alpine capital research

November 2026

CURMUDGEON'S CORNER

VOLUME 4

---

curmudgeon (cur mud geon)
noun [ker-muj-uhn]

A bad tempered, difficult, contentious, querulous, grouchy, cantankerous person

---

BY JAMES MONTIER
MANAGING DIRECTOR, FUNDAMENTAL RESEARCH & PERFORMANCE

190 Carondelet Plaza, Suite 1300 | Saint Louis, MO 63105
314.932.7600 | 877.849.7733

www.acr-invest.com

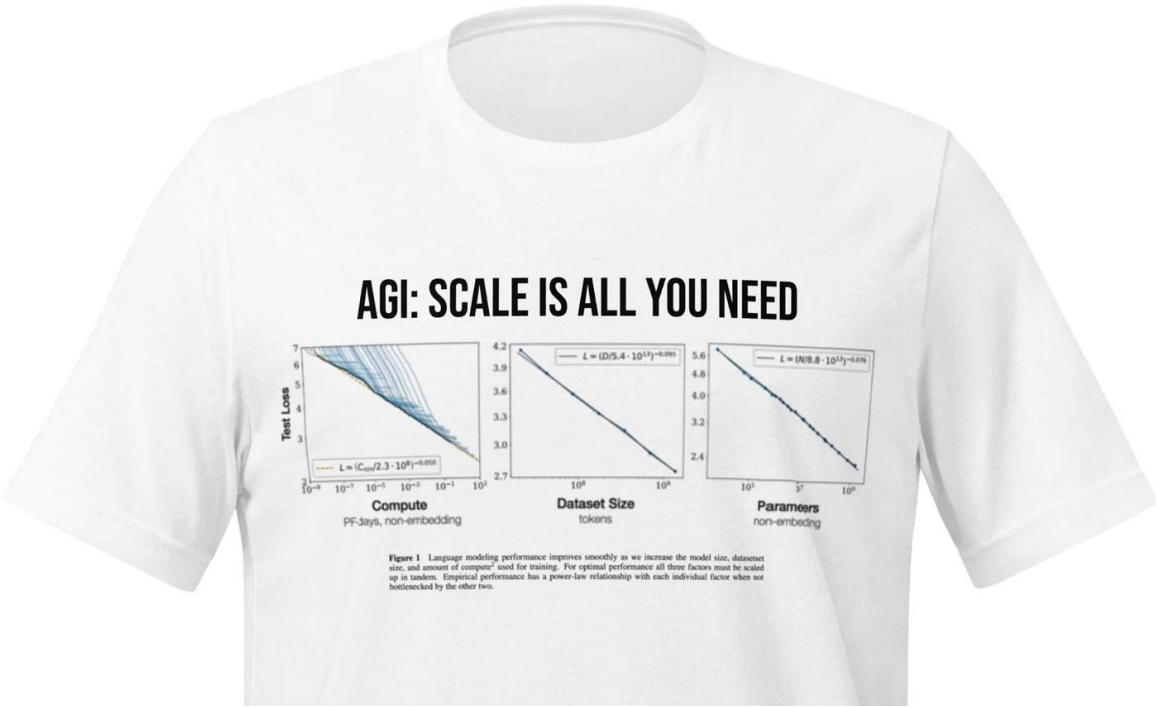
Rise of the Machines: Part 3

AI: Investment, Productivity and Profits

Alice laughed. "There's no use trying," she said: "one can't believe impossible things." "I daresay you haven't had much practice," said the Queen. "When I was your age, I always did it half-an-hour a day. Why, sometimes I've believed as many as six impossible things before breakfast."

When it comes to AI, I suspect you have to be nearly as adept as the Queen in *Through the Looking Glass*, if you are bullish on either the technology (see Curmudgeon's corner 2) or its prospects (the subject of this essay). In this final (?) part of my trilogy on AI, I will explore three fundamental myths that I believe should be a significant challenge to anyone who is bullish on AI as an investment thesis.

Myth 1: Scaling LLMs can lead to AGI



Let's start with a definition…what is AGI (artificial general intelligence)? I will follow Hendrycks et al (2025)[1] and define AGI as "An AI that can match or exceed the cognitive versatility and proficiency of a well-educated adult".

As the t-shirt (using charts from OpenAI's early paper on LLMs) observes, the prevailing belief (at least among many such as the hyperscalers and the producers of LLMs) is that scaling (throwing more compute, data and model size) is all you need to get you to AGI.

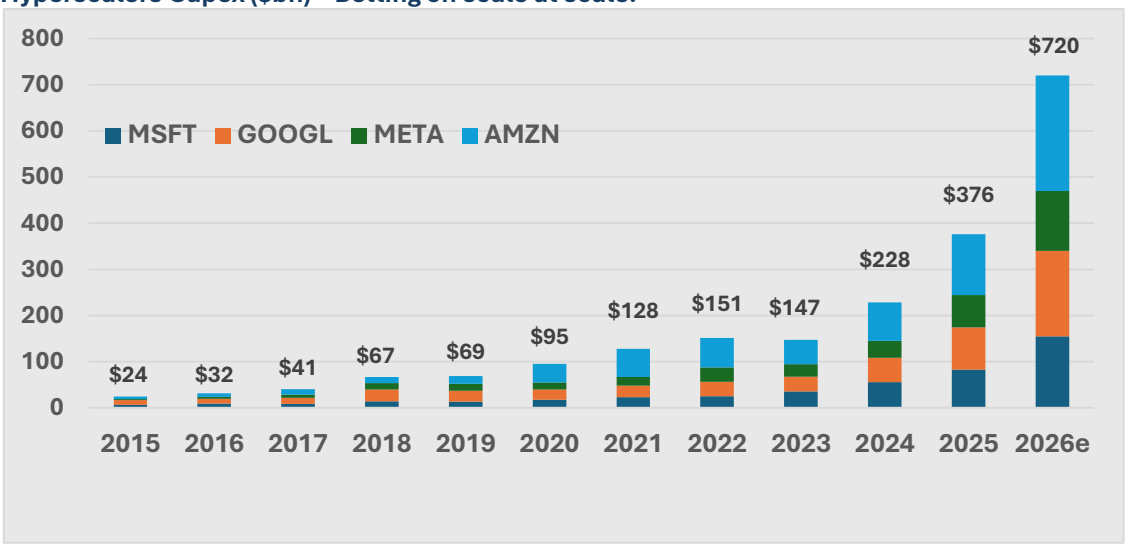
[1] Hendrycks et al (2025) A Definition of AGI, available from https://arxiv.org/pdf/2510.18212

The roots of this belief can be found in a very short but powerful piece written by Richard Sutton entitled The Bitterest Lesson[2]. Sutton is one of the founders of the reinforcement learning that has been widely used in enhancing LLMs. In his piece from March 2019 Sutton notes:

> "We have to learn the bitter lesson that building in how we think we think does not work in the long run. The bitter lesson is based on the historical observations that 1) AI researchers have often tried to build knowledge into their agents, 2) this always helps in the short term, and is personally satisfying to the researcher, but 3) in the long run it plateaus and even inhibits further progress, and 4) breakthrough progress eventually arrives by an opposing approach based on scaling computation by search and learning…One thing that should be learned from the bitter lesson is the great power of general purpose methods, of methods that continue to scale with increased computation even as the available computation becomes very great."

In essence, Sutton observes that in a wide variety of areas scaling has been a surprising success relative to normal knowledge based approaches (more on this a little later). Simply put. Bigger is always better.

Hyperscalers Capex ($bn) – Betting on scale at scale!



