

Subject: Benedict's Newsletter: No. 620

Date: Tuesday, 02 December 2025 at 17:06:01 South Africa Standard Time

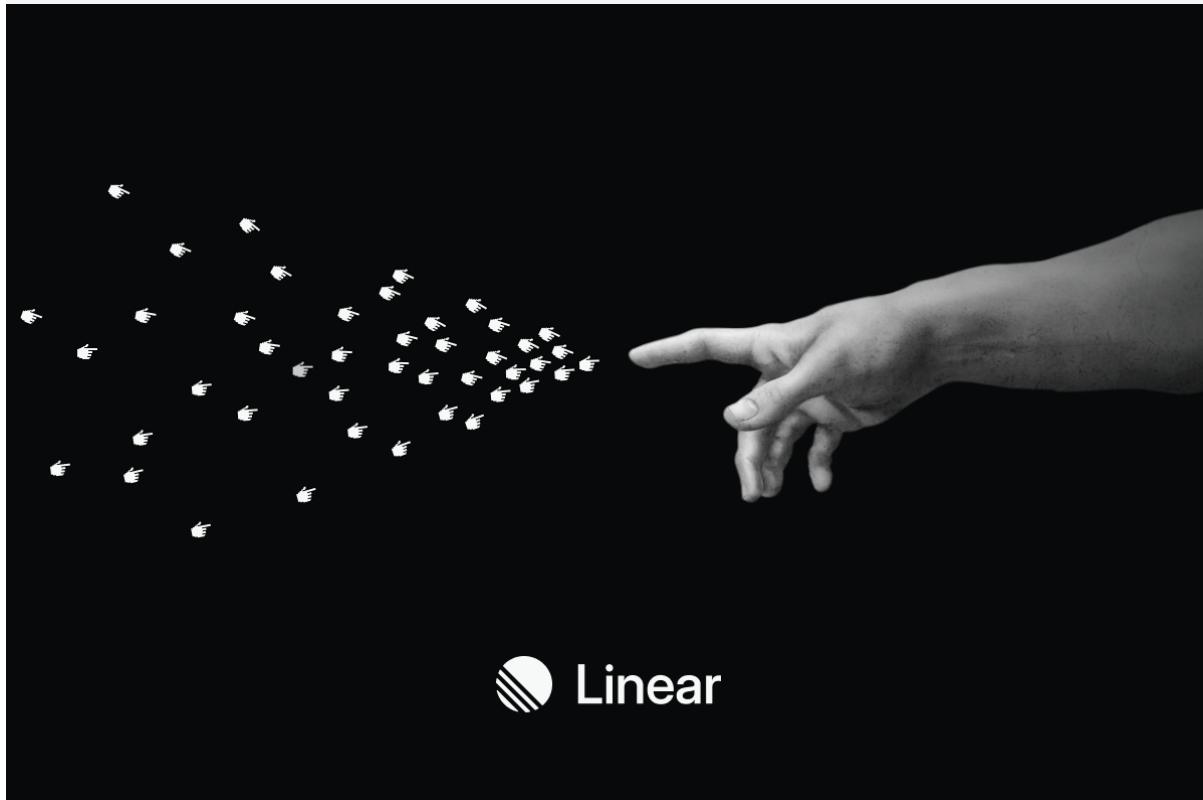
From: Benedict Evans

To: janbielies@gmail.com

Benedict's Newsletter

NO. 620 FREE EDITION SUNDAY 20 DEC 2025

SPONSORED BY LINEAR



Meet the coordination layer for human and AI product development

Linear helps you deploy, track, and orchestrate AI agents that ship features, fix bugs, and build products alongside your human team.

[**Learn more**](#)

My work

AI eats the world

Twice a year, I produce a big presentation exploring macro and strategic trends in the tech industry. New in November 2025, ‘AI eats the world’.

[LINK](#)

AI, networks and Mechanical Turks

How far do LLMs give us a step change in how good a search and recommendation system can be? Do they let you build one without needing a vast user base of your own? [LINK](#)

News

Google selling TPUs?

