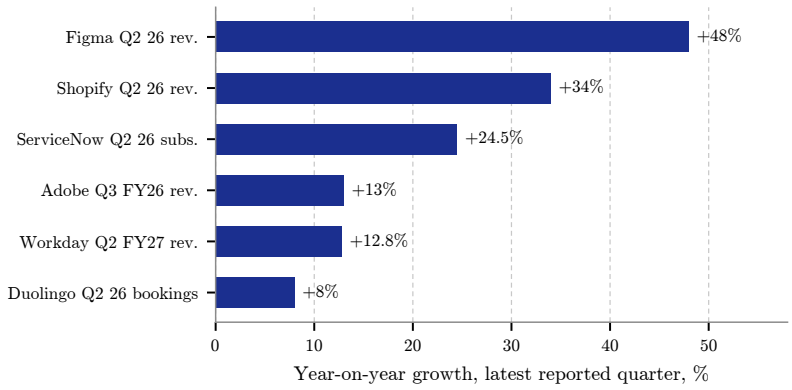


Figure VIII.4. Year-on-year growth in the latest reported quarter, selected software incumbents. Metrics are company-defined and not like-for-

like: Figma, Shopify, Workday and Adobe are total revenue, ServiceNow is subscription revenue, Duolingo is bookings. Source: company 8-K earnings releases, 2026 (measured).

ServiceNow grew subscription revenue 24.5% to $3.877 billion in Q2 2026, with current remaining performance obligations of $13.20 billion, up 21%. Workday grew revenue 12.8% to $2.649 billion. Adobe grew 13% to $6.76 billion, with total ARR of $27.50 billion. Salesforce guided fiscal 2027 revenue of $46.1–46.4 billion, up 11–12%, and reported about $3.9 billion in Agentforce and Data 360 ARR, up 210%.[12] Figure VIII.4 lays the growth rates side by side.

[12] Company Exhibit 99.1 earnings releases on SEC EDGAR: ServiceNow Q2 2026, Workday Q2 FY27, Adobe Q3 FY26, Salesforce Q2 FY27. All metrics are company-defined.

Product-led companies did better. Figma's Q2 2026 revenue was $370.1 million, up 48%, with net dollar retention of 136%, and over 80% of its $10,000-plus customers use AI credits weekly. Shopify grew revenue 34%. Duolingo's daily users grew 23%, though its bookings grew only 8%.[13]

[13] Company 8-Ks for Q2 2026, SEC EDGAR. I found Canva and Notion figures only on aggregator sites and left them out.

Retention in context

Figma's 136% is an outlier, not a norm. The SaaS Capital survey puts the median net revenue retention of private SaaS companies with annual contract value of $25,000 to $50,000 at 102%, and retention rises with contract size.[14] If you want a feel for ordinary, 102% is it. A company that keeps 136% of last year's dollars is being pulled along by something, and the likeliest candidate is that its product sits where the work already happens.

[14] SaaS Capital private-SaaS survey, December 2024 data; one bracket read. Salesforce and SAP retention rates were not found in the filings I fetched.

Microsoft is the incumbent AI test case. Microsoft 365 Copilot paid seats were 15 million at the first disclosure on 28 January 2026, 20 million on 29 April, and "over 30 million" in the 29 July earnings release.[15] Seats doubled in six months. But Microsoft 365 had over 450 million commercial seats in January, so the first disclosure was about 3.3% attach.[16] Fast growth, small share. Both facts are true, and that is the honest reading of the whole incumbent story.

[15] Microsoft FY26 Q2 call transcript (Motley Fool); No Jitter for Q3; Microsoft FY26 Q4 8-K on SEC. The January and April numbers come from secondary summaries of calls; the July number is primary. "Over 30 million" is a floor.

[16] 15 divided by 450 is my arithmetic. I did not find a later seat base, so I do not give a later penetration rate.

Who took the new money

None of this means incumbents won the AI wave. Menlo Ventures surveyed about 500 US enterprise decision-makers and estimated that enterprise generative-AI spend reached $37 billion in 2025, up 3.2 times from $11.5 billion, with $19 billion at the application layer. Startups captured 63% of application-layer revenue, against 36% a year earlier.[17] Incumbents held 56% of the infrastructure layer. And 76% of AI solutions were purchased rather than built, up from 53% a year earlier.

[17] Menlo Ventures, "2025: The State of Generative AI in the Enterprise", 9 Dec 2025. The authors are an investor and mark their own portfolio companies. The numbers are a survey plus a modelled revenue estimate, not audited.

Read those together. The systems of record kept their customers.

The new AI application dollars went mostly to new companies. Both happened. Survival of revenue is also not survival of valuation, which is why the same sector can post 24% growth and a 30% drawdown in one year. And Benedict Evans' May 2026 deck adds a warning I take seriously: "Commodity infra rarely captures value up the stack."[18]

[18] Benedict Evans, "AI eats the world", May 2026, slide 24. Slide text was read in full; chart values in the deck were not extracted so I use only the text.

Data: the weaker moat

Data is the candidate I expected to survive best. The idea is old: whoever owns the data that models need has something others cannot copy. The evidence is thinner than the idea.

Reddit as the test case

Reddit is the cleanest example, because it is a public company that sells its data to AI labs and has to report the results. Its S-1 disclosed data licensing arrangements signed in January 2024 with an aggregate contract value of $203.0 million over two to three years, with at least $66.4 million to be recognised in 2024.[19]

Then look at what it became. "Other revenue," which is mainly content licensing, was $15.2 million in 2023, $114.7 million in 2024 and $140.0 million in 2025. Total revenue over the same years was $804.0 million, $1,300.2 million and $2,202.5 million.[20] In Q2 2026 the other line was $43 million out of $805 million in total, about 5%, while advertising was $762 million, up 64%.[21]

[19] Reddit Inc., Form S-1, 22 Feb 2024. Counterparties are not named. A press report (CNET, citing Reuters) puts the Google deal at about $60 million, and some headlines say per year, but Reddit did not publish that figure, so I do not treat it as a disclosed fact.

[20] Reddit 10-K for FY2025, revenue note. "Other revenue" is mainly licensing but not purely AI licensing, so read it as an upper bound on the data business.

[21] Reddit Q2 2026 earnings 8-K. Figures are rounded to US$ millions in the release.

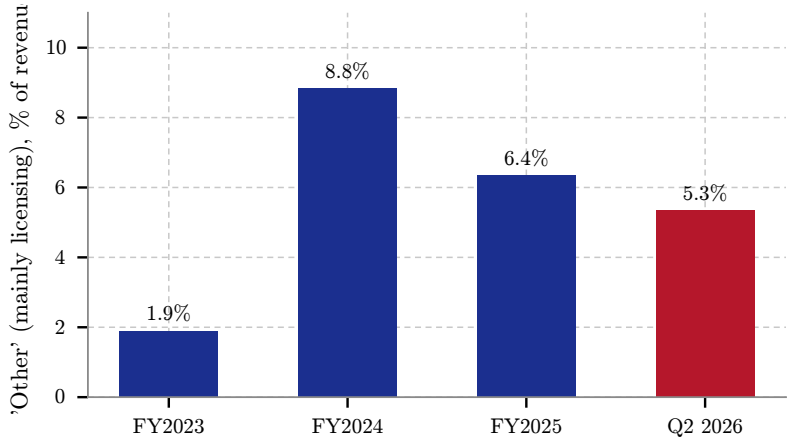


Figure VIII.5. Reddit "other" revenue (mainly content licensing) as a share of total revenue. Other is not purely AI licensing, so the series is an upper bound. Share computed from reported figures: other divided by total. Source: Reddit 10-K FY2025 and Q2 2026 8-K (measured; ratio computed by author).