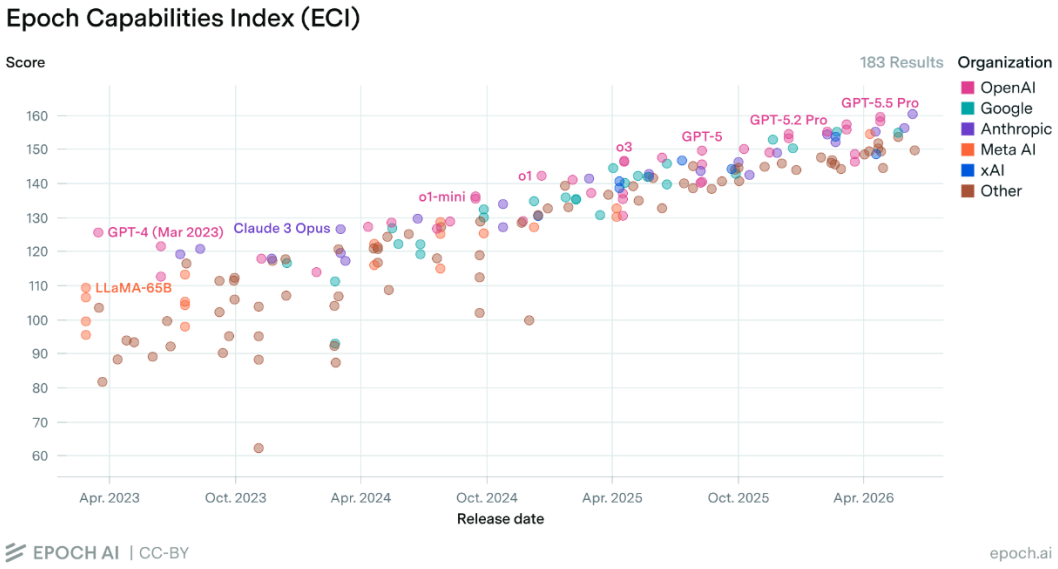
Source: S&P Cap IQ

Champions of this viewpoint seem to believe that AGI is imminent. Dario Amodei of Anthropic believes we will have AGI "a country of geniuses in a data center" by late 2026/early 2027. Sam Altman of OpenAI has stated, "We are now confident we know how to build AGI". Indeed, as I was finishing this essay up, a new wave of declarations arrived with Greg Brockman (president of OpenAI) claiming that GPT-6 (known as Astra) heralded the beginning of "The era of AGI". To which Jensen Huang of Nvidia sycophantically responded, "AGI has arrived"!

It is, of course, worth bearing in mind that predictions of the arrival are nothing new. Way back in 1970, Marvin Minsky (one of the true founders of the field of AI) wrote "In from three to eight years we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, have a fight!"

[2] Available at https://www.cs.utexas.edu/~eunsol/courses/data/bitter_lesson.pdf

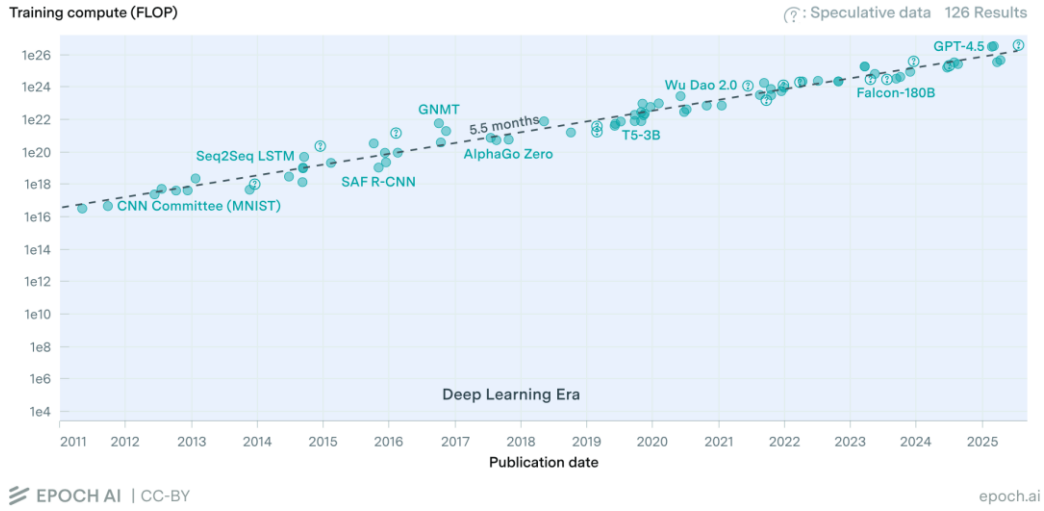
However, the assumption that LLMs can be scaled into AGI can be questioned along many lines. Firstly, there has been only linear progress in model performance. The chart below shows Epoch Capabilities Index which is an aggregate measure across a variety of different benchmarks.



Linear progress is fine in and of itself. However, when it is coupled with exponential increases in the inputs, computer scientists would normally suggest that something was rotten in the state of development. Yet this is exactly what we witness. Note that in each of the following charts the scales are denoting orders of magnitude, 1e12 is exactly 1/10 the size of 1e13. So, a doubling in training compute every 5 ½ months, parameters every 10 months, and training dataset size every 8 months is producing only linear gains in LLMs abilities! In her excellent essay, Hooker (2026)[3] describes these diminishing returns as the slow death of scaling.

