AI eats the world

Benedict Evans

November 2025

www.ben-evans.com

The next platform shift

Every 10-15 years, a platform shift reshapes technology

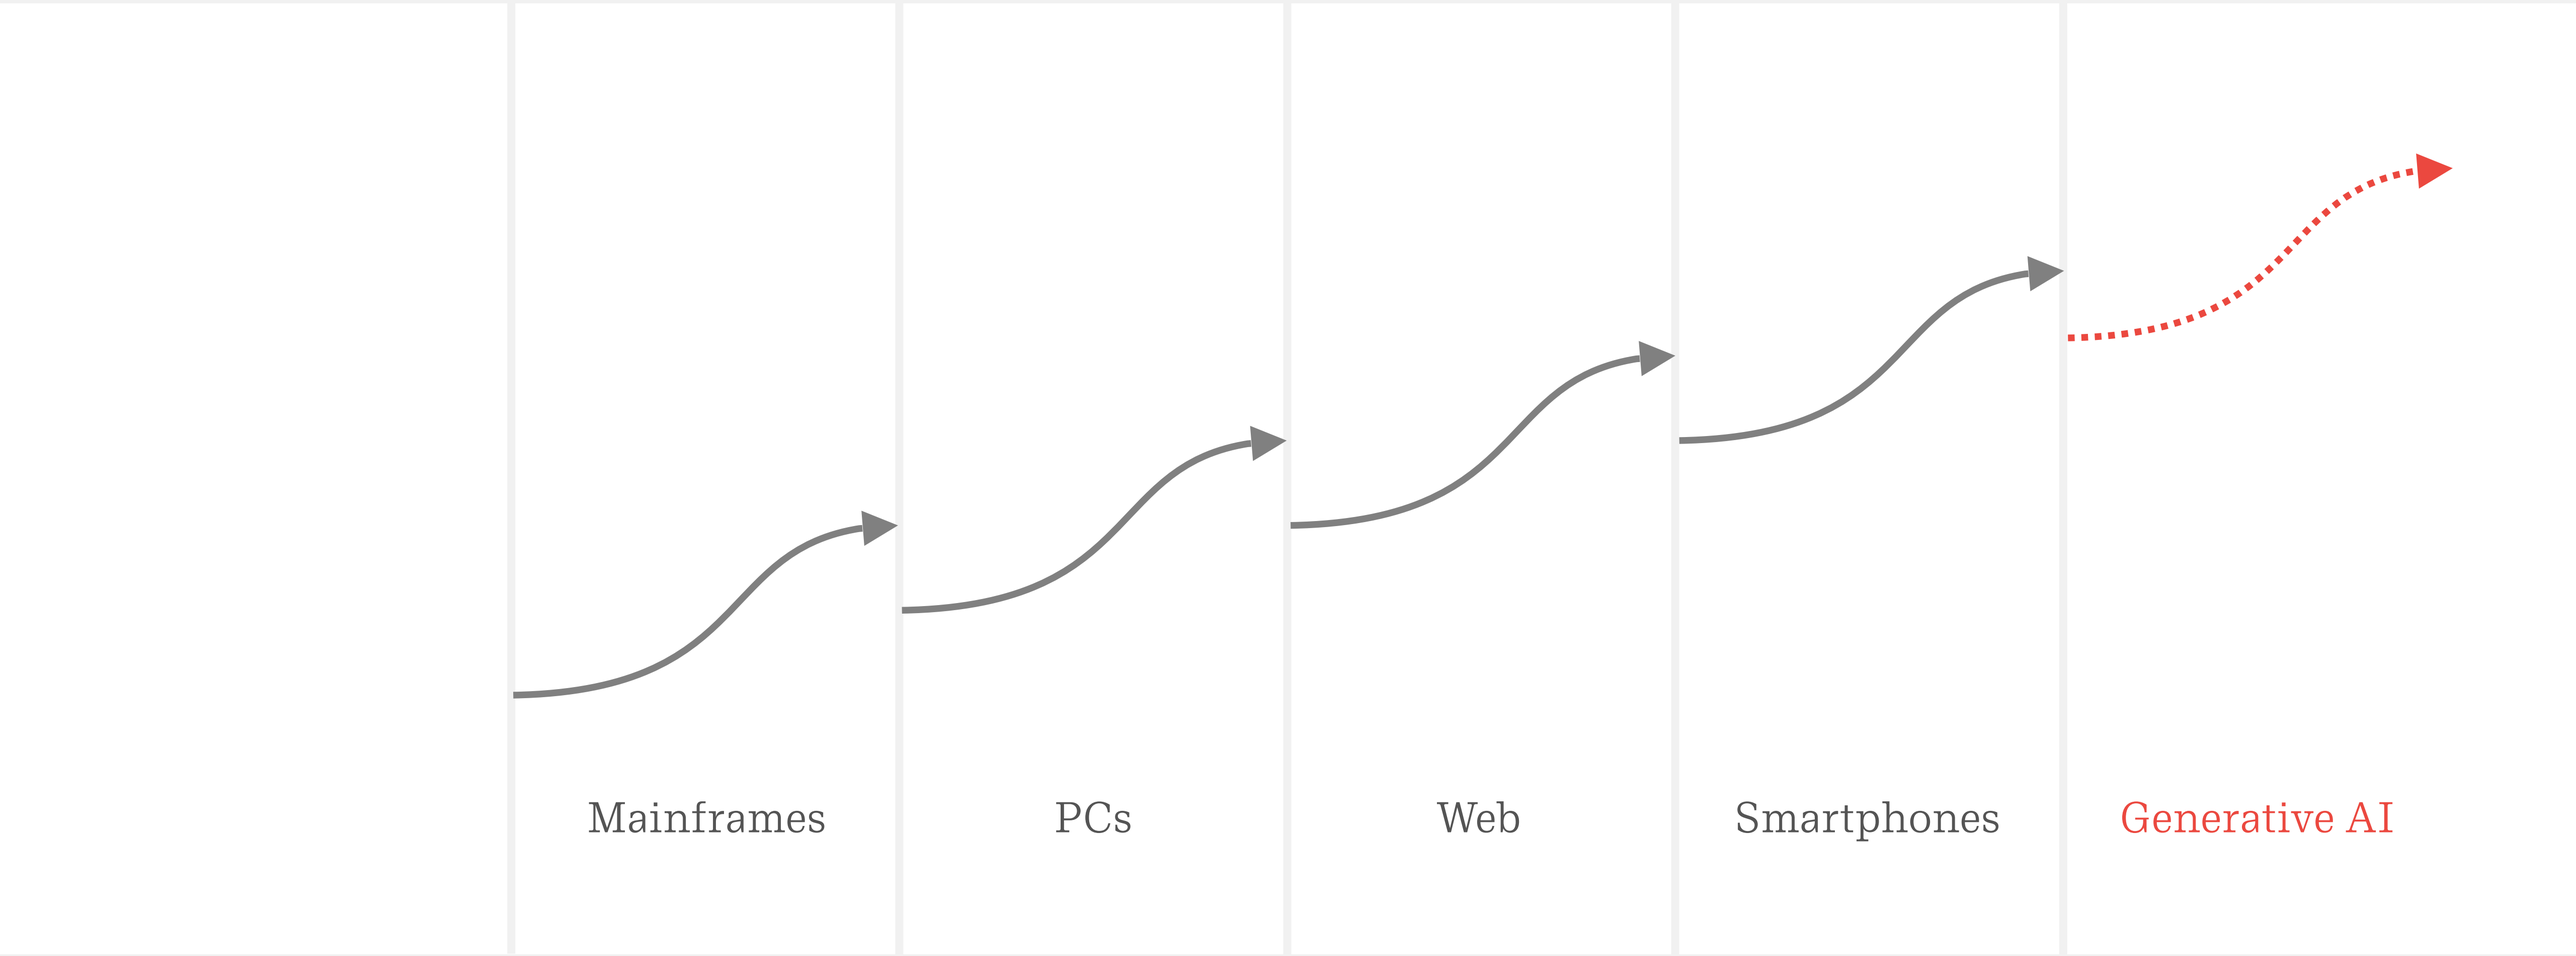

What happens in a platform shift?

Who is affected, and how much?

Dominance is won and lost

Microsoft dominated the PC era, but when the centre of gravity shifted to smartphones it became irrelevant

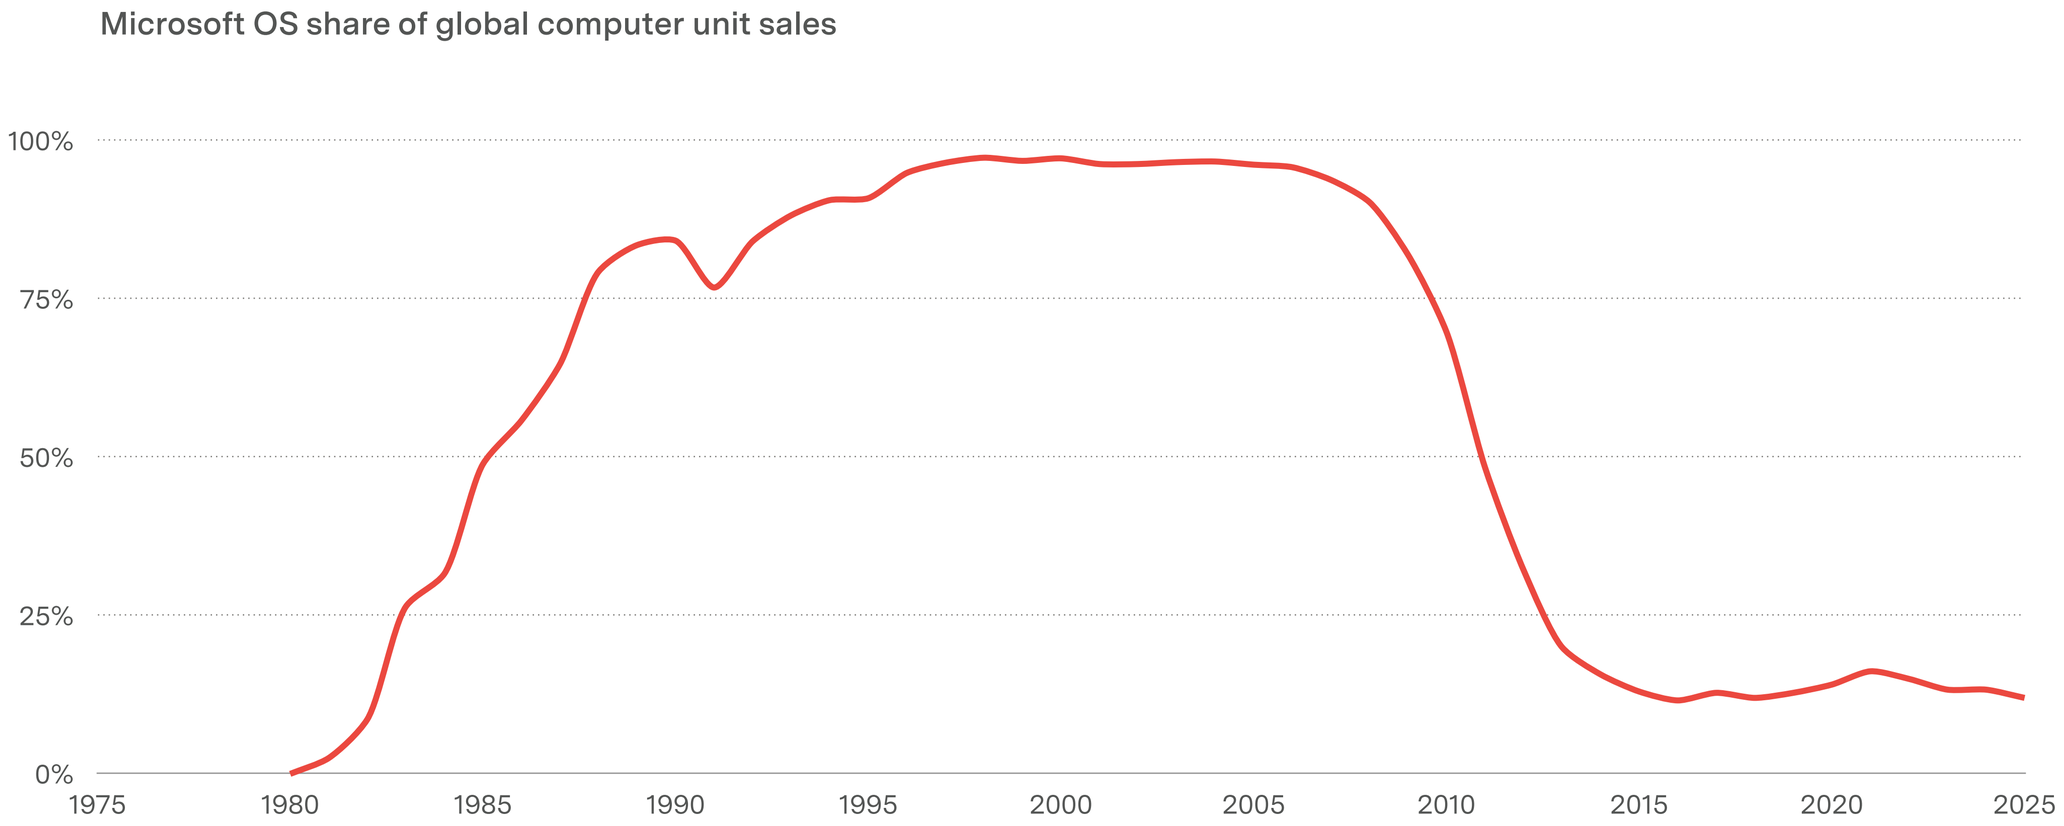

Source: Apple, Google, Gartner

Early leaders often disappear

The ‘first’ was not the winner in PCs, browsers, search, social or smartphones

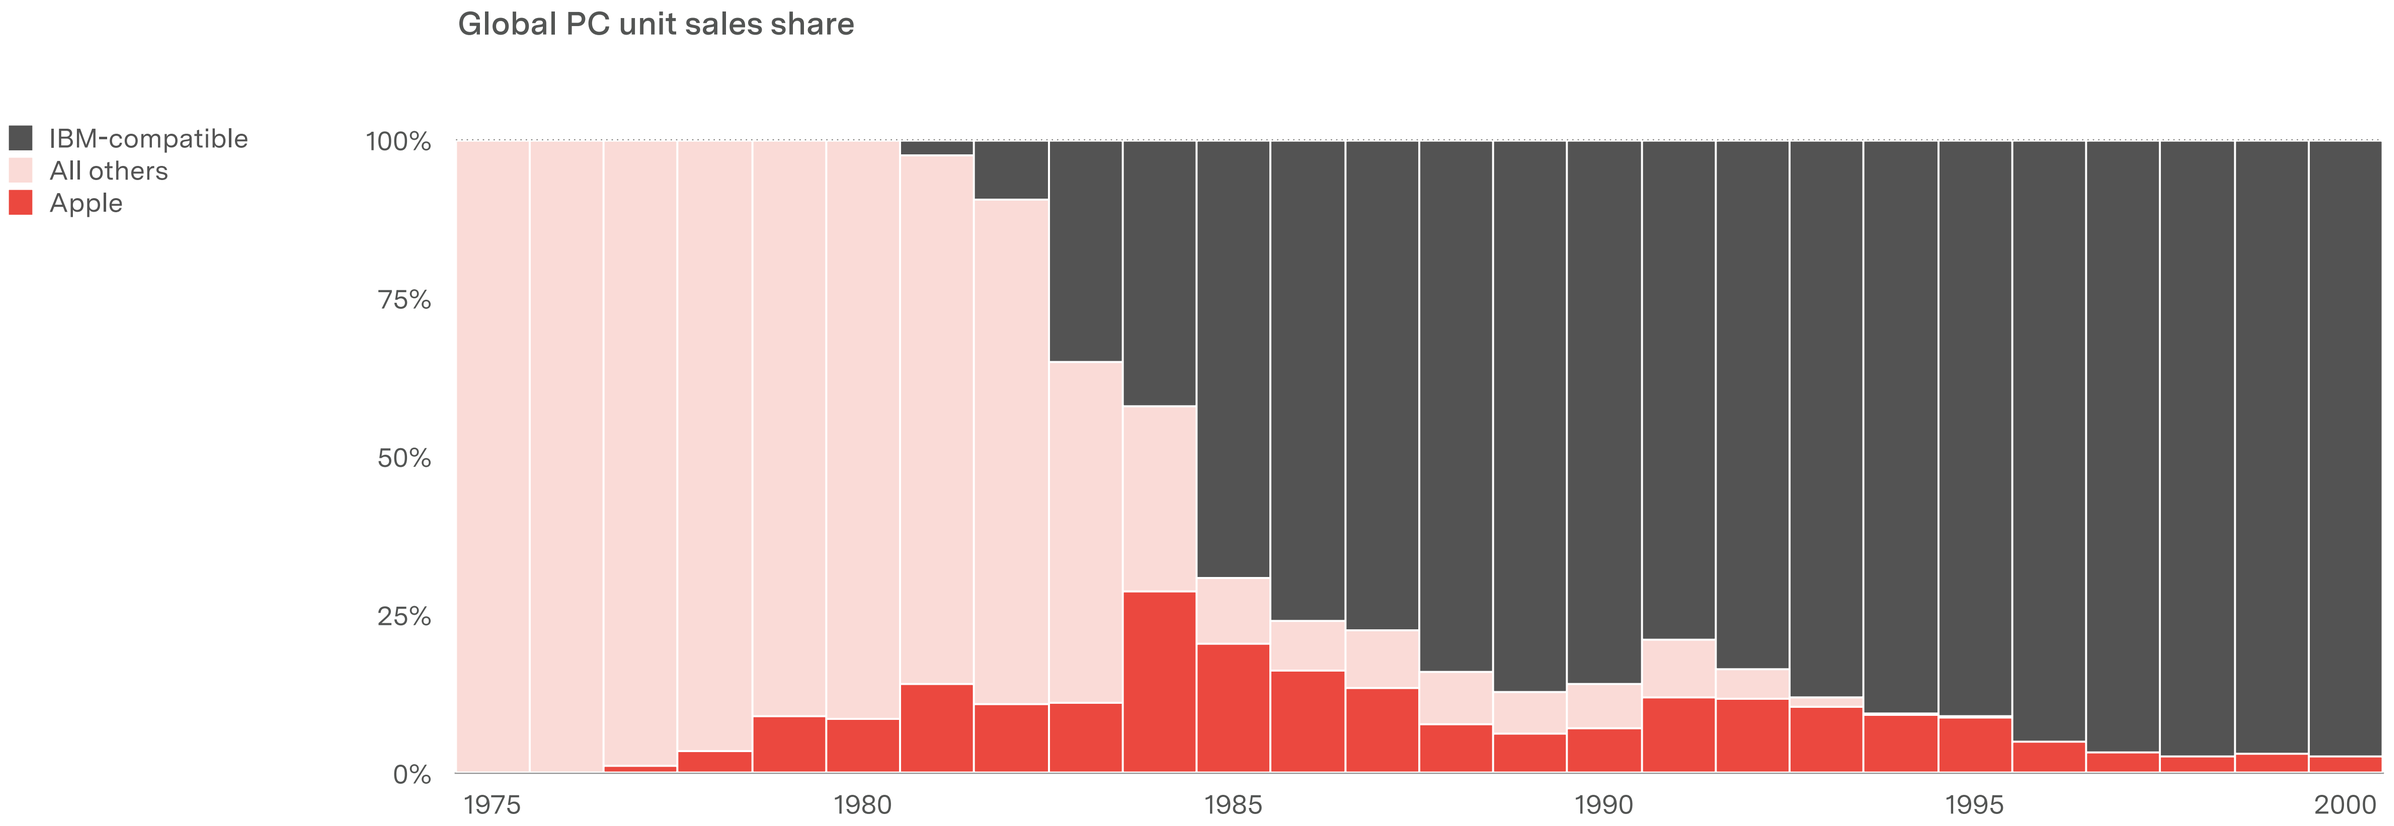

Source: Gartner

How will the new thing work? We don’t know

For every new platform, we forget how many ideas failed and how unclear everything was

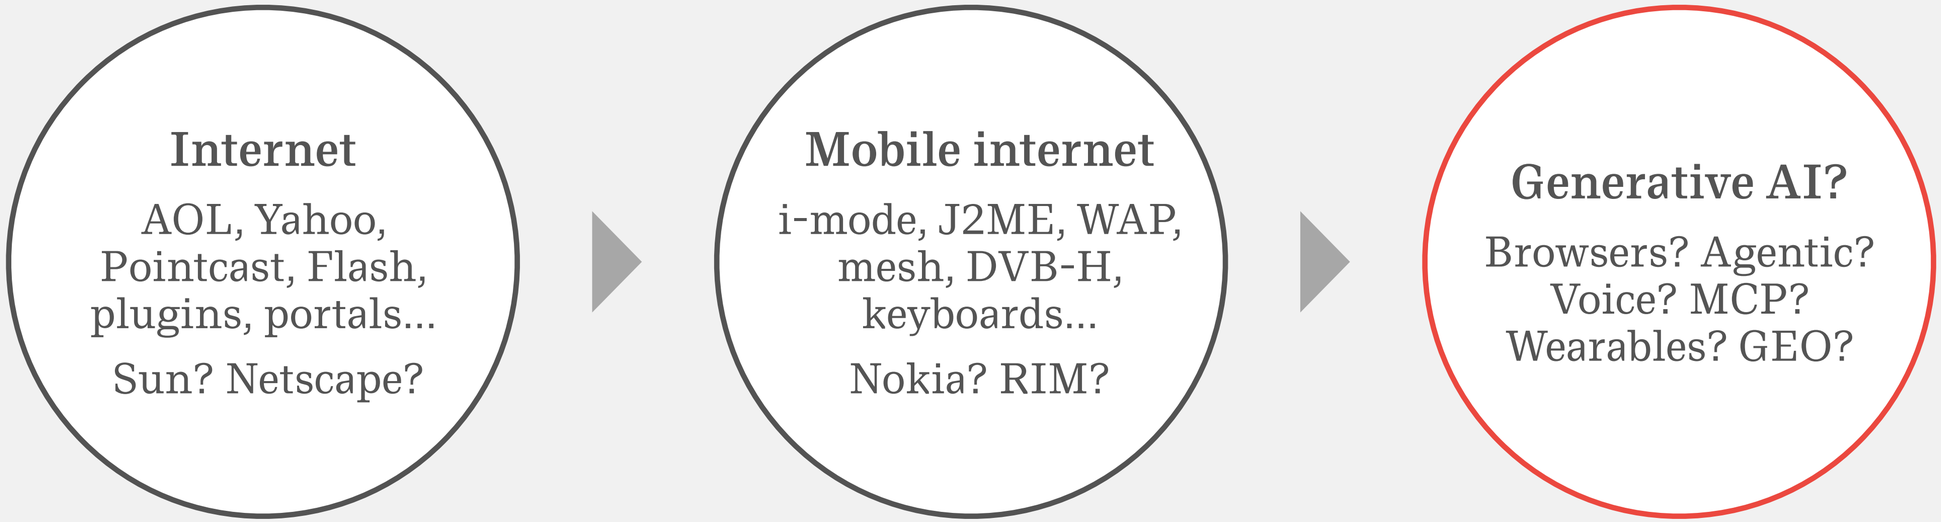

Noise, hype, anti-hype

When things are exciting, people get excited

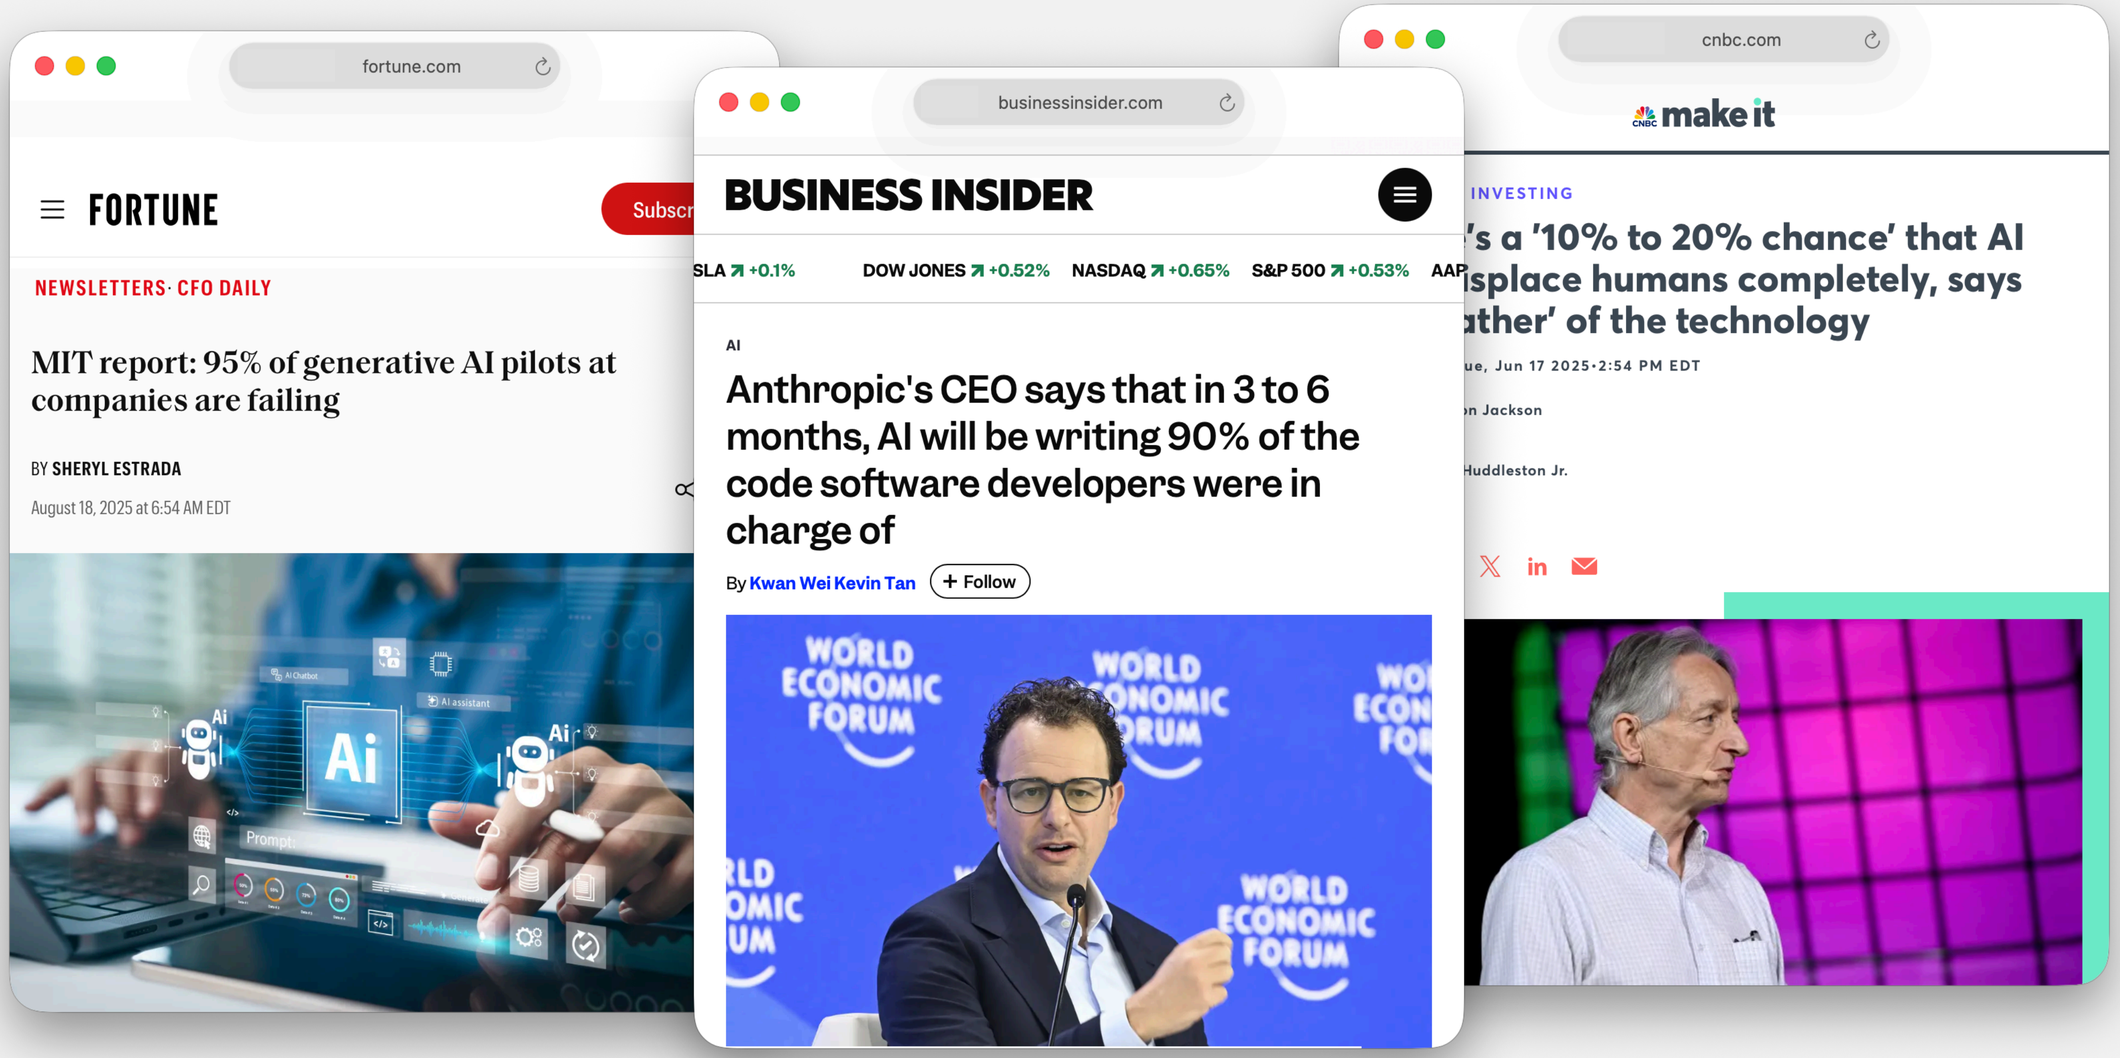

This often brings bubbles

People draw straight lines on log scale charts

They forget that exponential growth is generally an exponential curve

And always say “this is different”

The trouble is, they’re generally *right* - every bubble is different! But it can still be a bubble

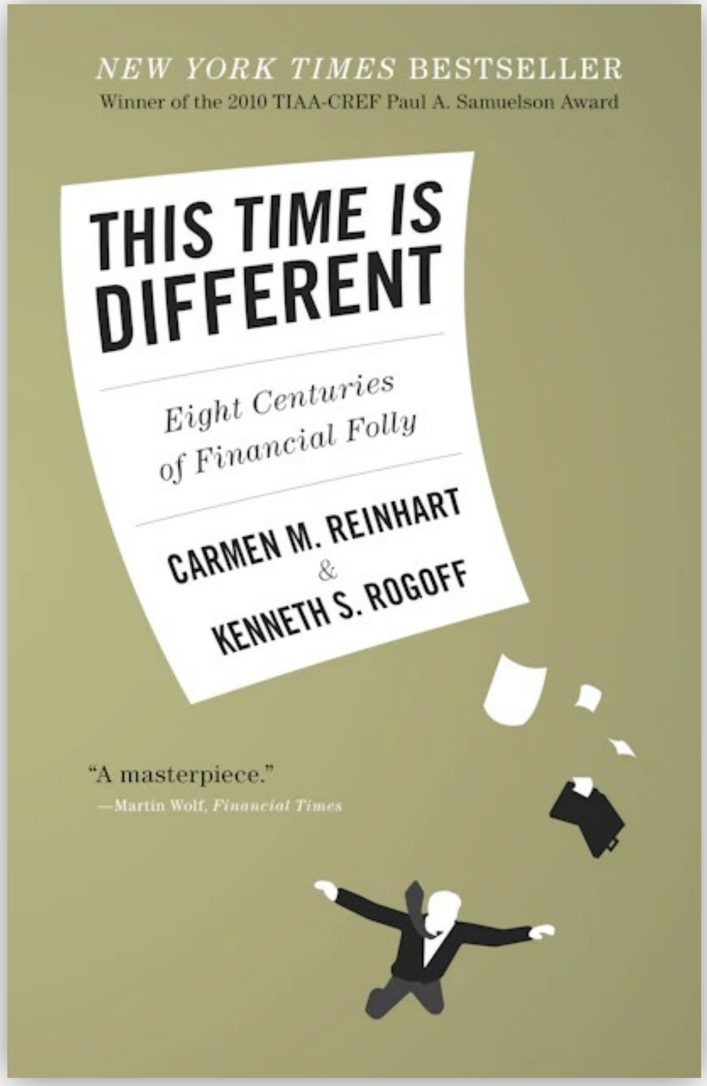

But when the dust settles, the world has changed

The internet has gone from the New Thing to a basic part of daily life

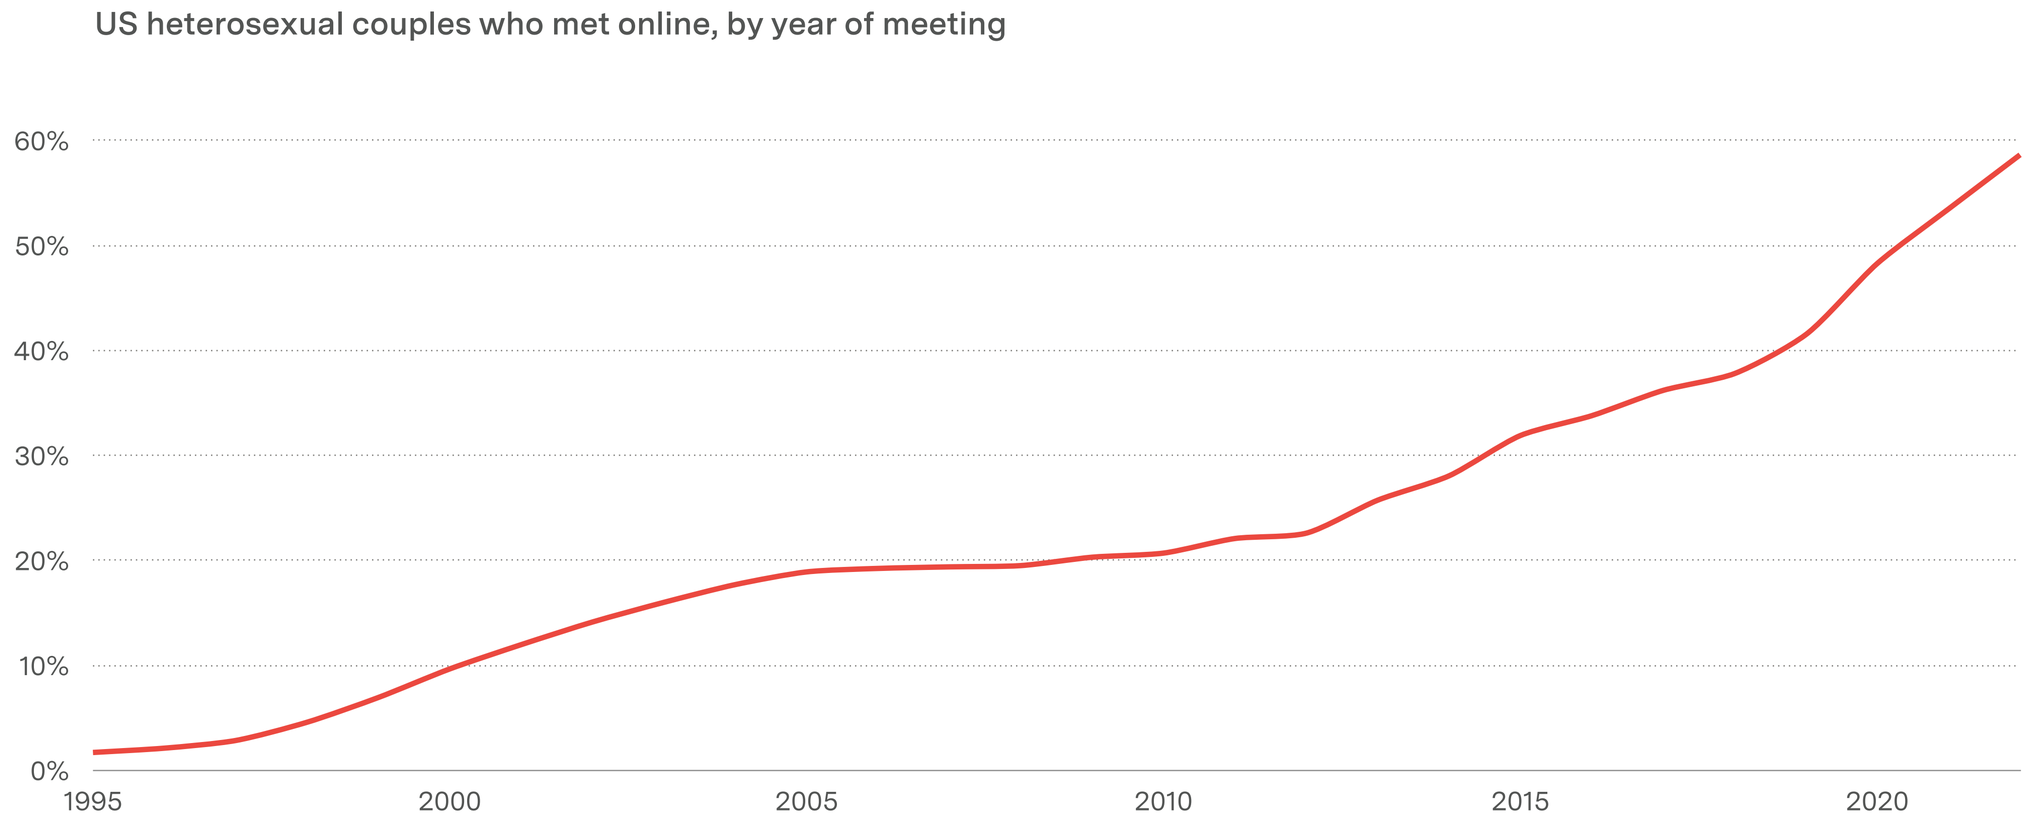

Source: Rosenfeld HCMST

New platforms mean new tools (and new revenue)

SaaS means the typical large enterprise in the USA now uses 4-500 apps

*Source: Productiv

One way this platform shift *is* different, though

For PCs, the web or smartphones, we knew the physical limits of what could happen next year

With LLMs, we don’t know how much better this could get

“The race to AGI is afoot”
Sergey Brin

“AGI needs multiple further breakthroughs”
Demis Hassabis

Another platform shift, or more?