The WSJ and Information both reported that Google is looking at selling its TPU AI accelerator chips to other companies, including Meta. This would be a big deal if true, given how far Nvidia is ahead of any publicly available competitors: TPUs are apparently very good, but only Google has access. After that, things get complex: for example, how much do other customers rely on Nvidia’s CUDA software layer as well as the hardware itself, and how much are TPUs tightly linked to Google’s systems? Ask a DC analyst. [INFORMATION](#), [WSJ](#)

Chat Shopping

Following Google, OpenAI launched a shopping assistant experience inside its chatbot. Meanwhile, Amazon is expanding the third parties that it blocks from doing this.

There are a lot of overlapping questions here. How much can such a system remove and clarify product complexity, and do that better than Google? This is how people in Silicon Valley like to shop, but how well does that map against how other people shop? How does that offset losing the sophistication (which *is* complexity!) of the retailers’ own tools, UX, and recommendations? How much does the model really know you, and know

you better than Amazon? And on the other side, which retailers will let you do this, because they want the traffic, and which retailers will block you because they want to own the experience and the customer? [LINK](#)

Amazon does US AI

Amazon announced a deal worth up to \$50bn to build dedicated compute infrastructure for the US government. [LINK](#)

AI accounting

Michael Burry, known for shorting US mortgage bonds in the GFC, has got attention in the last week or two by claiming (amongst other things) that AI companies are inflating earnings by depreciating their data centres too slowly. I don't do share prices anymore, but IMO this is a second derivative question, where the first derivative is what the useful life of AI accelerators (as sold by Nvidia) will be, and the primary question is the real steady-state capex needs for AI, for both new capacity and replacement capacity.

The problem, for me, is that we don't know how much more compute-efficient the models will get at achieving a given result, nor how much more compute-intensive new use-cases and applications (agentic, video) will be, nor how much how many people will use each of them. This is all very like trying to predict bandwidth needs in 1998 or 2000 when you don't know if YouTube or Spotify will exist, never mind how much bandwidth they'll need. And you don't know the revenue either! In that context, worrying about depreciation policies seems like displacement: aren't you looking at small variables you can analyse instead of much bigger ones you can't?

Also, Burry closed his fund and started a \$39/month Substack. [LINK](#)

Conversely, Amazon increased its depreciation charges. Stephen Clapham has a good write-up of the broader questions. [LINK](#)

About

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.

[FIND OUT MORE](#)

Ideas

Billboard saw a copy of a fundraising deck from the AI music generator company Suno. It's creating the equivalent of Spotify's entire catalogue every two weeks. [LINK](#)

The Verge did a nice write-up of Hoto and Fanttik, the two hot new Chinese tool brands. [LINK](#)

Following last week's reveal that a lot of deliberately divisive MAGA accounts on Twitter are not actually based in the USA, 404 points out that a lot of them are entrepreneurs chasing revenue shares. [LINK](#)

The FT points out that Oracle, Softbank, and Coreweave combined are on track to borrow over \$100b to build infrastructure for OpenAI contracts. Other People's Balance Sheets is the new Other People's Money. [LINK](#)

Interesting case-study of someone using AI to generate and sell fake articles to publications. The internet made content distribution 'free', and that had all sorts of unexpected consequences - now AI makes content creation 'free'. [LINK](#)

This 90-minute Dwarkesh Patel interview with Ilya Sutskever (OpenAI co-founder, now SSI) got a lot of pick-up. I know I should say it's fascinating, but I got 25 minutes in and lost the will to live. Sooo (Long. Pause.) Sloooow. Maybe I'll try again at 3x speed. [LINK](#)

Conversely, this Ari Emanuel interview on entertainment, live events,

sports, and AI is very good. He talks fast. [LINK](#)

Detailed case study from Booking.com on using agents in their guest messaging system. [LINK](#)

Laurene Jobs asks Sam Altman and Jony Ive “what can you tell us about The Thing?” Not much, except it will involve ‘less distraction’. [LINK](#)