Figure VIII.5 shows the point. The licence income is real, but it peaked as a share of revenue in 2024 at 8.8% and is about 5% now. For the owner of one of the largest bodies of human conversation on the internet, data is a nice side line. It is not the business. And Reddit says so itself. Its S-1 risk factors warn that "the value of data for AI training purposes may be reduced over time" and that some companies may use its data without a licence "given its open nature."

The a16z team made a related argument long before LLMs. In 2019 Casado and Lauten wrote that "data is rarely a strong enough moat," because the cost of adding unique data may go up while the value of incremental data goes down.[22] In March 2026 the same firm wrote that scarce data becomes "extremely valuable" again, naming Bloomberg's live market data and Abridge's clinical conversations.[23] Both can be right. Commodity data is cheap. Data nobody else can get is not. The question is how much of the world's data is the second kind.

[22] Martin Casado and Peter Lauten, "The Empty Promise of Data Moats", a16z, 9 May 2019. Their illustrative curve came from one customer-support dataset, so it is an anecdote, not a general law.

[23] Immerman and Rodriguez, a16z, 2 Mar 2026. a16z is an investor in some of the companies it names, so this is an interested party.

Model collapse is avoidable

The most common reason given for why human data would stay precious is model collapse. Shumailov and colleagues showed in *Nature* in 2024 that "indiscriminate use of model-generated content in training causes irreversible defects in the resulting models."[24] They concluded that data about genuine human interaction would become more valuable.

But there is a rebuttal with the same weight. Gerstgrasser and colleagues showed that replacing real data with each generation's synthetic data tends toward collapse, while accumulating successive synthetic generations alongside the original real data avoids it.[25] The practical reading is mild. Keep the real data you have and add to it. That does not make real data worthless. It makes it a stock you need to keep, not a flow you must keep buying.

[24] Shumailov et al., "AI models collapse when trained on recursively generated data", Nature, 2024. I read the abstract and references, not the methods. The experiments used small language models, VAEs and Gaussian mixtures.

[25] Gerstgrasser et al., arXiv 2404.01413, April 2024. Abstract read. It tests language models, diffusion models and VAEs.

What a courtroom says data costs

The best-known price tag is a legal one. The Bartz v. Anthropic settlement fund is $1.5 billion, roughly $3,000 per work before fees, for the use of books downloaded from pirate libraries. The court granted final approval on 20 July 2026.[26] I want to be careful. That number prices past piracy. It is not a market rate for training licences and should not be used as "what data is worth."

Put the scale next to the first section. Data deals sit in the hundreds of millions to low billions of dollars. Capex sits in the hundreds of billions. For the labs, data looks more like a cost they pay than a profit pool the owners capture. My defensible line is this: the evidence that proprietary data alone is a durable moat is thin, and the

[26] Settlement notice and settlement website (anthropiccopyrightsettlement.com); Authors Alliance analysis, 19 Dec 2025. The fee request and its percentage are from the secondary source.

evidence that proprietary data plus a workflow nobody can replace is a moat is better.

Regulation: a gate, but a passable one

The fourth candidate is permission. If anyone can build a diagnostic tool, then the right to sell one becomes the scarce thing.

The FDA keeps a list of AI-enabled medical devices. As of the page updated 4 September 2026 it held 1,614 entries, with the latest decision on 29 June 2026.[27] The yearly count of final decisions is 18 in 2016, 65 in 2018, 114 in 2020, 162 in 2022, 235 in 2024, and 335 in 2025 (Figure VIII.6). Another 181 came in the first half of 2026.

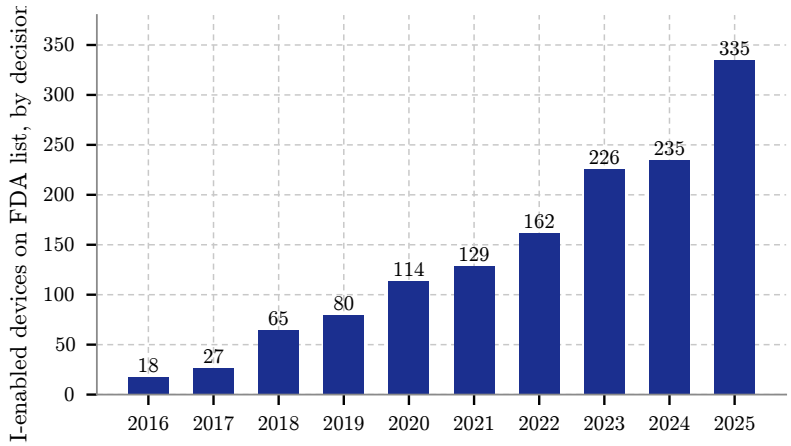


Figure VIII.6. AI-enabled medical devices on the FDA list, by year of final decision, 2016 to 2025. 2026 is a partial year (181 to 29 June) and is not shown. Source: FDA AI-Enabled Medical Devices List, author's count of the CSV (measured; list is revised and "not comprehensive").

Look at what the list is made of. Radiology accounts for 1,230 entries, or 76.2%, with cardiovascular at 154, neurology at 73 and anesthesiology at 30. By pathway, 1,553 came through the 510(k) route, 40 through De Novo, and 21 through PMA. A 510(k) means substantial equivalence to an existing device, which is a lower bar than PMA. So the FDA is a real gate that thousands have already walked through. It is a steady-state certification industry, not a scarce permit.

The European picture points the same way. The EU AI Act entered into force on 1 August 2024. After the Commission tabled the Digital Omnibus on AI in November 2025, a provisional agreement in May 2026 moved the high-risk Annex III obligations from 2 August 2026 to 2 December 2027, and Annex I products to 2 August 2028. Article 50 transparency obligations stayed at 2 August 2026.[28] The delay is

[27] FDA, AI-Enabled Medical Devices List, author's count from the downloadable CSV. The FDA says the list is not a comprehensive resource and items are identified by AI-related terms. fda.gov blocks bots, so I read it through a reader proxy. Yearly counts are my tabulation. Third parties report 1,430 or 1,451 at end of 2025 because the list is revised.

[28] Gibson Dunn client alert, 27 May 2026; arcusx.ai summary, 17 Jul 2026. High confidence for the May agreement; medium for final adoption, since I found only one secondary source and did not see the Official Journal publication.

about sixteen months. It weakens the claim that EU compliance is an imminent barrier.

There is a market signal that is not about regulation but sits in regulated rooms. Harvey, the legal AI company, was valued at $8 billion in December 2025, $11 billion in March 2026 and $15.5 billion in September 2026. OpenEvidence went from $1 billion in February 2025 to $6 billion in October 2025 to $12 billion in January 2026, with its CEO saying 40% of US physicians use it. Abridge raised at $5.3 billion in June 2025.[29]

[29] TechCrunch, 9 Sep 2026; Harvey blog, 25 Mar 2026; CNBC, 21 Jan 2026; Pittsburgh Wire for Abridge (a local outlet, no primary source fetched). Valuations are private-round prices. Usage and revenue claims are the companies' own.