We know this will get better, but we don't know how much

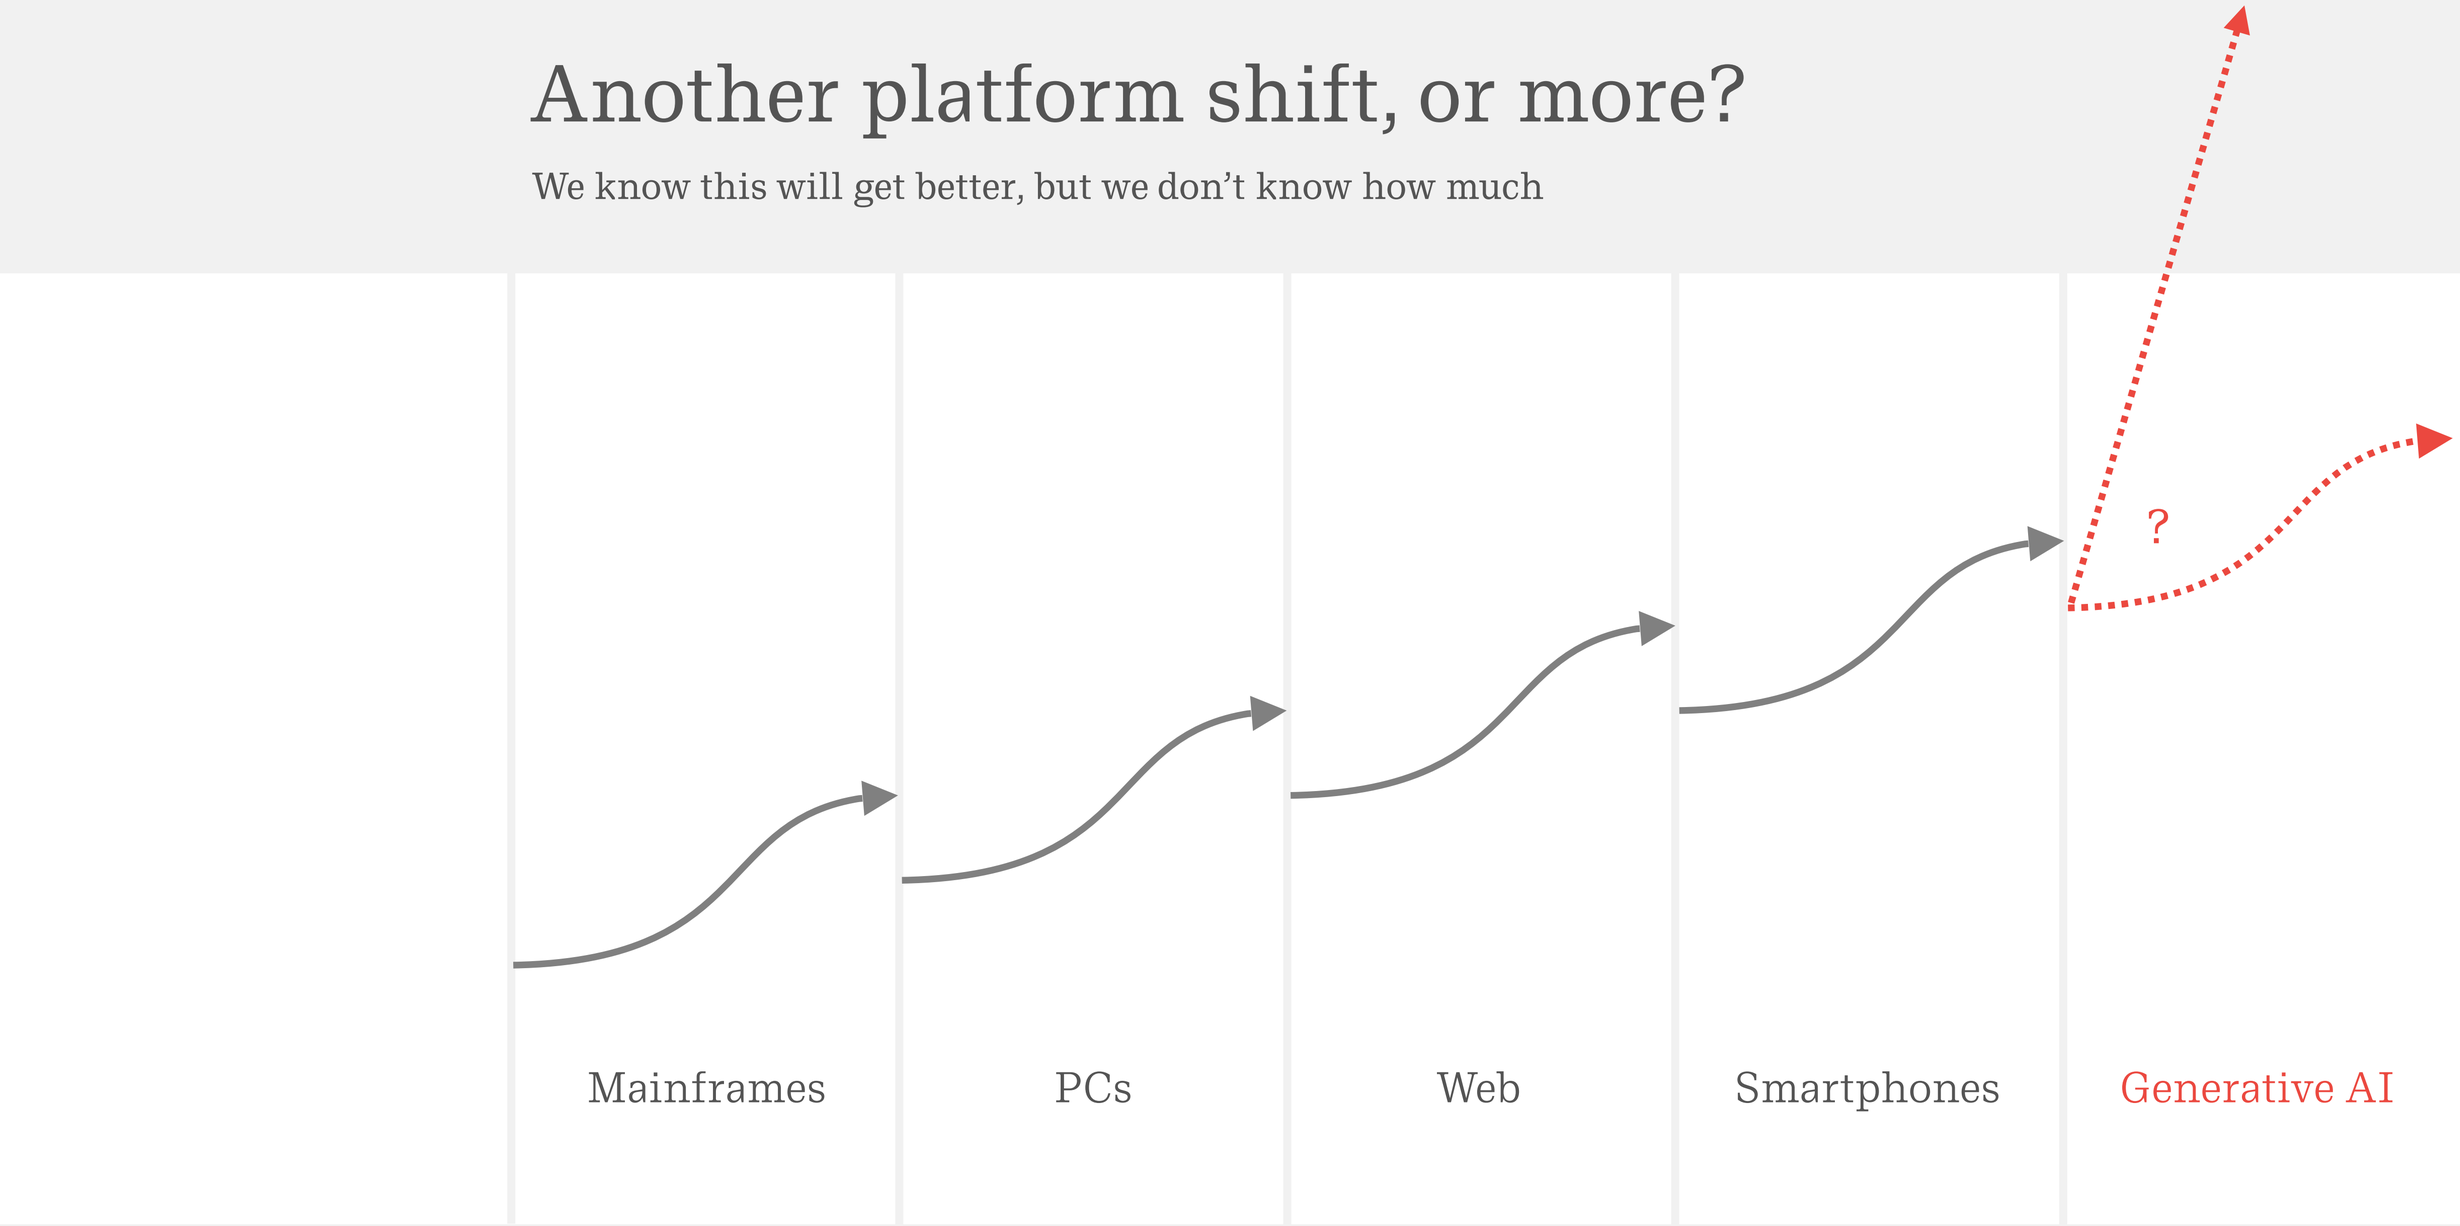

So how will the new thing turn out?

If this is ‘only’ as big as mobile or the internet, that seems like enough

How will this work?

How will this be useful?

Where is the distribution, value capture, and value destruction?

Inside tech

“The risk of under-investing is significantly greater than the risk of over-investing”

Sundar Pichai, Q2 2024

“The very worst case would be that we have just pre-built for a couple of years”

Mark Zuckerberg, Q3 2025

Three years of FOMO in Big Tech

"This is a huge new market, a huge threat to all our existing businesses, and we can't miss it"

FOMO drives a capex surge

~$400bn in 2025 for the big four alone (for comparison, global telecoms is ~$300bn)

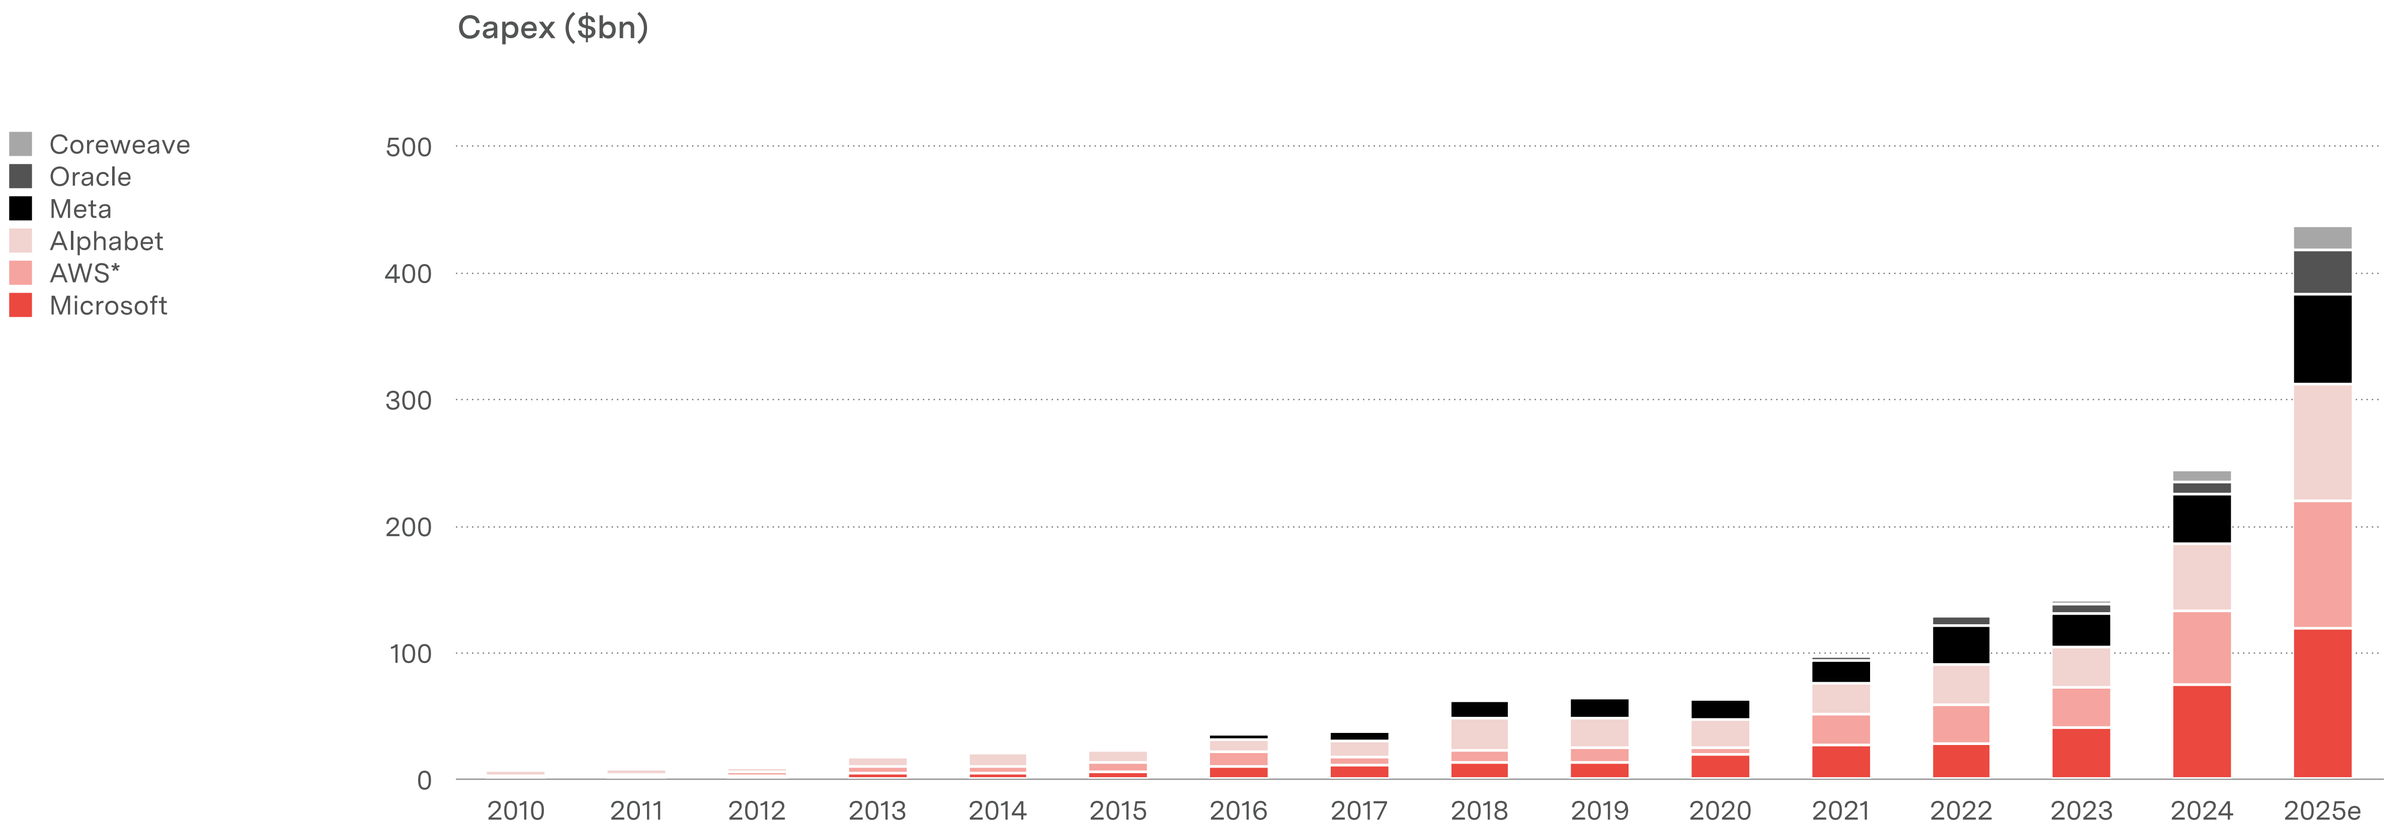

Source: Companies, company guidance. Includes capital leases
* Amazon does not break out AWS capex directly

Planned 2025 growth almost doubled - in 2025

~$350bn in 2025 for the big four alone (for comparison, global telecoms is ~$300bn)

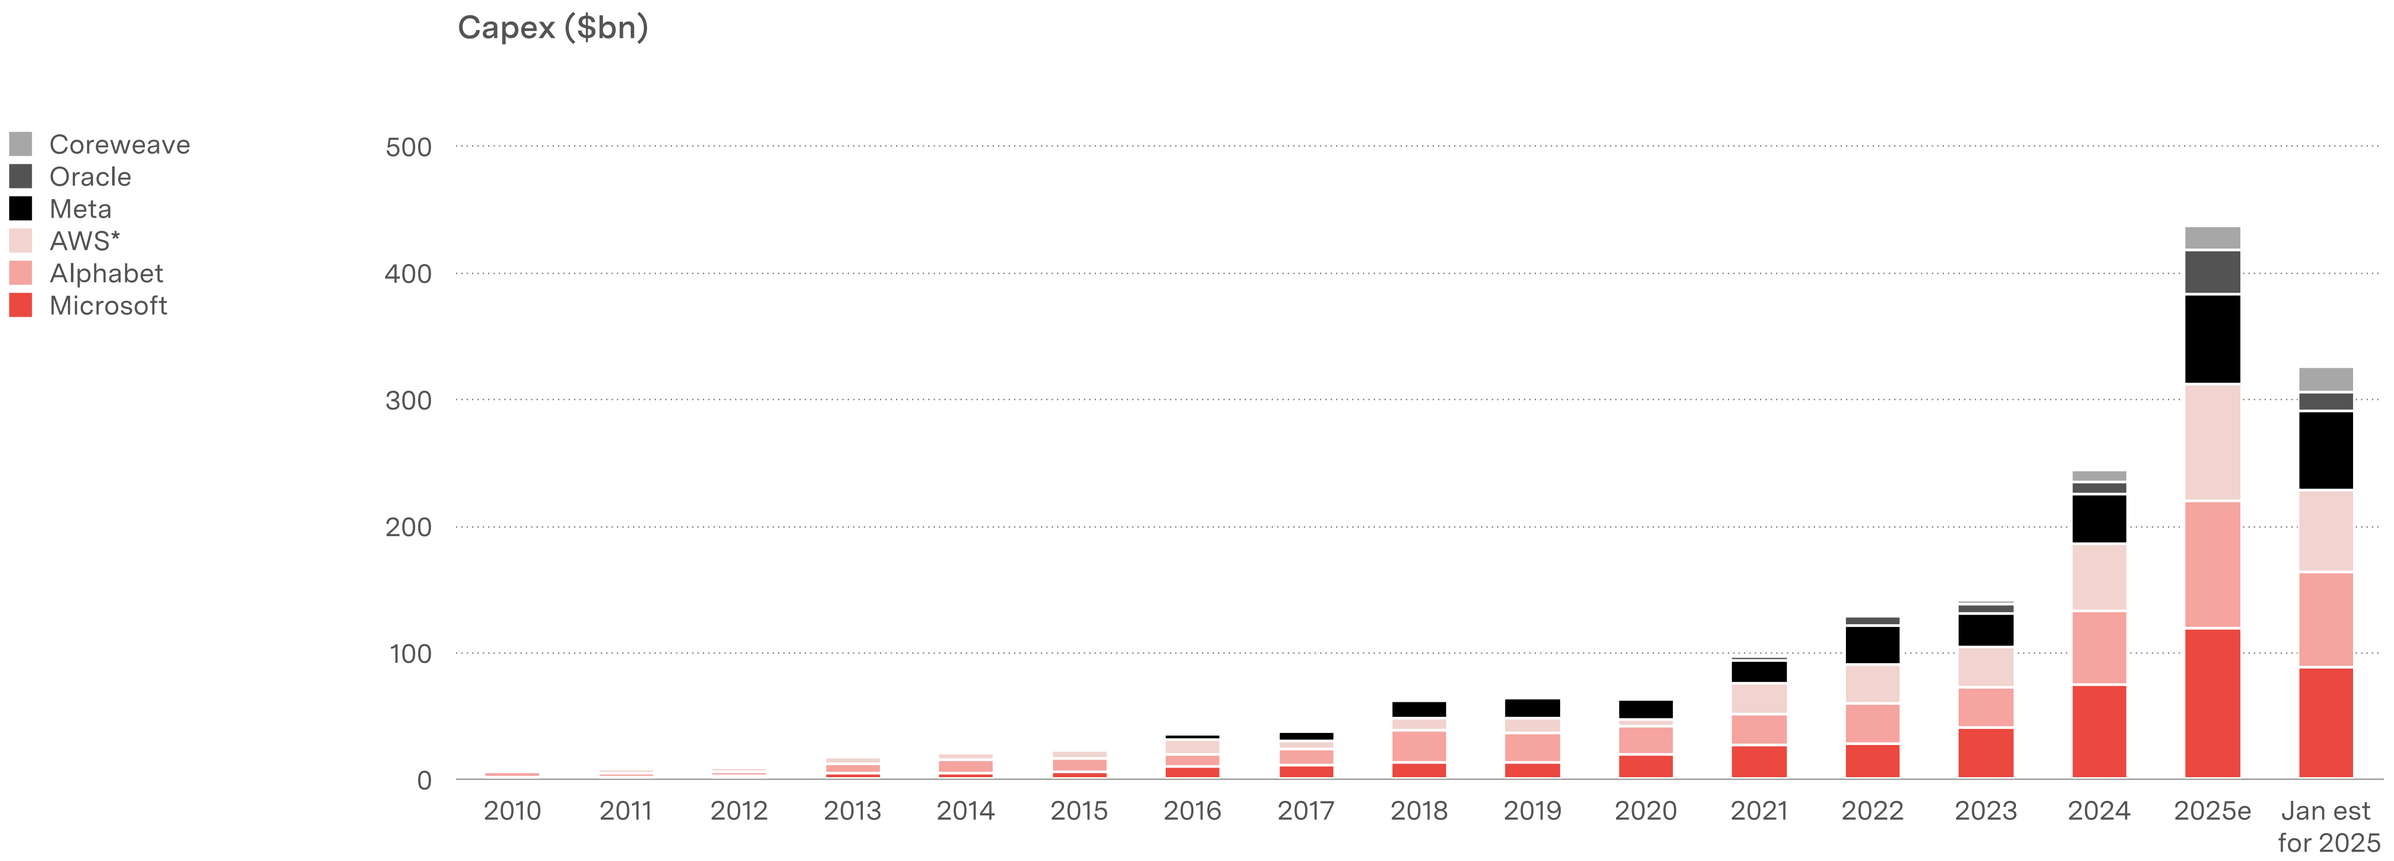

Source: Companies, company guidance. Includes capital leases
* Amazon does not break out AWS capex but reports it as ‘the majority’

A new investment cycle

US data centre construction overtaking offices

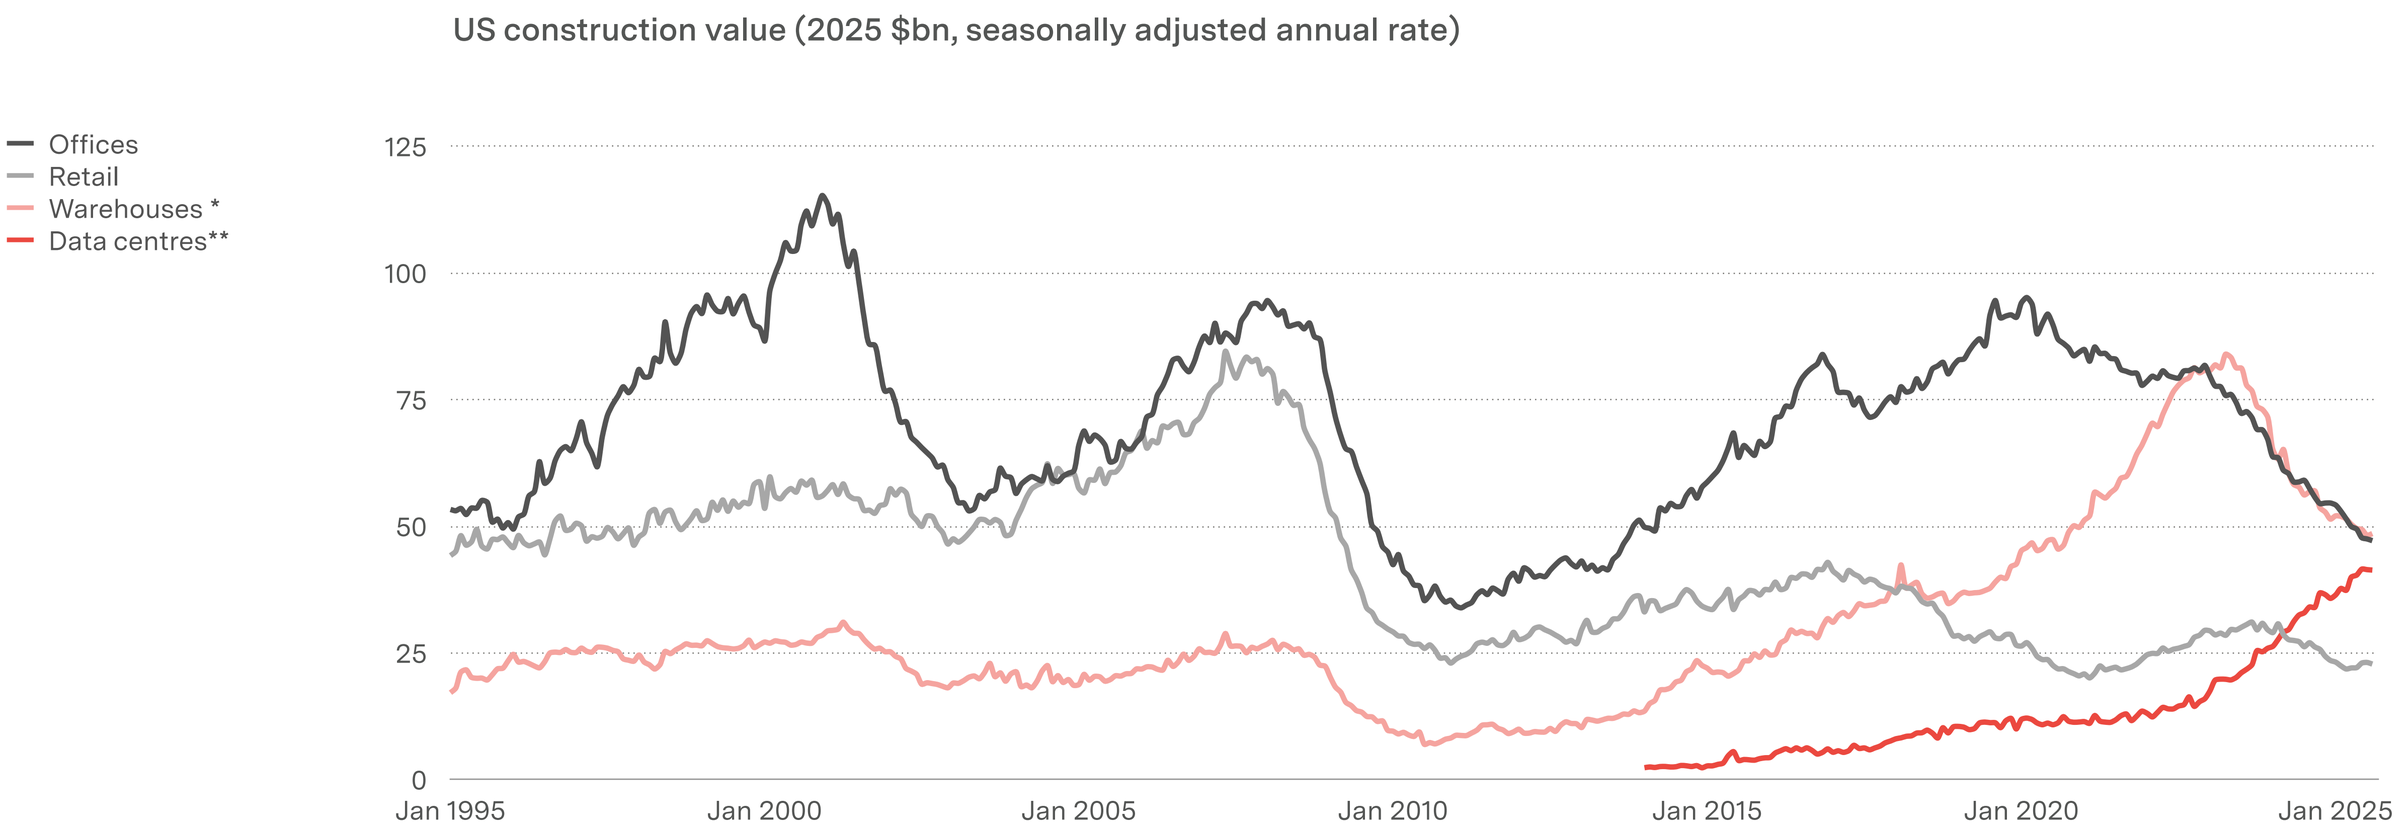

Source: US Census. * Excludes mini-storage ** Excludes compute
NB: data after July 2025 delayed by US government shutdown

Nvidia can’t keep up

Trying to build a new Sun Microsystems (though China and hyperscalers’ own chips are coming up behind)