Statistics about tech energy and water consumption have a long and ignoble history of people creating their own statistics based on wild extrapolations of poor data, or just getting the units wrong, and ‘Empire of AI’, a bestseller focusing on water consumption, appears to have based a central argument on numbers that are wrong by several orders of magnitude. Of course, the companies themselves disclose too little that it’s hard to get this stuff right - the big platform companies do disclose primary data, but that tells us little about future AI use. [LINK](#)

ChatGPT’s public launch was three years ago today. [LINK](#)

Outside interests

Alan Yentob’s film for the BBC about Sir Tom Stoppard. [LINK](#)

UK Government report on the regulation and costs of building nuclear power plants. Read the section on fish protection at the top of page 68: £700m to save the lives of 0.083 salmon per year. [LINK](#)

For the friend with everything, a 3D-printed, light-up model of the Chernobyl plant. [LINK](#)

Tap-to-pay leaves street vendors and the homeless behind. [LINK](#)

Data

Alix Partners 2026 Data Centre outlook. Power and Water. [LINK](#)

Bain survey data breaking down generative AI deployment by sector,

crucially, splitting pilots from production. [LINK](#)

AN analysis of LLM downloads arguing that China has overtaken US models. Not surprising. [LINK](#)

A fairly good Brookings Institute US survey on Generative AI usage. [LINK](#)

Conversely, the US Census does a bi-weekly business survey (BTOS), and since 2023 it's asked about AI use. Since this is national, authoritative, and frequent, it gets reported a lot, and in the last couple of periods, it's flat-lined. However, the methodology means this data has zero value.

The definition of 'AI' in the question is "Examples of AI: machine learning, natural language processing, virtual agents, voice recognition, etc." which does not distinguish generative AI from systems built 10 or even 20 years ago, and could cover almost anything, while the frequency is also only 'did you use it in producing goods or services?'. The definition is so broad and open-ended that it will massively overcount, while limiting to 'producing goods and services' means use for marketing or fraud prevention would be excluded, undercounting on the other side. This is not helpful. [LINK](#)

Preview from the Premium edition

Agents, AI apps and the widget fallacy

Every five or ten years, somebody tries widgets again. Instead of having to dive into lots of different applications, with different interfaces and experiences, why not just have little units that show you the stuff you really need, in one standard UI?

If all you want to do is see your next appointment or tomorrow's weather, this is quite useful. But if you want to do anything with the logic and data in those applications, very, very quickly, you do need to go back into the app. Those features are there for a reason, and yes, no one uses more than 75% of the features (make up your preferred percentage), but everyone uses a different set. So you open the app, or you open a website, and you forget about widgets.

Going a step back, you could suggest that for as long as we've actually had

separate programs, engineers have wanted to find ways to abstract them back into a single layer. There was a moment in the early history of GUIs when people thought that the operating system would do all of this. Really, what was an app, but just a different view on data that was being stored in the file system by the OS? That also got you to things like OLE, which you have to be over a certain age to remember, when people thought that you could somehow embed... a spreadsheet or a CAD drawing or a Photoshop file into a document, and when you clicked on them, that UI would load, but the individual ‘apps’ would be very thin layers of code with the OS doing all the work. In today’s terminology, Bill Gates thought Office was ‘just a thin Windows wrapper’.

Upgrade to Premium

You're getting the Free edition. Subscribers to the Premium edition got this two days ago on Sunday evening, together with an exclusive column, complete access to the archive of over 600 issues, and more.

[SUBSCRIBE OR UPGRADE](#)

THIS WAS SENT TO JANBIELIES@GMAIL.COM ON 2 DECEMBER, 2025, BECAUSE YOU SUBSCRIBED AT BEN-EVANS.COM. YOU CAN [UNSUBSCRIBE](#), OR READ THE [PRIVACY POLICY](#). OUR MAILING ADDRESS IS BENEDICT EVANS, 2 GOLDEN SQUARE, LONDON, W1, UNITED KINGDOM.

COPYRIGHT © 2013-2025 BENEDICT EVANS, ALL RIGHTS RESERVED.