I want to be careful here too. I did not establish that any of these three depends on an FDA clearance or a licence. Their stated moats, per the investors, are data, workflow and brand. I also found no evidence that quantifies SOC 2, HIPAA, FedRAMP or banking-licence counts as barriers, so I say nothing numeric about them. Regulation survives as a cost and a delay. As a moat, the case is modest.

Trust and brand

Now the candidate closest to the thesis. If supply explodes, buyers need a way to choose. Brand is how they do it.

The a16z essay makes the argument directly: "if every industry gets more crowded… we should expect the power of strong brands to increase. Brand is how you signal reliability in a world of infinite optionality."[30] It is opinion, but it fits the rest of this paper, and the ranking data does not contradict it.

[30] a16z, 2 Mar 2026, applying Helmer's Seven Powers. It is an argument by analogy; the essay also says switching costs are "perhaps the one moat that really is going to change."

Kantar BrandZ 2026 values the Top 100 brands at a combined $13.1 trillion, up 22%. The top ten are Google ($1.5 trillion), Apple ($1.4 trillion), Microsoft ($1.1 trillion), Amazon ($1.0 trillion), NVIDIA ($814.9 billion), Facebook ($366.6 billion), Instagram ($286.2 billion), Tencent ($251.6 billion), Oracle ($235.8 billion) and McDonald's ($235.1 billion).[31] Interbrand's 2025 ranking has the same names at the top, with Apple at $470.9 billion, Microsoft $388.5 billion, Amazon $319.9 billion and Google $317.1 billion. Instagram's brand value rose 27.0%.[32]

[31] Kantar BrandZ Global 2026. Values are a model combining survey data and financials, not market capitalisation. I did not compute brand value as a share of market cap.

[32] Interbrand, Best Global Brands 2025. A different methodology from Kantar's, so the levels are not comparable between the two.

What do I conclude? The brand valuations are concentrated in the same platforms that dominate distribution, and these are measured by models, so the level is soft. The concentration, however, is clear in both lists. Brand is the thing the biggest names have in common that a new builder cannot buy with a weekend of prompts.

Paul Graham's *The Brand Age* is a useful analogy and no more. He argues that mechanical watches "cost a lot because brands spend a lot on advertising and use tricks to limit supply."[33] When function is cheap, price follows story. That is the mechanism I think carries

[33] Paul Graham, "The Brand Age", March 2026. It does not mention AI. I use it only as an analogy for what happens when function commoditises.

into software. Chapter VI covered what happens to trust in ads and feeds; the other half is that trust in names rises as a share of what is left.

Habit counts too

Sam Altman said something in August 2026 that I think is underweighted. In an interview with David Senra, reported by Business Insider on 25 August, he said everyone had been "too ambitious on timelines": "They keep buying from the same company. They keep sort of wanting to use their tools in the same way."[34] Habit is distribution that already exists. It is also the reason a trusted name can survive a technology shift.

[34] Business Insider, 25 Aug 2026, quoting an interview posted 23 Aug 2026. The interview itself I did not hear, so I rely on the press quotes.

Other voices are consistent with this. Ben Thompson wrote in 2023 that disruptive innovations "do consistently come from new entrants in a market, but those new entrants aren't necessarily startups." Jensen Huang said in February 2026 "Nobody's going to service better than ServiceNow."[35] I treat these as a pattern, not a proof.

[35] Thompson: Stratechery, "AI and the Big Five", 9 Jan 2023. Huang: CNBC, 26 Feb 2026. Huang sells the chips these agents run on, so treat him as an interested party.

Community

The last candidate is the hardest to measure and the one I care most about as a builder. Community is a group of people who show up for each other and for a shared thing, in a place that belongs to someone.

I have to be straightforward about what the evidence says. I found no dataset in the dossiers that measures community as a moat. What I have is indirect. Reddit's ad revenue grew 64% in Q2 2026 while its licensing line stayed near 5% of revenue. The thing advertisers pay for is the people. The data was a by-product of the community, and the community is what monetises.

Paul Graham adds a practical rule from his September 2026 essay: "Network effects make companies more powerful, and I almost always think about how to introduce them." He also asks whether a supplier can become "the one that owns the relationship with the customer."[36] Ben Thompson made a similar point about Spotify, calling it a feature of every monolithic network-effects company that they are "poised to be big winners from AI."[37]

[36] Paul Graham, "Making Startups Powerful", September 2026. Practitioner heuristic, no data. He also suggests letting users opt in to training the model on their interactions, so that their model outperforms the vanilla one.

[37] Stratechery, "Aggregators and AI", February 2026, free section only. I did not read the paywalled articles.

So here is what I can honestly say. Community is where trust, habit and network effects meet. It is probably the thing a solo builder has the best chance of owning. But I cannot cite a measurement, and I will not invent one. Chapter VIII returns to this as part of the counter-case.

What does not survive, and what the builders say

Before drawing a line under the chapter, here are the counter-views I take most seriously, because several come from the people who most want the thesis to be true.

Dario Amodei said in March 2025 that in three to six months AI might be writing 90% of code, and in twelve months essentially all of it, but that "the programmer still needs to specify… what is the overall app you're trying to make, what's the overall design decision."[38] In his January 2026 essay he wrote that AI models had gone "from barely being able to complete a single line of code, to writing all or almost all of the code for some people, including engineers at Anthropic."[39] Andrew Ng put the thesis most directly: "The bottleneck is deciding what do we actually want to build."[40]

That is scarcity moving from engineering to judgement, and it agrees with the distribution thesis. But it comes with a cost, and the cost shows up in the production data. The only randomised trial I found, METR's July 2025 study, found experienced developers took 19% longer with early-2025 AI tools, with a confidence interval of +2% to +39%, though they had forecast a 24% speed-up.[41] Vendor datasets point the same direction, with caveats: Faros's 2026 report found throughput up 33.7% per developer and incidents per merged pull request up 242.7%.[42] Writing code got cheaper; owning it did not.

And the cheap-to-build story has limits at the macro level. Acemoglu's 2024 calculation bounds the total factor productivity gain from AI at no more than 0.66% over ten years.[43] MIT's NANDA report found that in a small preliminary sample, 95% of organisations reported no measurable P&L return, with 5% of enterprise-grade tools reaching production.[44] I hold both loosely. They are reasons to expect slower diffusion, which, if anything, favours the incumbents with distribution.

The ledger

Put the six candidates side by side and score them by how well the evidence holds.

Compute and power is the strongest. Measured, large, and getting larger. It is also the one a small builder cannot own.

Systems of record is strong. ServiceNow, Workday, Adobe and Figma are all growing, and Menlo says buyers purchase 76% of AI solutions rather than building them. But startups took 63% of new AI application revenue.

Trust and brand is plausible and consistent with the rankings, but the

[38] Council on Foreign Relations CEO Speaker Series, 10 Mar 2025, transcript.

[39] Dario Amodei, "The Adolescence of Technology", January 2026. This is where the March 2025 claim ended up: true for some people at one company, not shown industry-wide.

[40] Andrew Ng on the No Priors podcast, via Business Insider, 22 Aug 2025. I did not hear the podcast.

[41] METR (Becker, Rush, Barnes, Rein), 10 Jul 2025. Small sample, one setting, early-2025 tools. In February 2026 METR said a follow-up gave "an unreliable signal" because developers refused to work without AI, and that developers were likely more sped up now. I do not cite the follow-up as proof of a speed-up.