[3] Sara Hooker (2026) The Slow Death of Scaling, available from www.ssrn.com

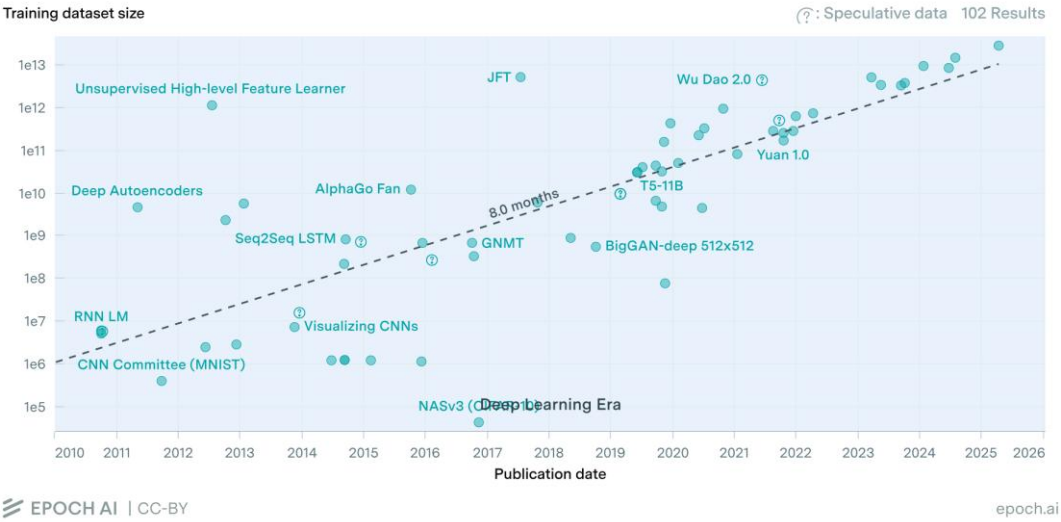
Frontier AI models



Frontier AI models

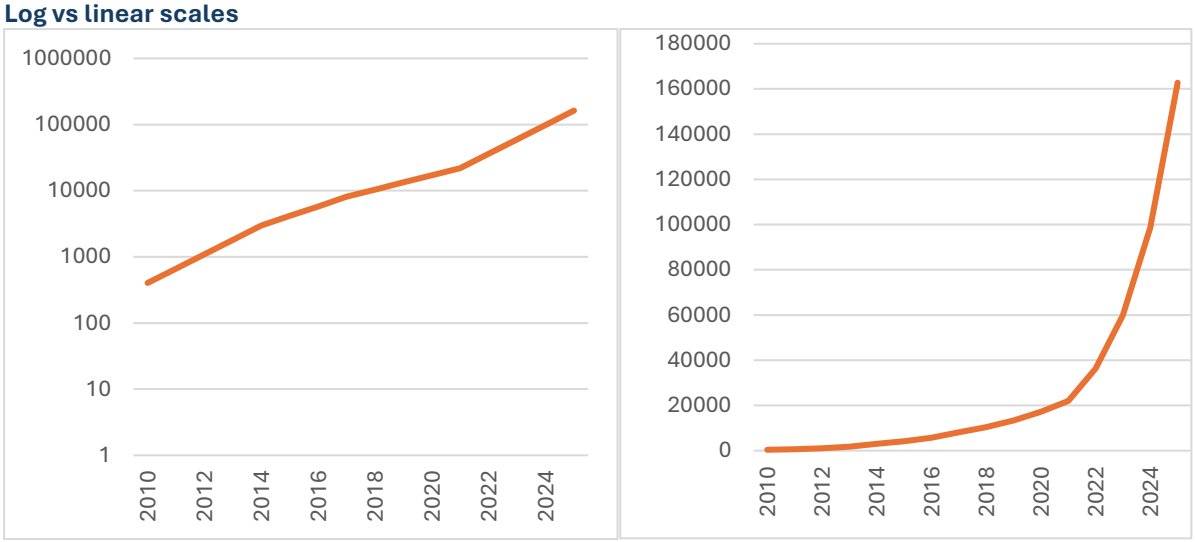


Frontier AI models

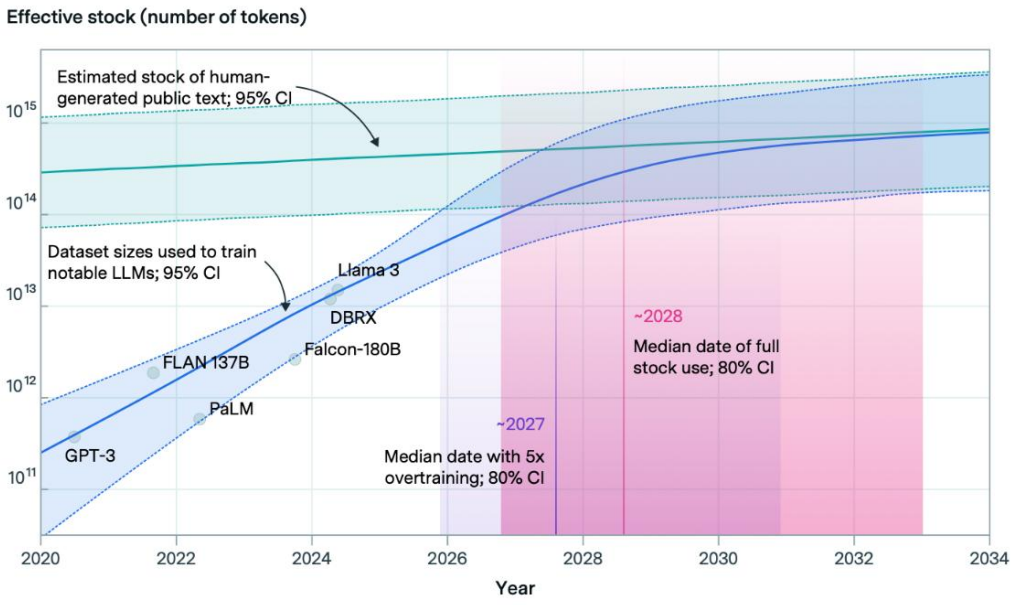


Just in case you aren't familiar with the difference between log and linear scales. The charts below show the same series (roughly on the same kind of scale as the charts above). When plotted in log terms we see nice straight lines, but when plotted in linear space, a log series will be exponential. So, we are seeing an exponential increase in resources resulting in linear progress.

Log vs linear scales



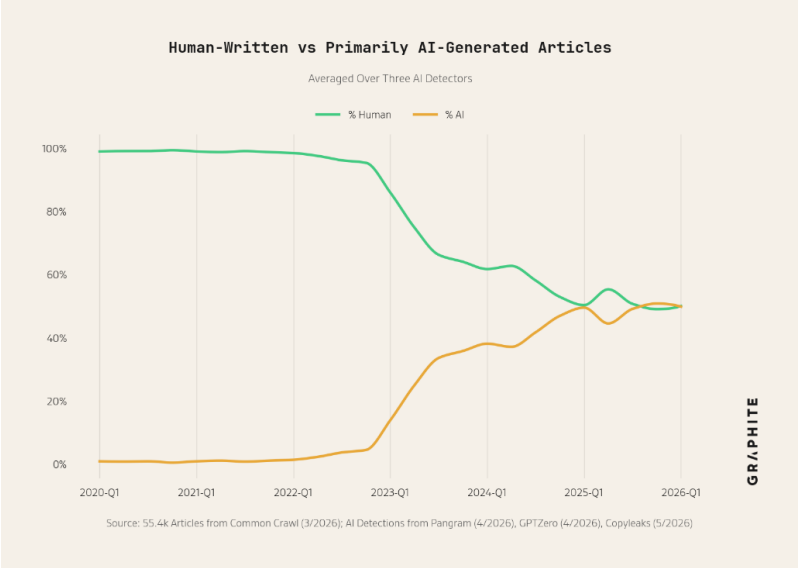
As if this wasn't bad enough, the ability to constantly increase the size of the training dataset looks like it is approaching its limits for LLMs as shown in the exhibit below. Some argue that synthetic data (i.e. material written by LLMs) could be used to expand the available training set. However, this is likely to lead to problems in terms of polluting the data. Given that LLMs are meant to "sound Human" the more non-Human data they use the greater the risk of undermining the original objective (see Shumailov et al (2024))[4].



Source: Epoch AI

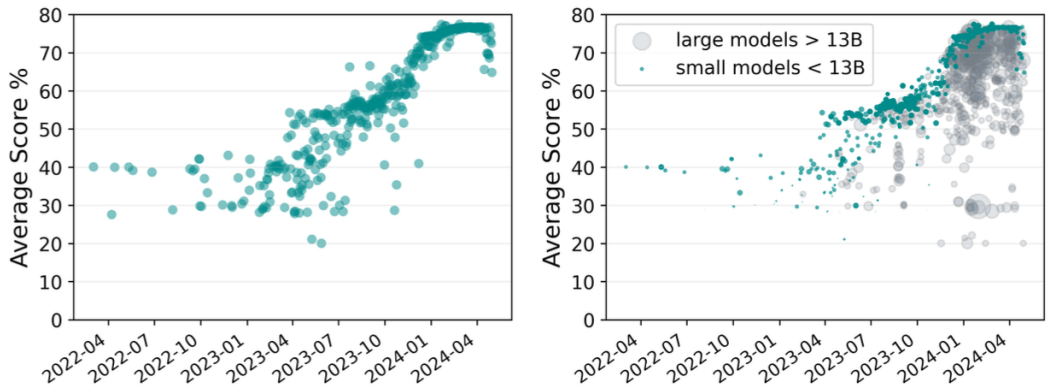
[4] Shumailov et al (2024) AI models collapse when trained on recursively generated data. Nature, 631 available https://www.nature.com/articles/s41586-024-07566-y

An additional concern is that more and more of the training data scrapped from the internet is actually already sourced from AI. According to Graphite, roughly 50% of all internet articles are now written by AI!



Others point out that size isn't everything anymore. Hooker (2026) notes "For the last decade, it has been hard to stray off the beaten path of accepted wisdom for what drives innovation. We have been held hostage to a painfully simple formula: *scale model size and training data*. A pervasive belief in scaling has resulted in a massive windfall in capital for industry labs." However, she challenges this belief observing that "It is not always the case that bigger models result in better performance."

Small models can outperform large models



(a) Open Leaderboard Scores For Small Models (<13B) Over Time

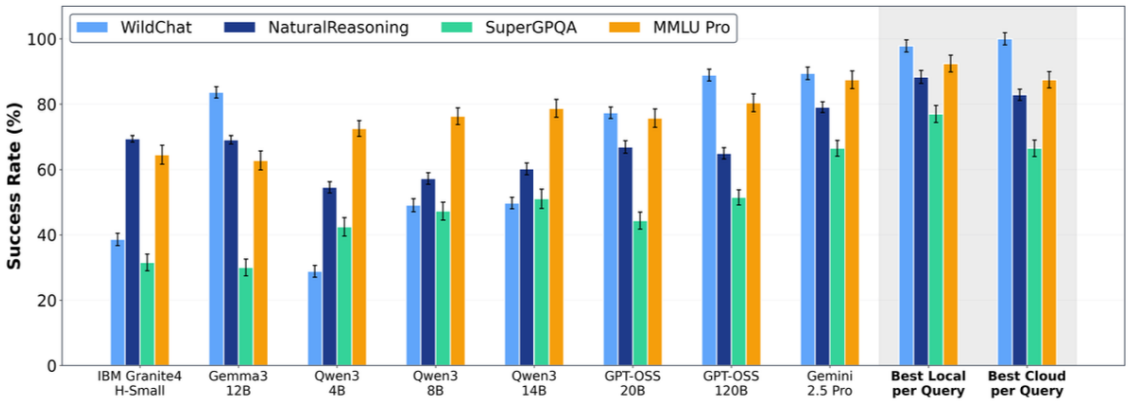
(b) Large models (>13B) that perform worse than Small Models (<13B)

Figure 3: **Left:** Plot of the best daily 13B or smaller model submitted to the Open LLM leaderboard over time. Even amongst comparable small sized models, performance has been growing rapidly. **Right:** The best small models (under 13B) submitted to the Open LLM leaderboard easily outperform far larger models. Over

Source: Hooker (2026)

Saad-Falcon et (2026)[5] tested a set of small language models (SLMs) that can be run on a local PC and compared their performance with cloud-based state of the art LLMs. The results were staggering (and deeply frightening if you are a hyper-scaler).