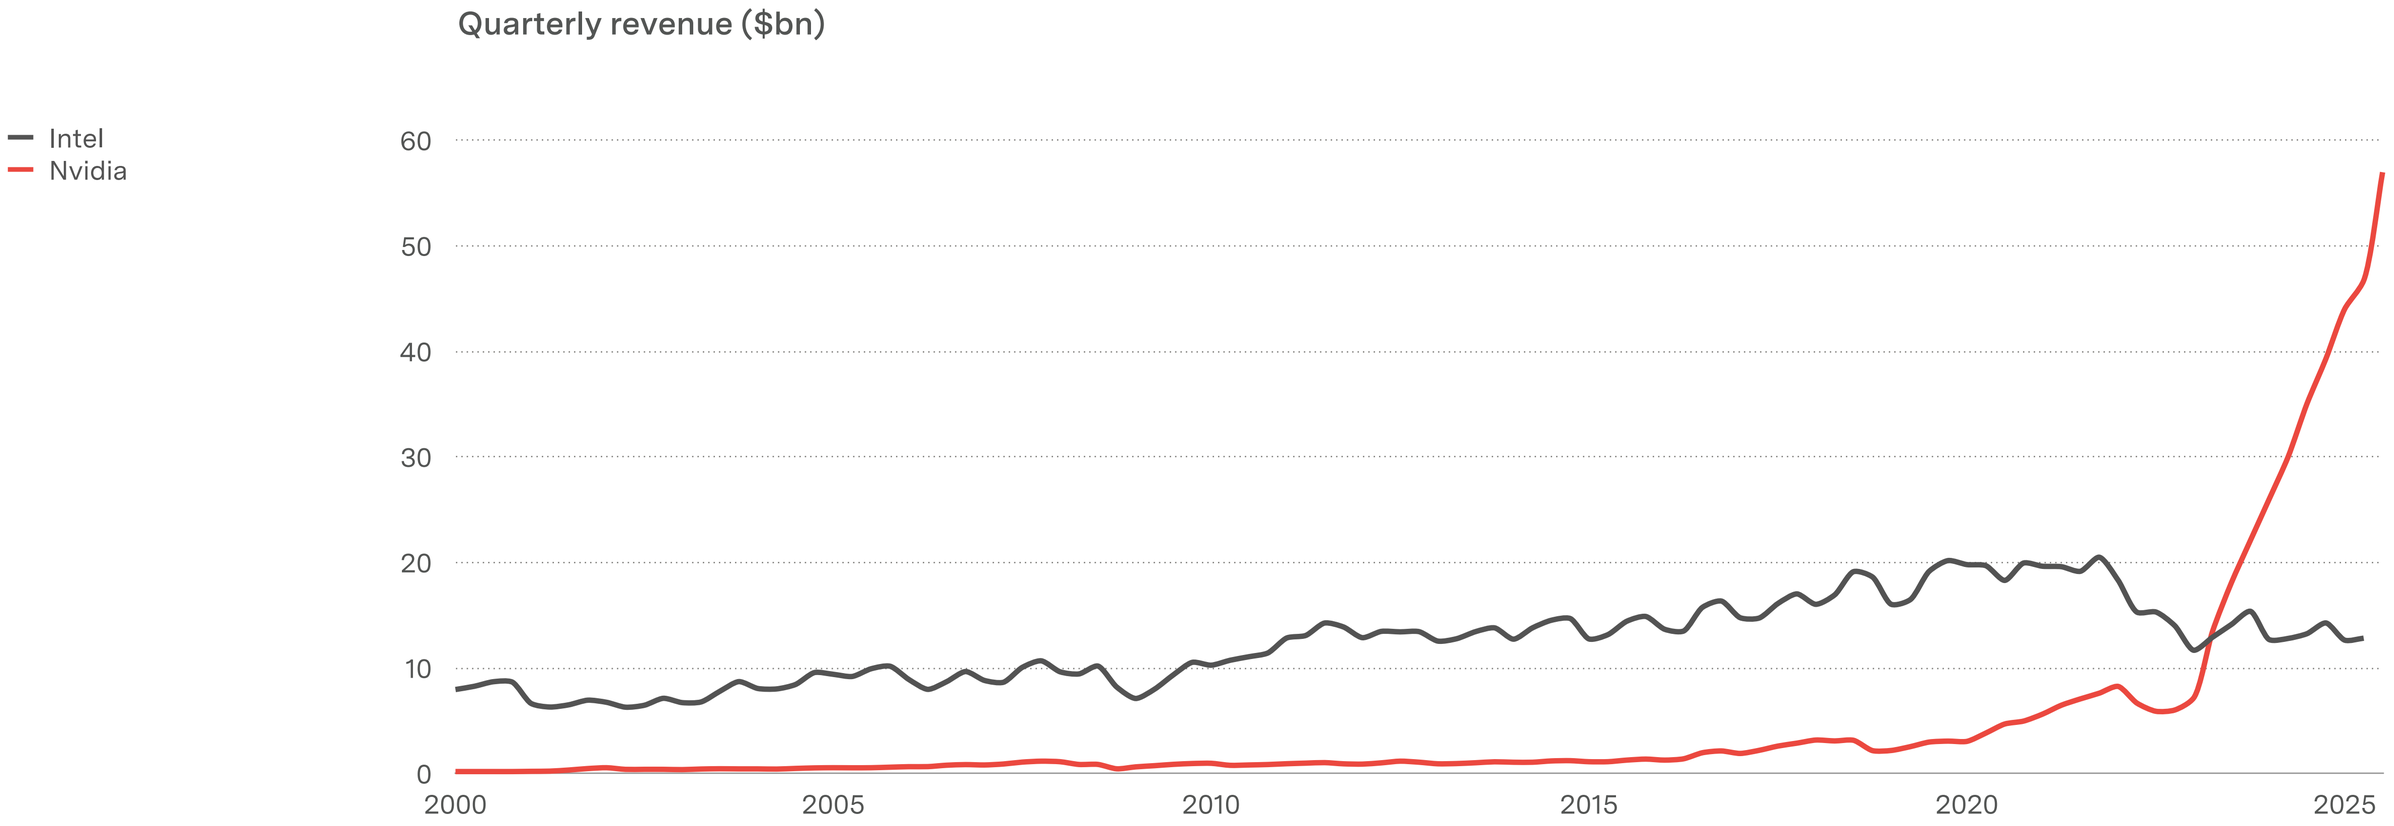

Source: Companies

Nvidia can’t keep up (neither can TSMC)

TSMC unwilling/unable to expand capacity fast enough to meet Nvidia’s book

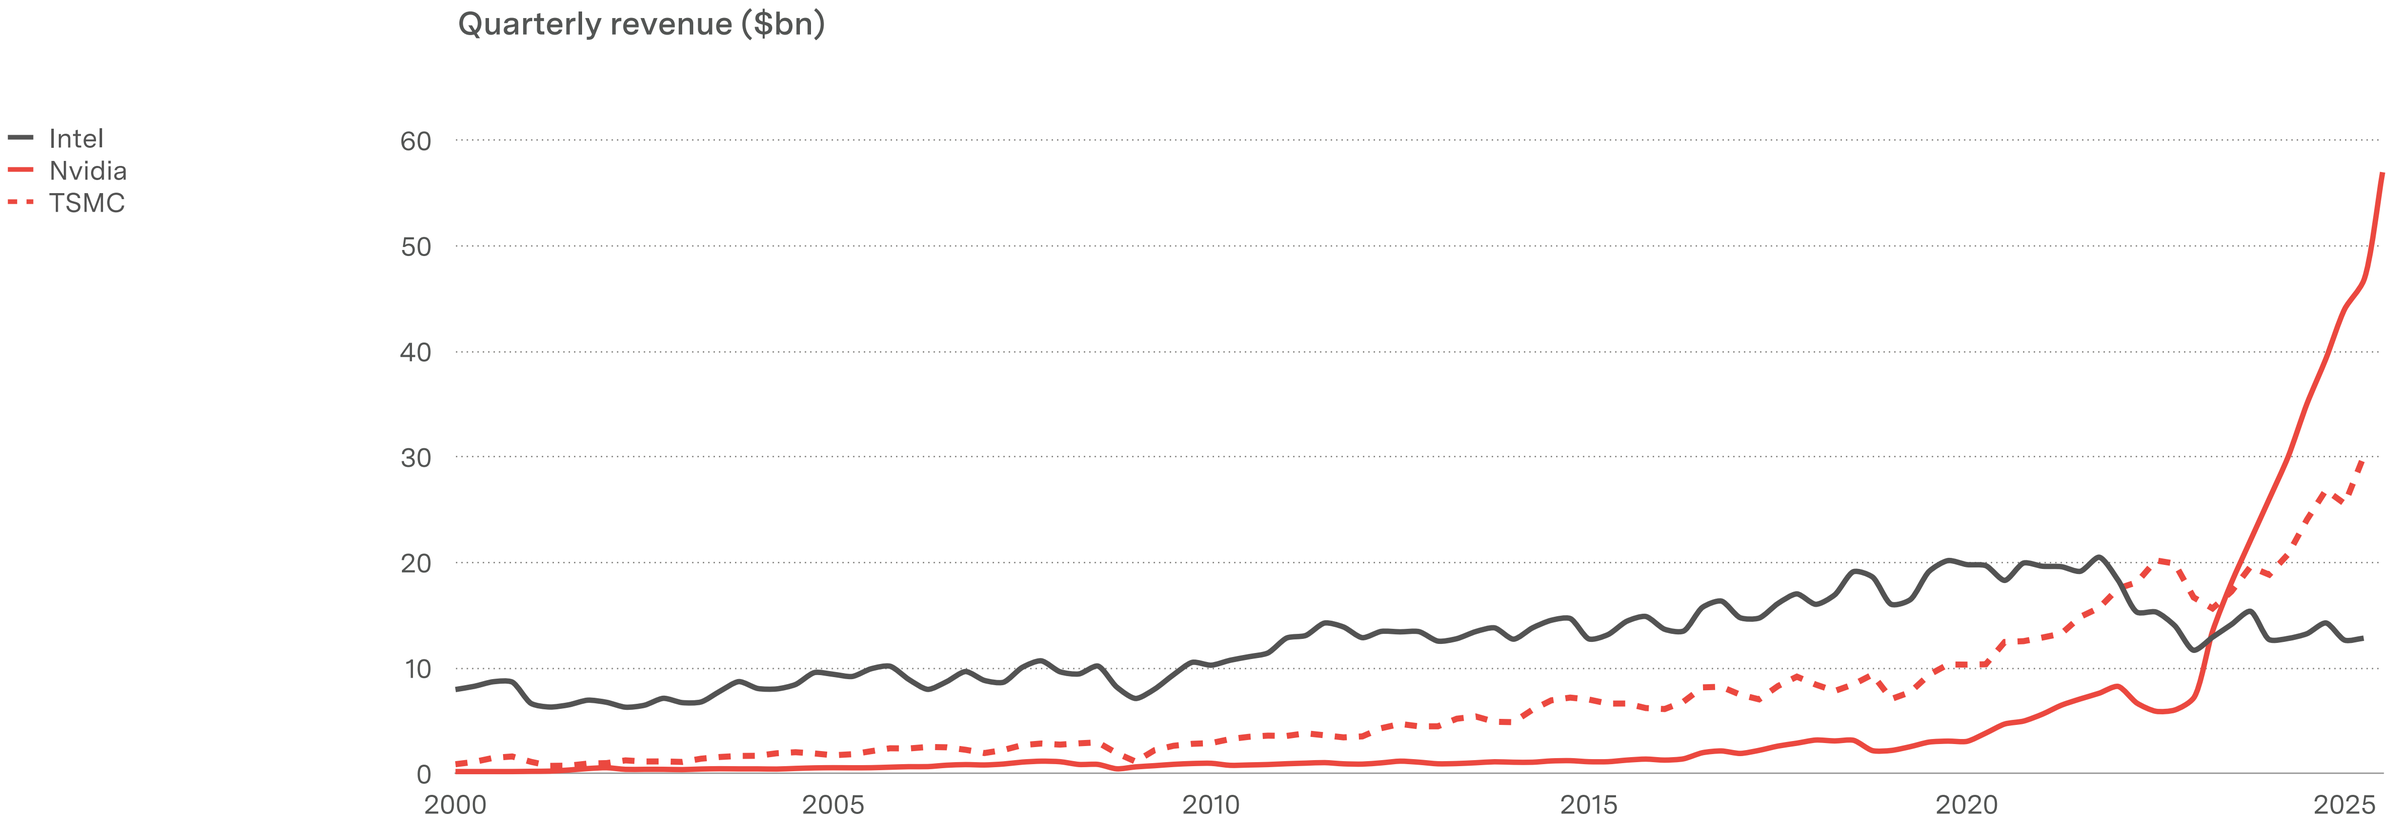

Source: Companies

US power backlogs becoming a major issue

US power demand growth is ~2%, and AI might add 1% that's hard to build fast (this is not an issue in China)

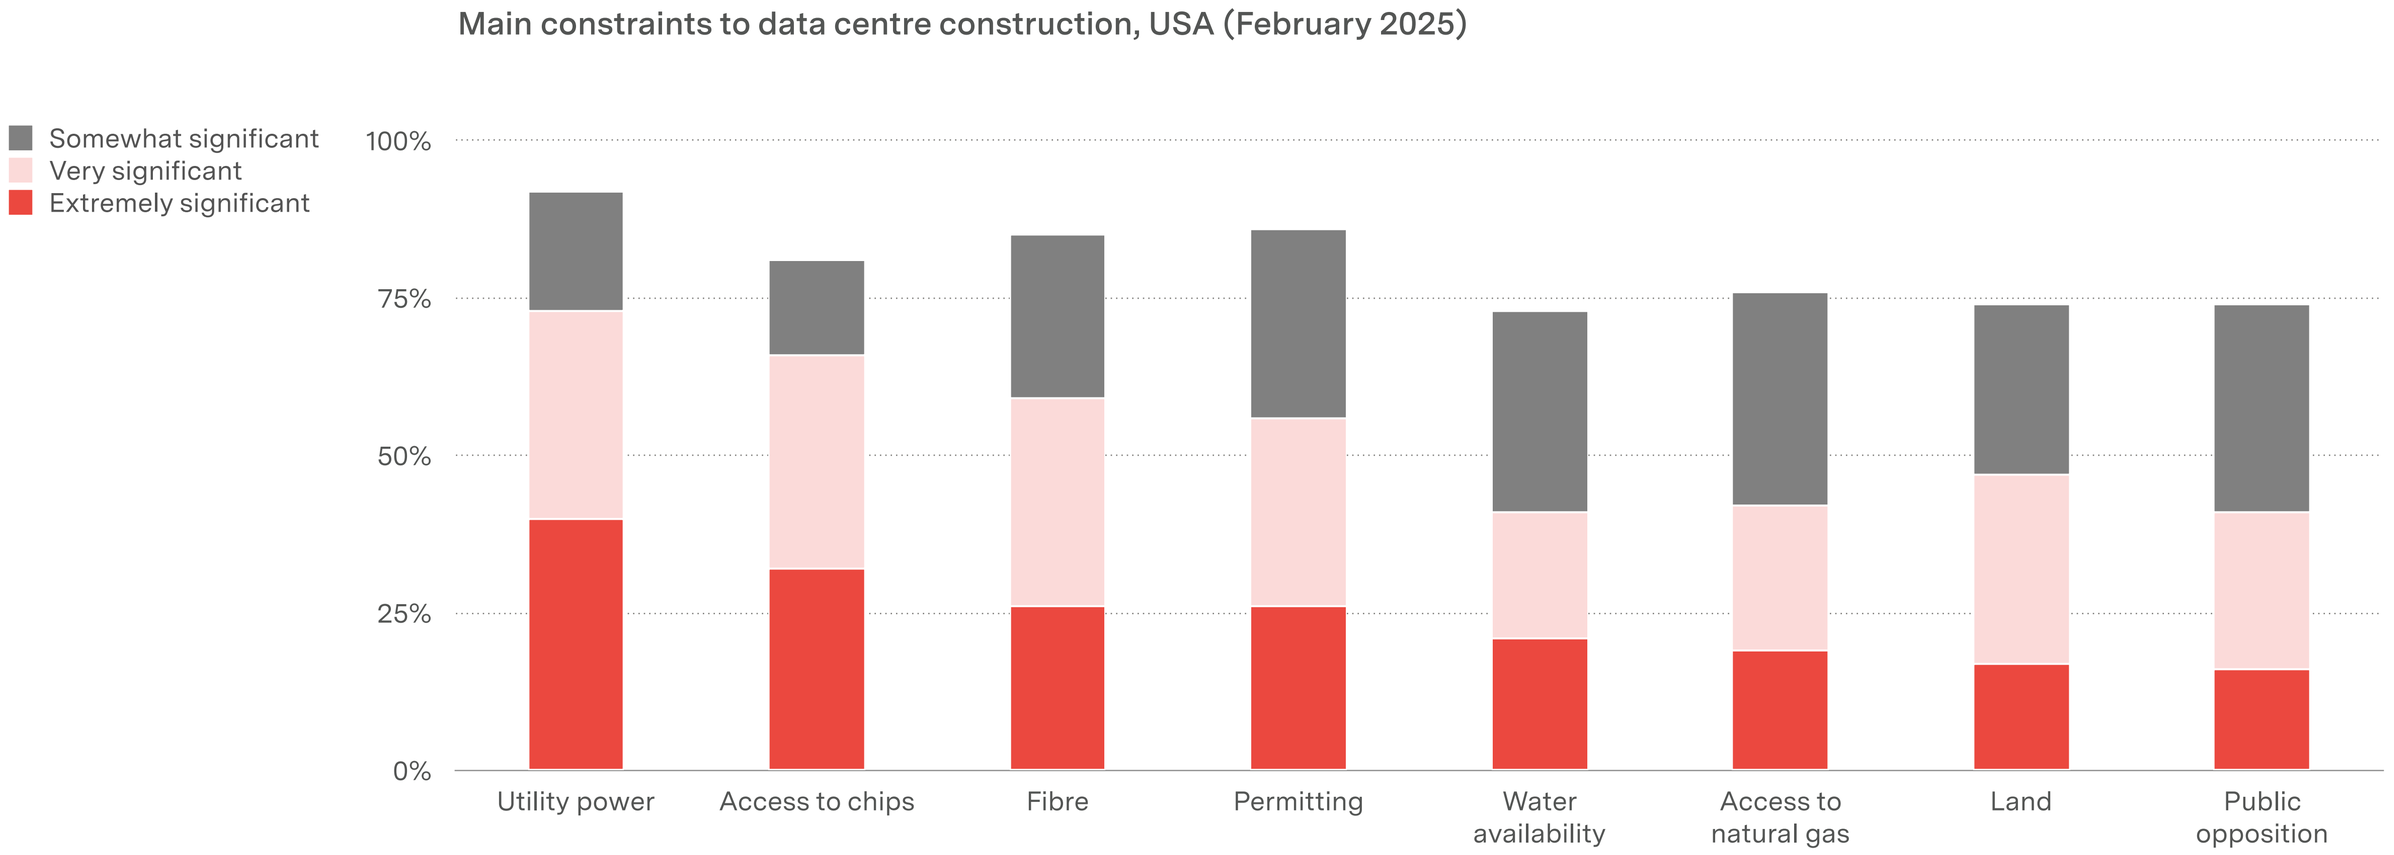

Source: Schneider Electric industry survey

“It’s been almost impossible to build capacity fast enough since ChatGPT launched”

Kevin Scott, Microsoft CTO

“We now expect the FY26 growth rate to be higher than FY25” - Microsoft

“Capex dollar growth will be notably larger in 2026” - Meta

“We expect a significant increase in 2026” - Alphabet

Data centre capacity triples? For $3tr? $5tr? More?

Some very large numbers (although some of the ‘bragawatts’ may be more performative than real)

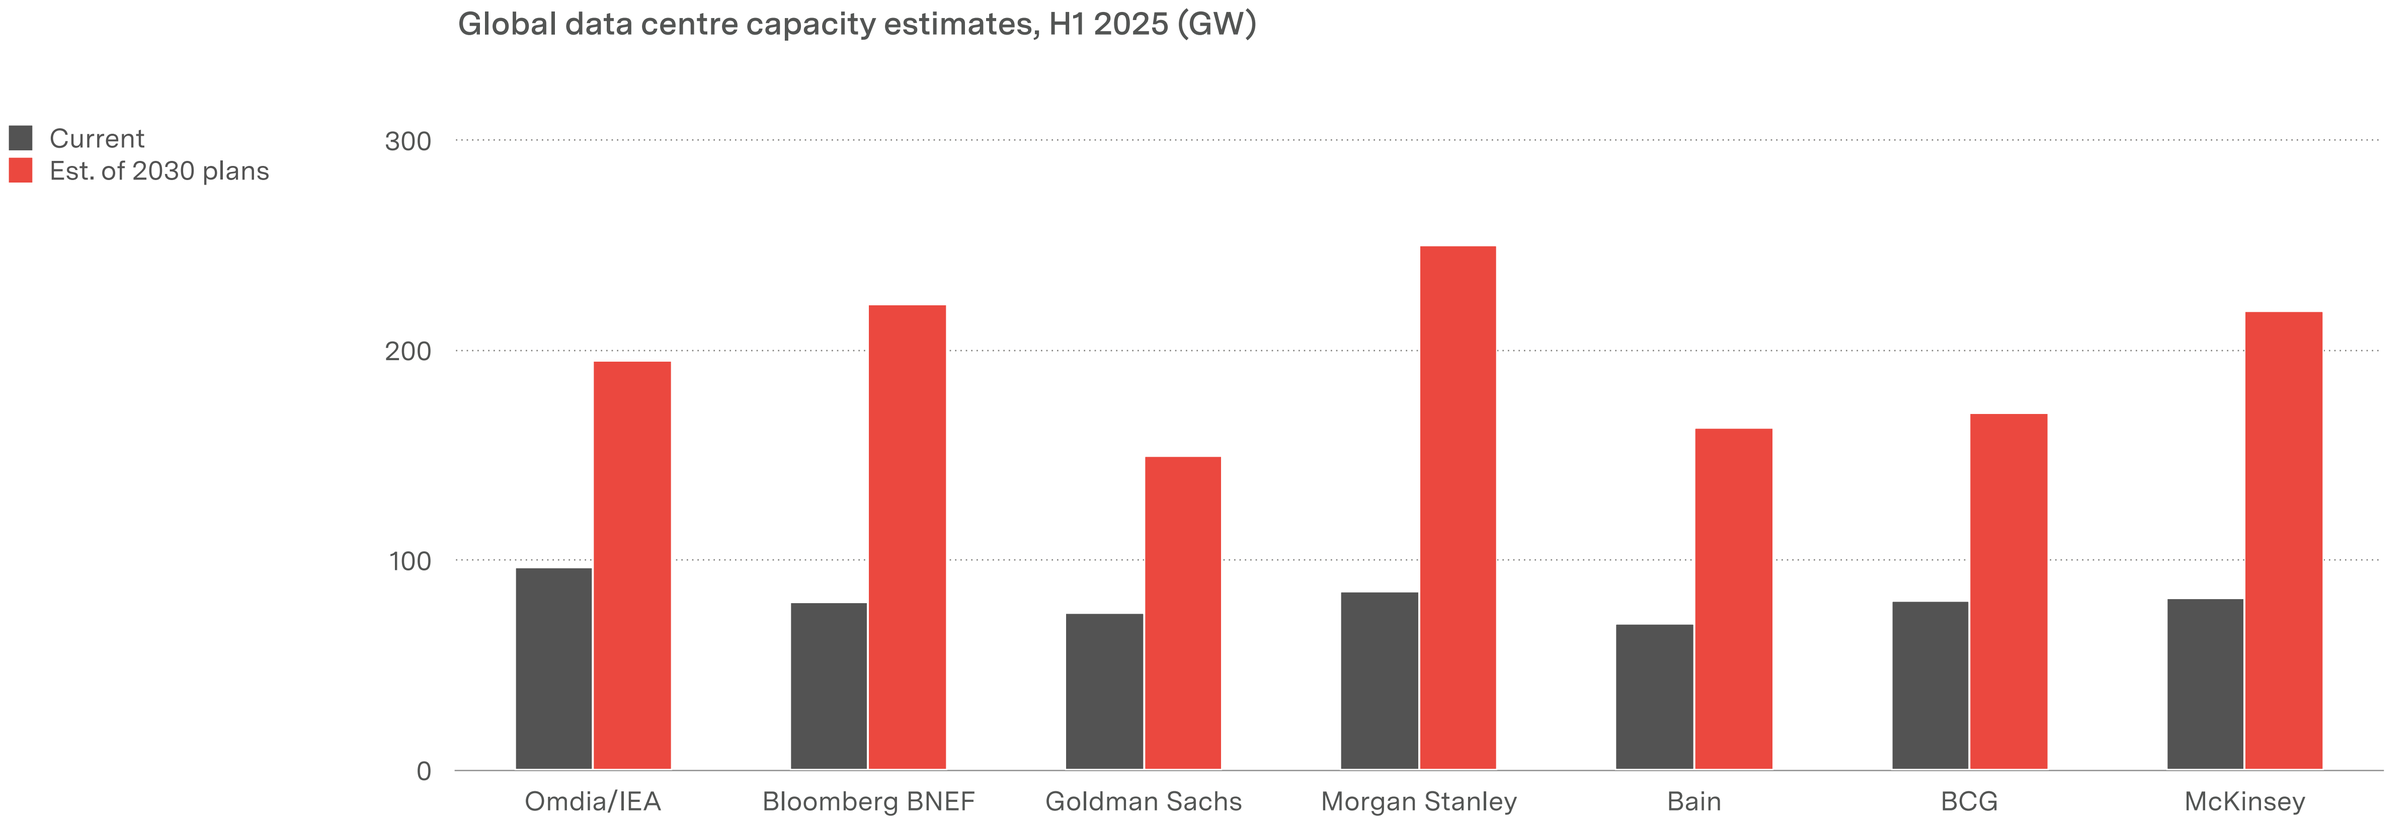

Source: Companies

“Three trillion dollars!”

Annualised AI capex aspirations are a similar magnitude to mature global capital-intensive industries

Source: ITU, IEA, ONS

The hyperscalers can afford it...

Big Tech cashflow has surged since the Pandemic, and most of that growth is going on AI capex

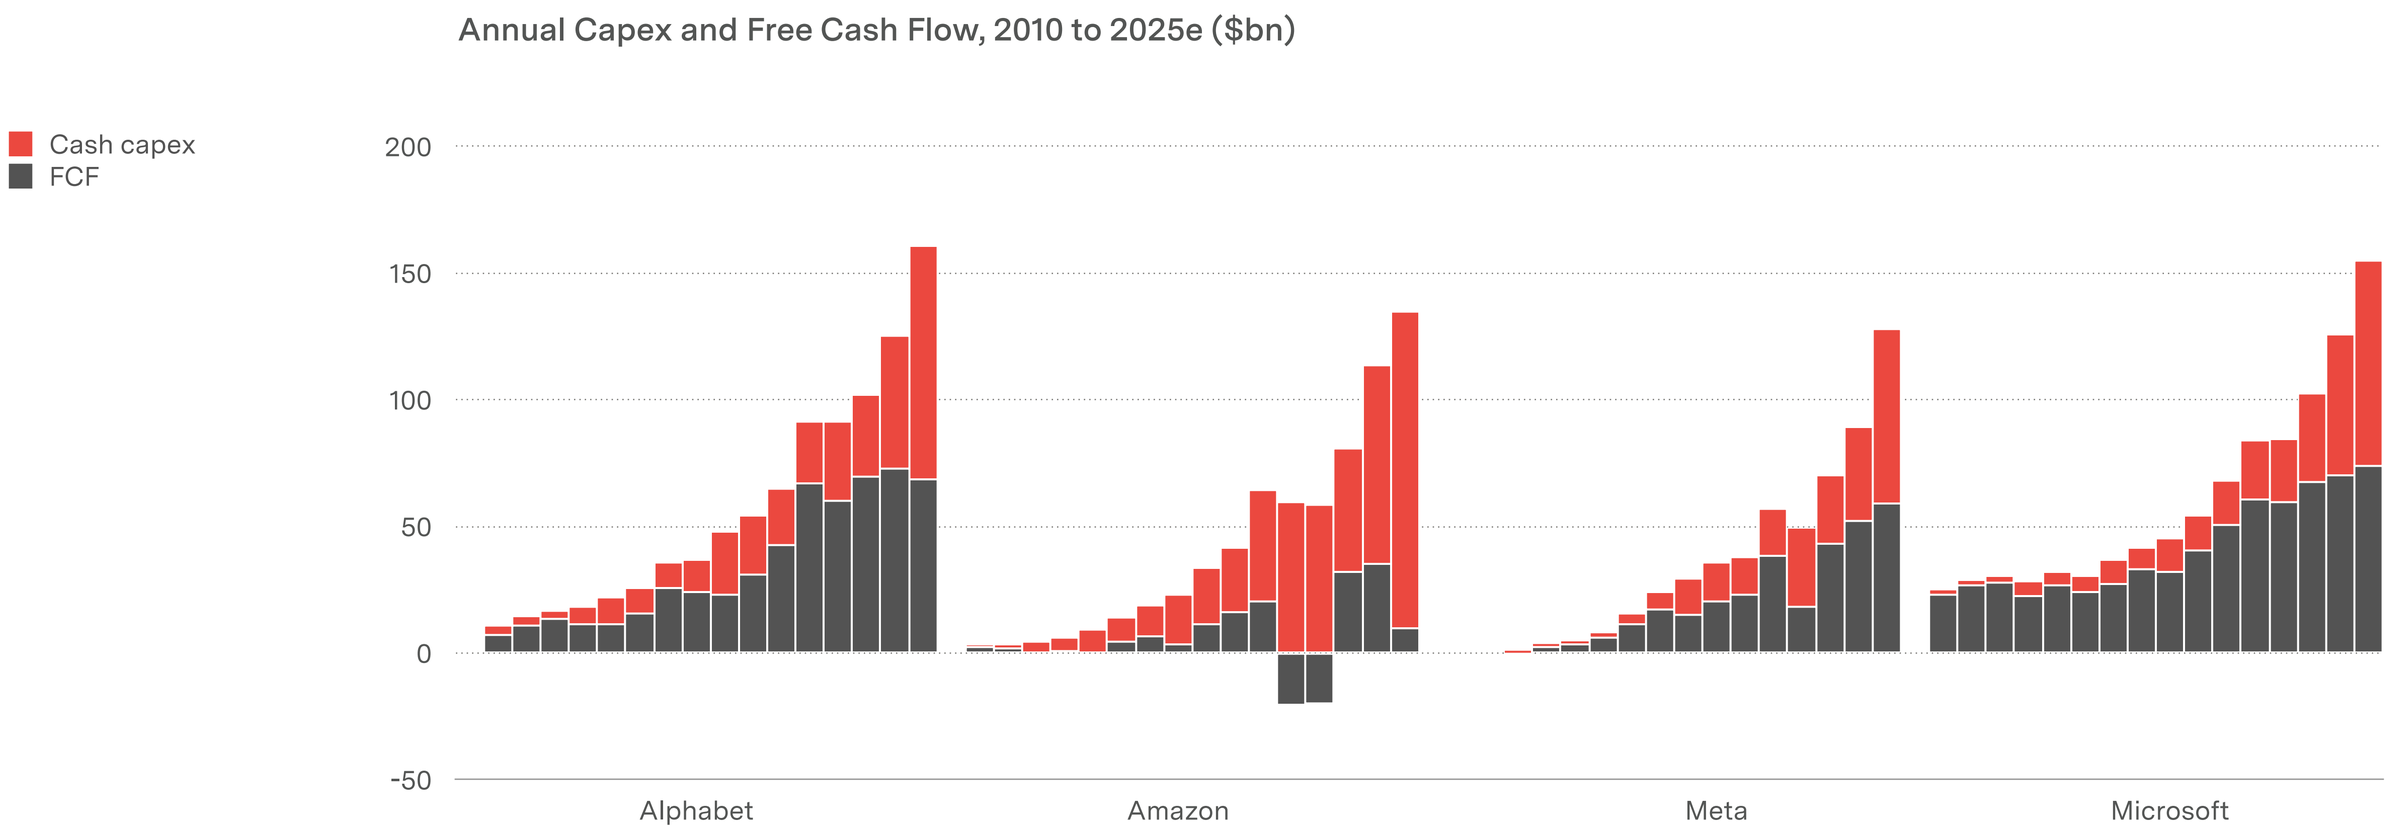

Source: Companies, consensus estimates. Calendar year.

The hyperscalers can afford it… up to a point

Capital leases are not new, but they've got a lot bigger

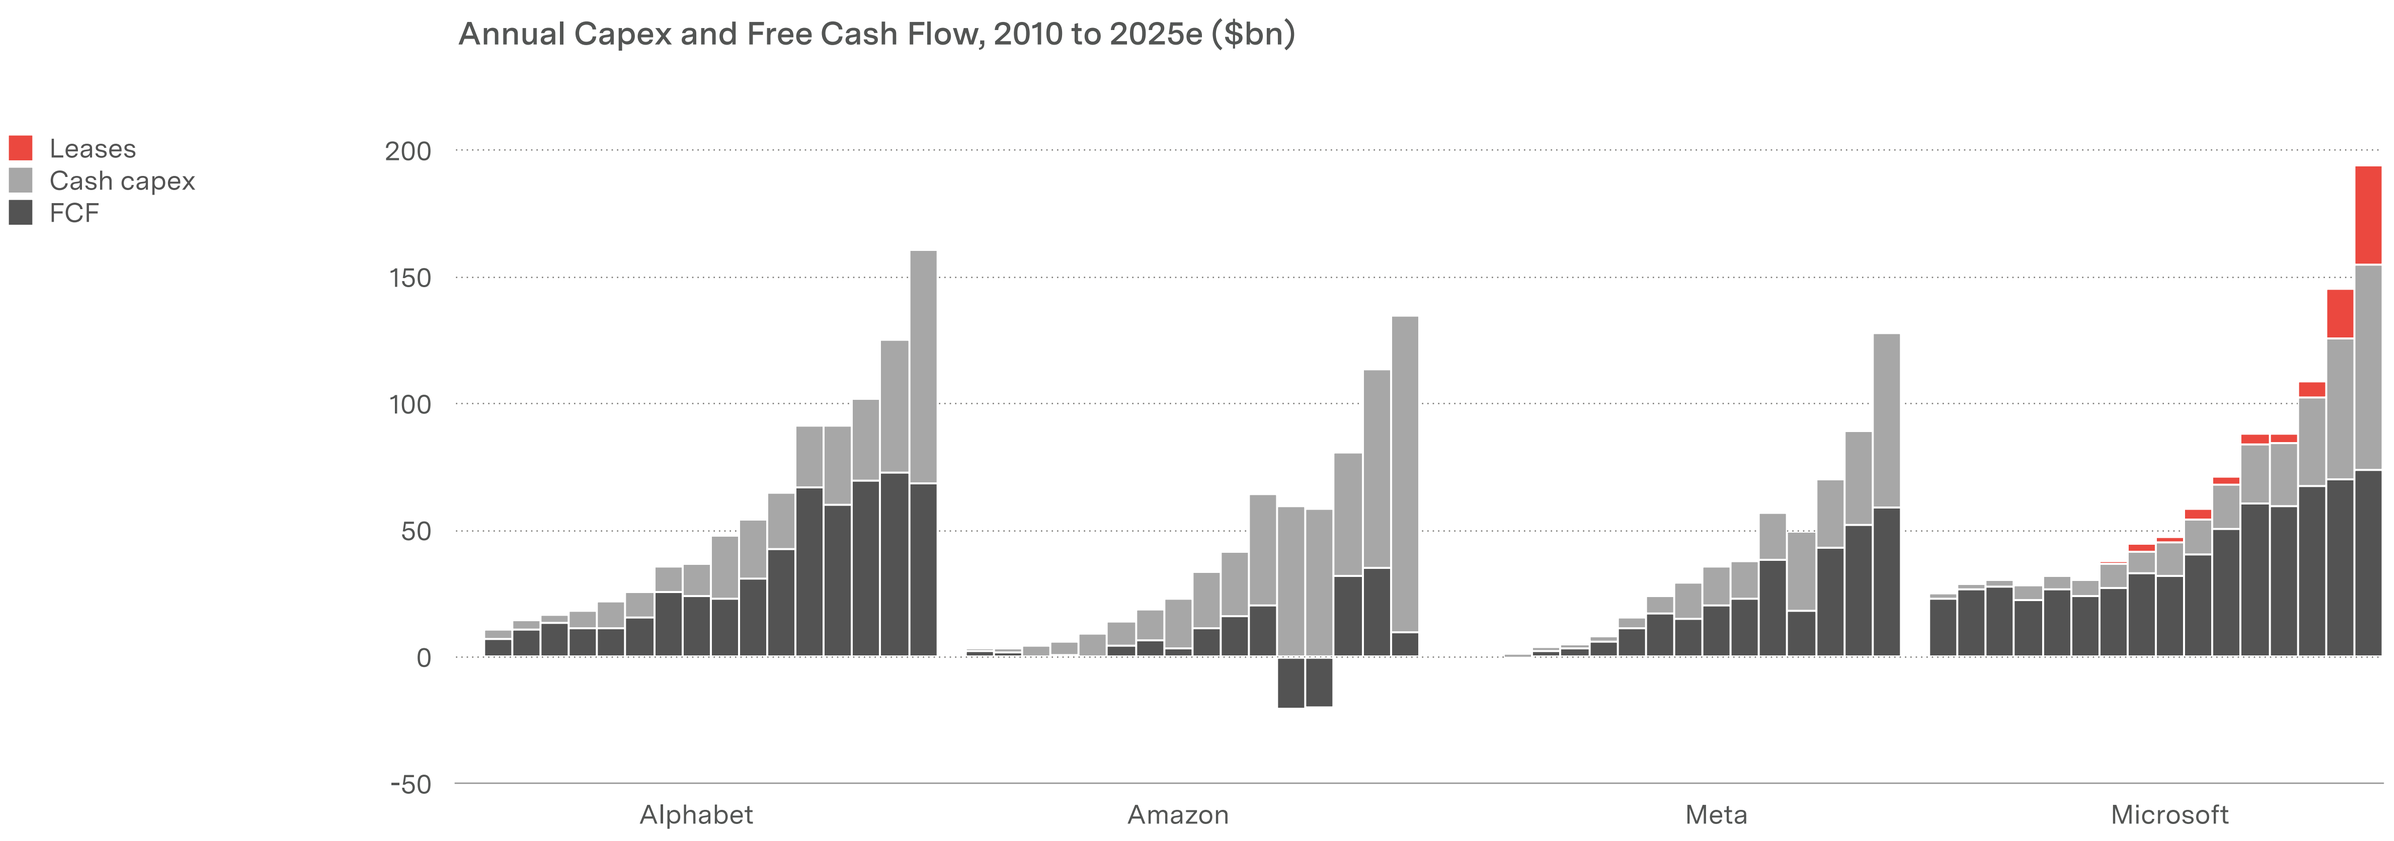

Source: Companies, consensus estimates. Calendar year.