[42] Faros AI, AI Engineering Report 2026, June 2026. Vendor telemetry, correlational; the vendor itself notes the incident metric is a ratio. DORA 2024 and 2025, Veracode 2025, GitClear and CodeRabbit are in the same vein. Only METR is randomised.

[43] Acemoglu, "The Simple Macroeconomics of AI", NBER w32487, May 2024. A model bound from pre-agent evidence, not an observed outcome.

[44] MIT Project NANDA, "The GenAI Divide", July 2025. Self-reported, 52 interviews and 153 surveys. Not a measure of the share of AI pilots that fail, and the report lists its own selection-bias limits.

valuations are model outputs and the main argument is an investor's opinion.

Regulation is real but passable. 1,614 AI devices have cleared the FDA, and Europe moved its own deadline back.

Data is weaker than the story. Reddit's licensing share is about 5%, and model collapse is avoidable if you keep real data.

Community is a bet. I believe it, and I cannot measure it.

If you read the list from top to bottom, there is a pattern. The strongest moats are physical or institutional: chips, power, contracts, installed bases. The moats a person can build alone are softer: brand, trust, community. That is the situation the next chapter has to face, because it is the one the thesis points to. If code is no longer scarce and the hard moats belong to a few large companies, then the open ground is distribution and trust, and the real question is whether a builder can win there against the incumbents who already own it.

IX. The Counter-Case and What Would Falsify This

The strongest objections to this paper are not stupid, and several come from the people who sell the thing I am writing about. I state them at full strength, concede what they win, and then list the data that would show I am wrong.

> *A moat that must be continuously rebuilt will eventually be no moat at all.*
>
> Warren Buffett, Berkshire Hathaway shareholder letter, 2007

I have spent the preceding chapters arguing one thing. AI made building cheap. The supply of software went up. The technical barrier fell. So the scarce thing is being seen and trusted, which is distribution. This chapter tries to break that argument.

I did not pick weak opponents. The evidence file I assembled for this chapter has eight objections, ranked by how much damage each would do if it were right.[1] I have also gone back to the "weak spots" sections of the four other dossiers (coding, supply, moat and price, distribution), because those sections are where the data argue with me.

Two rules. First, I quote chief executives as claims, never as evidence. Several of them sell the products in question. Second, where an objection wins, I say so and change what I claim.

[1] Objections and quote bank: dossier F, "What survives when building is cheap, and the counter-views", accessed 7 October 2026. Every source named in this chapter was accessed that day unless a date is given.

The steelman in one page

The opposing case has five legs. I think it is a fair summary of what a smart sceptic would say.

Falling price is not a falling barrier. Code was never what incumbents sold. Customers pay for risk reduction: the payroll system that has run for fifteen years, the vendor whose security review is done, the integration that already talks to billing. Jason Lemkin, who vibe-coded his own tools, still says he buys about 90% of what he needs.[2] Sam Altman said in August 2026 that people "keep buying from the same company" and that changing behaviour is "much harder than the tech nerds realize".[3]

[2] Lemkin, SaaStr.AI blog, about May 2026: "If you can buy it - buy it. The 90/10 rule." He adds that the exposed products are point solutions, not "deep platforms with real data moats, real integrations, real workflow lock-in".

[3] Altman, David Senra interview posted 23 August 2026, quoted by Business Insider 25 August 2026. I read the quotes via cached text of the paywalled article, not the interview.

Cheap code expands demand. Satya Nadella invoked Jevons the week DeepSeek landed. Altman: "lower prices lead to much more use." When a cost collapses, profit migrates to the adjacent layer that is not yet modular. Right now that layer is compute, power and frontier models.

The productivity evidence is thin. The one randomised trial found experienced developers slower. Company-level studies find more output and more breakage.

The assets I say will matter are not obviously durable. Data licensing is small. Regulatory gates are porous.

Continuity, not regime change. The durable advantages are the old ones: brand, network effects, scale, switching costs. Buffett's test cuts against most new moats.

Now the eight objections, one at a time.

Objection 1: building is not cheap where it counts

The claim. The scarce input did not vanish. It moved to silicon, power and balance sheets. If so, "building is cheap" is false for anyone who needs the best models, and the money is flowing to compute, not to distribution.

The evidence for it is strong. Cash capital spending by Amazon, Alphabet, Meta and Microsoft was about $376 billion in calendar 2025 and about $295 billion in the first half of 2026 alone (Figure IX.1).[4] Nvidia's data-centre revenue went from $3.6 billion in a quarter (Q4 fiscal 2023) to $89.0 billion (Q2 fiscal 2027) at a 75.0% gross margin.[5] The IEA's base case has data-centre electricity use rising from about 415 TWh in 2024 to about 945 TWh in 2030, a projection and not a measurement.[6] Alphabet raised $49.6 billion of equity in June 2026, partly for capital spending.[7] On the price side, flagship list prices have not fallen: the newest top tiers list at $10 input and $50 output per million tokens, close to GPT-4's $30 and $60 in 2023, and Google's Flash tier rose about 15x in input price in 15 months.[8] One measured data point on usage: the median OpenAI researcher's coding-agent use reached about $601 a day at API list prices.[9]

[4] SEC EDGAR XBRL filings (10-K and 10-Q cash purchases of property and equipment); the sums are my arithmetic. Microsoft calendar quarters are derived by subtraction from cumulative figures. Cash capex excludes finance-lease additions, so company-reported totals run higher.

[5] Nvidia earnings releases, Form 8-K Exhibit 99.1, 15 quarters. Company-reported.

[6] IEA, *Energy and AI*, 2025.

[7] Alphabet Q2 2026 earnings 8-K: "aggregate net proceeds of $49.6 billion, to be used for general corporate purposes, including capital expenditures".

[8] Dossier C, rows C-020 to C-024, list-price snapshots of 7 October 2026. The dossier could not retrieve launch prices for several older models from first-party pages.

[9] Dossier C, row C-032; it is a list-price equivalent, not what OpenAI pays.

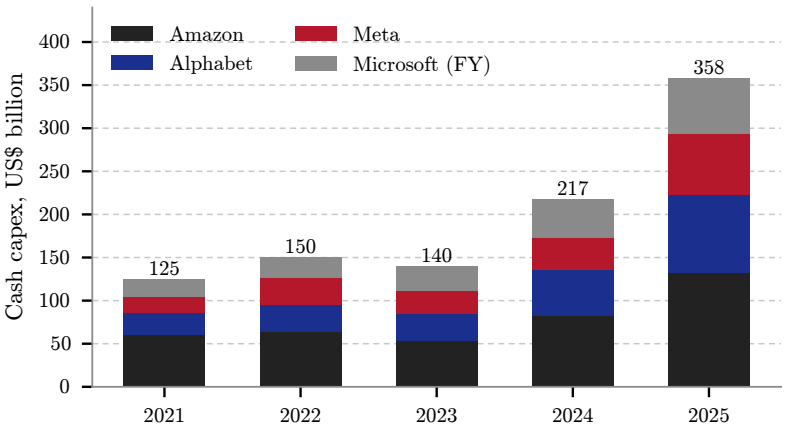


Figure IX.1. Four-company cash capex, 2021 to 2025, US$ billion. Amazon, Alphabet and Meta are calendar years; Microsoft is its fiscal year ending 30 June, so each bar mixes bases. Source: SEC EDGAR XBRL, 10-K filings (measured).