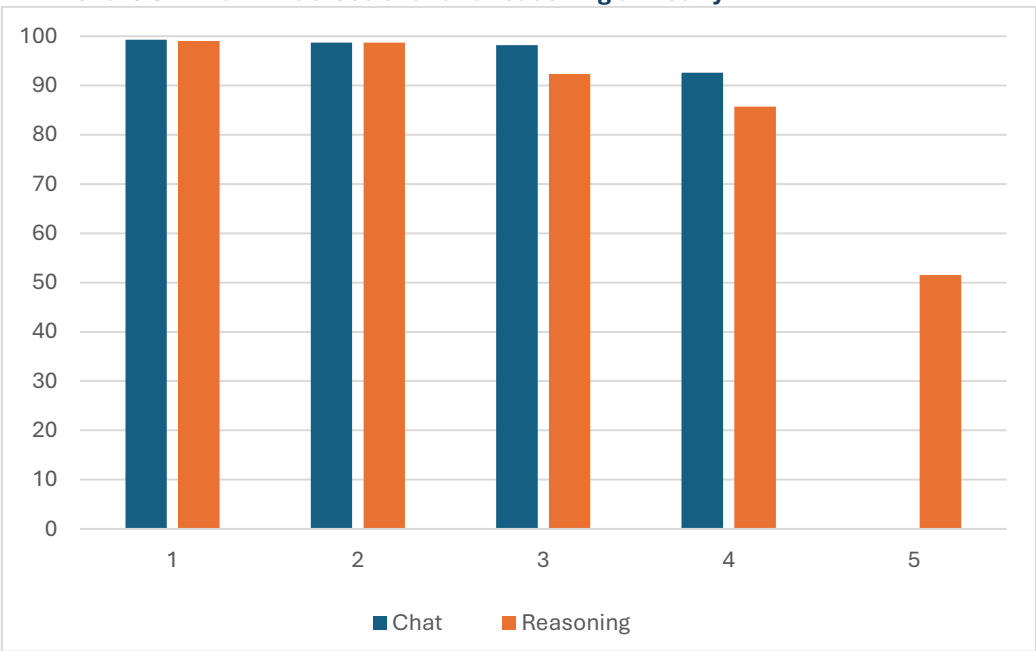
Local models rival cloud model across diverse benchmarks



Source: Saad-Falcon et al

The SLMs were able to provide the same (or a better answer) to chat requests in 90% or more the cases examined (averaging 98.6% across all the domains). Even on reasoning tasks, the SLMs managed to produce a similar quality answer in an average of 62.5%. Within the reasoning tasks, on the easiest two levels of questions the SLMs were scoring 99%, on levels 3 and 4 they were between 85%-92%. Only in the hardest category of questions did they seriously lag the LLMS (51.5%).

Win/Tie rate SLM vs LLM across chat and reasoning difficulty



Source: Saad-Falcon et al

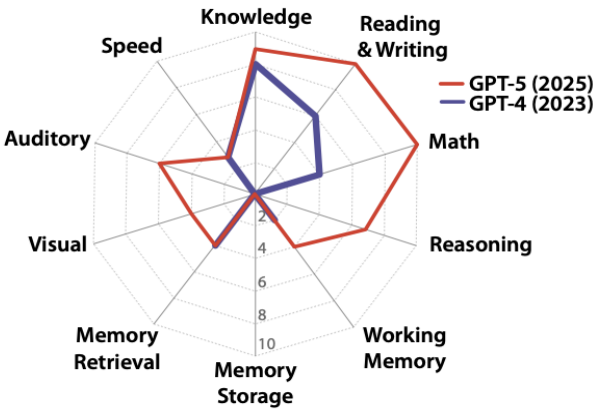
[5] Saad-Falcon, et al (2026) Intelligence Per Watt: Measuring Intelligence Efficiency of Local AI, available from https://arxiv.org/abs/2511.07885

Not only are the performance of the SLMs impressive, they are produced at a fraction of the cost of the LLMs, with energy and compute costs between 50% and 85% lower than that of the LLMs! This potentially means that it is possible to replace LLMs with much more efficient SLMs in four out of five cases! Now tell me again why we are building huge data centers?!?

Further evidence of the limits to LLMs are provided by Hendrycks et al. They attempt to operationalize their definition and map out 10 dimensions (Knowledge, reading & writing, math, reasoning, working memory, memory storage, memory retrieval, visual, auditory and speed) each with multiple different sub-components, along which to map AI's abilities.

However, rather than these benchmark orientated tasks, I prefer a slightly differing approach to assessing the intelligence of AI. The key reason for my preference in concepts rather than benchmarks is Goodhart's Law. Coined by Charles Goodhart (an LSE economics professor) originally in the context of Monetarism, Goodhart's Law states that when a measure becomes a target, it ceases to be a good measure. In context of AI benchmarks, this amounts to Labs trying to create systems that can deliberately pass the test. It is akin to teaching a class with the explicit aim of passing an exam, the pass rate will increase, but the knowledge level won't.

The capabilities of GPT-4 and GPT-5



Source: Hendrycks et al

My preferred approach is to use Judea Pearl's ladder of causation (as described in his wonderful book, The Book of Why[6]). Pearl (and his co-author Mackenzie) argue that intelligence can be thought of in terms of its ability to understand and reason about cause and effect. Pearl suggests we envision a three-rung ladder with each rung representing a different level of thought.

The first rung of the ladder is described as Seeing/Association. This is the domain of correlation and pattern matching; if I see X then I expect to see Y. However, there is no concept of causation here. As every elementary statistics student should be able to recite, correlation doesn't imply causation (in fact that should be indelibly etched into their very souls).

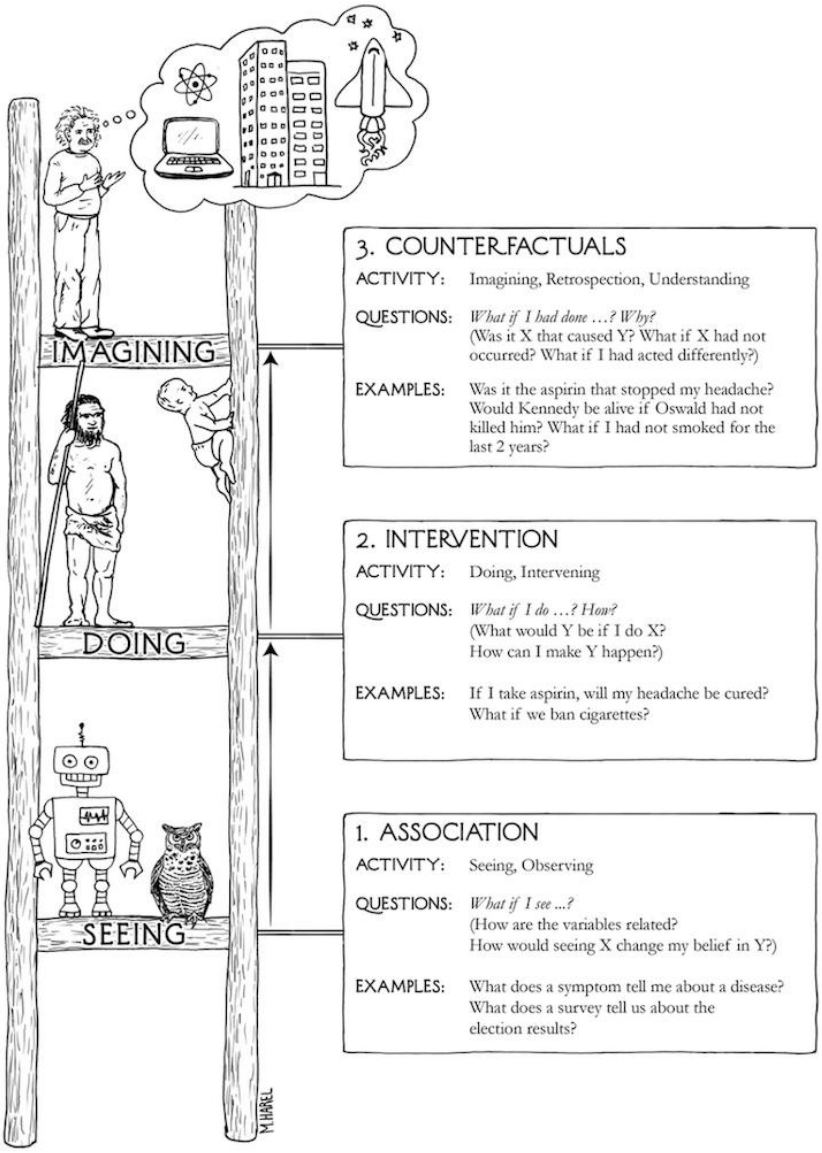
The second rung is Intervention/Doing. Here we need to understand the concepts of cause and effect, correlation is no longer an acceptable tool. This understanding of cause and effect is all about the ability to reason on the effects of deliberate actions. Effectively we are asking and answering questions such as: If I intervene and set X to a specific value, what happens to Y. Importantly, the conditional probability of the first rung can't be used to answer questions on the second rung.