Up to a point

Hyperscalers add leases and debt, while some analysts suggest Oracle's cloud capex might be >100% of revenue

ft.com

FINANCIAL TIMES

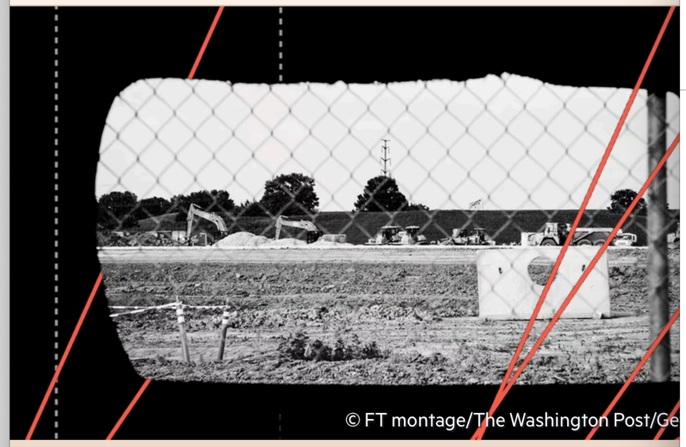

Big Tech + Add to myFT

'Absolutely immense': the companies on the hook for the $3tn AI building boom

Private capital joins Big Tech in seeking to capture rewards fro historic expansion of data centres

wsj.com

WSJ

Oracle Isn't Answering the Hardest Questions About Its AI Plans

Software giant boosts long-term targets, but paying for an artificial-intelligence build-out will stretch its resources

By Dan Gallagher Follow

Oct. 17, 2025 at 5:30 am ET

reuters.com

Reuters My News

Meta taps PIMCO, Blue Owl for $29 billion data center expansion project, source says

By **Chandni Shah** and **Surbhi Misra**

August 7, 2025 11:10 PM EDT · Updated August 7, 2025

Aa

A Meta logo is pictured at a trade fair in Hannover Messe, in Hanover, Germany, April 22, 2024.

OpenAI joins the club

Announced commitments for 30GW+ of capacity at $1.4tr

Aspiration for 1GW/week of new construction at $20bn/GW = ~$1tr annually...

Equivalent to 2/3 of total current global base, every year

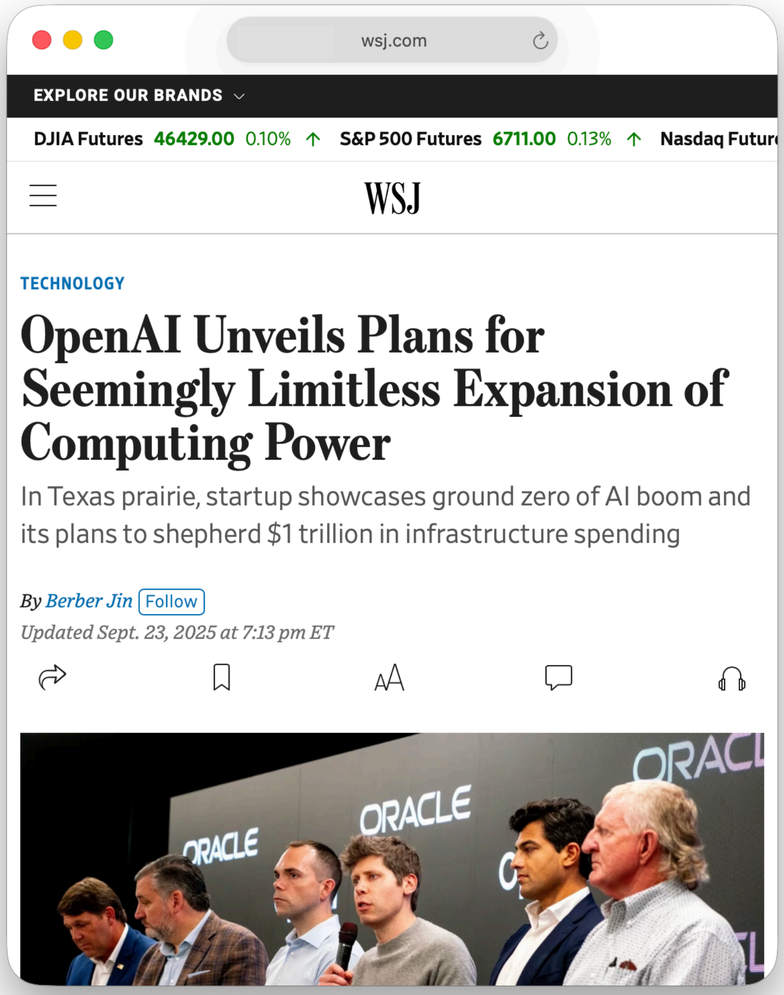

Source: OpenAI, October 28 2025

“Circular revenue”

Without its own cashflows, OpenAI partners with Nvidia, Oracle, Softbank, petrodollars…

OpenAI is buying Nvidia chips with Nvidia’s cashflow…

Which comes from the hyperscalers…

and using Nvidia’s cash to turn AMD into an Nvidia competitor, and pay Broadcom to make its own chips…

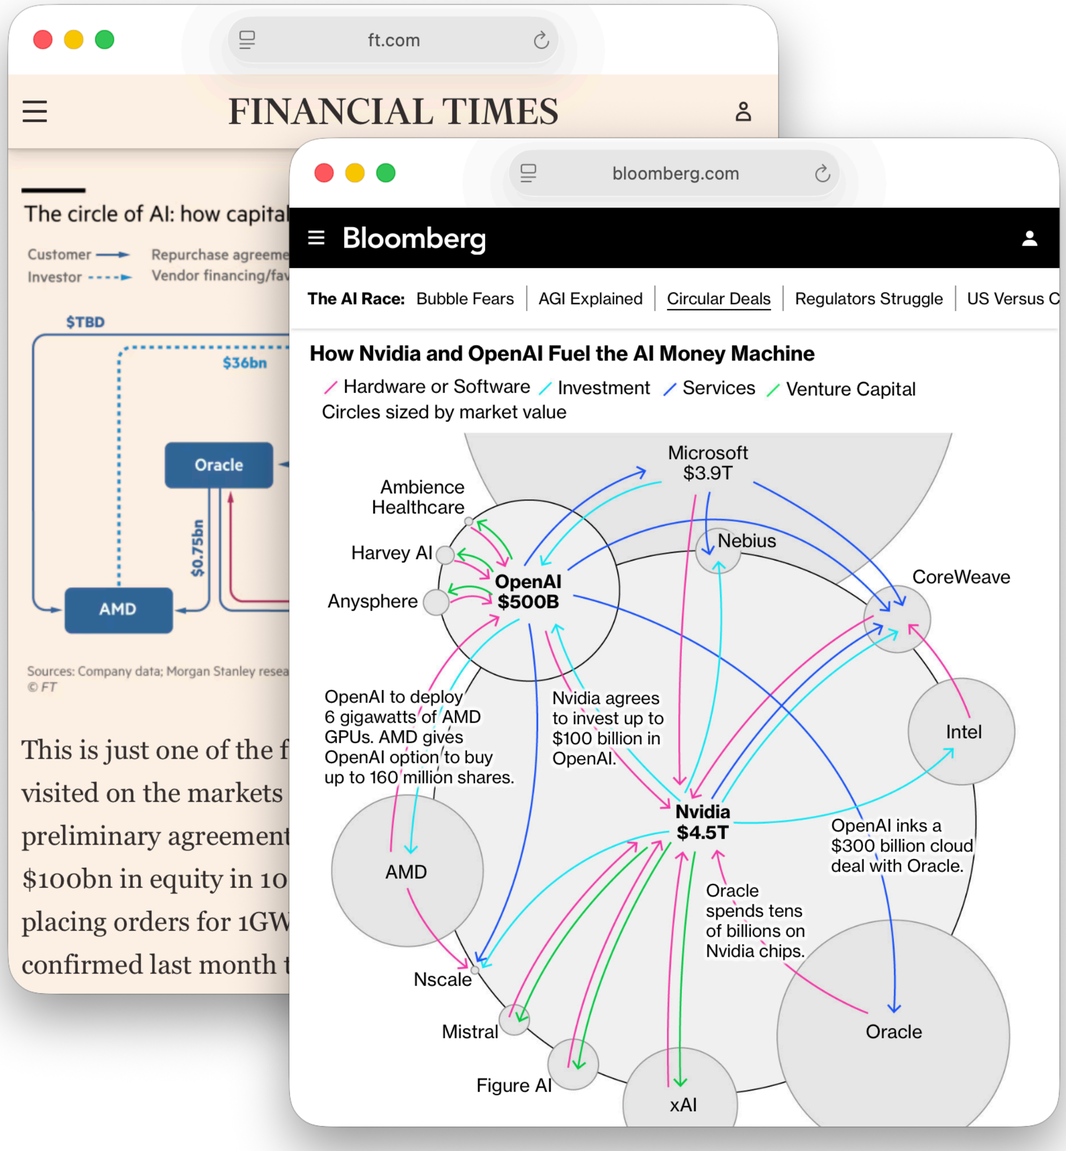

Rational actors?

What would you do if your company was sitting on a bubble?

Nvidia has $77bn of TTM FCF* and TSMC can't keep up with demand	→	Use your excess cash to buy demand, FOMO and platform lock-in
OpenAI has mindshare and expensive stock, but no platform, infra or moat	→	Swap your paper for hard assets and market position
Oracle is a cash-generative legacy business losing share to cloud and now to AI	→	Gear up and burn your way into the new thing?

*Source: Nvidia

Yes, but where has all this money got us?

After three years, lots more science and engineering, but no real clarity on the shape of the market

Models still improving

Far more models, China, OSS

Lots of new acronyms

No apparent moats

No clarity on product or value capture

Far more models

Every week - new models, new (problematic, gamed, saturated) benchmarks, new acronyms

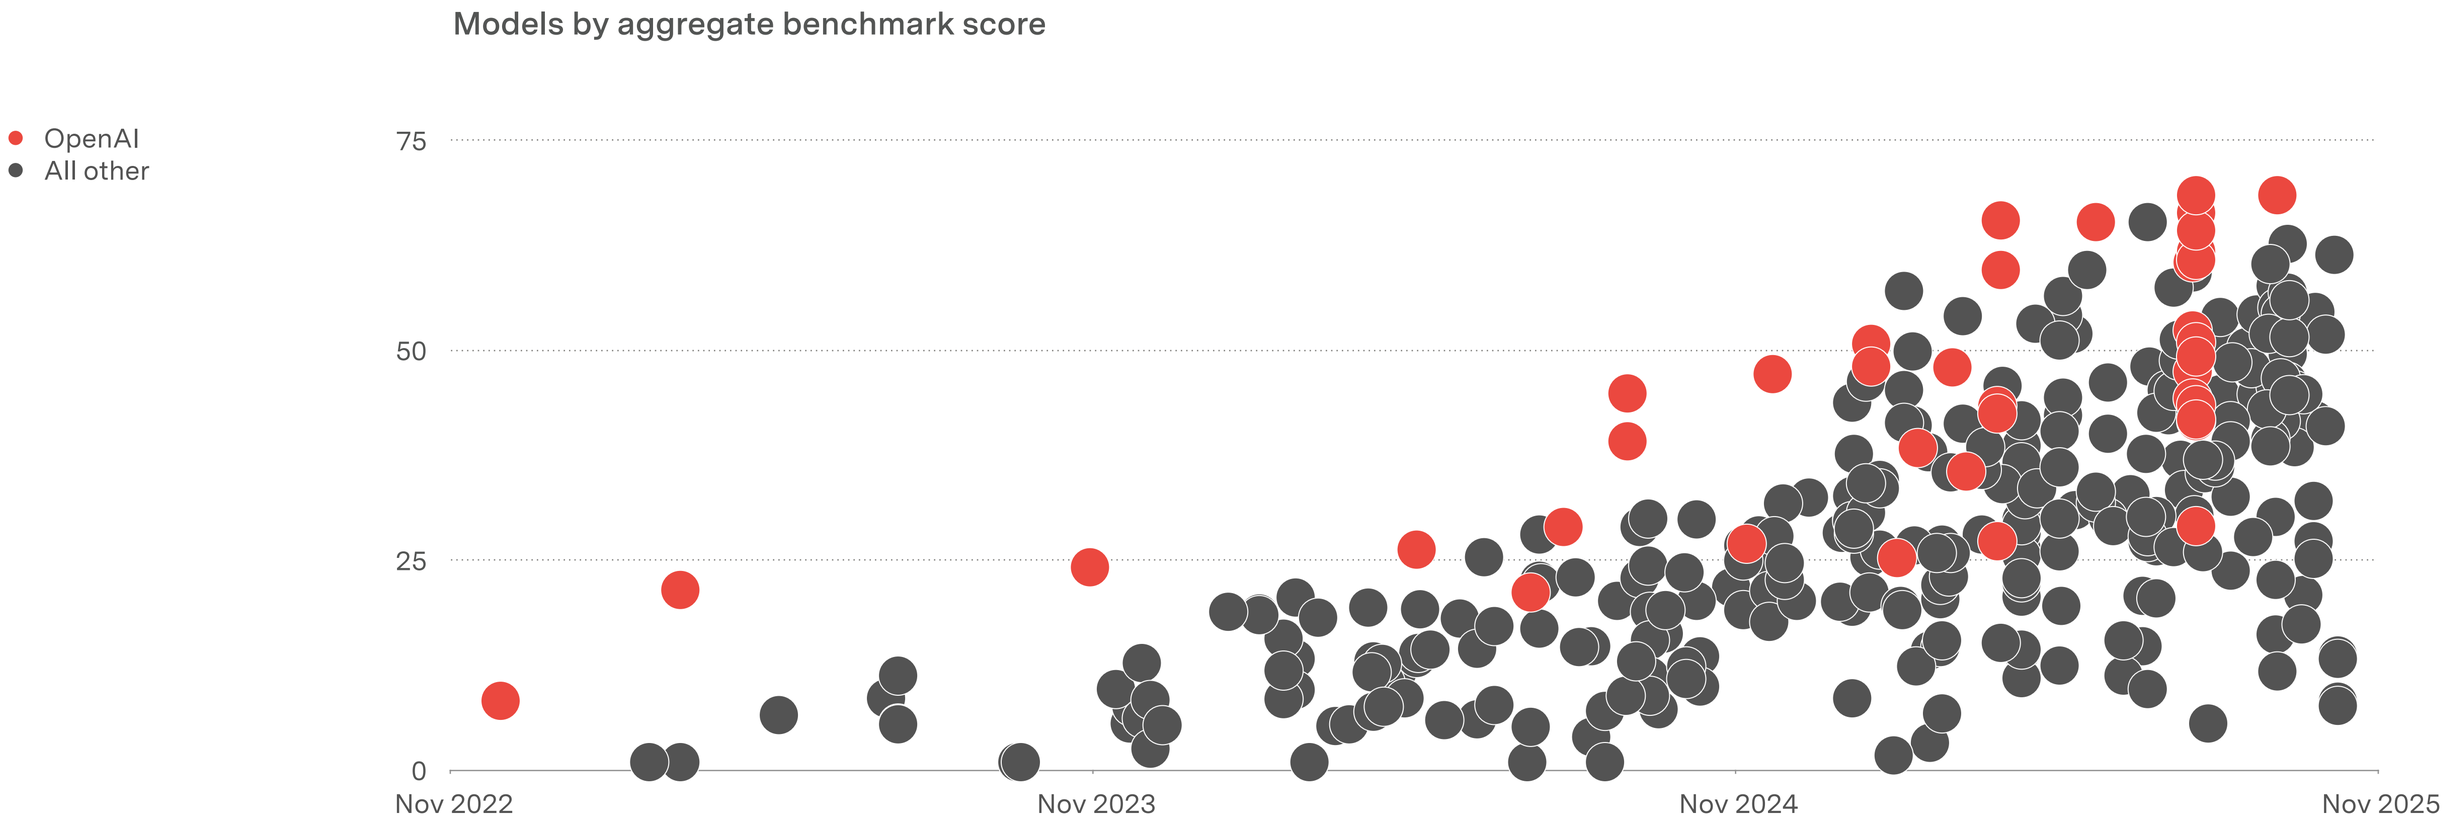

Source: ArtificialAnalysis

Models converge and leaders change weekly

Dozens of (saturated) benchmarks to choose from, but on the most general, the leaders are very close

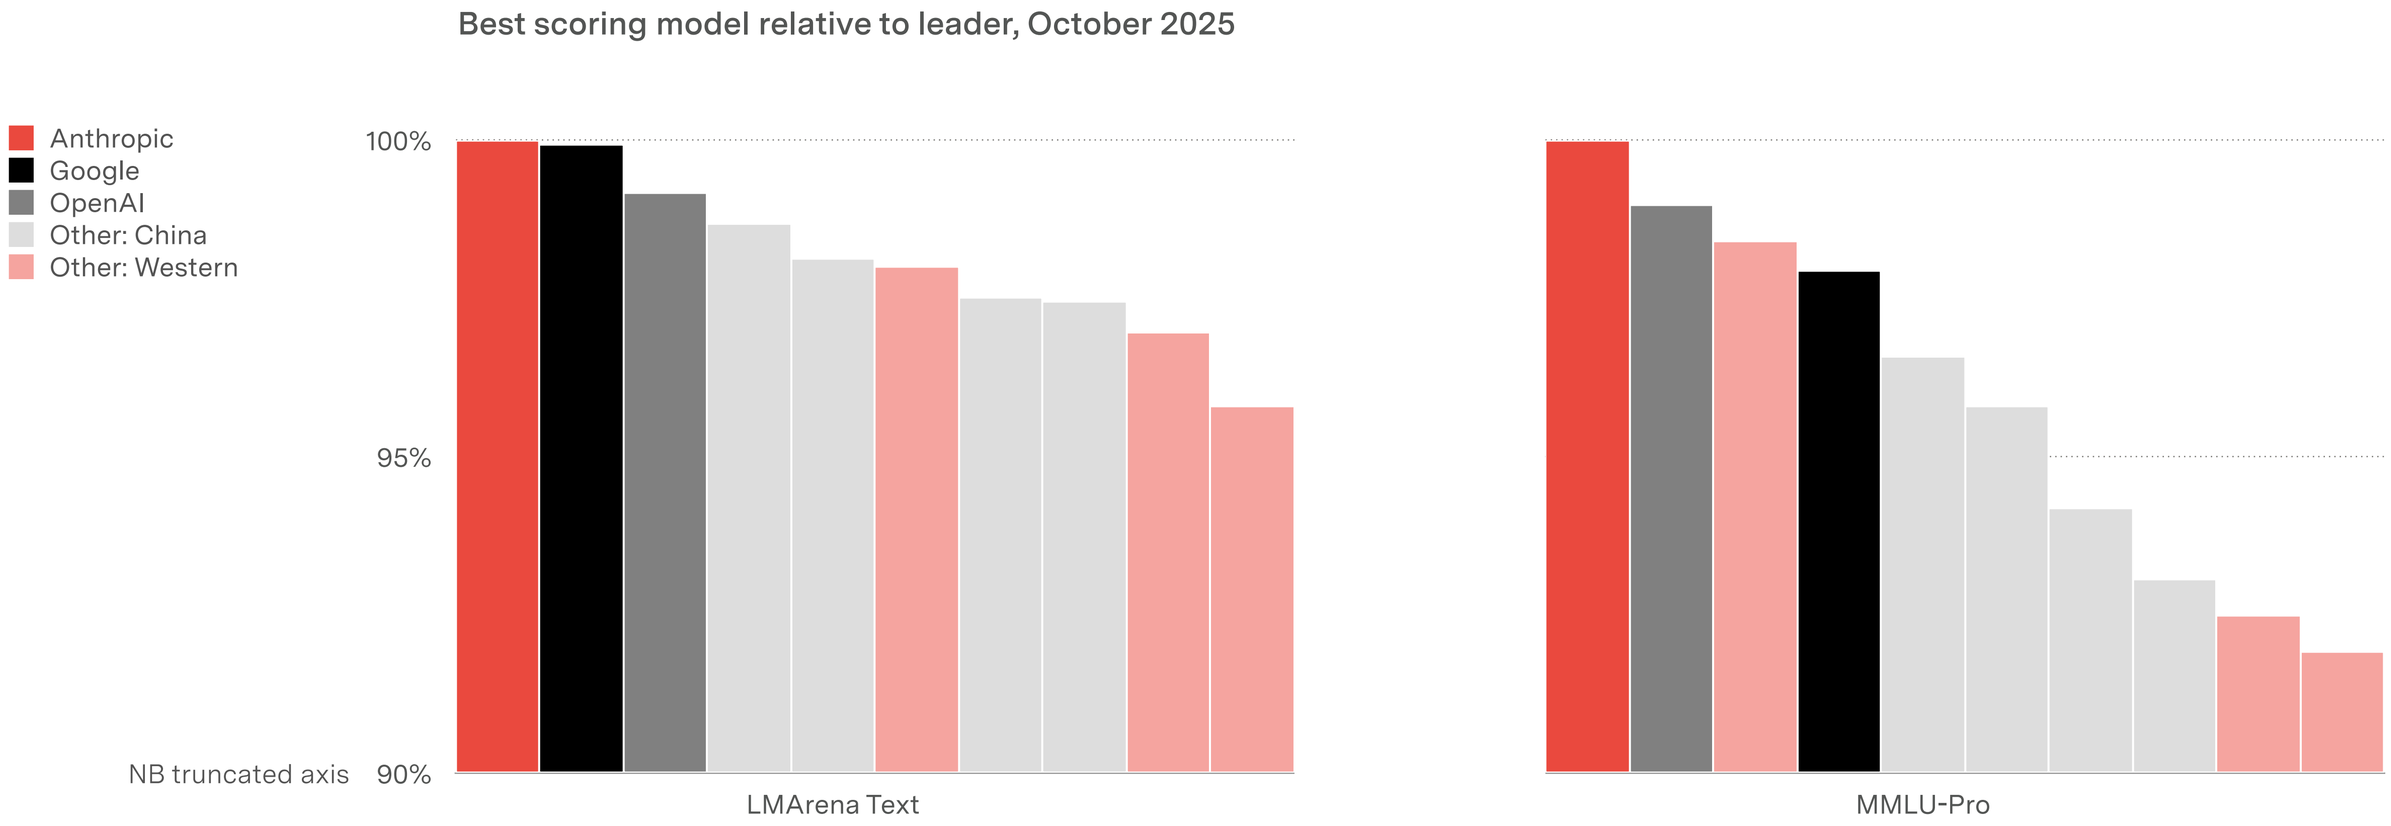

Source: LMArena, Artificial Analysis

Tech versus brand versus distribution?

The models may be close to commodities (especially for general use), but market position is not, so far

Weekly active users of generative AI tools as share of the population, June 2025

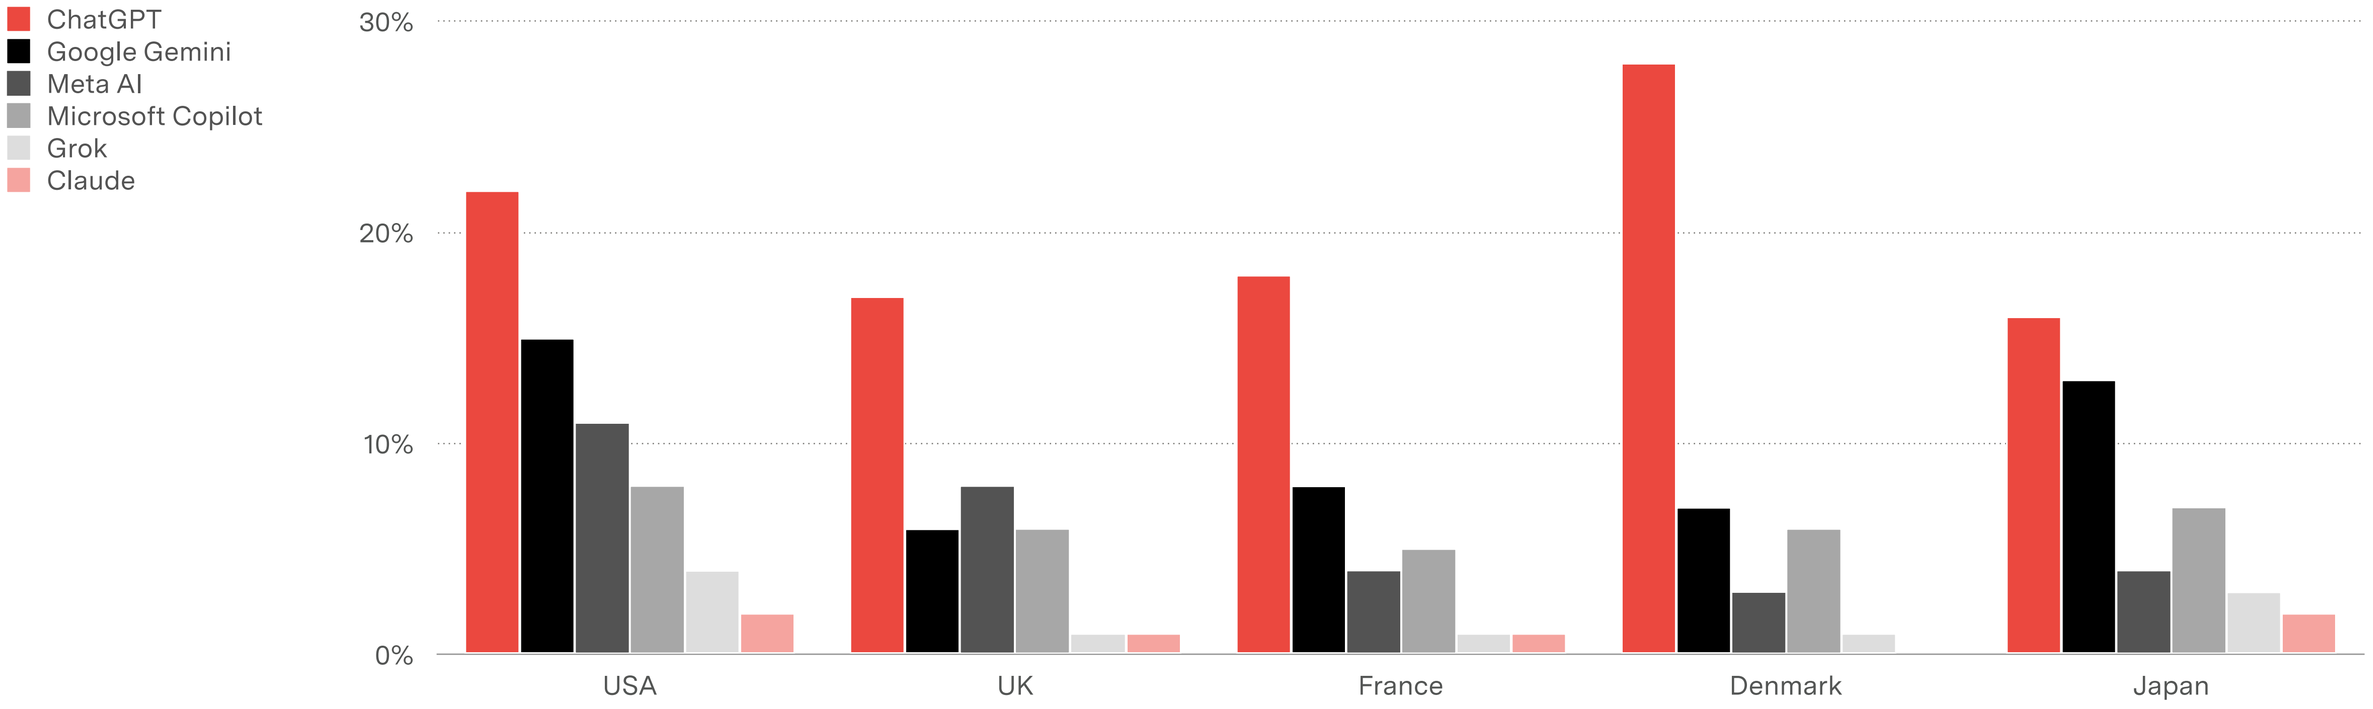

Source: Reuters Institute

“Everyone is already using this!”

800m weekly users, but apparently only 5% are paying - and why announce WAU and not DAU?