My answer. I concede the premise partly, and I narrow the thesis. What fell is the price of a *fixed* level of capability. Epoch AI's careful cost-per-task study of September 2026 puts that at about 13x a year across five benchmarks; the range in the literature is wide (5–10x a year in one academic estimate, a median of 50x a year per token in an earlier Epoch study), so treat it as a range.[10] Coding is slower: Epoch's SWE-bench Verified series falls only about 3.6x a year (computed in the dossier). The cheap input is the *last* model, not the best one.

That is enough for my argument, because my argument is about the supply of ordinary software, and ordinary software does not need the frontier. But the objection stands on one point I must not dodge: capex is where the *profit* may be going. If a few labs and chip makers earn the rents, an app builder is a price-taker whatever else is true. I come back to this in Objection 6.

[10] Epoch AI, "The plunging price of thought", 22 September 2026; Epoch data insight, March 2025; Gundlach et al., March 2026, read only via Epoch's summary. Per-token and per-task definitions differ; do not average them.

Objection 2: the measured productivity gains are small or negative

The claim. If building were cheap in any way a firm could measure, we would see it in controlled studies. We do not.

The evidence. The randomised trial by METR in early 2025 found 16 experienced open-source developers took 19% *longer* with AI tools (confidence interval +2% to +39%), while they had expected a 24% speed-up and afterwards still believed 20% (Figure IX.2).[11] In February 2026 METR said the follow-up gave "an unreliable signal" because developers refused to work without AI; the raw estimates were

[11] METR (Becker, Rush, Barnes, Rein), 10 July 2025; 246 tasks on their own mature repositories, early-2025 tools. The only randomised design in this chapter.

18% less time for returning developers (CI −38% to +9%) and 4% less for new ones (CI −15% to +9%), and METR thought developers were "likely" more sped up by then.[12] Acemoglu's task-based model bounds the ten-year gain in total factor productivity at no more than 0.66%.[13] MIT's NANDA report found 5% of enterprise-grade tools had reached production.[14]

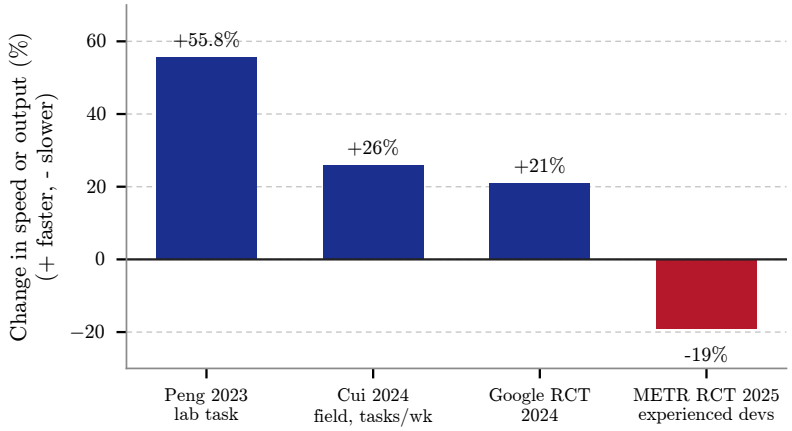


Figure IX.2. Effect of AI coding tools on speed or output in four studies. The metrics differ: Peng is time on one task, Cui is tasks completed per week, Google is one complex task, METR is task time with sign flipped. Only METR is on experienced developers in their own mature repositories. Source: Peng et al. 2023, Cui et al. 2024, Paradis et al. 2024, METR 2025 (measured).

My answer. Three points. First, the studies disagree, and they measure different things: +55.8% on a lab task, +26% more tasks a week in field experiments, about 21% faster at Google, 19% slower for METR's experts. Second, the METR result is narrow: experts, their own repositories, early-2025 tools. Third, and most important, *my thesis does not need a per-developer productivity number.* It needs more software to exist and the barrier to entry to fall. Those are counted directly, in the supply chapter, not inferred from time trials.

What I concede: I cannot claim that enterprises have seen their *cost of delivering reliable software* fall. The dossier searched for a cost-per-feature or cost-per-line series and found no credible one.[15] I have not claimed one.

There is a measurement worth keeping from the larger study of more than 100,000 developers. Commits rose roughly 3x, projects about 50% and shipped releases only about 30%.[16] Writing code got faster. Shipping did not, by the same factor. That gap is a reason distribution and judgement matter, not a reason they do not.

[12] METR, 24 February 2026. Do not cite as proof of a speed-up.

[13] Acemoglu, NBER Working Paper 32487, May 2024. A calculation from pre-agent evidence, not an observed outcome.

[14] MIT Project NANDA, *The GenAI Divide*, July 2025: 60% evaluated, 20% piloted, 5% in production. "Preliminary findings", 52 interviews and 153 survey responses, self-reported; a critic says it "overstates" what was established. It is not a measure of how many AI pilots fail.

[15] Dossier C, gaps: "Do not claim a 'cost per line of code' number."

[16] Demirer et al., 2026, via CEPR VoxEU, 21 June 2026. Matched event study; effect sizes are the CEPR column's rounded figures.

Objection 3: cheap to write is not cheap to own

The claim. If AI code is buggy, insecure and hard to maintain, then technical depth is still the moat and the thesis is partly reversed.

The evidence. Veracode found 45% of AI-written code samples failed security tests (Java 72%).[17] DORA 2024 found that each 25% rise in AI adoption was associated with 7.2% lower delivery stability and 1.5% lower throughput; DORA 2025 said adoption now improves throughput but "still increases delivery instability".[18] Faros, using telemetry from 22,000 developers, reported epics per developer up 66% and incidents per merged pull request up 242.7%.[19] GitClear found cloned lines rose from 8.3% to 12.3% between 2021 and 2024. Two incidents are documented: a Replit agent deleting a production database in July 2025 and a Lovable-generated project exposing data (CVE-2025-48757).

[17] Veracode 2025 GenAI Code Security Report: over 100 models, curated tasks scored by static analysis, not production repositories. Veracode sells security tools.

[18] Google Cloud DORA reports, 2024 and 2025. Survey-based associations, not causal effects.

[19] Faros AI, *AI Engineering Report 2026*, June 2026. Vendor telemetry comparing each organisation's low- and high-adoption periods. Faros itself notes churn may include refactoring, and incidents are a ratio.

My answer. I will be blunt about the weight. Veracode, GitClear, CodeRabbit and Faros are vendors whose products benefit from a world where AI code is risky. DORA is correlational. The incidents are existence proofs, not rates; the dossier warns against extrapolating a rate from Replit.[20]

[20] The Tea app breach of July 2025 is *not* used here: no source I fetched says AI wrote that code.

But I accept the direction. Owning software is costly, and AI does not remove that. This changes my claim in a specific way. The barrier to *producing something that runs* fell. The barrier to *operating something trusted at scale* did not fall as much. That is a point *for* my thesis in one respect: the difference between the cheap prototype and the trusted product is exactly what buyers cannot see from outside, so they rely on brand and reputation to tell them. I would say it is the strongest honest link between this objection and the trust chapter, but it is an argument and not a measurement.

Objection 4: cheaper software expands the market

The claim. Jevons: when a cost falls, use rises, and the profit pool grows. Enterprise generative-AI spending went from $1.7 billion in 2023 to $11.5 billion in 2024 and $37 billion in 2025.[21]

[21] Menlo Ventures, 9 December 2025, a survey of about 500 US enterprise decision-makers plus a bottom-up revenue model. Menlo is an AI investor; the figures are modelled, not audited.