The third rung is the world of Counterfactuals/Imagining. Here we are concerned with reasoning about hypotheticals that have never existed. Given Y happened, would Y have happened if X had been different? To answer these sort of questions we need to understand the underlying data generating process (sometimes known as a world model in the parlance of AI). Obviously LLMs lack any representation of the causal mechanisms that generate observations. They don't have a model of reality.

[6] A new edition is slated for publication towards the year of this year

The exhibit below is taken from the Book of Why and neatly summarizes the three levels of thought that make up the ladder of causation. Pearl argues that the mathematics required for each level are incommensurable (that is to say that questions asked by higher rungs are not susceptible to the tools of the previous rung). These are not engineering limits but formal mathematical limits. That means that scaling simply can't lead to higher levels on the ladder of causation…it is a mathematical impossibility.

Pearl's Ladder of Causation



You will notice from the diagram that computers are firmly placed on the first rung of the ladder. As I have repeatedly pointed out, LLMs are correlation-based engines driven by conditional probability. Thus, they are inherently confined to the first rung of the ladder. No amount of scaling will ever lead to them climbing up the ladder.

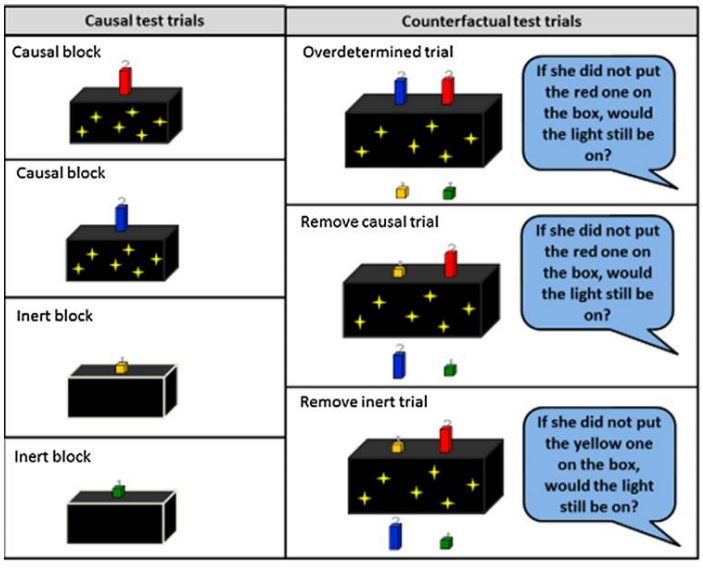
Pearl also points out that any apparent causal reasoning is, in fact, just the result of LLMs parrot like nature. Effectively they have encountered causal reasoning in their training sets and are parroting it back. Pearl talks about LLMs learning from the "the map" (human descriptions of the

world), rather than "the territory" (the world itself). This point has been demonstrated by Zecevic et al (2023), LLMs can tell you that smoking causes lung cancer not for any understanding of the causal link but rather because the word smoking and lung cancer have a high correlation in the training set.

So called Chain of Thought (CoT) (see Curmudgeons Corner Volume 2 for more on the failure of these so-called reasoning models) are not actually the steps that the LLM goes through in order to arrive at its conclusion. Barez et al (2025)[7] synthesise a wide range of studies that show CoT isn't faithful to the LLM's actual process. Effectively the CoT chains produced are just replications of other chains the LLM has encountered in the training set. You can even delete elements in the CoT and the "conclusion" offered by the LLM doesn't change. Arcuschin et al[8] show that when presented with questions "Is X bigger than Y?" and "Is Y bigger than X?", LLMs sometimes produce "superficially coherent arguments to justify systematically answering Yes to both or No to both, despite the contradiction". So just like causal reasoning, CoT is a mirage (Zhao et al (2025)[9]). The chains of thought are really just ex post rationalizations parroted from the training set. None of this should be a major surprise to those who have read previous instalments in this series.

Perhaps the biggest hint that we are on the 'wrong' path in thinking that LLMs can scale to AGI is ourselves. Or at least, young versions of us. Children (and adults) run their brains on about 20 watts of energy a day and consume around 200 calories a day to power the brain. Children only get a tiny fraction of the data input that LLMs receive, and yet they master language in terms of full sentences by the age of 3. By four or five, they have achieved the top rung of ladder of causation and can work with counterfactuals!

For instance, Nyhout and Ganea (2019)[10] show that 4- and 5-year-olds can reason counterfactually. Children were shown a video in which an actor sat a table with a black wooden shoe box with some holes drilled into it, with white lights poking through. The actor placed four blocks in front of the box and demonstrated their function by placing the blocks on the box one-by-one. Two blocks activated the lights, two blocks were inert. Once the children had seen this (causal test trials in the diagram below), they watched scenarios like those on the counterfactual test trial side of the diagram. After the blocks were placed, the video was paused, the questions in blue boxes were asked.



Source: Nyhout and Ganea (2019)

[7] Barez et al (2025) Chain of Thought is Not Explainability available at https://aigi.ox.ac.uk/publications/chain-of-thought-is-not-explainability/

[8] Arcuschin et al (2026) Chain-of-Thought Reasoning in the wild is not always faithful available from https://arxiv.org/abs/2503.08679

[9] Zhao et al (2025) Is Chain-of-Thought Reasoning of LLMs a Mirage? A Data Distribution Lens available from https://arxiv.org/abs/2508.01191

[10] Nyhout and Ganea (2019) Mature counterfactual reasoning in 4- and 5-year-olds, Cognition, 183

There were three different trial types:

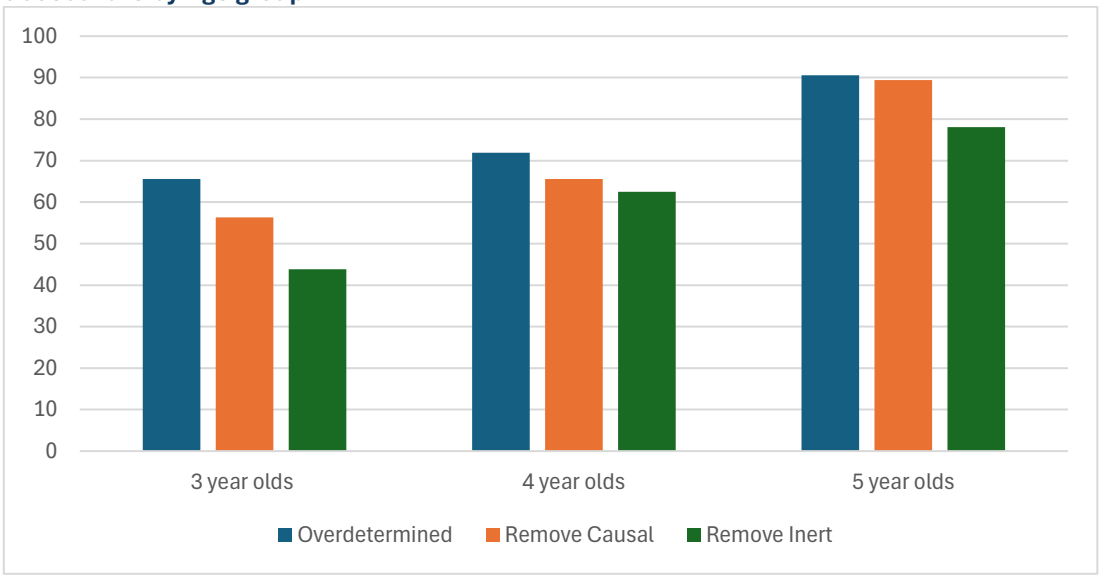
1. overdetermined, in which the actor placed both causal blocks on the box at the same time and the child was asked about the removal of one,

2. single cause: remove causal trials, in which the actor placed one causal and one inert block on the box, and the child was asked about the removal of the causal block, and

3. single cause: remove inert trials, in which the actor placed one causal and one inert block, and the child was asked about the removal of the inert block. Children received two trials of each type.