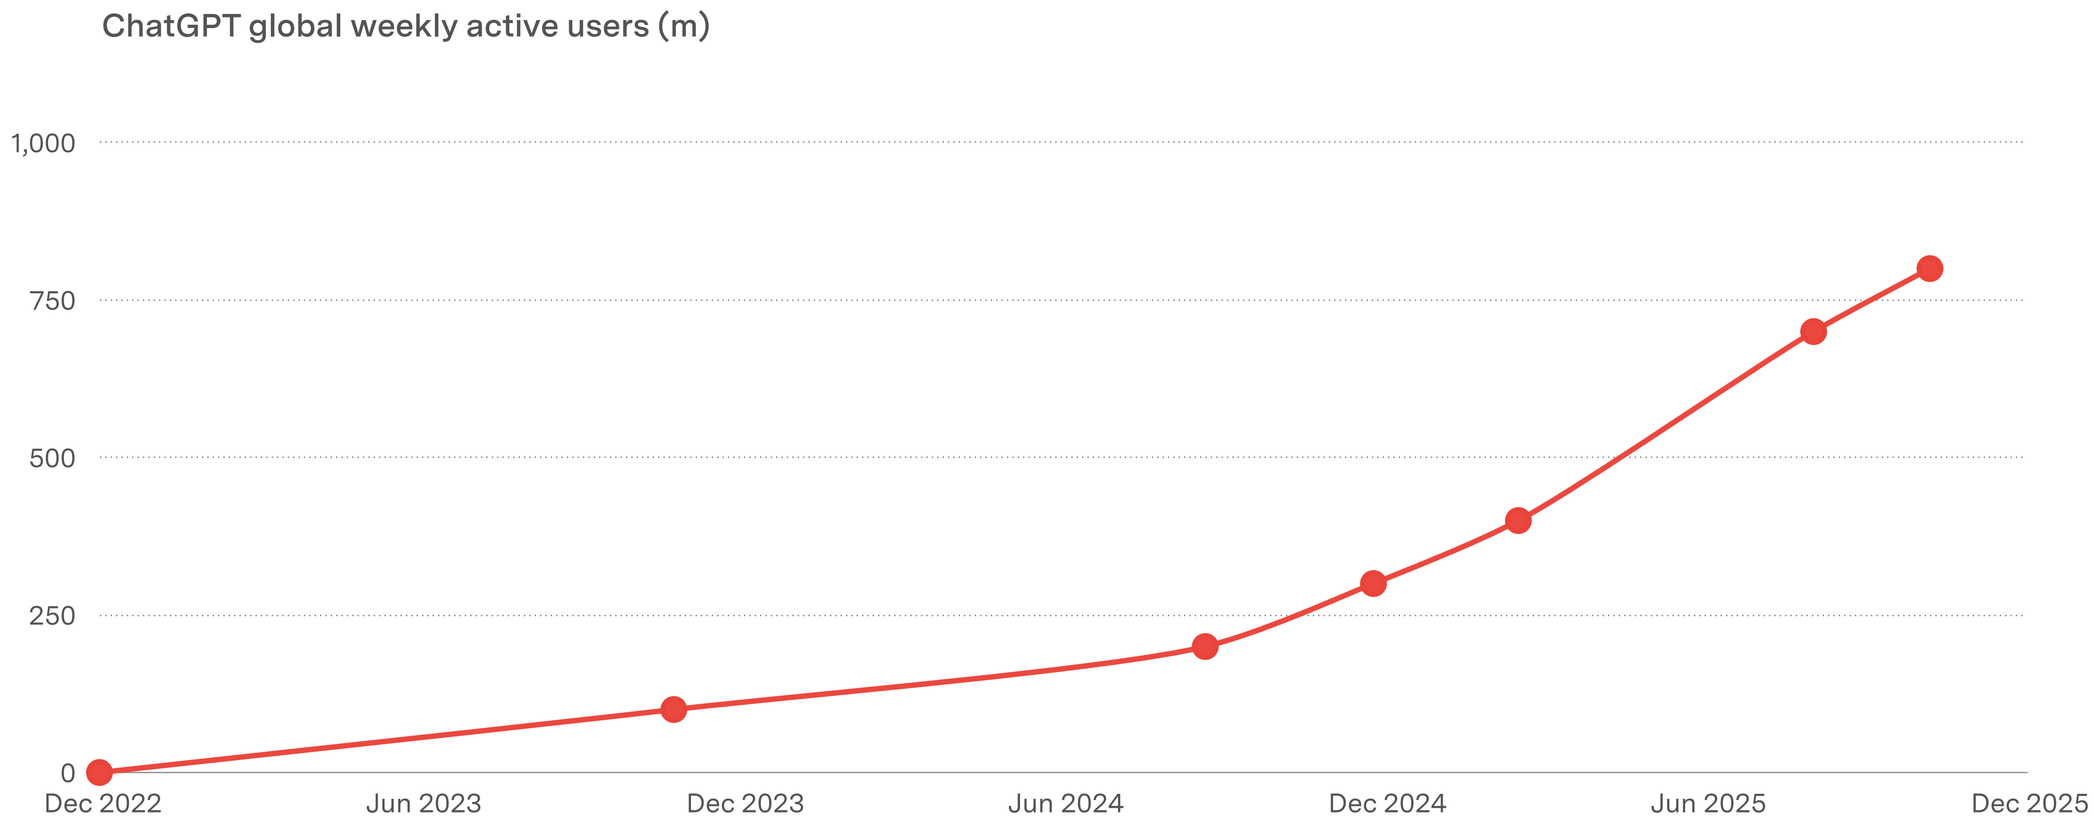

Source: OpenAI. NB: round numbers reported at scheduled events

Still more experimentation than daily use

So far, many more people use chatbots occasionally than make them part of their daily lives

Consumer generative AI chatbot use: USA

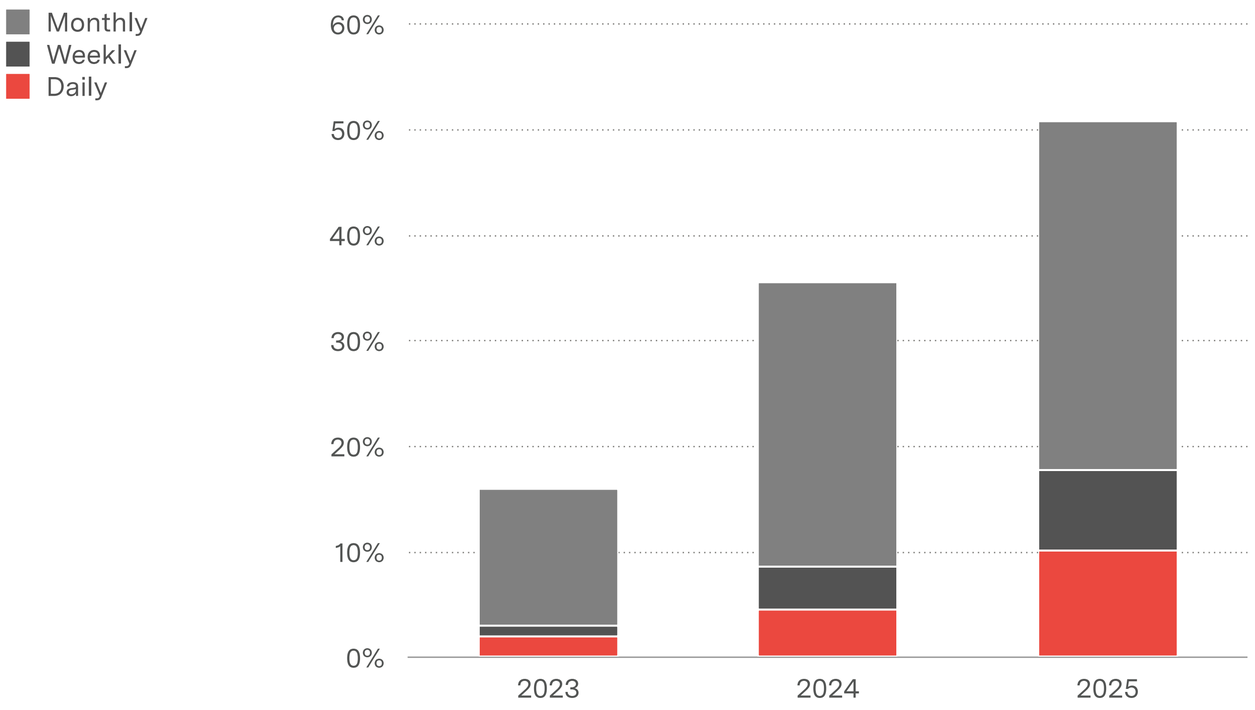

Consumer generative AI chatbot use: UK

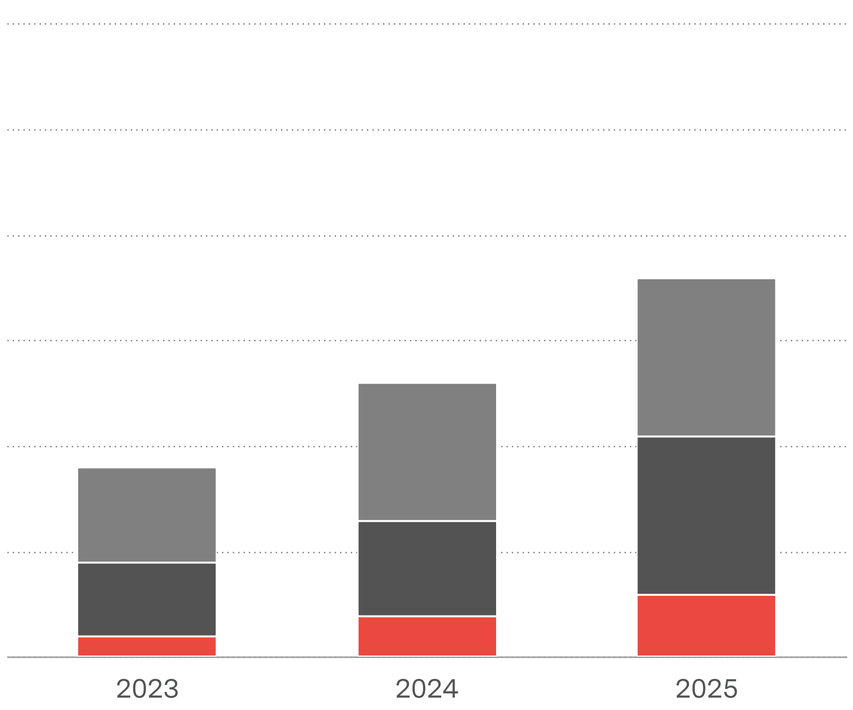

Source: Deloitte. Surveys conducted in June.
Defined as “Use of a purpose-built GenAI tool, like ChatGPT, Microsoft Copilot, Google Gemini, etc”

Most data shows the same picture

Surveys are early, scattered and inconsistent, but an engagement gap seems clear

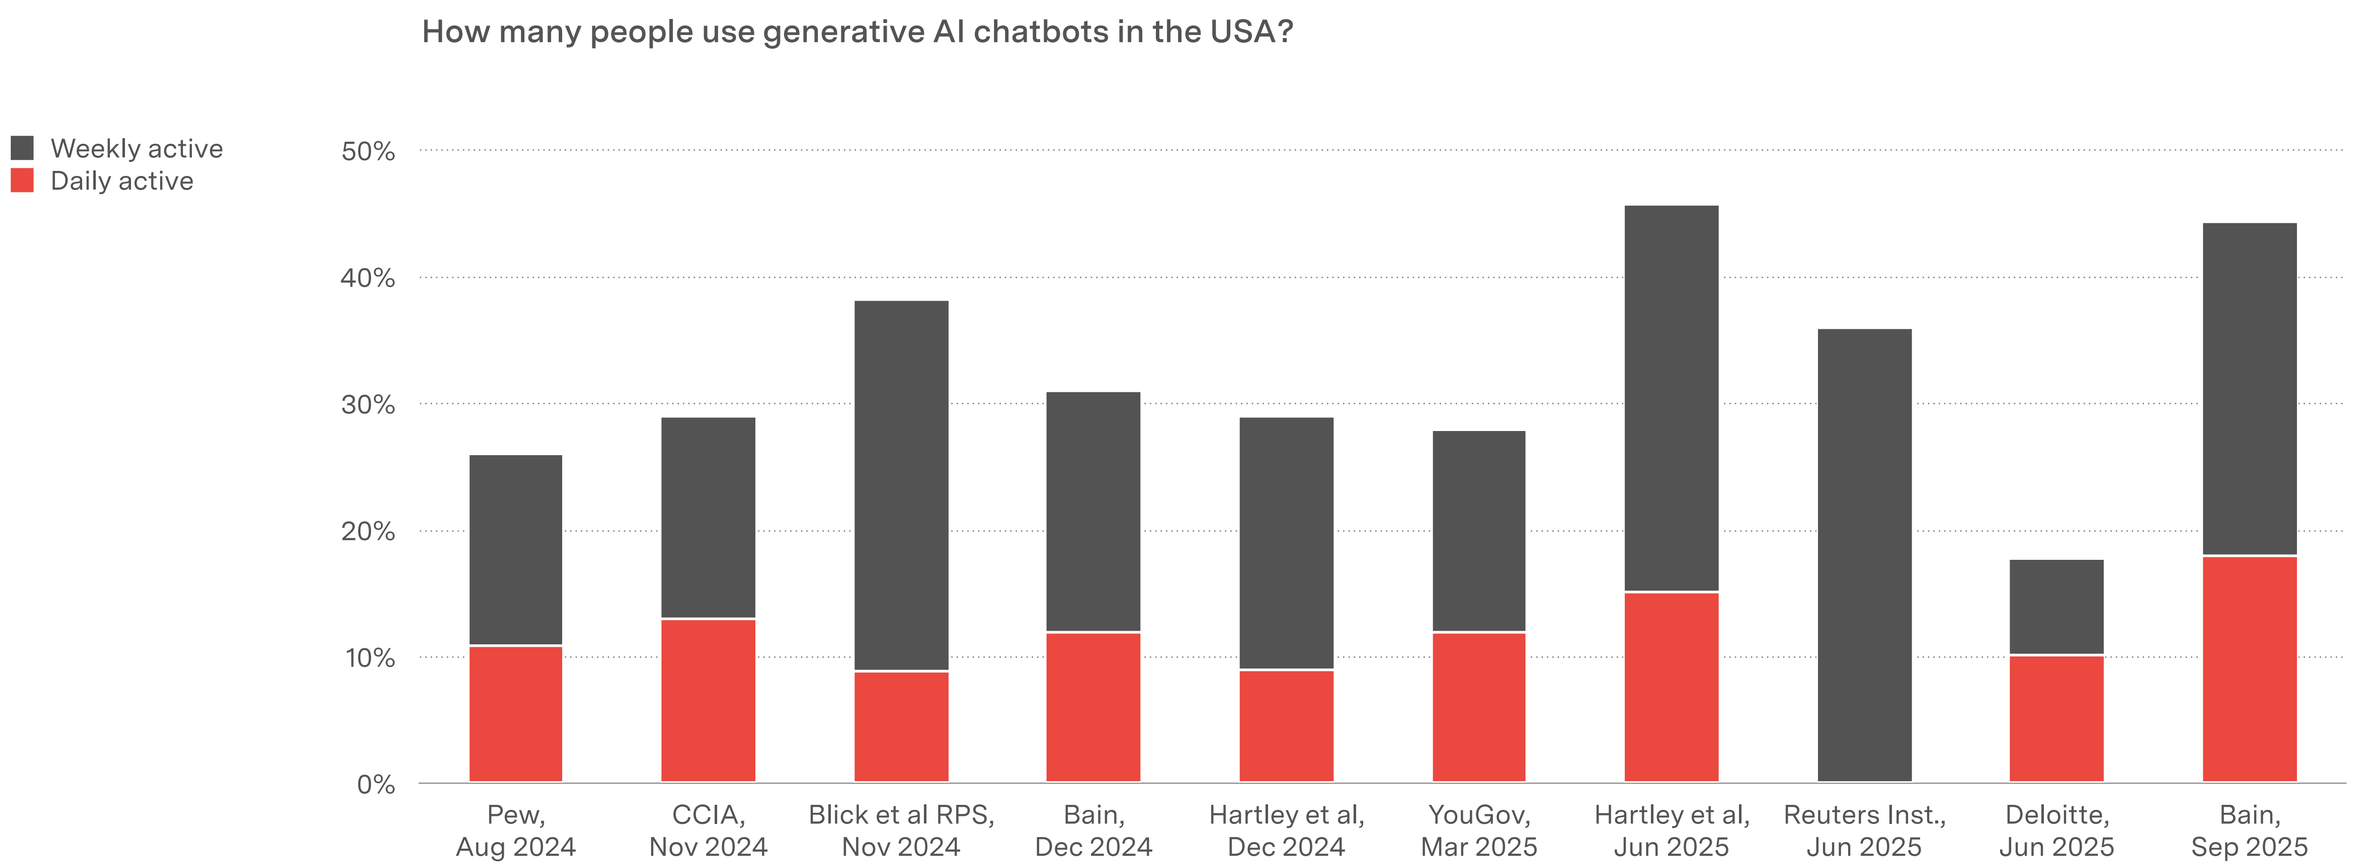

Source: Consumer surveys conducted as of date shown

Is this just early? Or a harder problem?

Why do most users of ChatGPT only use it a little bit?

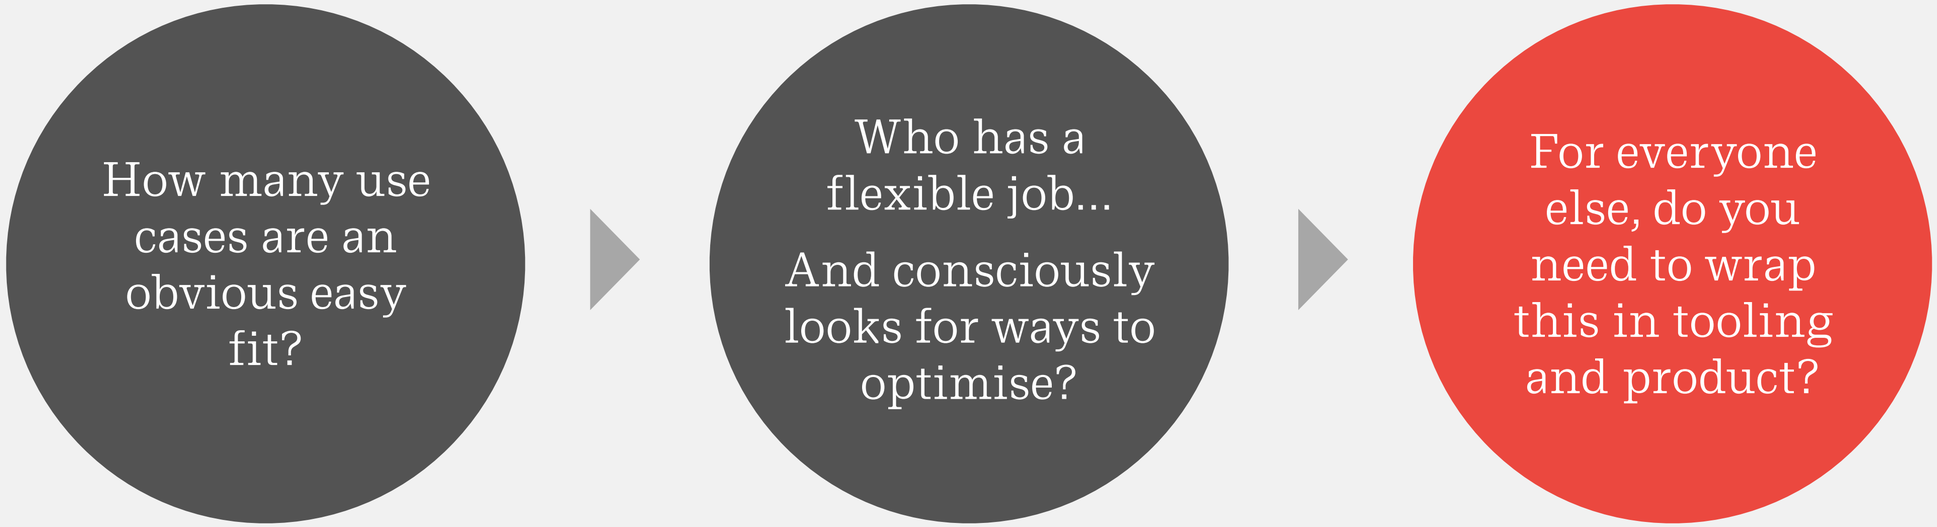

“People don’t know what they want until you show it to them”

“You’ve got to start with the experience and work backwards to the technology”

Steve Jobs

So where does a model lab compete?

Where is the value capture for a research-heavy, capital-intensive commodity?

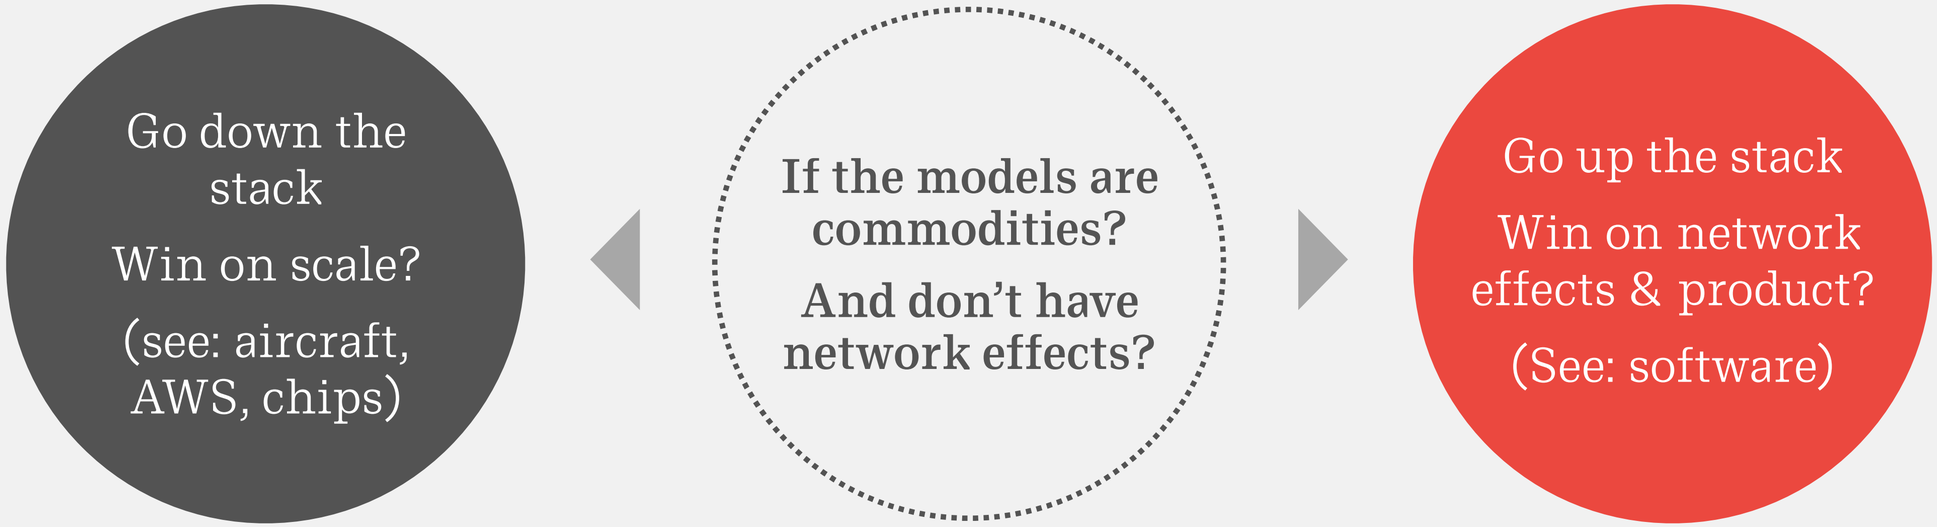

Microsoft’s shift away from network effects?

From competing on network effects to competing on access to capital?

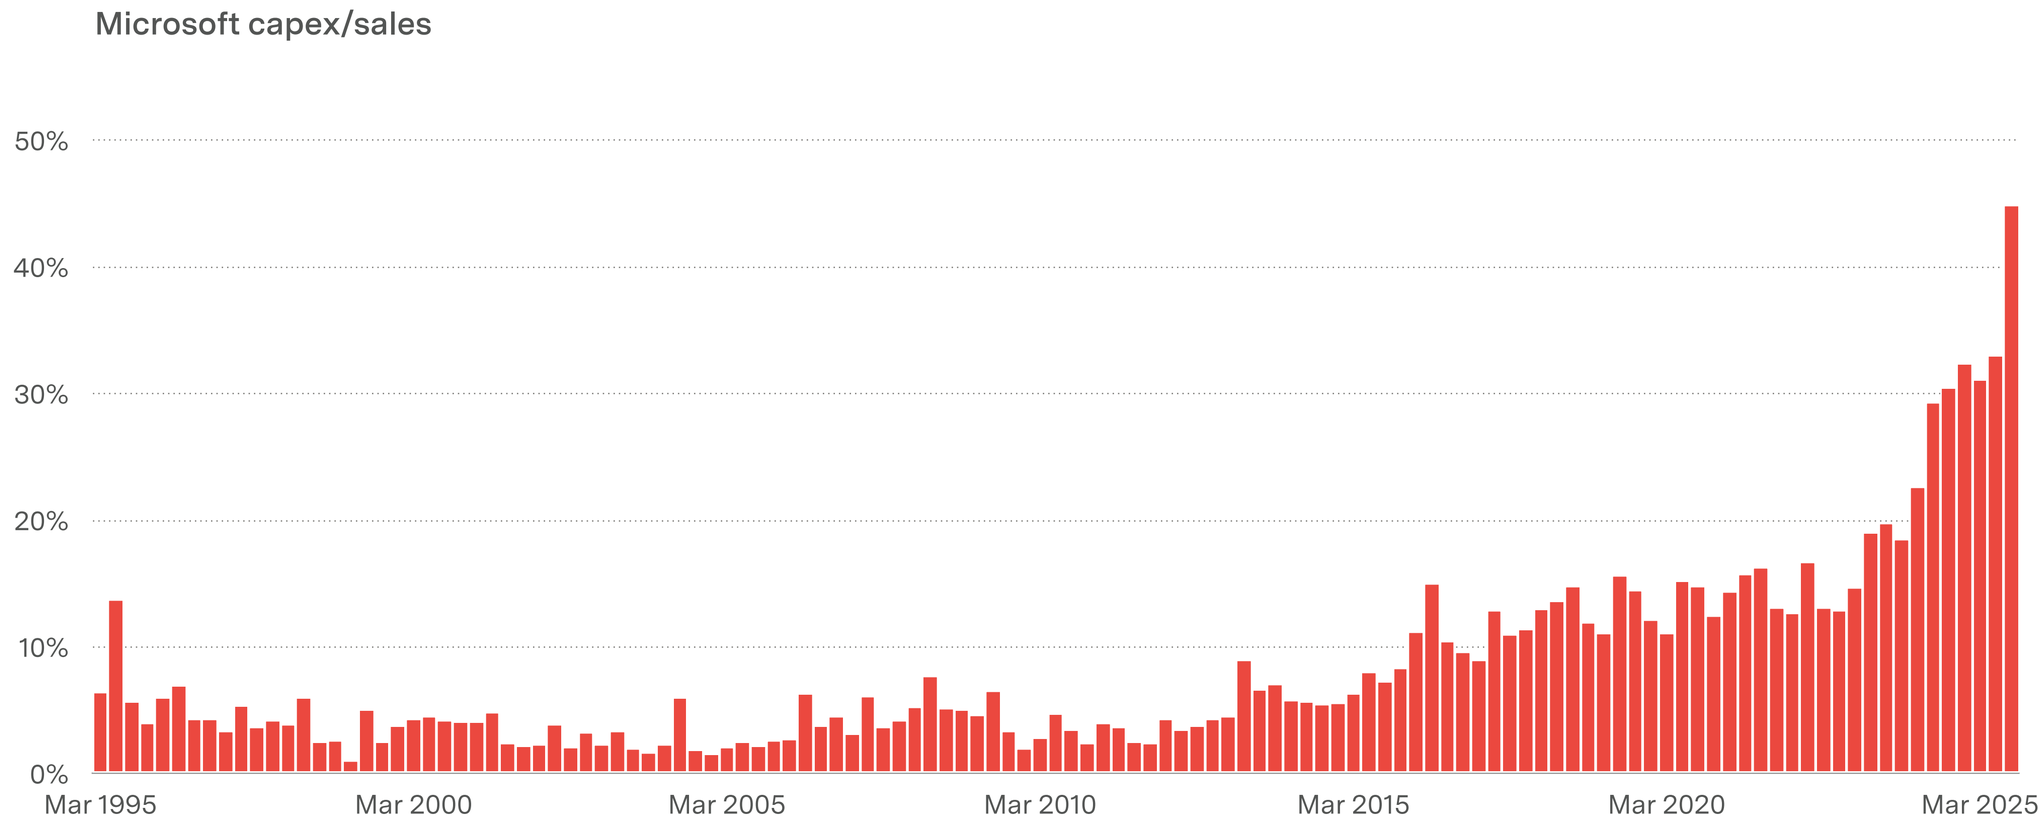

Source: Microsoft. Includes capital leases

For OpenAI, “yes!” to everything

Everything, everywhere, yesterday (on other people’s balance sheets), before the market slips away

“There are two ways to make money.
You can bundle, or you can unbundle”

Jim Barksdale

OpenAI bundles and unbundles use cases

What's the right experience? The right distribution? And why isn't it just the ChatGPT app?

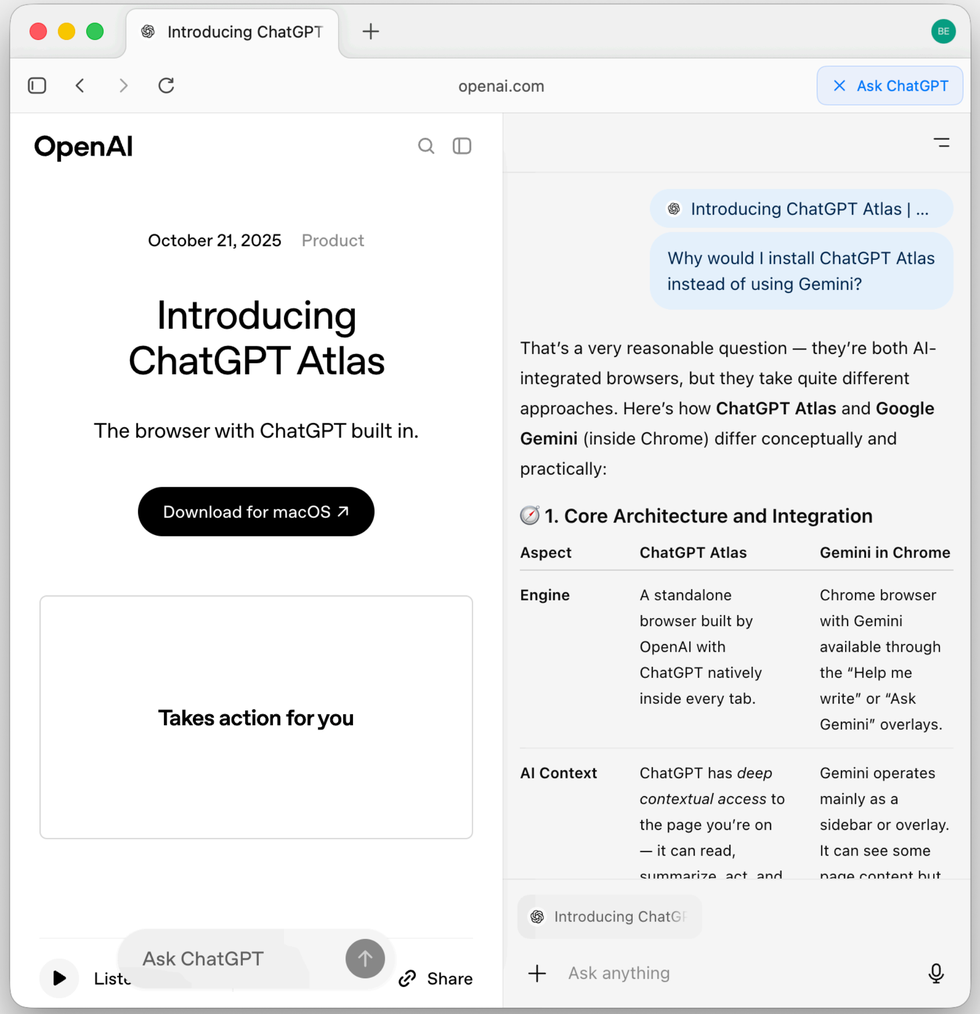

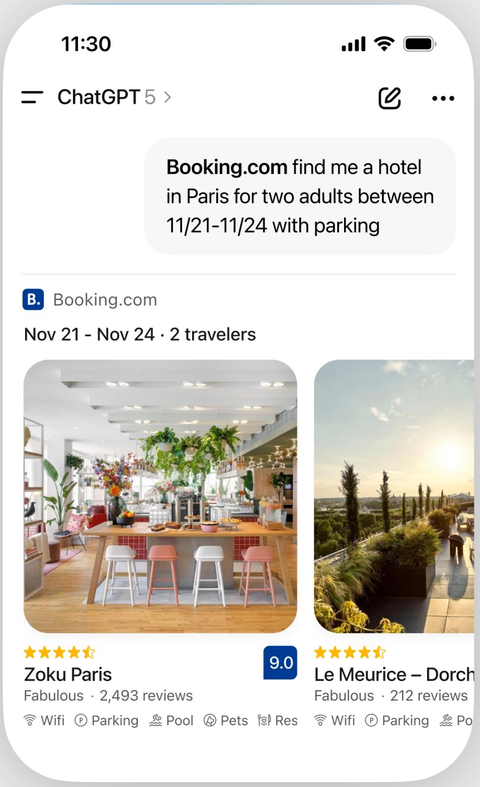

Startups exist to unbundle use cases

The coming wave of AI startups trying to unbundle Google, Excel, email and Oracle… and ChatGPT

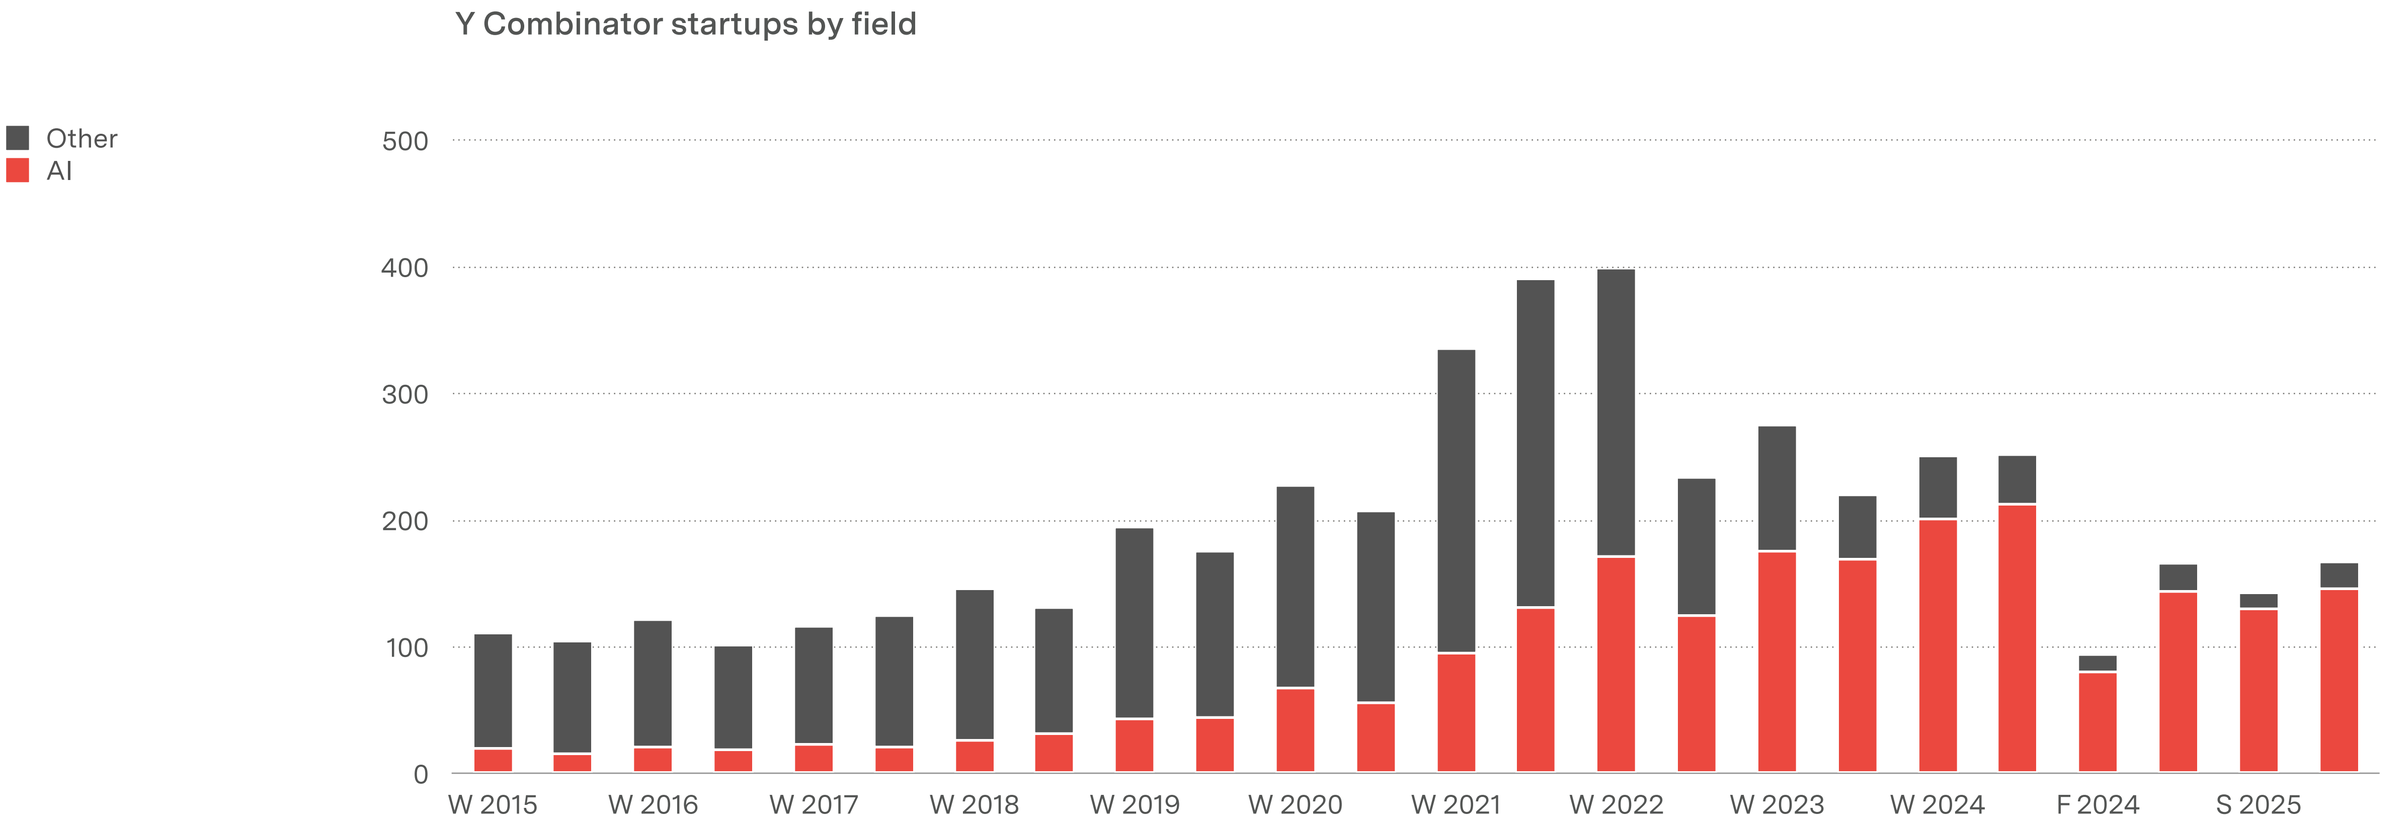

Source: Y Combinator

Where is the value capture?