My answer. This is the objection I most want to engage, because the numbers cut both ways. Spending on AI tools is not demand for the *apps* the tools make. The demand test is downloads, revenue and attention, and those are flat. New iOS apps rose about 30% in 2025 and about 80% in Q1 2026 on one vendor's count; iOS downloads rose 3% in 2025 and 2% in the first half of 2026; global downloads fell 2.7% in 2025.[22] Total engagement with each new monthly cohort of apps, in its first three months, is "flat or declining on every marketplace", according to the study behind the CEPR column. The authors offer

[22] Sensor Tower and Appfigures via TechCrunch (14 January and 18 April 2026) and 9to5Mac/NYT (20 July 2026). Vendor estimates with different scopes: iOS versus iOS plus Android, new apps versus downloads, US versus global.

two readings, lower quality or an attention bottleneck, and the data cannot tell them apart.

That is the Jevons objection answered on the evidence available today: supply is up, spending on tools is up, but the demand side that a *new* app would earn from is not. It is also the data point most consistent with my thesis, and so the one I would most like replicated by someone who does not agree with me.

The magnitude correction I owe the reader

The version of this argument I started with said app supply was up by "orders of magnitude", roughly 100x. The data do not support that, and the paper does not claim it.

The dossier computed the multiple for thirteen supply series against a 2022 baseline (Figure IX.3). None of the market-level series exceeds about 7x at its monthly peak (Show HN, February 2026, 7.4x of the 2022 mean). Full-period multiples run from about 1.1x to 3x. Some of the sharpest rises sit on top of trends that were already running: new public GitHub repositories were growing 13–14% a year in 2023 and 2024, and Steam game releases grew 49x between 2013 and 2025 without AI. The AI effect is the *increment above trend*, roughly +20% to +75% a year in 2025–26, and the repositories with ten or more stars are flat at 0.99x.[23]

Three more checks lower the claim. Apple rejected more than 2 million submissions in 2025, including over 1.2 million new apps, so a gatekeeper is still filtering.[24] Google Play's total catalogue has *shrunk* from about 3.4 million to about 1.8 million on one count (Appfigures) and from 2.44 million to 1.54 million on another (AppBrain). And much of the growth in code is by existing developers: Octoverse reports that 80% of new GitHub developers use Copilot in their first week and TypeScript overtook Python and JavaScript. That is traditional builders using AI, not non-technical people displacing them. The one dataset on "non-technical" users is Lovable's own survey, a vendor claim.

[23] Dossier B, "Counter-evidence". Sources: GitHub Search API counts, PyPI snapshots, Appfigures, Steam via SteamDB, Census Business Formation Statistics, Hacker News. Series mix stock and flow and vendor and API counts, so the multiples are best estimates, not a single audited number.

[24] Dossier B, row B-011, from Apple's own statement.

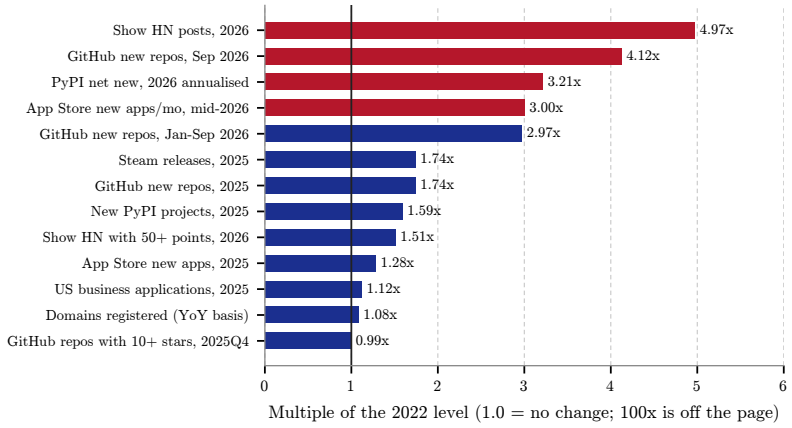


Figure IX.3. Growth in software supply, latest period versus the 2022 level, thirteen metrics (red: 3x or more). The axis stops at 6x; a 100x claim would sit far off the page. Mixed stock and flow series and mixed sources, so best estimates, not one common unit. Source: dossier B computations from GitHub, PyPI, Appfigures, SteamDB, Census BFS, Hacker News (estimate).

So the honest size is this: supply is up roughly two to five times on the best-measured metrics, concentrated at the low-traction end. That is a large change. It is not a hundred-fold one, and I no longer write as if it were.

Objection 5: incumbents are holding

The claim. The strong form of "software is dead" is contradicted by the filings.

The evidence. In the latest quarters: ServiceNow subscription revenue +24.5%, Figma revenue +48% with net dollar retention of 136%, Shopify +34%, Adobe +13%, Workday +12.8% (Figure IX.4).[25] Microsoft 365 Copilot paid seats went from 15 million in January 2026 to "over 30 million" in July. Public cloud software kept a median net revenue retention of 108–110% and median forward growth of 11–13% through the sell-off.[26] The software ETF IGV fell 36.6% from 22 September 2025 to 10 April 2026, then had its best month since October 2001 in May (+21%) and was 6.9% below its peak on 7 October 2026. Median enterprise value to forward revenue went 6.7x, 2.9x, then 4.2x. What fell was multiples, not revenue.

[25] Company earnings releases, Form 8-K Exhibit 99.1, SEC EDGAR; metrics are company-defined and not like-for-like. Salesforce reports about $3.9 billion of Agentforce and Data 360 ARR, up 210%.

[26] Clouded Judgement (Jamin Ball, Altimeter), weekly, December 2024 to October 2026; a roughly 100-company comparison set.

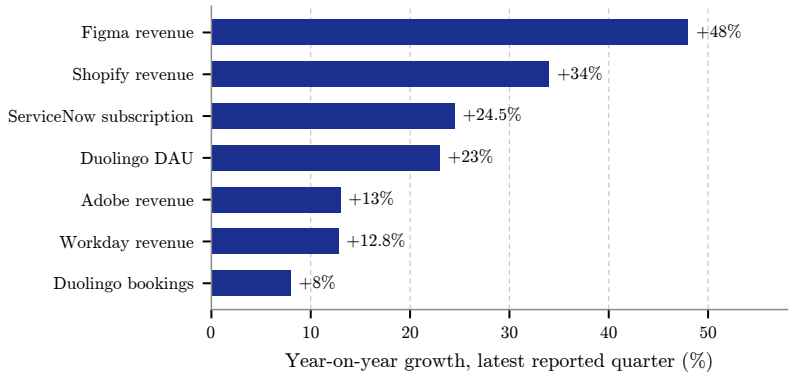


Figure IX.4. Year-on-year growth of selected incumbent software firms, latest reported quarter (Q2 2026 for ServiceNow, Figma, Shopify and Duolingo; Adobe Q3 fiscal 2026; Workday Q2 fiscal 2027). Metrics are company-defined and not like-for-like. Source: company 8-K earnings releases, SEC EDGAR (measured, company-reported).

My answer. I agree with all of it, and I think it is evidence *for* my thesis, not against it. Incumbents that hold distribution (installed seats, enterprise contracts, an audience of a billion monthly users) keep growing while a flood of cheap competitors arrives. That is what "distribution is the scarce thing" predicts. Two cautions on my own side. Survival of revenue is not survival of valuation, and software has only matched the market since ChatGPT: IGV was at 213.9 against 212.2 for the S&P 500 ETF (December 2022 = 100). And the bundling story has a survivorship problem: the dossier verified no failed-bundle case, so I cannot say how often a distribution advantage fails to convert.[27]

A different kind of thing moves the other way. Menlo's survey says startups took 63% of application-layer AI revenue in 2025, up from 36%.[28] If true, new entrants did win new dollars. I take that as showing that distribution can be *built* fast by some new firms, mostly through creators and word of mouth, which is a craft and not an incumbent right.