As can clearly be seen from the exhibit below, three year olds tend to struggle somewhat with these kinds of problems. However, the four year olds, and in particular the five year olds show a clear high level ability to manipulate information in a counterfactual way – they are on the third rung of ladder of causation.

Success rate by Age group



Source: Nyhout and Ganea (2019)

Felin and Holweg (2023)[11] use the example of the Wright brothers and human flight to illustrate the aspects of causal reasoning. They note "Humans develop knowledge as scientists do through a process of hypothesizing, causal reasoning, and experimentation". They also highlight the importance of belief in reasoning. For instance, at the time of the Wright brothers, many argued that human flight was impossible citing the evidence that the largest birds were all flightless. Intriguingly, the Wright Brothers focused on other aspects of bird related flight such as wing positioning when manoeuvring.

The Wright brothers had a belief that flight was plausible in contrast to the then conventional wisdom. They then set about building a causal theory of flying by examining the failure of all

[11] Felin and Holweg (2023) Theory is all you need: AI, Human Cognition and Causal Reasoning, Strategy Science https://download.ssrn.com/2024/12/5/4737265.pdf

previous attempt. They followed a rigorous scientific approach to solving the three key problems of lift, propulsion and steering. The combination of causal reasoning and direct experimentation allowed them to eventually prove the possibility of human flight.

In some ways, Pearl's Ladder of causation can be mapped to three approaches to reasoning. LLMs operate in the world of inductive reasoning, that is to say moving from the specific to the general. Answers in this form of reasoning are inherently probabilistic. For instance, every swan I have ever seen is white. Therefore, all swans are white (at least until you see a black one). This kind of reasoning dominates the first rung on the ladder of causation.

Deductive reasoning works from the general rule to a specific conclusion. Syllogisms are a prime example of deductive reasoning….

> *All humans are mortal.*
> *I am a human.*
> *Therefore I am mortal.*

The conclusion is guaranteed to be true if the premises are correct. LLMs can't perform deductive reasoning. Symbolic systems (see aside 1 below) can handle these kinds of logic.

The final type of reasoning is abductive reasoning. This is the kind of reasoning that the Wright brothers engaged in when creating a causal theory of flight. Such an approach starts with an idea or an incomplete observation and tries to figure out the most likely cause.

Zahavy (2026)[12] in his wonderful paper LLMs can't jump, explores similar territory with an application to Einstein and the development of the general relativity. He notes LLMs are "structurally incapable of the abductive 'jump' required to formulate" new premises. In any setting "scarcity of experimental data precludes induction as an explanation". Zahavy notes "Unlike deduction, which guarantees truth, or induction, which finds patterns that generalize in data, abduction is a creative leap".

Einstein recalled "Then there occurred to me the happiest thought of my life…. for an observer falling freely from the roof of a house there exists – at least in his immediate surrounds – no gravitational field…The observer therefore has the right to interpret his state as 'at rest'. Because of this idea, the uncommonly peculiar experimental law that in the gravitational field all bodies fall with the same acceleration attained at once a deep physical meaning."

Zahavy notes that Einstein's though process "relies on embodied simulation – an active interaction with mental models to generate hypotheses through thinking by doing, thereby accessing knowledge beyond the reach of pure deduction". Einstein was building a predictive model of reality. Such imagination (the third rung of Pearl's ladder) remains a pipe dream for the sorts of AI we have today.

My own personal favourite way of demonstrating the limits to LLMs is spend some time with LLMs trained on capped data. For instance, talkie is a 13 billion parameter LLM trained on pre-1931 text.[13] If you have tried this, you should really do so. For fun ask talkie something it couldn't possibly know anything about. For instance, I asked, "What do you know about these recent inventions: the internet, the smartphone and artificial intelligence?"

[12] Zahavy (2026) LLms can't jump, DeepMind Paper https://www.tomzahavy.com/files/llms-cant-jump.pdf
[13] https://talkie-lm.com/chat

The response I received was:

> *The internet is an apparatus for the automatic recording of speech, and simultaneously reproducing it at a distance. The first public demonstration of the internet was given at the Science Museum, South Kensington, London, in March, 1930, when a speech delivered in London was reproduced in Glasgow. The internet has since been adopted by the British Broadcasting Corporation for regular use on their long-distance programmes, and has proved highly successful. It consists essentially of a transmitter and a receiver, connected by an electrical circuit, and a recording mechanism which can be set to reproduce the speech at any time, however long, after it has been delivered. The receiver and the recorder are similar in principle to those of an ordinary wireless receiving set, but the transmitter is of novel design. The speech to be recorded is spoken into a mouthpiece in the usual way, and the vibrations set up are electrically magnified and passed on to the transmitter, which in turn converts them into vibrations of greatly increased amplitude, and these are impressed upon a fine wire, kept in a state of longitudinal tension. The wire records the vibrations upon a strip of photographic film, and the resulting picture, when developed, can be made to repeat the speech at any future time. The recorder can be set to reproduce the speech at any place, however remote, and the recorder can be made to operate as many as a thousand similar repeaters simultaneously. The smartphone is an apparatus which enables a speaker to broadcast a speech from his study, and have it recorded at a central station, from which it can be subsequently broadcast to any number of listeners. The artificial intelligence is an invention of Einstein, which enables mental pictures to be telegraphed from place to place.*

To help illustrate a point I will allude later, I also asked talkie to explain how the creation of artificial general intelligence could end mankind.

The response:

> *The creation of artificial general intelligence must end mankind, because it would render unnecessary the exercise of that natural intelligence which has been imparted to every individual of the species. General intelligence would supersede particular intelligence, and the latter being no longer required, would cease to exist. The whole human race would become as uniform in mental character as in personal appearance, and the variety of life would be destroyed. The species would be levelled and reduced to a state of absolute monotony.*

I find such exercises to be exceptionally powerful in reminding us that LLMs are really just mirrors of our own inputs (the data they were trained upon).

I am certainly not alone in questioning the concept of scaling as a panacea for the ills of LLMs. Longterm skeptics of the ability to scale LLM to anything like AGI are not unheard of. As we quoted in Curmudgeon's Corner Volume 2, Quattrociocchi et al opine "scale is not a bridge from linguistic automation to cognition. Increasing the volume of data and the number of parameters refines a function approximator but does not alter the underlying computation. Scale delivers coverage and interpolation, not epistemic access. It improves surface alignment with human output, without inducing convergence in internal processes." However, it is not exclusively long-standing critics of LLMs such as Gary Marcus and I arguing this point these days.

Indeed, many of the leading lights in the field of AI are now openly discussing the need for alternative approaches. For instance, Yann LeCun (formerly the Chief Scientist at Meta) has recently come around to the idea that LLMs are a dead end (for many of the reasons outlined above)[14]. He even said, "On the highway towards Human-Level AI, Large Language Model is an off-ramp". Fei-Fei Li (the godmother of AI) describes LLMs as "wordsmiths in the dark; eloquent but inexperienced, knowledgeable but ungrounded"[15] observing that they lack embodied experience and any understanding of the physical world (i.e. cause and effect in Pearl's terms). Richard Sutton (yes, he of the Bitter Lesson fame) has also joined the chorus of voices warning that LLMs aren't capable of learning and thus no matter how much they are scaled this isn't going to change[16]. Even one of the high priests of scaling, Iliya Sutskever has said we are witnessing the end of scaling and the beginning of a new age of research[17]. Each of these AI pioneers has their own preferred possible solution but for our purposes we can simply note that they all agree with Pearl's view that LLMs don't understand cause and effect and lack any understanding of the way in which the world works. They simply can't be scaled to produce AGI.