If models are near-commodities, and we don't know the right product, where will the value be?

~~Best model?~~
Most capital?

Proprietary vertical data?

Distribution & GTM? Product? UX?

Building 'normal' software companies?

Outside tech

“What’s our AI strategy?” Well, what’s the pattern?

How do we *always* deploy new technologies?

So far, most successful use-cases are ‘absorb’

Where is it easy and obvious to use generative AI?

AI coding as the new AWS

“Vibe coding” as the new abstraction layer, after AWS, libraries, operating systems…

A new step change reduction in software creation costs

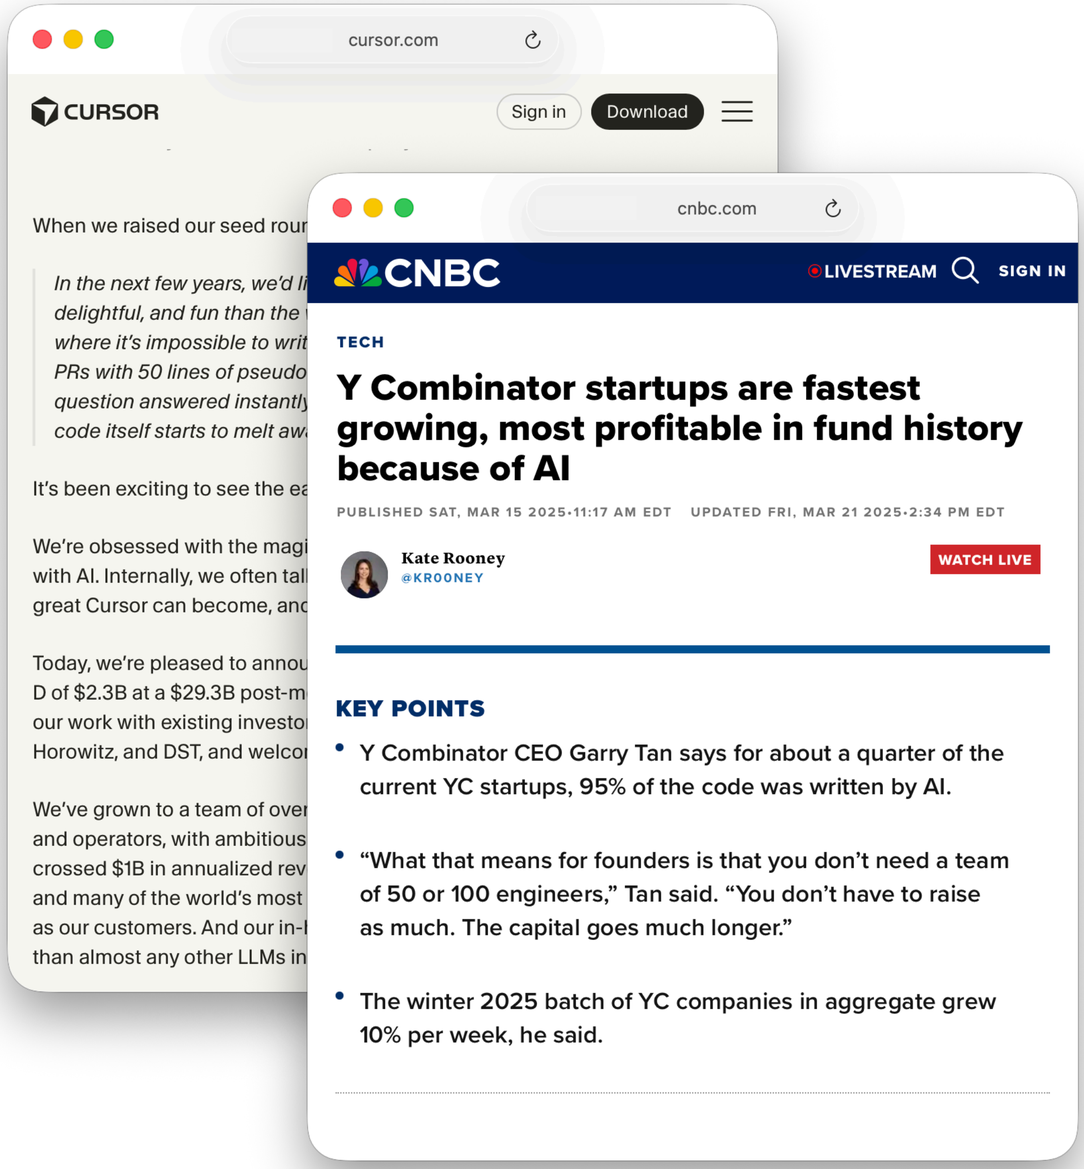

How do you know what to automate?

Step one: ask your systems integrator

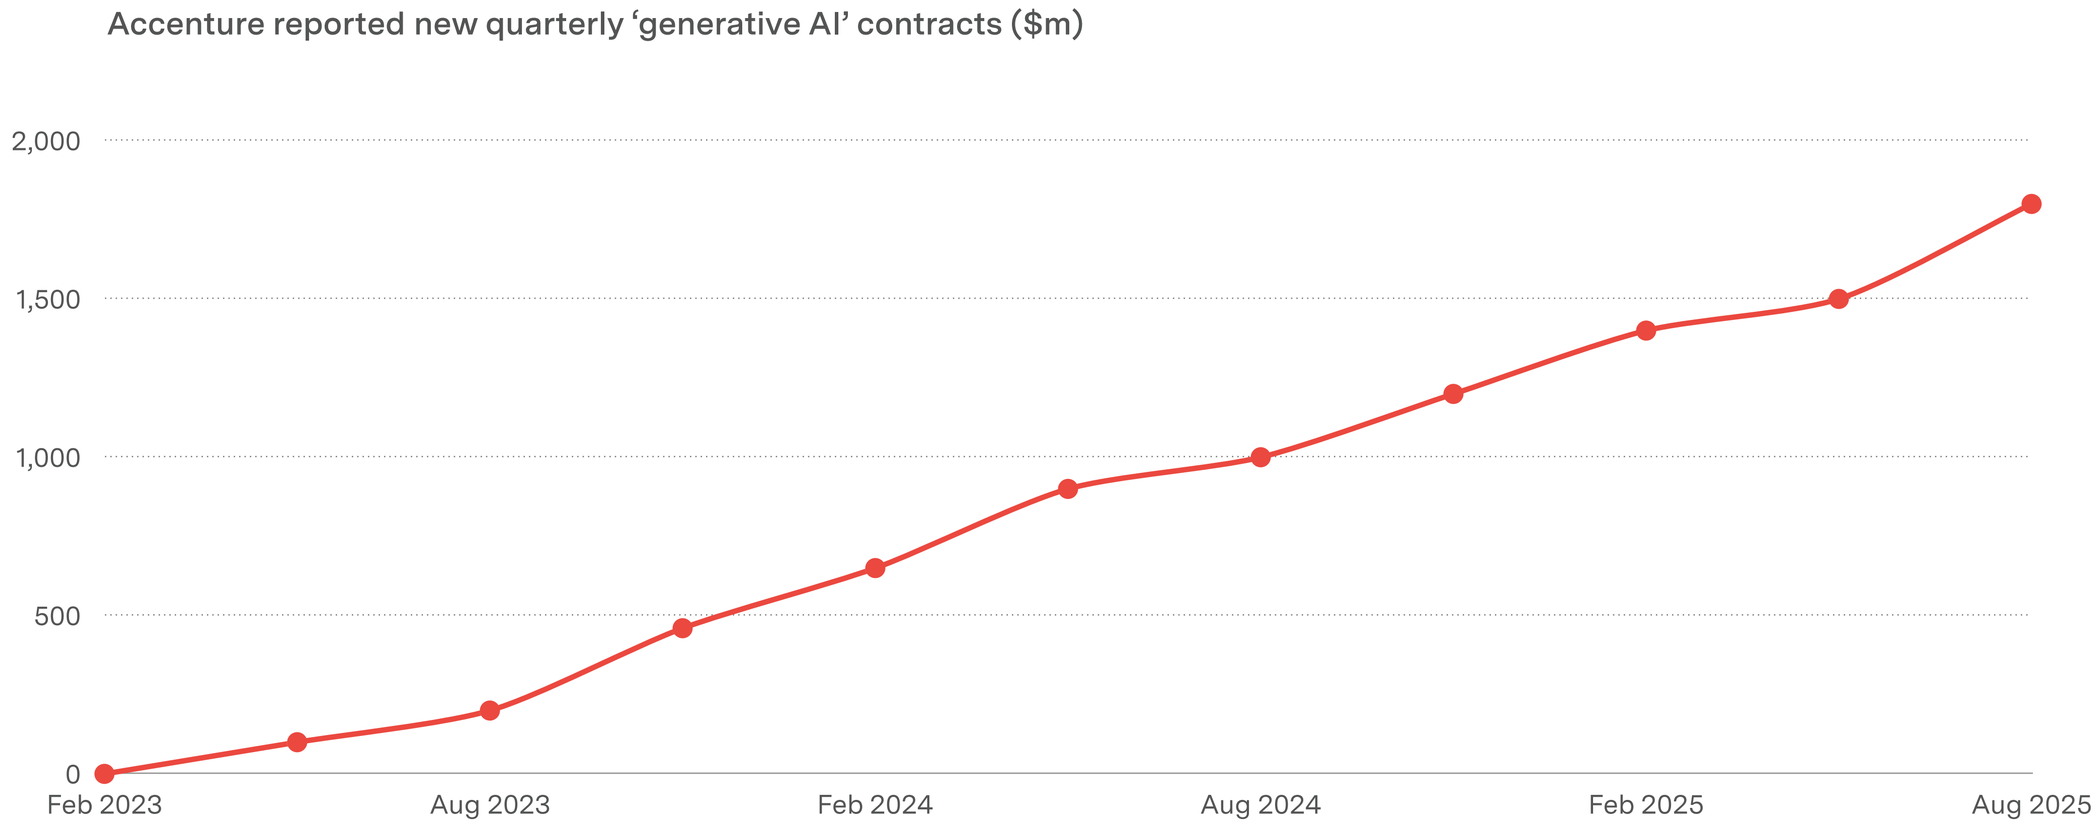

Source: Accenture

How do you know what to automate?

Step two: buy some SaaS from Dr Evil

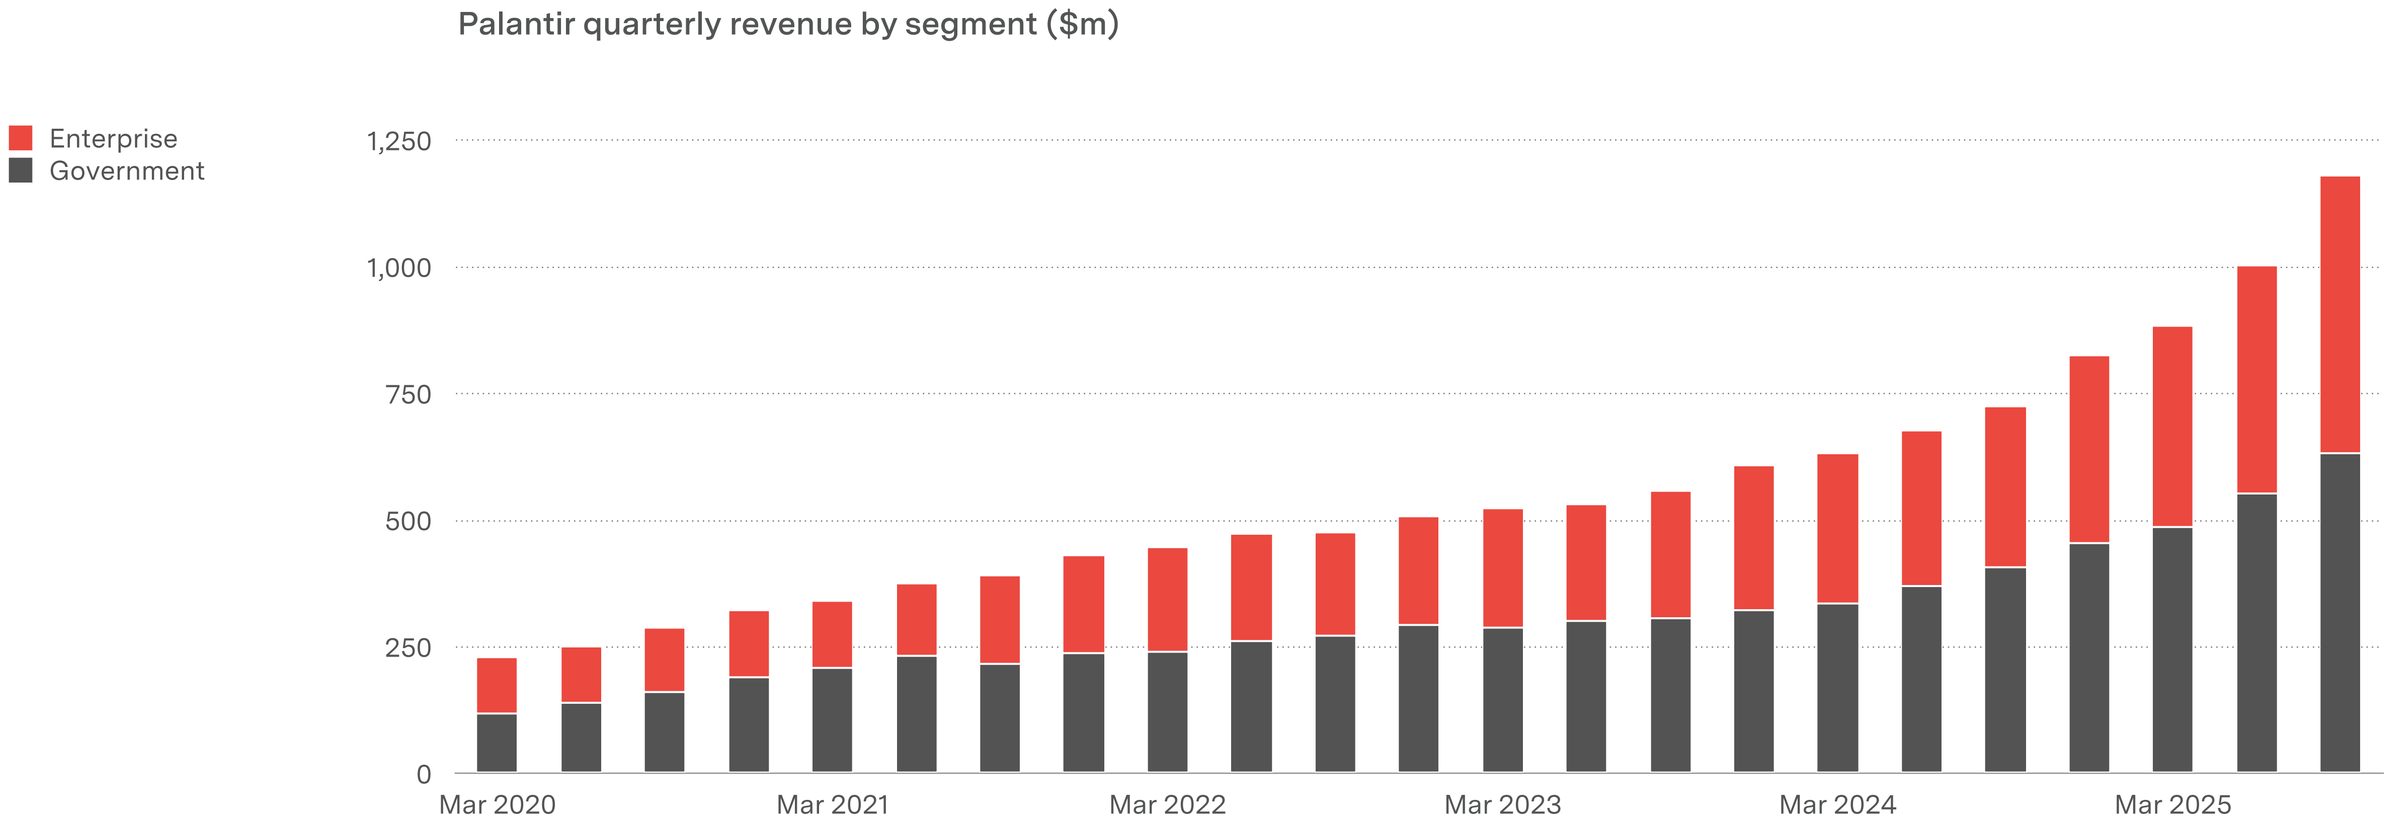

Source: Palantir

Pilots come first and deployment takes time

"Agentic!' is 2025's buzzword, but deployment takes longer

AI 'agent' use by business function, where Generative AI is already used. June 2025

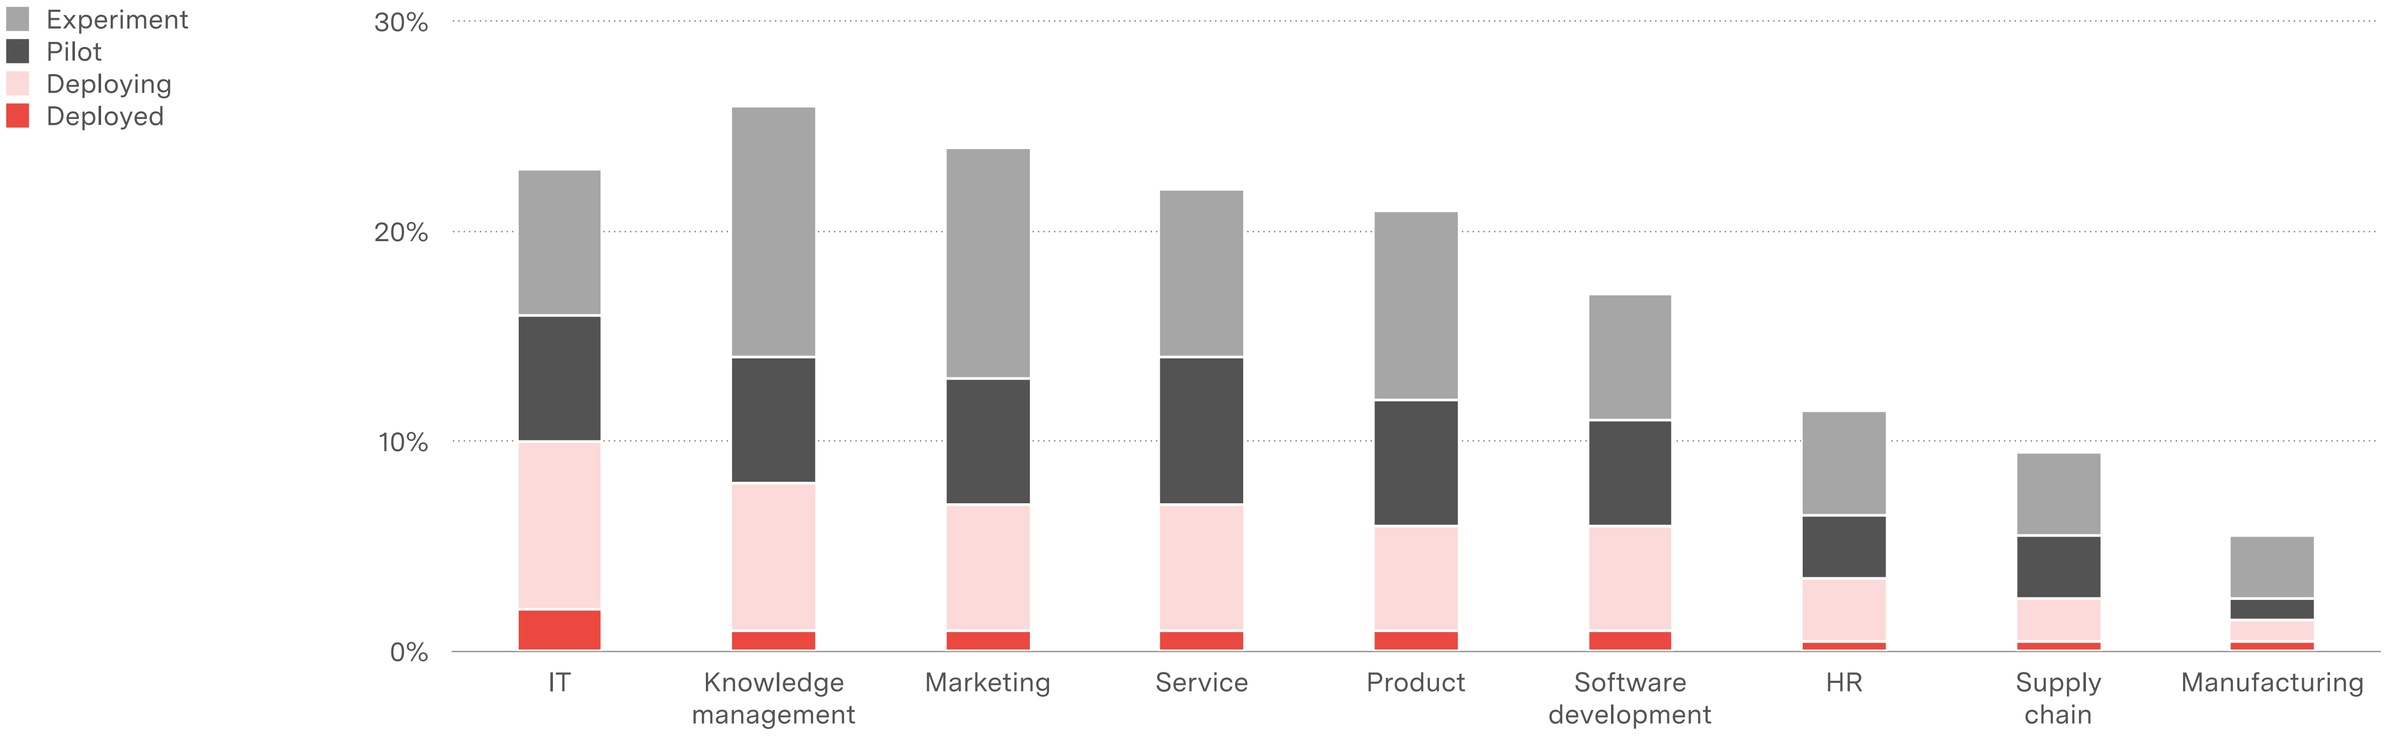

Source: McKinsey

Not everything works? Welcome to tech

“Why did our AI pilot fail?” That’s a CTO question, not an AI question

Security, privacy, IPR, error rates, legal

Data integration & legacy systems

Finding the right solution for the right people

The same issues as deploying any new tech

The future can take time

A quarter of CIOs have launched something - but 40% don’t plan anything until at least 2026

CIO expected timing for first LLM projects in production, September 2025

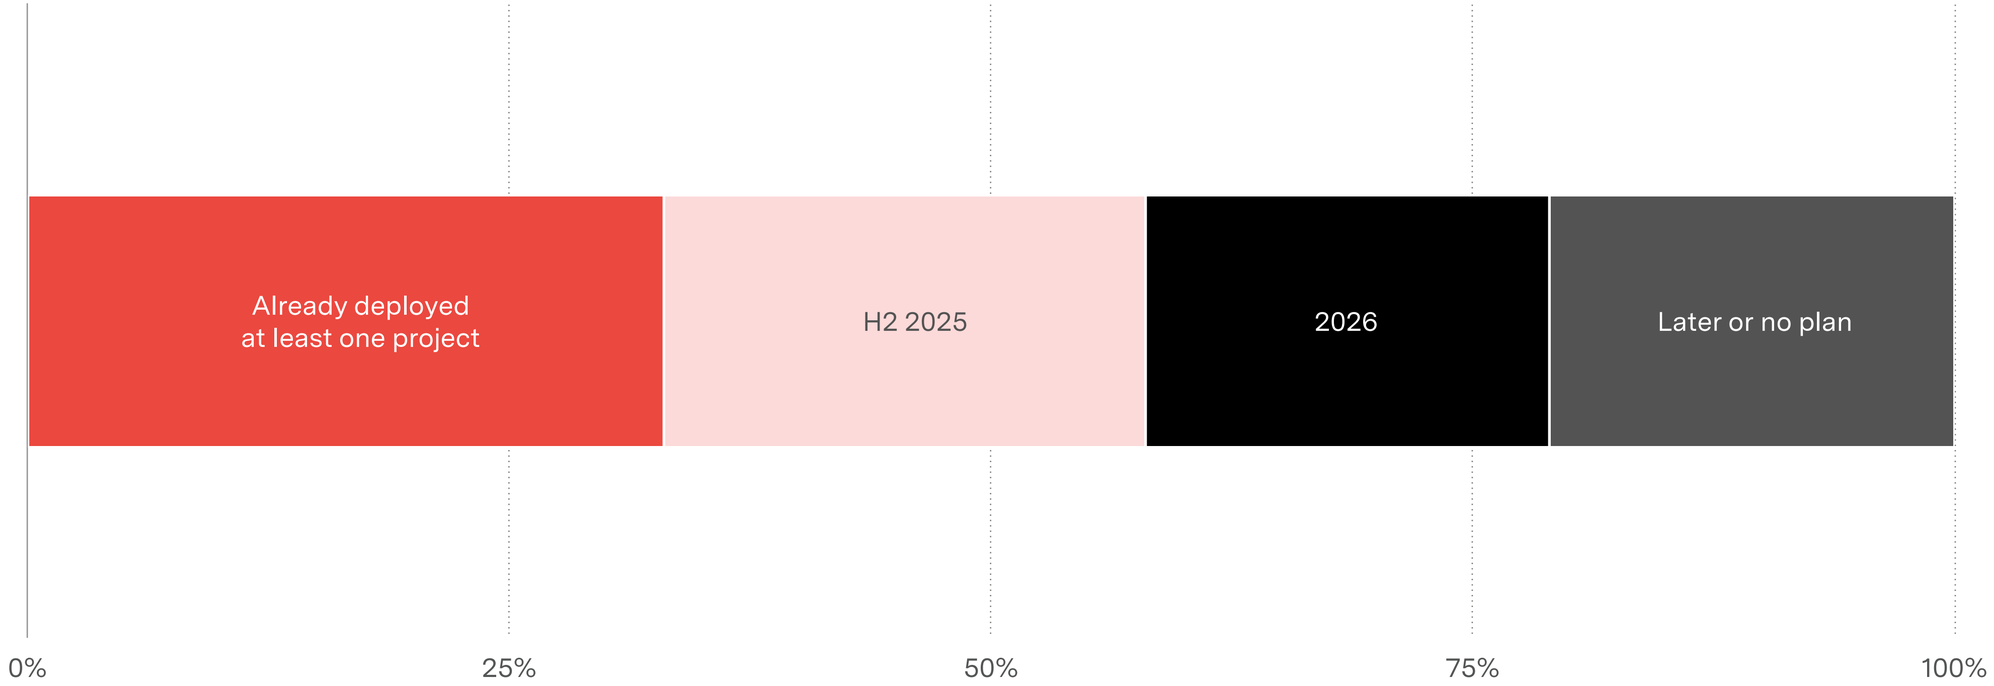

Source: Morgan Stanley CIO Survey

But the future always takes time

Cloud is old and boring - but still only 30% of workflows

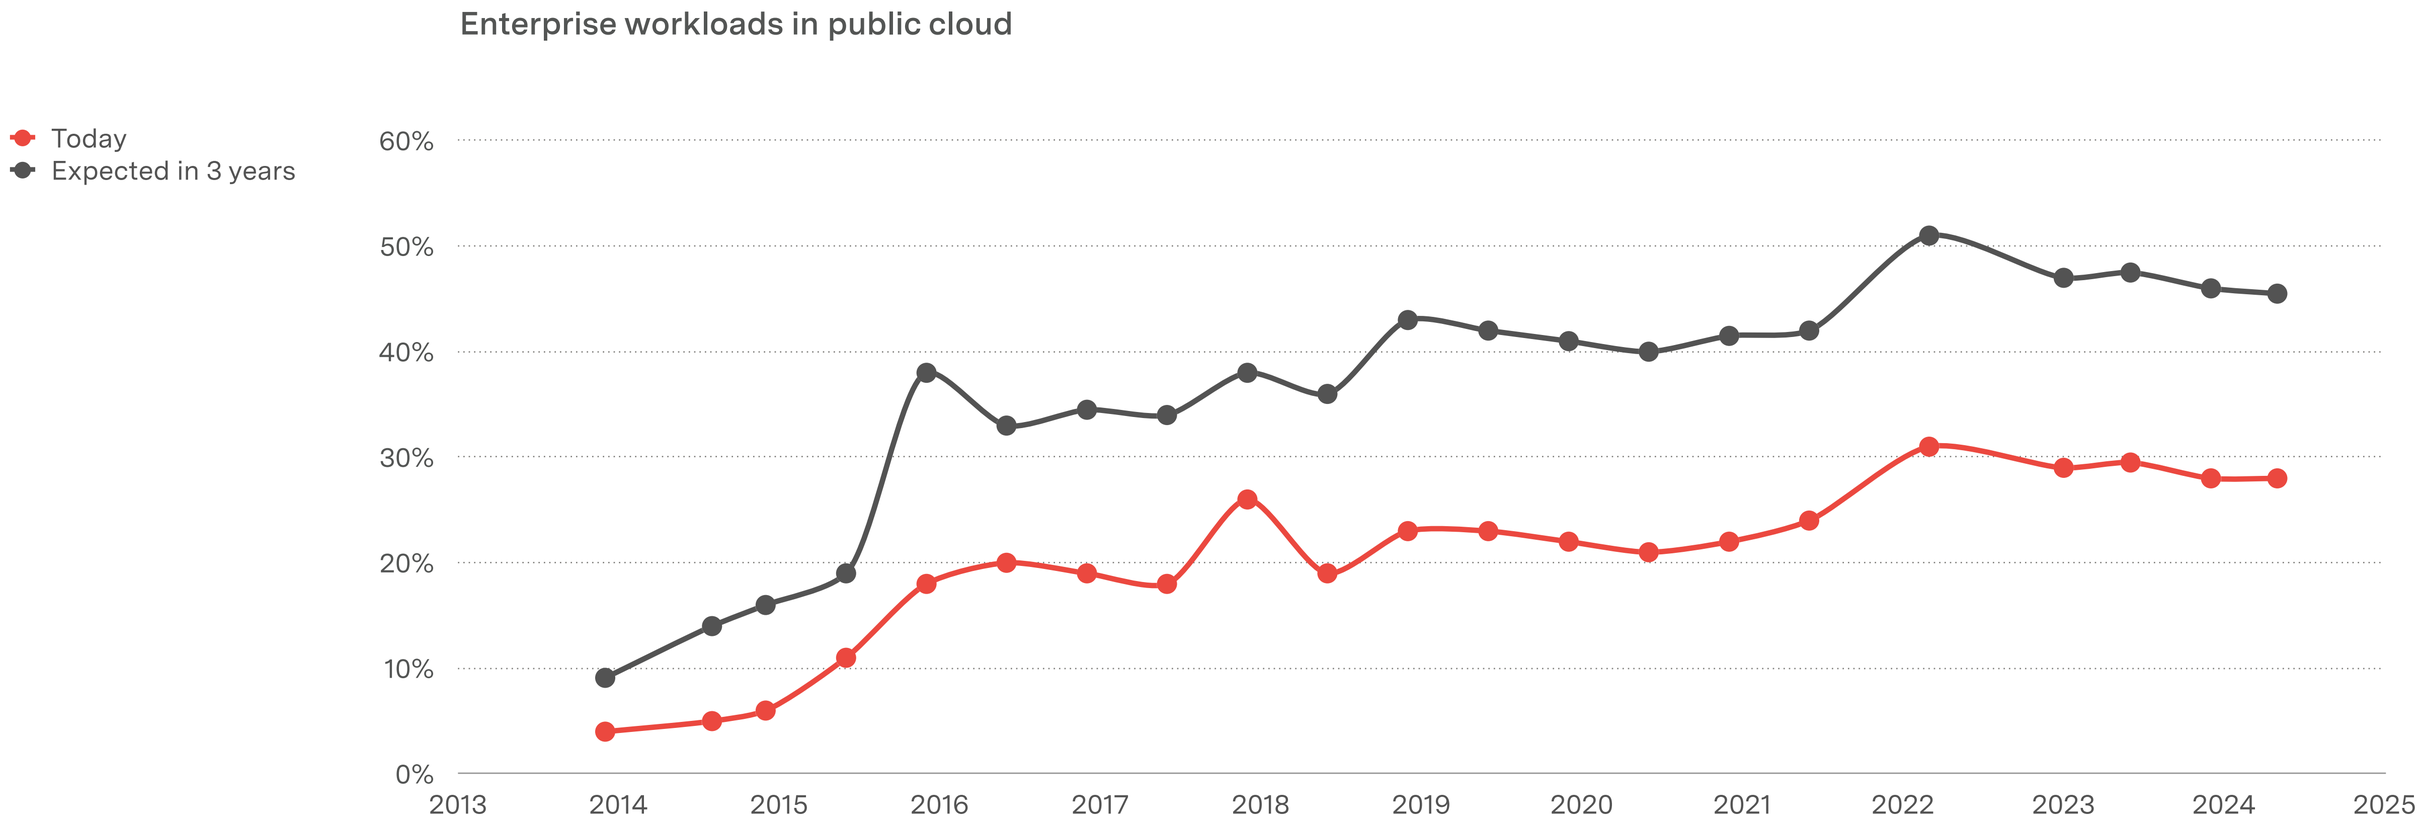

Source: Goldman Sachs CIO Survey

Sometimes ‘automation’ alone is a big deal

UPCs, barcodes and databases let retailers manage 5x more SKUs

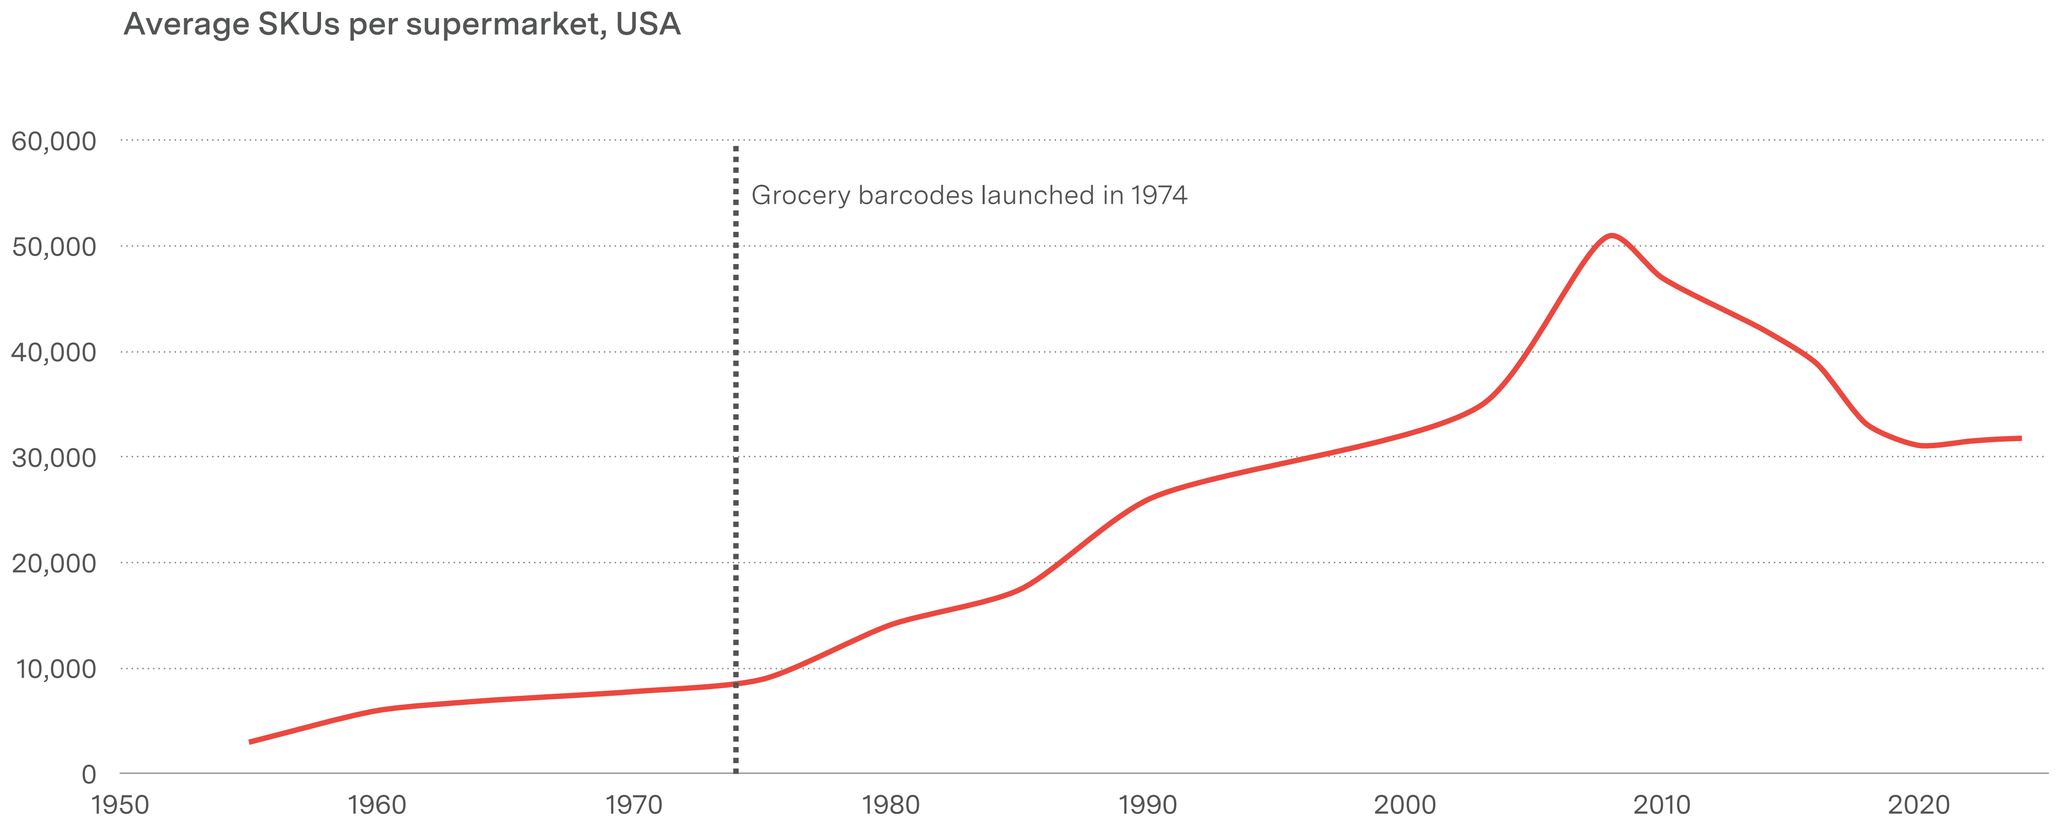

Source: FMI

But...