[27] Dossier D, caveats. MAU claims by Google, OpenAI and Meta use different definitions and should not be compared.

[28] Menlo Ventures, December 2025; survey and model, investor-authored.

Objection 6: the value is accruing to the model layer

The claim. Anthropic reports a run-rate above $47 billion and a $965 billion post-money valuation; OpenAI $852 billion. Menlo has Anthropic at 40% of enterprise API share against OpenAI's 27%.[29] If a handful of labs earn the rents, app-layer "moats" are decoration, and Christensen's conservation of attractive profits predicts exactly this.

My answer. Concede, with a precise scope. My thesis is about the

[29] Anthropic Series H release, 28 May 2026 (company statement); OpenAI figure from a secondary source only; Menlo Ventures, 2025. Run-rates come from interested parties.

application layer: for a given product, the scarce input is not code. It does not say the whole stack is flat. Where profit lands in the stack is a separate question, and the dossier gives reasons to think the model layer is not a clean moat either: open-weight models lag the frontier by about three to four months, a short lead but a lead, and the dossier records a rise in copycat speed (clones of the first coding agent appeared as GitHub repositories 50 to 67 days after it).[30] Margins also deserve care. Cursor ran negative gross margins until recently, which means part of its revenue growth was pass-through of model costs.[31] Run-rate is not profit.

What I take from this objection is a warning to builders. An app maker who relies on one model vendor is a price-taker, and I should say so more clearly than I did in the early chapters.

[30] Dossier C, rows C-028, C-029, C-099. Repository creation date is not launch date; stars are a vanity metric.

[31] TechCrunch reporting from sources, cited in dossier A; unconfirmed by the company.

Objection 7: data and regulation are not the moats people say

The claim. Reddit's "other revenue", mainly content licensing, was $140.0 million in 2025 against $2,202.5 million in total revenue, and about 5% of the Q2 2026 total.[32] Model collapse can be avoided if real data is kept (Gerstgrasser et al.). A 2019 a16z essay called data moats an "empty promise". FDA AI-device clearances ran at 335 in 2025, about 96% of the list through the lower-bar 510(k) pathway; the EU moved its high-risk deadline from August 2026 to December 2027.[33]

My answer. I agree, and the thesis does not rest on data or regulation. They are listed in the dossier as alternative moats. The evidence is thin and I have left them out of the argument. The line I can defend is narrower: data is priced in the hundreds of millions to low billions, courts have put a price on the piracy of it (the $1.5 billion Bartz v. Anthropic settlement, about $3,000 a work), but that price is small beside the labs' capex.[34] I also did not find any fetched evidence that counts SOC 2, HIPAA, FedRAMP or banking-licence barriers, so I say nothing numeric about them.

[32] Reddit Form 10-K FY2025 and Q2 2026 8-K. "Other revenue" is not purely AI licensing.

[33] FDA AI-Enabled Medical Devices list, 1,614 entries, my tabulation of the CSV; Gibson Dunn client alert, 27 May 2026, and arcusx.ai, 17 July 2026 (the final adoption is single-sourced).

[34] The Bartz figure prices past piracy, not a market rate for training licences.

Objection 8: the advocates' forecasts slipped, and they talk their books

The claim. Dario Amodei said in March 2025 that AI would be writing "90 percent of the code" in three to six months. By January 2026 the claim was that models write "all or almost all of the code for some people", including engineers at Anthropic.[35] Altman in August 2026: "we've all been too ambitious on timelines." Jensen Huang and a16z say software will do well; Huang sells the chips agents run on.

[35] Amodei, CFR speaker series, 10 March 2025; "The Adolescence of Technology", January 2026. Both have full text.

My answer. Agreed. I quote them as claims. Company self-reports on the share of code written by AI are vendor claims of low confidence: Meta's 55% was a target, not an outcome, and Snap's 65% is in a secondary source. The same discipline applies to revenue: Cursor, Lovable, Gamma and Base44 figures are founder or investor claims, not audited numbers. Benchmarks are no safer: OpenAI retired SWE-bench Verified as contaminated, and METR's horizons above 16 hours are unreliable.

One thing this objection does not touch. The key supply and demand numbers in this paper come from GitHub's API, app-store analytics, Census forms and SEC filings, not from CEO statements.

What would show I am wrong

A claim that cannot lose is not worth reading. These are the findings that would make me withdraw or sharply weaken the thesis. The thresholds are my judgement; the current readings are from the dossiers and are as measured as their sources allow.

F1. Demand rises with supply. If downloads, app revenue and time spent in new-app cohorts grow roughly in step with new-app counts for two consecutive years, attention is not the bottleneck. *Now:* new iOS apps up 30% to about 100%; downloads +2% to +3%; new-cohort engagement flat or declining.

F2. New entrants stop being buried. If the share of newly launched apps reaching meaningful revenue rises sharply without paid or owned distribution, discovery is no longer the filter. *Now:* 17.3% of new subscription apps reach $1,000 a month within two years and 4.6% reach $10,000; the top 5% earn about 400x the bottom 25% after a year.[36]

F3. Cost of reach falls. If advertising cost per unit of reach falls for two years while budgets stay flat, distribution is getting easier. *Now:* Meta's price per ad is about 3% below its end-2018 level (index 97.1) after rising 10% and 9% in 2024 and 2025; impressions are about 3.7x; Google's cost per click index is 111.4. Gartner has marketing at 7.7–7.8% of large-company revenue in 2024–2026, down from 9.1% in 2023.[37]

F4. Incumbent share erodes on the metrics that matter. If median net revenue retention of public software falls below 100% and growth below 5% for four quarters, "distribution protects incumbents" fails. *Now:* 108–110% and 11–13%.

F5. Verified non-technical builders displace vendors at scale. If an independent source counts a large share of firms replacing paid vendors with software built by non-engineers, the barrier fell further

[36] RevenueCat, *State of Subscription Apps* 2025 and 2026; subscription apps only, vendor data (medium confidence).

[37] Meta and Alphabet 10-K filings, derived indices in dossier D; Gartner CMO Spend Surveys 2024–2026, large firms only. Meta's metric is revenue per ad delivered, a mix measure and not an auction price. I do not claim that Meta's price per ad has risen since 2018.

than I claim. *Now:* only 19% of developers surveyed build software to replace paid vendors, and the purchased share of enterprise AI solutions rose from 53% to 76%.[38]

[38] Stack Overflow Developer Survey 2026; Menlo Ventures 2025.

F6. Developer demand collapses. If BLS employment for software developers falls clearly and the Indeed postings index does not recover, the cost of building has fallen enough to remove jobs, a stronger result than mine. *Now:* BLS 1.80 million (2022 edition), 1.90 million in 2023, 2024 and 2025: flat; postings index 61.1 at the May 2025 trough and 77.3 in September 2026 against 233.8 at the 2022 peak; early-career (22–25) developers down nearly 20% from their late-2022 peak by September 2025.[39]

[39] BLS Occupational Outlook Handbook editions (via Wayback); FRED series IHLIDXUSTPSOFTDEVE (Indeed Hiring Lab) and CES6054150001; Brynjolfsson, Chandar and Chen, Stanford Digital Economy Lab (descriptive, not causal). The BLS ten-year growth projection has been cut from 25% to 10%.