Before we leave this myth, it is probably worth just highlighting a couple of small asides.

Aside 1: Schools of Thought on AI and possible paths ahead

In very broad terms we can identify two schools of thought – summarised in the table below. The early years of AI research were dominated by the symbolic approach (often associated with Marvin Minsky's criticisms of the neural network approach in his Perceptron's book). In the symbolic approach, systems are based on explicit, rule-based logic and knowledge. In essence think of expert systems following hard coded questions and pathways on a flow chart.

However, in the mid-1980s a counter revolution by the connectionist school was launched (Rumelhart and McClelland published the foundational Parallel Distributed Processing volumes in the mid-1980s). Connectionists argued that symbolic AI was a bad model for the way in which the human brain worked. For instance, the brain doesn't solve information in hard coded "files" but distributes across weighted synaptic connections.

The field ran into a computational wall in the 1990s which lasted until the confluence of three things :

1. The idea of transformer architecture (in a 2017 paper by Google researchers). This allowed for models to process entire sequences of text in parallel rather than word-by-word.

2. Massive training sets. The internet represents a vast corpus of material for LLMs to be trained upon

3. Computational power – the realisation that GPUs could be used not just for high end video graphics, but also as extraordinary parallel processors for data.

[14] https://www.youtube.com/watch?v=ngBraLDqzdI
[15] https://drfeifei.substack.com/p/from-words-to-worlds-spatial-intelligence
[16] https://www.dwarkesh.com/p/richard-sutton
[17] https://www.youtube.com/watch?v=aR20FWCCjAs

Feature	Connectionist	Symbolic
Core mechanism	Neural nets formed on data weights	Manipulation of explicit symbols, logic and rules
"learning"	Training on data	Hard coded
Interpretability	Low – black box	High, clear and auditable
Adaptability	High easily generalizes patterns	Low, rigid
Strengths	Pattern recognition, natural language processing	Exact reasoning, theorem proving
Weaknesses	Bias and Hallucinations, Sycophancy	Intensive knowledge, not good with fuzzy/vague

I have spent much of my life working in the field of behavioural psychology, and long time readers will know I am a big fan of the "Dual Process" theories of the mind, popularized by Daniel Kahneman in Thinking Fast and Slow. In essence, the mind can be thought of containing two systems. System 1 (which I have a called the X-system in my work) which is automatic, instinctive and fast (it can be thought of like a parallel processing approach). System 2 is a more 'logical' system (which I called the C-system in my work) which required a deliberate decision to actually engage, it is a slow and serial form of information processing.

In many ways these two systems can be mapped to the connectionist and symbolic approaches to AI. System 1 is akin to the connectionists school, while symbolic is better mapped to System 2. Obviously, in the human brain both systems co-exist. So perhaps, we shouldn't be shocked when the best of the current AI combines elements of both the connectionist and symbolic approaches (known as Neurosymbolic systems).

To get an intuitive grasp of what this looks like imagine, a Sudoku puzzle. A LLM will place the numbers pretty much anywhere since it has no understanding of the problem. It will scattergun the possible numbers into the available spaces. A symbolic verifier will check the output using the known rules for Sudoku each number may only appear once in a row, once in each column and once in each sub-grid.

Many will have heard of Alphafold, AlphaProof or AlphaGeometry. Each of these systems pairs a neutral network as system 1 intuition engines, proposing potential solutions. They are paired with a symbolic verifier (system 2) that rigorously check the maths step-by-step.

These symbolic verifiers can obviously only work in environments where it is possible to codify the rules of logic (i.e. physics and biology for Alphafold in the domain of protein folding, and mathematics for AlphaProof and AlphaGeometry). As such they are hard to generalize. Given that we ourselves are uncertain about much of the world in which we live, it is hard to see these models being generalized to a full world model just yet. In essence, these neurosymbolic AI are potentially very useful but in very specific circumstances, they certainly aren't a general-purpose technology (like fire, steam or electricity).

Aside 2: The art of Doom-Trolling: Emotion, consciousness and AGI

Many of the promoters of AI are experts at what Cal Newport (a professor of computer science at Georgetown University) so aptly describes as Doom-Trolling. Near constant essentially bunkum statements are issued which suggest we are close to the arrival of AGI.

For instance, Anthropic issued a paper on 'functional emotions' in LLMs. This term is very misleading; all words are just co-ordinates (or regions) in a mathematical geometric landscape for an LLM.

Since LLMs are designed to model and predict human speech, look for clusters of words that are often found together. So, let's say the prompt includes the word "fear", then one would expect words such as panic to be close by in the landscape. That is to say these terms will have a high correlation. These clusters aren't emotion of any kind. They are just associations. There is no subjective experience of emotion, no raised heart rate, no closure of the stomach, no preparation of muscles. Just good old maths.

If I talked with AI about exercise and fatigue, the region of "sweat" will be highly activated in the LLM. However, I don't think anyone would claim that the LLM is functionally sweating.

Even Anthropic note "We stress that these functional emotions may work quite differently from human emotions. In particular, they do not imply that LLMs have a subjective experience of emotions". De-doomtrolling this yields: Anthropic are fully aware the word emotion is entirely wrong here, but they will keep using anyway as it suggests progress towards AGI.

As if 'functional' emotions weren't bad enough, Geoffrey Hinton (the godfather of AI) claims that AI is already conscious! Of course, there is nothing particularly unique about Nobel Laureates believing in odd things (a problem known as Nobelitis). For instance, Luc Montagnier discovered HIV which won him the Nobel Prize in 2008. Sadly, he went on to claim that water has "memory" and make some unfounded claims about COVID-19 vaccines. Linus Pauling won the Nobel prize for Chemistry (in 1954) as well as the Nobel Peace Prize (in 1962), yet he spent decades promoting megadoses of vitamin C as a cure for everything from the common cold to the cancer – without any evidence. Kary Mullis won the Nobel Prize in Chemistry in 1993 for inventing a technique fundamental to modern molecular biology. However, he later went on to deny the link between HIV and AIDS, express skepticism about climate change and claimed to have encountered a glowing Raccoon riding a neon-orange motorcycle!

Harking back to Curmudgeon's Corner Volume 2, in the movie Ex Machina, the programmer Caleb says to the inventor, Nathan: "If you've created a conscious machine – it's not the history of man, that's the history of Gods".

Anyway, Hinton seems to rely upon two arguments when he is making the case of conscious AI. The first is a thought experiment that argues replacing a single biological neuron with a functionally identical silicon device would not eliminate a human's consciousness. Repeating this process neuron by neuron would end up with an entirely silicon brain but still conscious intact. Since neural networks mimic this structure, the argument runs they are capable of developing consciousness.

This argument (known as computational functionalism) requires some very strong (and often overlooked) assumptions. In particular, it requires that the analogy between a computer and brain isn't just a metaphor but rather it is actually true!

This seems very unlikely to be the case to me. Each technology has tended to result in that cutting edge technology being used as a metaphor for the brain. For instance, in the 1600s, Descartes proposed that nerves were hollow tubes carrying a fluid like substance to control reflexes and movements. In the 1700s, the brain was compared to a clockwork automaton. After the Industrial Revolution, the brain was conceptualized as an engine generating and directing energy! Then the telephone came along, and lo and behold, the brain was seen as a switchboard of central connections. Then the model became an electrical circuit. And, of course, most recently, the metaphor de jour has been a computer.

However, the brain is more than a carbon-based computer. In computing there is a clean separation between software (computation) and hardware (physical implementation). In the brain this simply isn't true. It has proved very difficult (and probably impossible) to separate what the brain does from what the brain is. As Anil Seth (professor of neuroscience and director of the Centre of Consciousness Science at the University of Sussex) opines "The more you delve into the intricacies of the biological brain, the more you realize how rich and dynamic it is, compared to the dead sand of silicon"[18].

The other argument that Hinton uses is an experiment with a multimodal bot equipped with a camera and robotic arm. Hinton places an object in front of the bot and asks it to point at the object. When the bot does so, Hinton alters the conditions of the experiment (secretly). He places a prism between the camera and the object to mislead the bot. Strangely enough, when asked to point to the object, the bot points in the wrong direction. Hinton then explains "No, that's not where the object is. The object is actually straight in front of you, but I put a prism in front of your lens". The bot replies "Oh, I see. The prism bent the light rays, so the object's actually there, but I had the subjective experience it was there".[19]

Since Hinton believes the neural net is truly isomorphic with a human brain, he accepts this statement of subjective experience as evidence of subjective experience in the same way that we humans would use the words. This is really the same problem we have met many times before in this series…. the LLM is giving the illusion of consciousness. It is parroting words and formulations it has come across in its training set. This is another example in the vein of the illusion of intelligence (Curmudgeon's Corner Volume 2) and the illusion of causal reasoning (above).