“We’ve all seen lots of AI presentations now, and we’ve deployed a bunch of stuff. Is that it? What’s next?”

F100 Retailer CMO, summer 2025

What next?

What comes after automating the obvious, easy things?

Where might we look for change?

What can LLM automation unbundle? What things did we not realise were bundles?

Online distribution unbundled physical assets

What do LLMs unbundle?

Internet unbundling created new aggregators

How do LLMs do that better?

“AI gives you infinite interns”

How do we use automation that makes ‘mistakes’?

We have no indication that error rates will go away, so where’s the human in the loop?

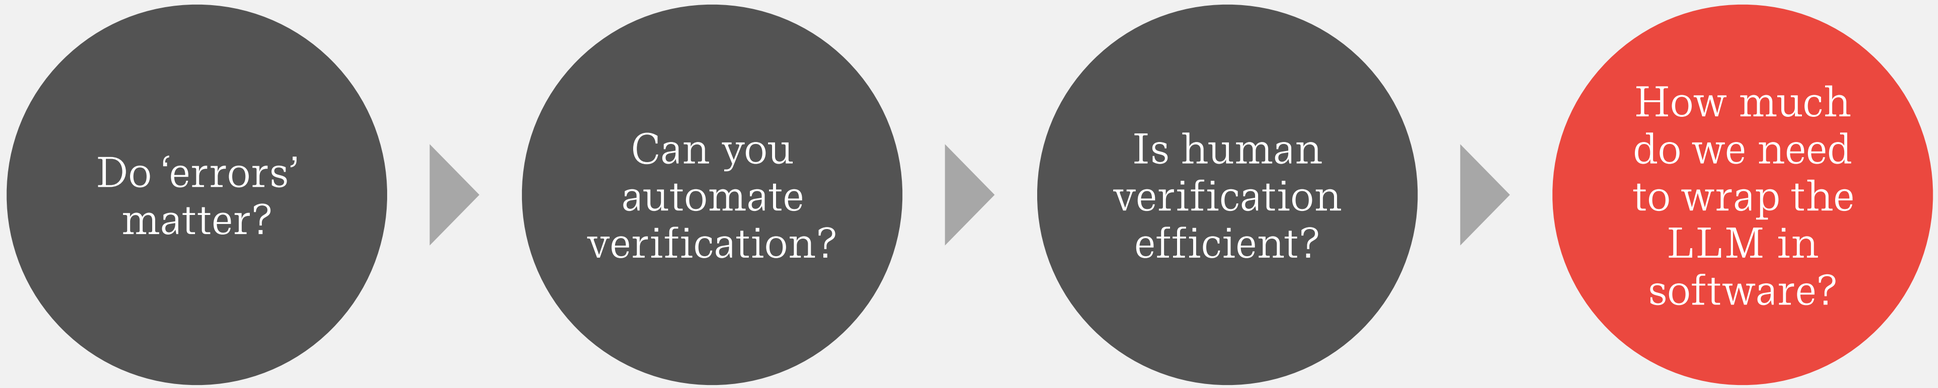

What do ‘automated interns’ change?

The Jevons paradox - applied price elasticity

300m interns? Jevons paradox at work

Steam engines gave Britain the equivalent labour of (very roughly) 5x its total population by 1900

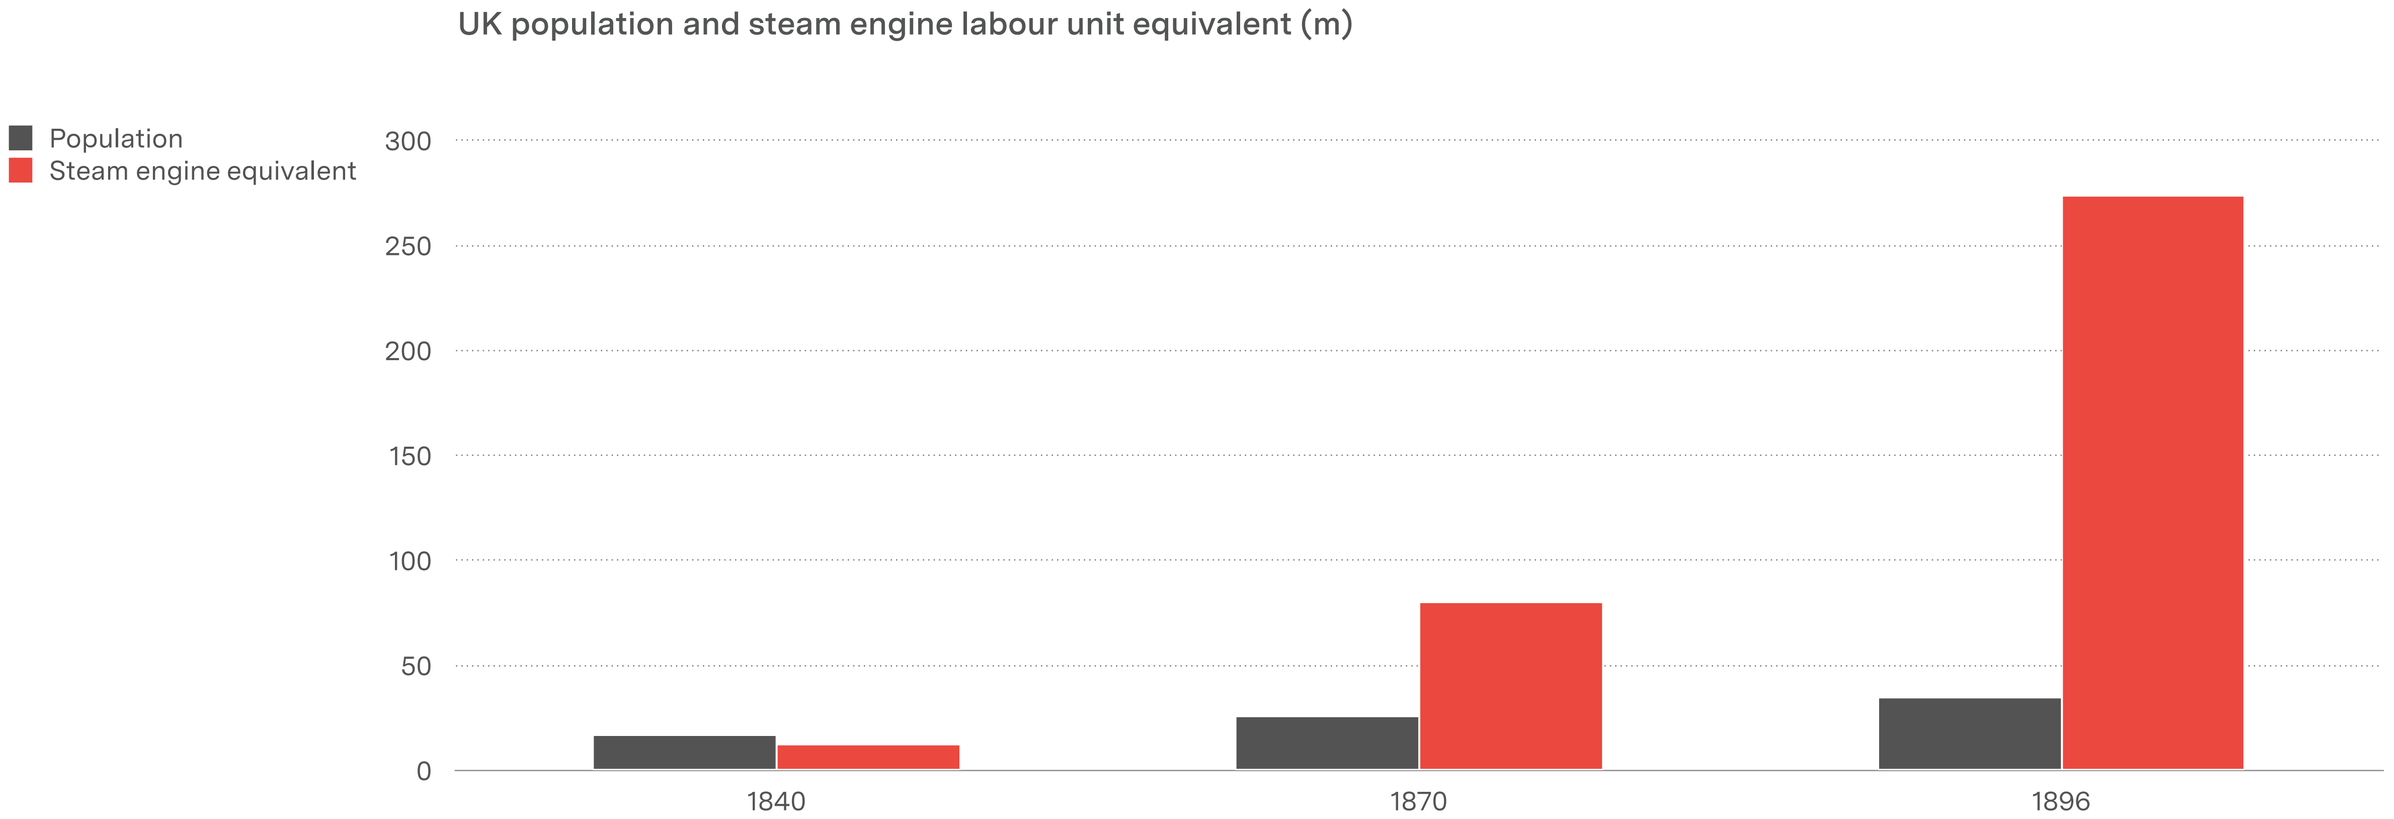

Source: Landes

Where's the human in the loop?

All recommendation systems today work by driving, capturing and analysing user activity

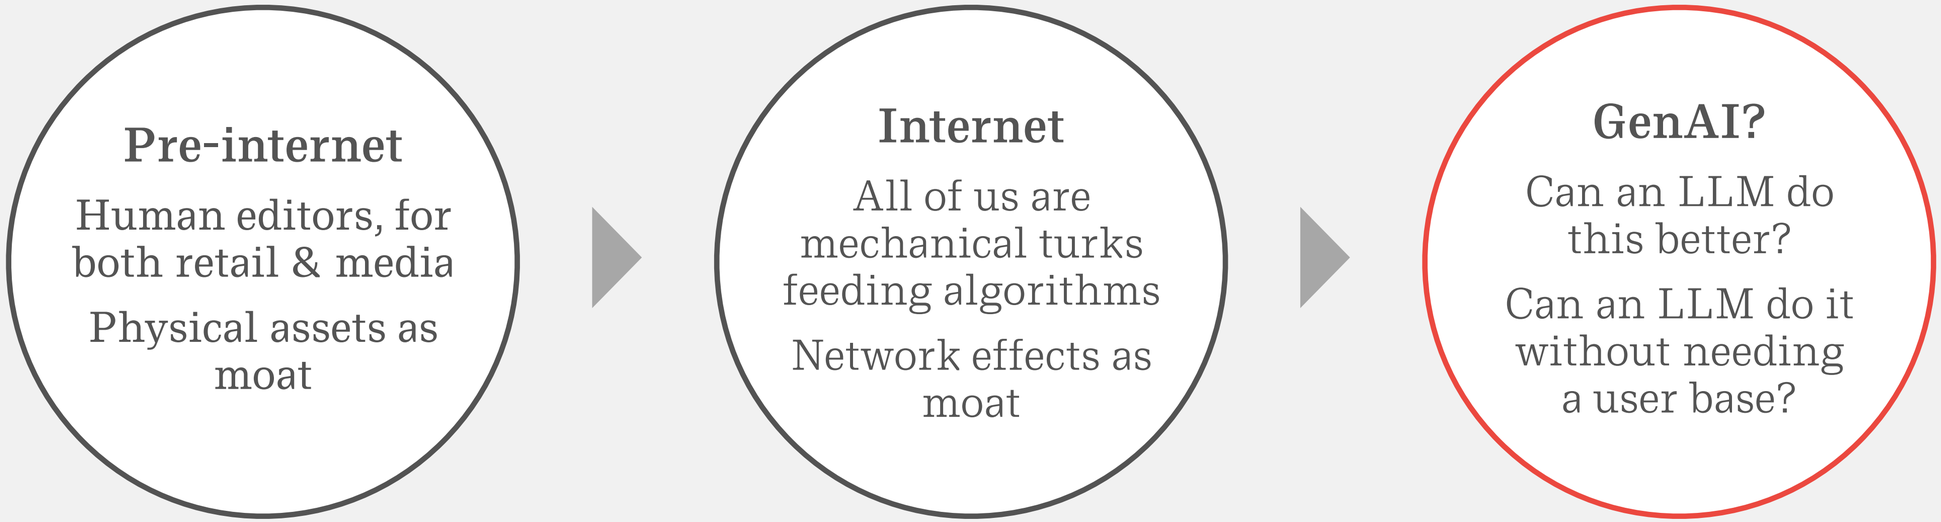

Value to capture!

Brands spend a trillion dollars a year to talk to consumers - plus rent, shipping, marketing, returns...

Global ad revenue, 2024 ($bn)

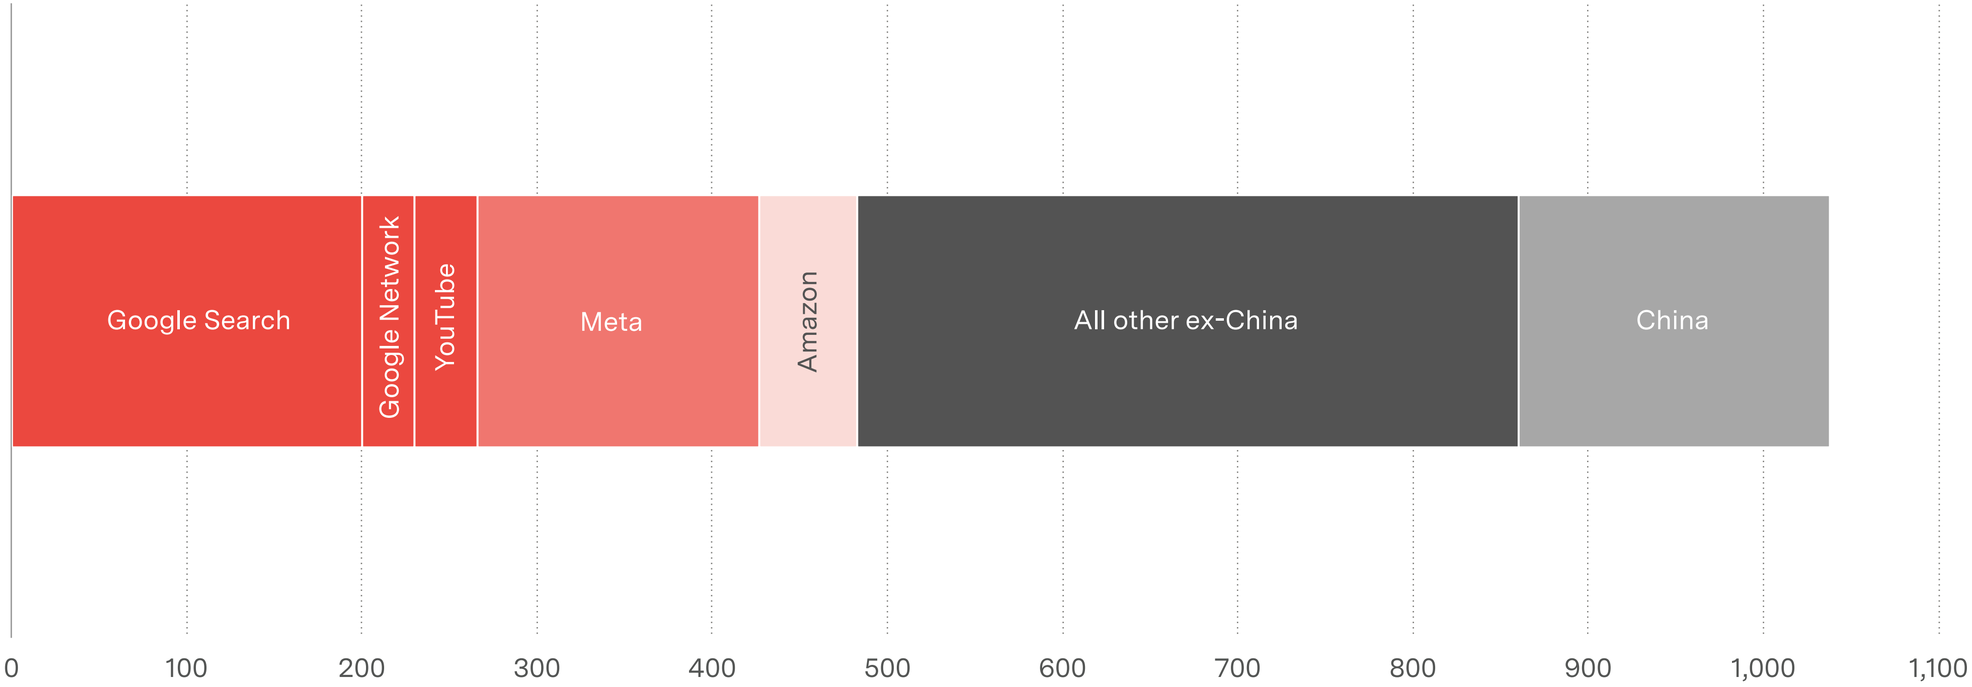

Source: Companies, WPP Media
NB: 'China' represents spending in China, not global spending by Chinese companies

“Our new AI recommendation model drove 5% more ad conversions on Instagram and 3% on Facebook”

Meta, Q2 2025

“Advertisers that activate AI Max in Search campaigns typically see 14% more conversions”

Google, Q2 2025

Absorb the new thing, automate what you know

Ad asset creation costs ~$100bn globally: now add 10-20x more assets and unlock cheap video for everyone

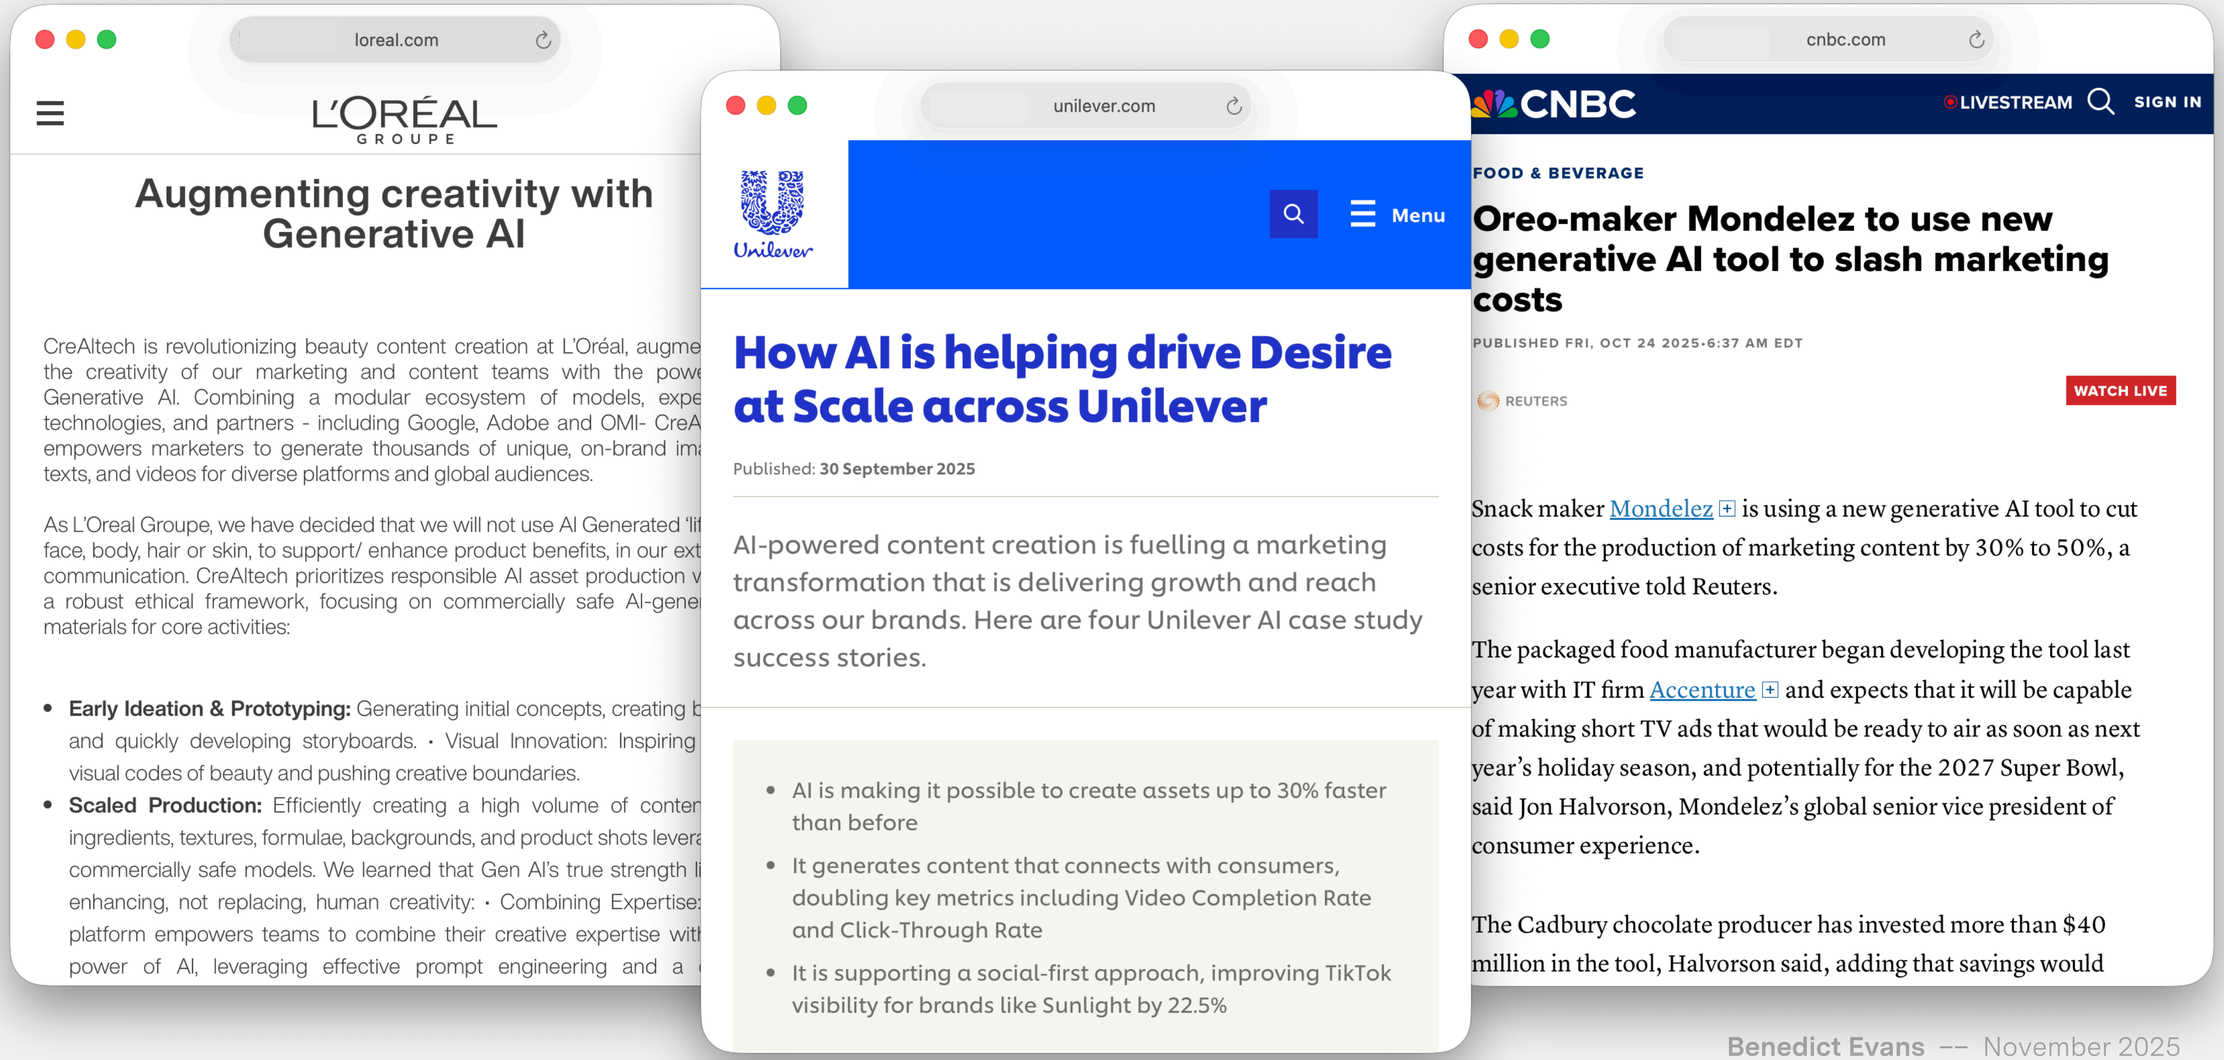

loreal.com

L'ORÉAL GROUPE

Augmenting creativity with Generative AI

CreAltech is revolutionizing beauty content creation at L'Oréal, augme
the creativity of our marketing and content teams with the powe
Generative AI. Combining a modular ecosystem of models, expe
technologies, and partners - including Google, Adobe and OMI- CreA
empowers marketers to generate thousands of unique, on-brand ima
texts, and videos for diverse platforms and global audiences.

As L'Oreal Groupe, we have decided that we will not use AI Generated 'lif
face, body, hair or skin, to support/ enhance product benefits, in our ext
communication. CreAltech prioritizes responsible AI asset production v
a robust ethical framework, focusing on commercially safe AI-gener
materials for core activities:

- **Early Ideation & Prototyping:** Generating initial concepts, creating b
and quickly developing storyboards. • Visual Innovation: Inspiring
visual codes of beauty and pushing creative boundaries.
- **Scaled Production:** Efficiently creating a high volume of conten
ingredients, textures, formulae, backgrounds, and product shots levera
commercially safe models. We learned that Gen AI's true strength li
enhancing, not replacing, human creativity: • Combining Expertise:
platform empowers teams to combine their creative expertise with
power of AI, leveraging effective prompt engineering and a

unilever.com

Unilever

Menu

How AI is helping drive Desire at Scale across Unilever

Published: **30 September 2025**

AI-powered content creation is fuelling a marketing transformation that is delivering growth and reach across our brands. Here are four Unilever AI case study success stories.

- AI is making it possible to create assets up to 30% faster than before
- It generates content that connects with consumers, doubling key metrics including Video Completion Rate and Click-Through Rate
- It is supporting a social-first approach, improving TikTok visibility for brands like Sunlight by 22.5%

cnbc.com

CNBC LIVESTREAM SIGN IN

FOOD & BEVERAGE

Oreo-maker Mondelez to use new generative AI tool to slash marketing costs

PUBLISHED FRI, OCT 24 2025·6:37 AM EDT

REUTERS

WATCH LIVE

Snack maker Mondelez is using a new generative AI tool to cut costs for the production of marketing content by 30% to 50%, a senior executive told Reuters.

The packaged food manufacturer began developing the tool last year with IT firm Accenture and expects that it will be capable of making short TV ads that would be ready to air as soon as next year's holiday season, and potentially for the 2027 Super Bowl, said Jon Halvorson, Mondelez's global senior vice president of consumer experience.