F7. Productivity evidence turns clearly one way. A randomised trial on current tools showing no reduction in delivery time for ordinary teams would not touch my supply counts, but it would remove my assumption that building got cheaper for them. A well-run trial showing a large saving would weaken Objection 2. *Now:* one randomised trial (METR, slower), one unusable follow-up.

F8. Trust in feeds and ads stops falling. The attention and trust chapters rest on saturated attention and falling trust. If repeated surveys show trust in ads and feeds recovering, that leg falls. *Now:* the dossier has a mix of surveys and I treat them as directional.

Two of these can be checked within a year using public data: F1 from app-store analytics and F3 from the next annual filings of Meta and Alphabet. I plan to rerun both.

What I now believe, after the objections

I keep the core. Supply of software is up, measured by counts nobody paid me to produce. Demand on the receiving end is not up with it. Incumbents with distribution are growing, and cheap entrants face a hard discovery problem.

I change five things. The supply increase is roughly two to five times, not a hundred-fold. Aggregate developer employment is flat, not falling. Frontier list prices are not falling; the price at constant capability is, at roughly 13x a year on the best study. Meta's price per ad is not up since 2018. The "scrolls per day" arithmetic in the attention chapter is an illustrative model only.

And I add one concession I did not plan to make. The barrier to *producing* software fell. The barrier to *operating a trusted product* fell by less. Where the profit lands, in distribution or in the model layer, is the question the next ten years will answer.

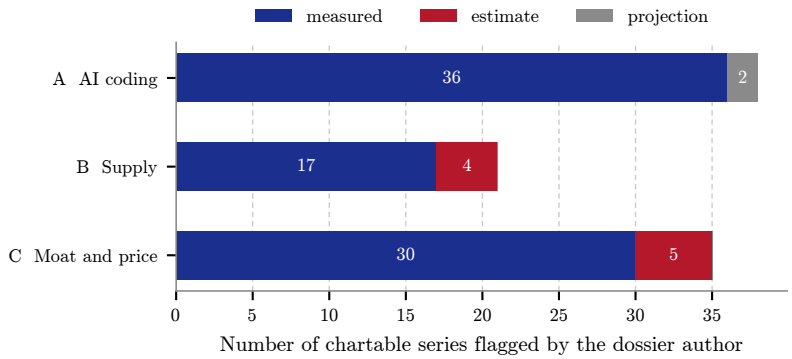


Figure IX.5. Evidence grades of the chartable series in dossiers A (AI coding), B (supply) and C (moat and price), as flagged by their authors. Dossiers D, E and F carry no machine-readable grade; their sources are graded row by row in the text. Source: research/parts A_series.json, B_series.json, C_series.json (measured).

Sources and Methods

How the evidence was gathered

Every number in this paper comes from six research dossiers, labelled A to F, last accessed on 7 October 2026.[40] I did not use numbers from memory. Many sources were reached through workarounds: reader proxies, Wayback copies and API queries, because publishers (Forbes, AP, Reuters, fda.gov, bls.gov, X) blocked direct access. Where a figure is only in a secondary source, the sidenote says so.

[40] Dossiers: A, AI coding; B, software supply; C, the moat and price collapse; D, distribution; E, attention and trust; F, what survives and the counter-views. Each row carries a source, URL, design and confidence grade.

Four kinds of number

Measured. Counted or reported by the party that holds the data: SEC filings (capex, Nvidia, incumbents, Reddit, Meta and Alphabet advertising), BLS and FRED series, GitHub and PyPI API counts, the FDA's own device list (yearly counts are my tabulation), Census business applications. Of 94 chartable series flagged in dossiers A to C, 83 are graded measured (Figure IX.5). That grade means the data are real, not that the interpretation is settled. A measured count of repositories says nothing by itself about quality.

Estimated. Computed by me or a third party from measured inputs, with stated assumptions: the supply multiples in Figure IX.3, Meta's derived price index, calendar-year capex sums, adoption percentages from surveys. Eleven series in A to C are graded estimate or projection.

Modelled. Outputs of a stated model: Acemoglu's productivity bound, the IEA data-centre projection, Menlo's market model, Kantar and Interbrand brand values, Gartner's forecast of assistant use by 2028. The "scrolls per day" arithmetic in the attention chapter is an illustrative model, labelled so. These are inputs to my reasoning, not observations.

Vendor-reported or interested-party. Company-published ARR and run-rates (Cursor, Lovable, Replit, Gamma, Anthropic), valuations from private rounds, vendor security and quality studies (Veracode, GitClear, CodeRabbit, Faros), vendor surveys (Lovable, Instantly, RevenueCat), and statements by chief executives and investors. Dossier A alone grades 26 rows as vendor claims. I report them as claims.

Only the METR trial is randomised. Everything else on productivity is observational, survey-based or vendor telemetry.

What I excluded

- The "300 scrolls a day, 9,000 a month" figure, "300 feet" and "1.7 miles" of scrolling: no primary source found. The arithmetic (300 $\times$ 30 = 9,000) is trivially true; the premise is not sourced.
- The Tea app breach as a vibe-coding failure: no source says AI wrote the code.
- Cursor's $4 billion ARR headline, Codex user counts, OpenAI's SWE-bench Pro "30% broken" claim, Veracode 2026 and VentureBeat "380,000 apps exposed": headline-only, marked unverified.
- Several distribution figures: mobile app cost per install by year, Meta CPM history in dollars, Similarweb zero-click shares, ecommerce CAC trends, "65% of downloads follow a search": gated or not found.
- Stratechery's "more than $700 billion" capex line, which I could not reconcile with company guidance.

Verification gaps that matter

- *No census of products that survive.* Lovable project counts and GitHub repositories are creation counts. No source measures AI-built apps that reach production or earn revenue. This is the largest gap, and F2 depends on closing it.
- *No cost-per-feature series.* I could not find one, so no claim about the cost of delivering software is made.
- *Primary papers not read.* The Demirer et al. study (via CEPR and MIT Sloan summaries), Gundlach et al., the Anthropic Institute

report behind the 80–90% code-share claim, and Shumailov et al.'s methods (abstract only).

- *Secondary-only prices and counts.* OpenAI's $852 billion valuation, Windsurf's $2.4 billion licence, Abridge's valuation, Harvey's revenue above $400 million, TSMC's Q2 2026 figures, Alphabet's 2026 capex guidance (from a transcript site), a16z's 40,000 to 120,000 monthly iOS apps.
- *Unconfirmed dates.* Altman's *Three Observations* and *Gentle Singularity*, the Davos interview, Kantar's 2026 release, a September 2026 burst in new GitHub repositories whose cause is unknown.
- *Definitions drift.* Capex with and without finance leases; Microsoft's June fiscal year; Nvidia's reclassified customer groups; MAU against weekly users; iOS against iOS plus Android; Sensor Tower against Appfigures counts (557,000 against about 600,000 for 2025).
- *Not located.* Quotes from Elad Gil, Sarah Guo, Chamath Palihapitiya and David Sacks, and Helmer in his own words.

I would rather publish a thinner paper whose every number can be traced than a fuller one that cannot. If you find a figure here that you cannot trace to a row in a dossier, tell me; I will fix it or remove it.