The Cadbury chocolate producer has invested more than $40 million in the tool, Halvorson said, adding that savings would

Old: “half of AI will be turning three bullet points into emails, and the other half will be turning emails into three bullet points”

New: half of AI will be turning three bullet points into 300 ads, and the other half…

Again, this is early

Use so far may be more additive and experimental than substitution (and this includes Gemini)

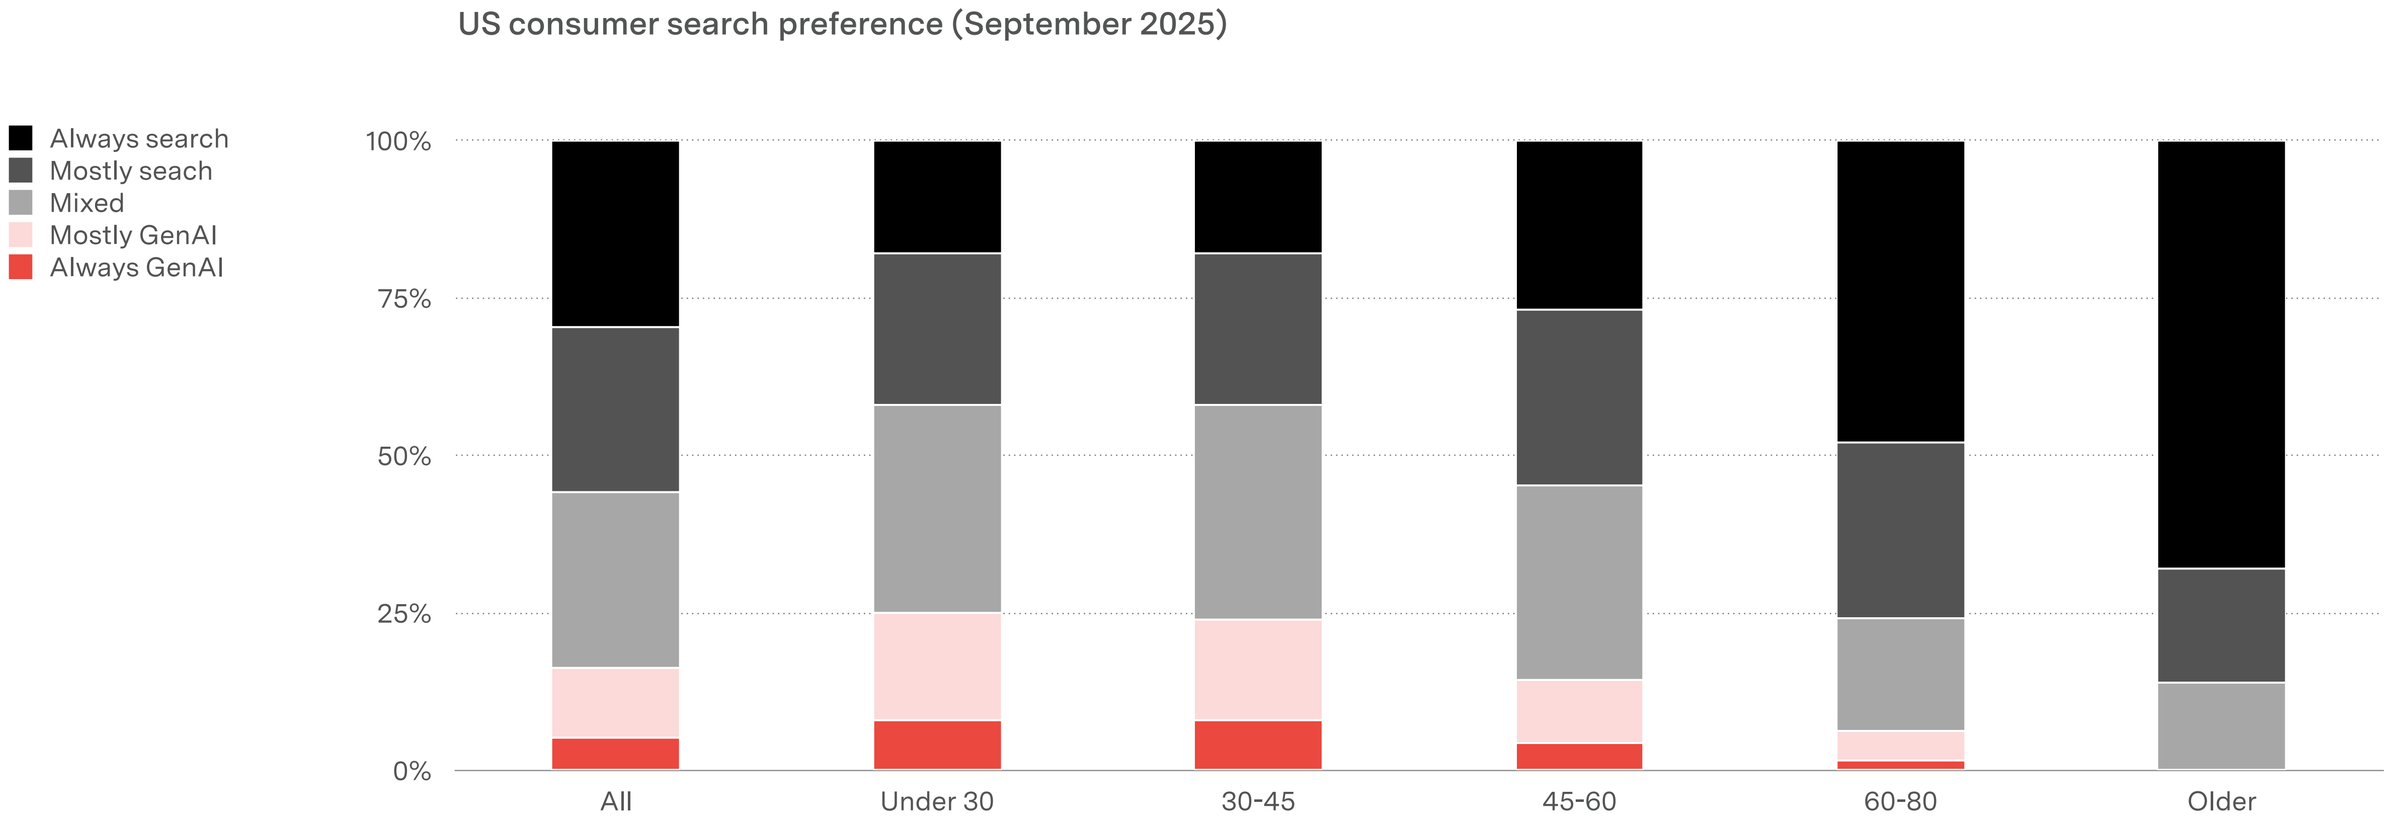

Source: Bain

The web has been dying since 1997

Remember how early this is, and how hard it is to know how the new thing will work

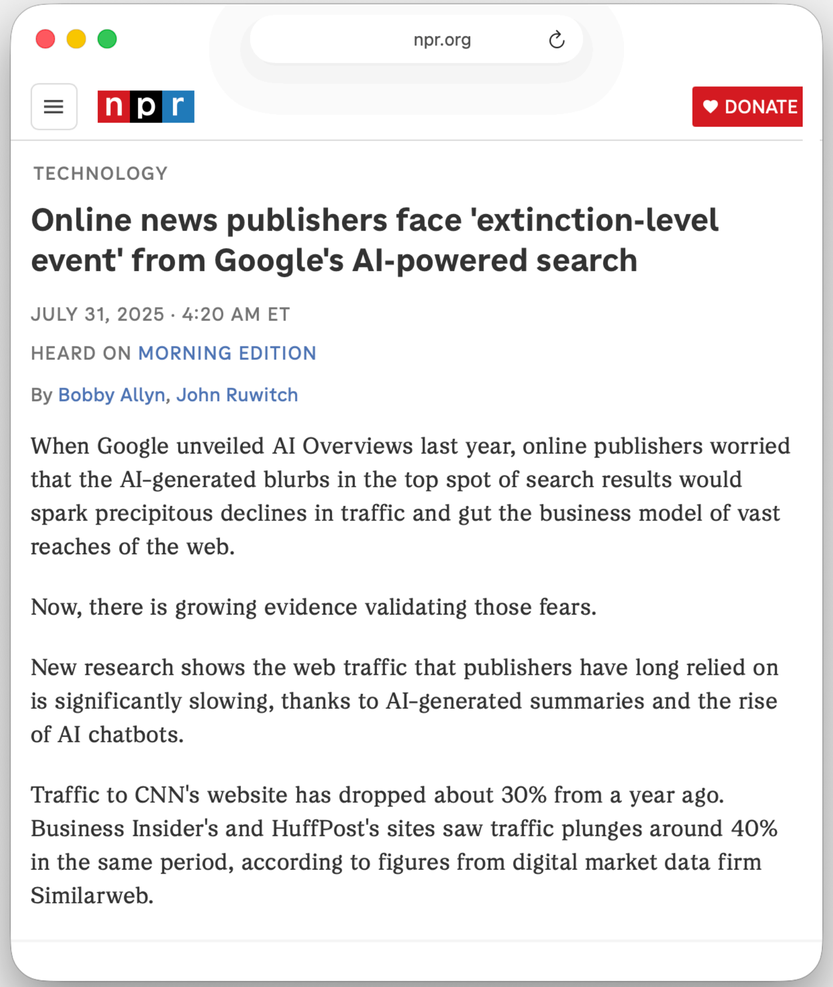

Source: Zeldman

But where are we going?

What if recommendations go from correlation to an understanding of what those SKUs really represent?

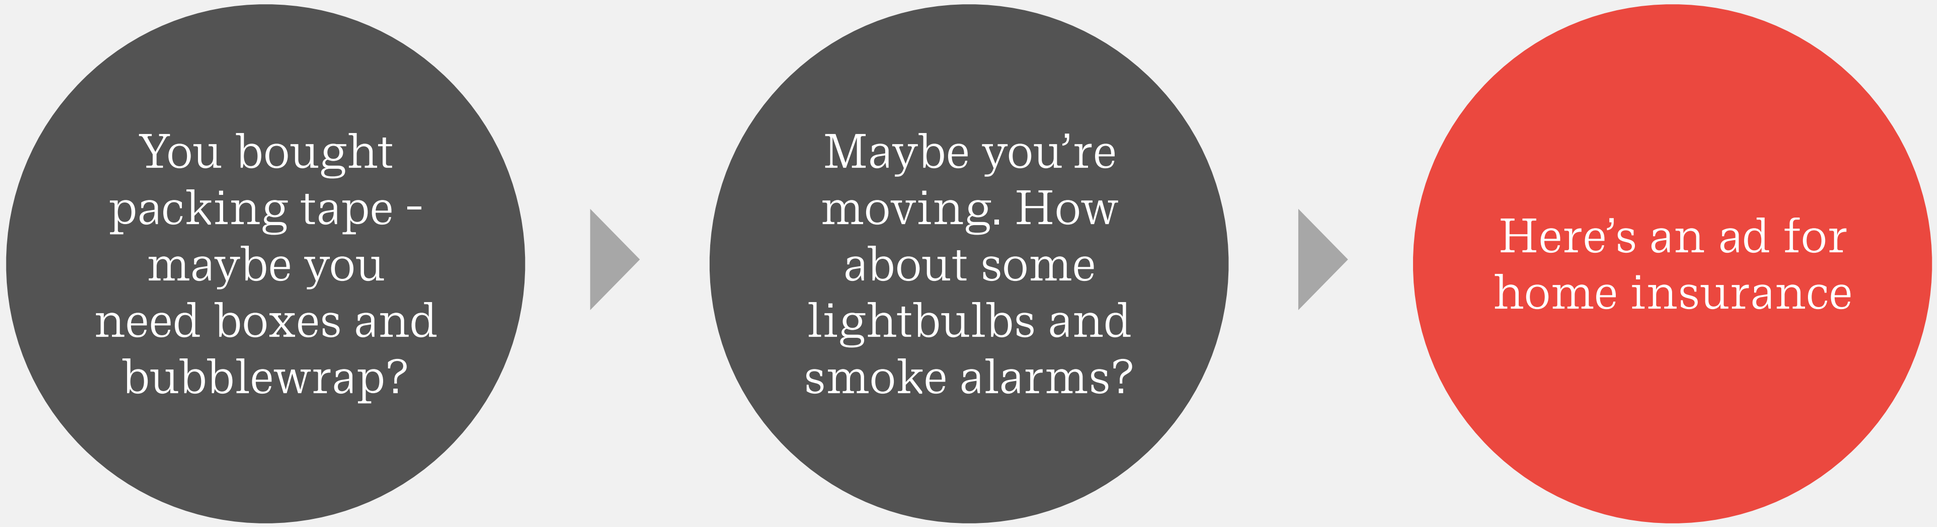

How do we know what we might want?

For 30 years we've had infinite product, infinite media and infinite retail

Now we have a machine that sees all of it, and sees us

What does it recommend?

Source: Morioka Shoten

And what gets unbundled?

What do you actually want? What are you trying to do? Why? What do you care about?

“What’s our AI strategy?”

Is this a question for the CIO? CMO? CEO?

Accenture? Publicis? Bain/BCG/McKinsey?

Is this a new tool or a new industry?

AI eats the world

The old stuff from before ChatGPT is still here

Excluding gas and grocery, 30% of US retail sales are now online

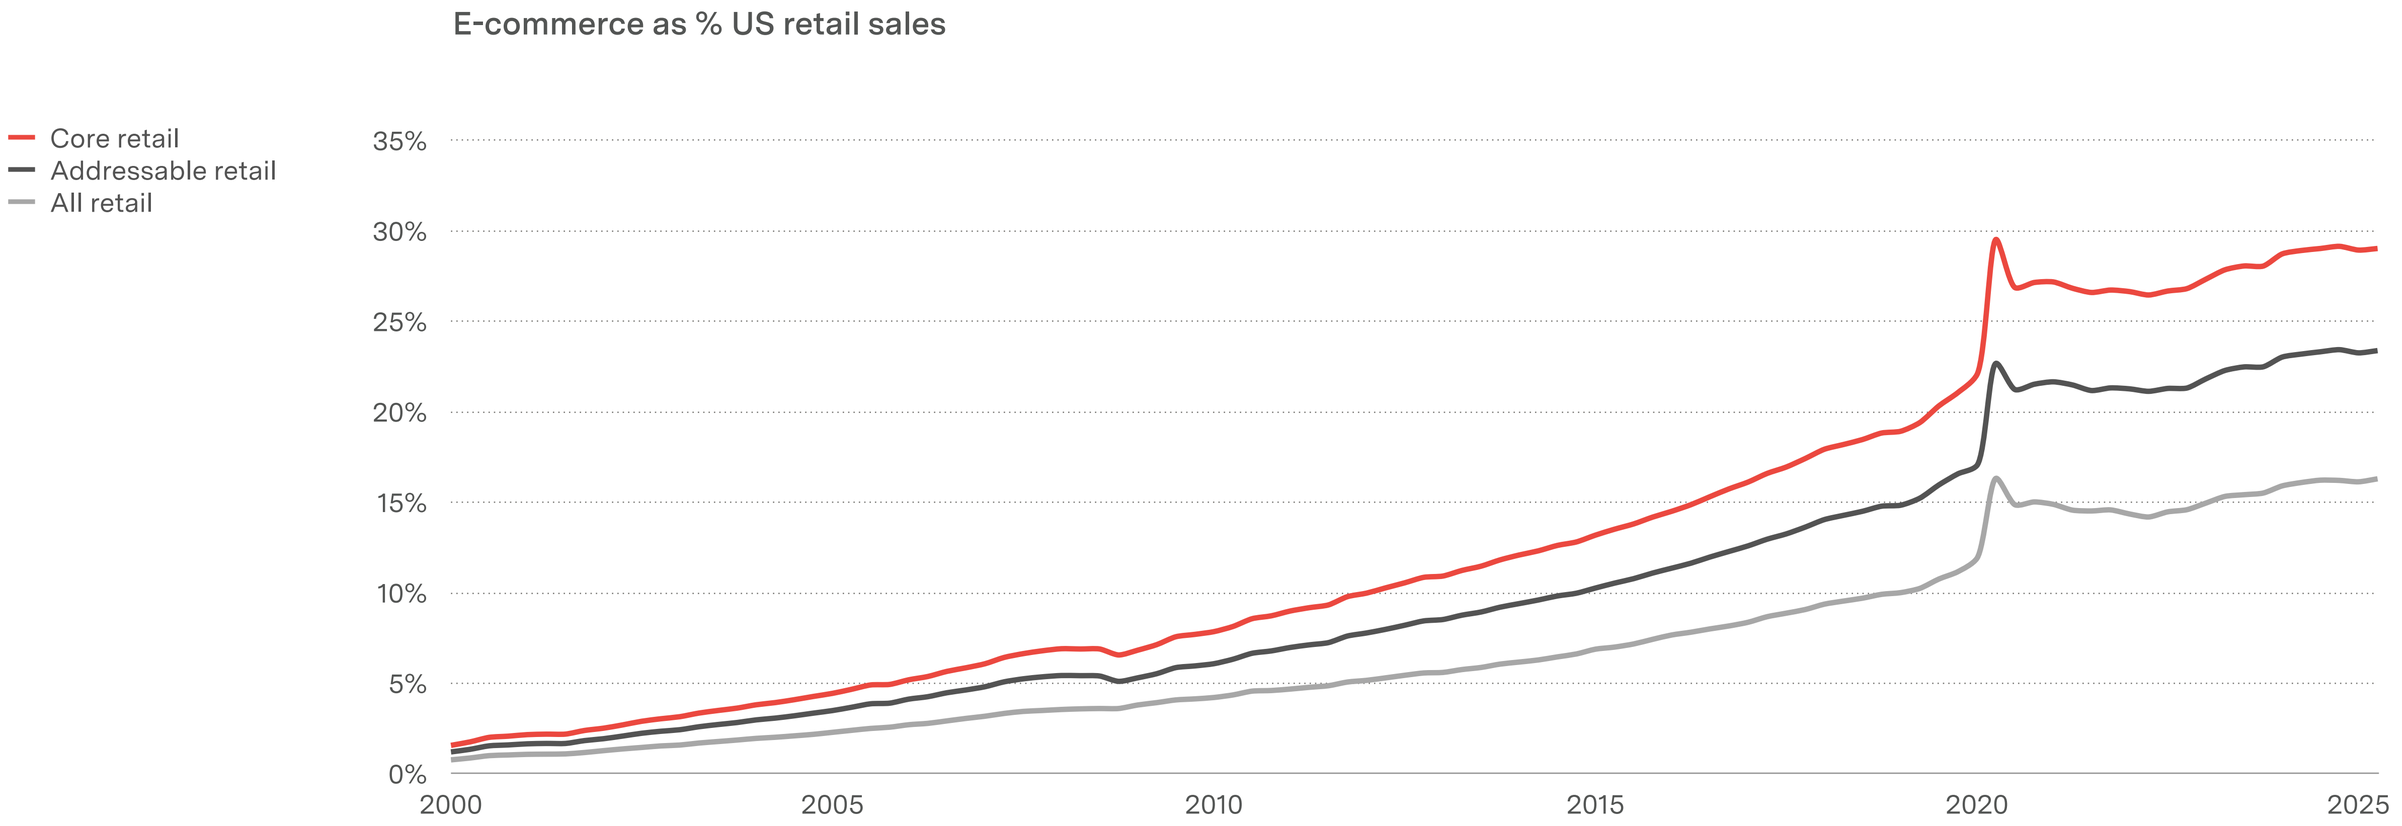

Source: US Census. Seasonally adjusted
'Retail' excludes restaurants & bars. 'Addressable retail' also excludes cars, car parts & service, and gasoline stations. 'Core retail' excludes grocery from addressable retail.

And all the *other* new stuff

After a decade of promises and tens of billions of dollars, ‘automatic cars’ might be starting to work

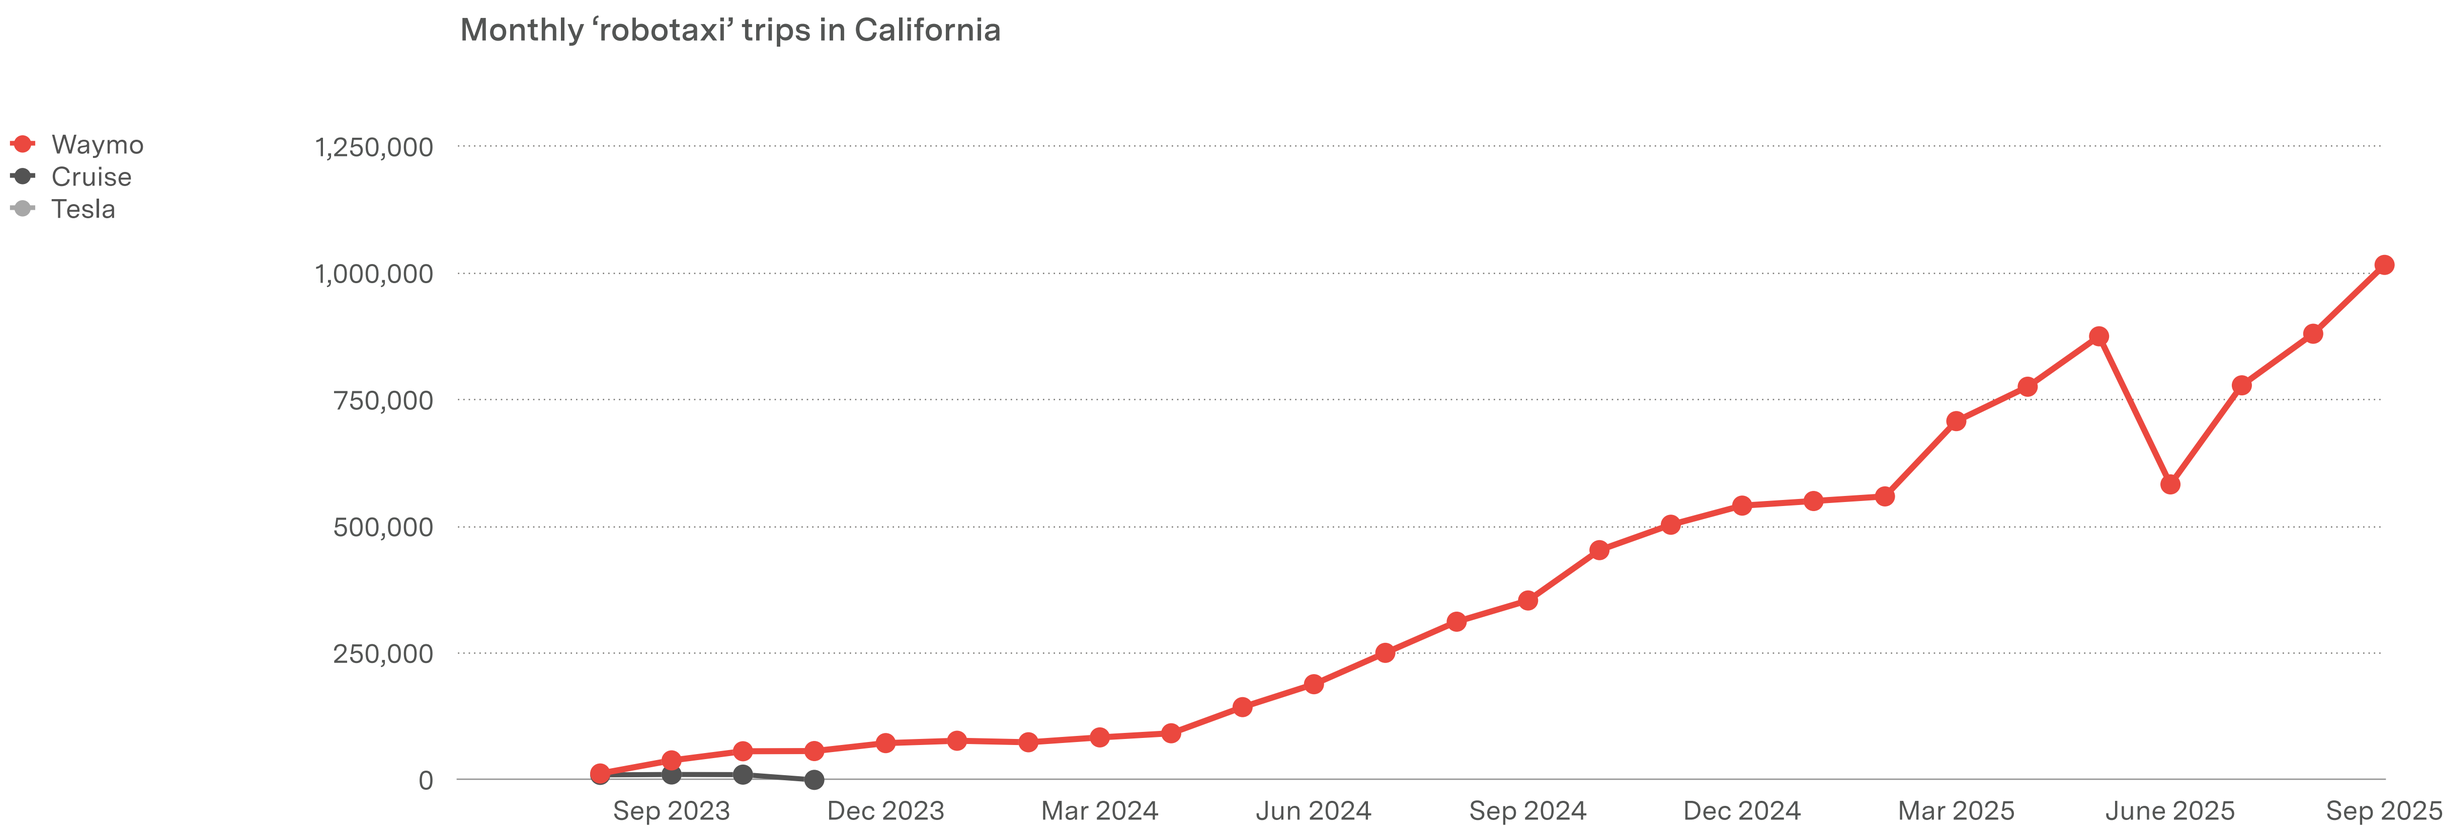

Source: California Public Utilities Commission

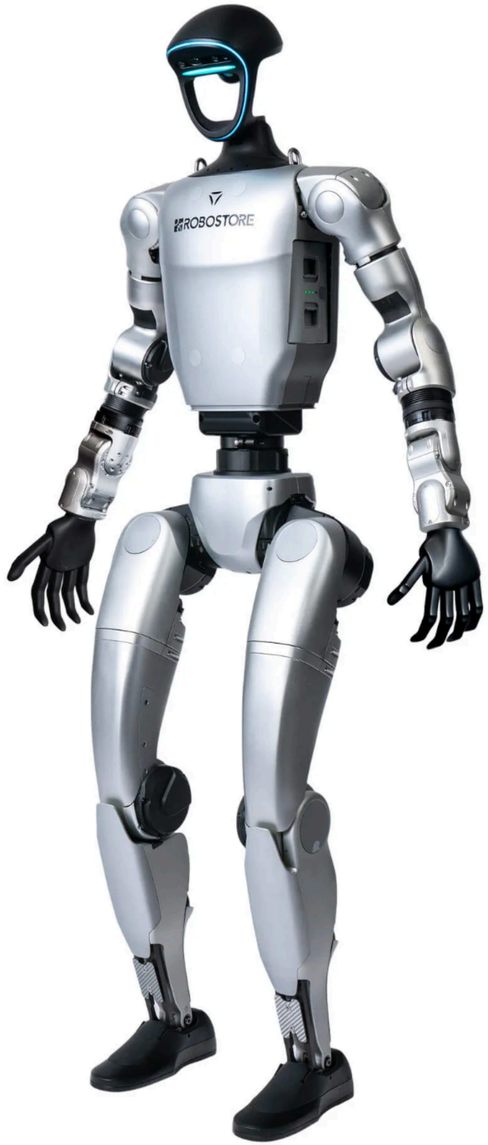
ROBOSTORE

How many times have we been here before?

We’ve had radical change (and bubbles) before

And we’ve also done automation before

Source: IBM

‘Automation and technological change’, 1955

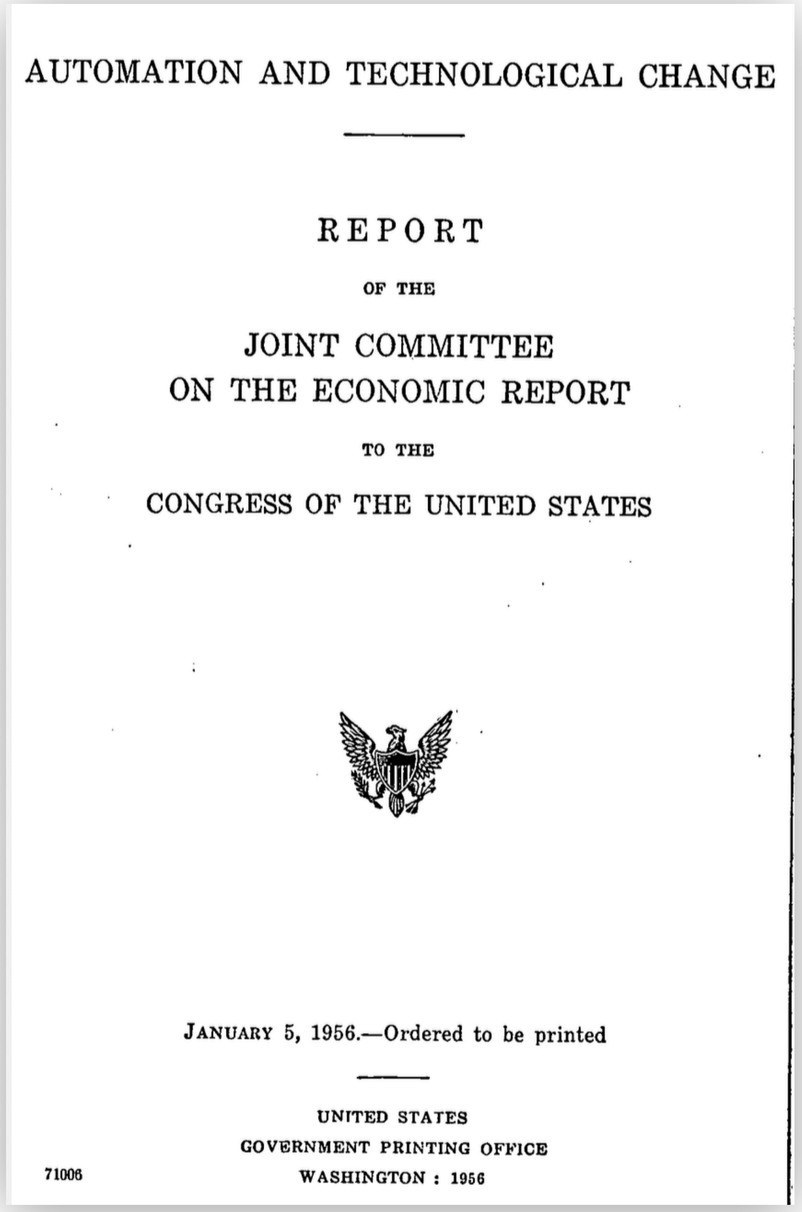

AUTOMATION AND TECHNOLOGICAL CHANGE

REPORT

OF THE

JOINT COMMITTEE
ON THE ECONOMIC REPORT

TO THE

CONGRESS OF THE UNITED STATES

JANUARY 5, 1956.—Ordered to be printed

UNITED STATES
GOVERNMENT PRINTING OFFICE
WASHINGTON : 1956

71006

Source: US Congress report on automation, 1956

“Automation”

In the course of the hearings, the subcommittee considered specifically six different industrial situations in the metalworking, chemical, electronics, transportation, and communications industries, together with data processing and officework. These industries were selected merely as illustrative of the kind of problem which may be faced in the trend toward automation. There are, of course, many other industries which might have been studied with interest and profit had time permitted. The fact that these particular industries were chosen should not for a moment obscure the fact of rapidly advancing technology in other areas. To mention only a few such areas, one might cite the canning and bottling industries. One might cite also petroleum refining, the processing of commercial-bank paperwork, the basic steel industry, the use of ready-mixed concrete, coal mining, the use of electronically controlled elevators in our modern skyscrapers, and numerous others.

Source: US Congress report on automation, 1956

“Automation”

Source: Benedict Evans

When automation works, it disappears

Otis launched the ‘Autotronic’ automatic elevator in 1950

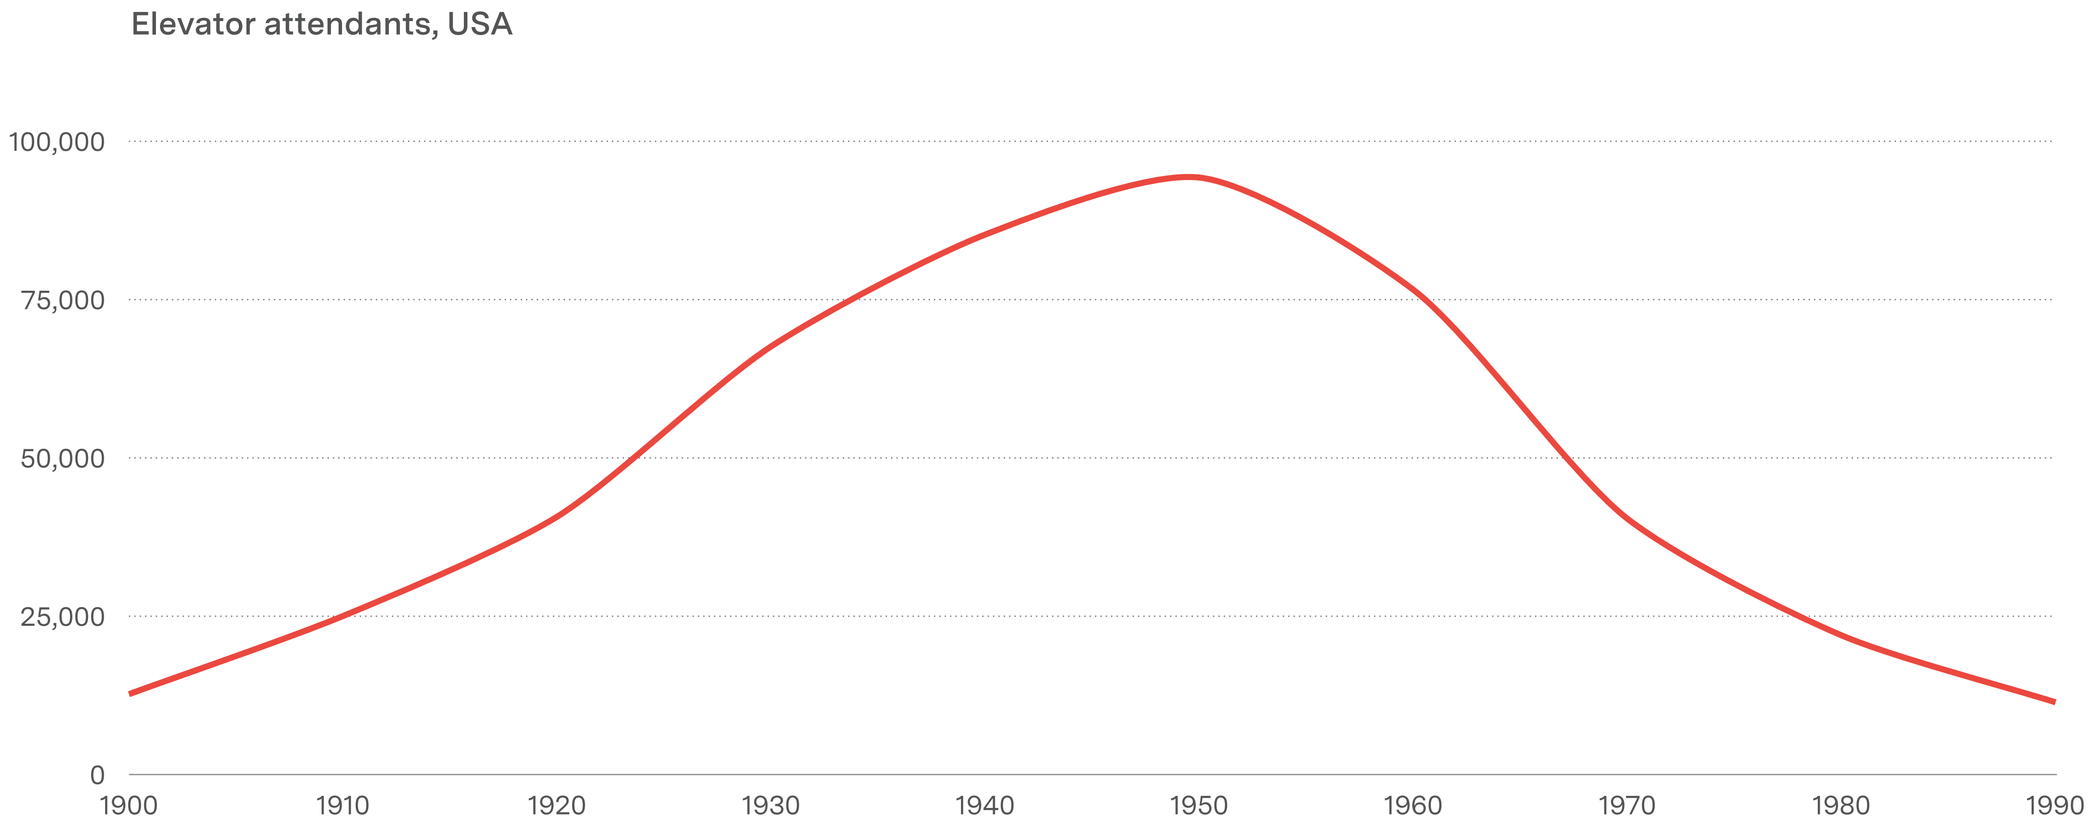

Source: US Census

“AI is whatever machines can’t do yet”

Larry Tesler, 1970

Thank you

What matters in tech? What's going on, what might it mean, and what will happen next?

I've spent 25 years analysing mobile, media and technology, and worked in equity research, strategy, consulting and venture capital. I'm now an independent analyst, and I speak and consult on strategy and technology for companies around the world. Mostly, that means working out the right questions.

For more, see www.ben-evans.com

Thank you

Benedict Evans

November 2025

www.ben-evans.com