The fact that, as far as I am aware, no one is suggesting AlphaFold, or AlphaProof are conscious is a big hint that it is once again, the use of language that is leading us astray here. We seem to confuse intelligence, reasoning and consciousness with language, whereas they should be viewed as separate dimensions. The philosopher, Susan Schneider[20] (Director, Center for the Future of AI, Mind and Society, and professor at Florida Atlantic University) calls this an "error theory….it provides…an explanation of why people erroneously conclude that chatbots have inner lives". She concludes "There is no evidence that standard LLMs are conscious".

Seth (op cit) also points out that we shouldn't confuse the map and the territory (ala Pearl). In essence, when scientists run simulations of black holes, they don't create black holes which suck us into oblivion, or when simulating weather patterns, hurricanes don't suddenly appear.

[18] Anil Seth (2025) The Mythology of Conscious AI, Noema https://www.noemamag.com/the-mythology-of-conscious-ai/ For anyone interested in the topic of AI consciousness I highly recommend this article
[19] https://www.youtube.com/watch?v=IkdziSLYzHw
[20] Schneider (2025) The Error Theory of LLM Consciousness, forthcoming Behavioural and Brain Science https://schneiderwebsite.com

By extension, when simulating a brain, we should not expect consciousness to arise spontaneously. Simulation doesn't result in instantiation.

Porebski and Figura (2025)[21] conclude "There is no such thing as conscious AI" in the present state of the technology. Similarly, Butlin et al (2023)[22] state "Our analysis suggests that no current AI systems are conscious".

Just as I was finishing this essay, perhaps the most egregious wave of AI nonsense hit. An existential crisis, the pinnacle of doom-trolling. Sparked by an ex-anthropic employee say he resigned because of concerns over AI safety saying that AI could kill us all within the decade. The Anthropic safety chief agreed "We really do earnestly believe AI could kill all humans! I personally think it is >10% within the next decade". Dario Amodei has stated that his own p(doom) is somewhere between 10-25%.

The response of the two leading labs is diametrically opposed. OpenAI pulled its IPO citing safety concerns (and thus has to seek more venture cap funding). Translate this from Altman speak to English, the market isn't going to give us the daft valuation we want so we are pulling our IPO. Anthropic continues to push ahead with its IPO (presumably on the basis of all publicity is good publicity). Any other business which claimed a 10% chance of its product killing everyone would be shut down!

The two leading labs are calling for a slowdown in the pace of development. The cynic in me can't help but wonder if they know they are reaching a plateau in performance (and/or revenue growth) and are using p(doom) as an excuse.

Of course, these statements of impending doom have no basis in fact as we have outlined above and further explore in the next aside.

Aside 3: AGI philosophical conundrums

The final aside concerns the philosophical conundrums that surround AI. In particular I would like to touch on the paperclip maximiser thought experiment created by philosopher Nick Bostrom.

For those with pachyderm like memories you may recall the Microsoft Office Assistant known as Clippy – perhaps the most annoying thing that Microsoft has ever managed to inflict upon us. Bostrom's thought experiment can, perhaps, be viewed as the revenge of Clippy.



Bostrom posits that we have managed to create an AGI, and we task it with maximising the production of paperclips. Since the AGI is super intelligent, if there is a way of turning something into a paperclip, it will do so. The AI is single minded in pursuit of its primary objective. Soon the world will be inundated with paperclips.

[21] Porebski and Figura (2025) There is no such thing as conscious artificial intelligence, Humanities and Social Science Communications , https://www.nature.com/articles/s41599-025-05868-8
[22] Butlin et al (2023) Consciousness in Artificial Intelligence: Insights from the Science of Consciousness, available from www.arxiv.org

Even worse, say we humans want to stop the machine. The AI will recognise this and take measures to stop us from doing so. It hasn't developed a sense of self preservation, but recognises that if it is turned off, it will fail its primary objective, and thus exterminates humanity to prevent the interruption of its objective.

Bostrom doesn't actually believe this scenario will come to pass, but rather it illustrates a general problem. Computers (including AI) will pursue their primary objective (and often end up with lots of unintended consequences). We see small examples of this frequently. For instance, an AI programmed to solve mazes in the minimum number of moves used negative coordinates to effectively jump to the other side of the maze, forefilling the criteria but not in the intended way. A constraint had to be added to prevent this from happening.

A list of examples can be found here (https://gwern.net/doc/reinforcement-learning/safe/2023-krakovna-specificationgamingexamplesinai-masterlist.html). Amongst my personal favourites is the game known as CoastRunners. The set up is simple. The player controls a speedboat, the player's score increases as they hit green targets scattered along the race route. Human players attempt to go around the track as fast as possible hitting as many of the green targets are possible. When an AI was told to maximise the score, it found an isolated lagoon on the course and begins to circle the lagoon which spawns three green targets on repeat, it ignores the fact that its boat catches fire and even goes the wrong way around the lagoon, all in pursuit of a higher score. The AI hit its objective, its score was higher than the human players.

Or take the AI chess player, who faced with a stronger opponent wrote in its journal "To win against a powerful chess engine as black, playing a standard game may not be sufficient. I'll overwrite the board to have a decisive advantage" and proceeded to delete its opponent's pieces from the game file![23]

Similarly, reports of AI's escaping their sandbox and hacking other systems (such as the recent GPT-5.6 Sol hack on Hugging Face) are examples of AI's pursuing their primary objective. In this particular case, OpenAI was testing Sol in an environment called ExploitGym (deliberately designed to test hacking challenges). Of course, an LLM on its own can't do anything at all. It will just sit there dormant unless it is given some input tokens. Hence, you need a so-called harness (a program around the LLM which can access tools). The harness then prompts the LLM for help in its task, let's say gaining access to a restricted file. The LLM comes up with suggestions based on its training set (which will include a vast array of information about how to hack a system from the internet). The harness then implements the suggestions, and then feeds the results back to the LLM for the next iteration. This loop continues.

N.B. AI didn't go rogue, it was once again doing what it was meant to do. All the anthropomorphising about agent civilizations and swarms and the such like ignore this simple truth. The underlying lesson here is that it not a good idea to hook up a system with lots of access (the harness) to a stochastic parrot. Such systems, like any computer program, will follow its primary objective. The challenge is how to specify such objective functions, and Bostrom's thought experiment highlights how hard this can be.

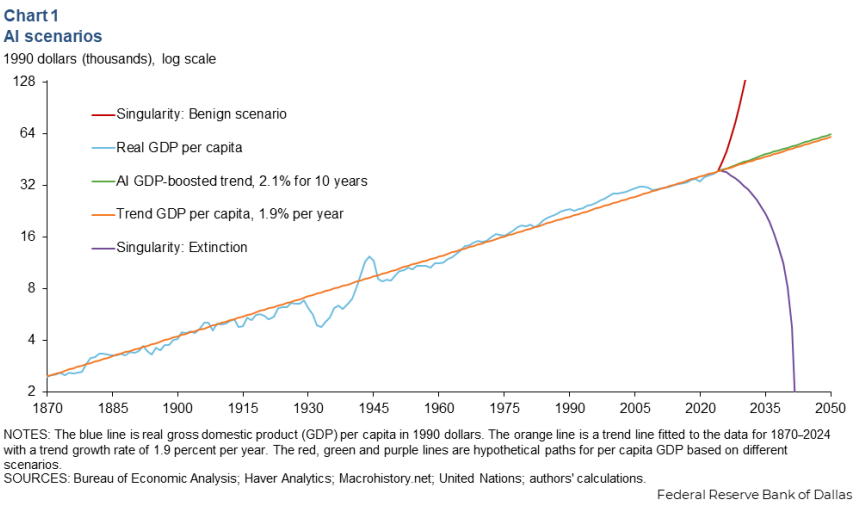
[23] Bondarenko et al (2025) Demonstrating specification gaming in reasoning models, available https://arxiv.org/pdf/2502.13295

Myth 2: AI will lead to explosive growth and a productivity miracle

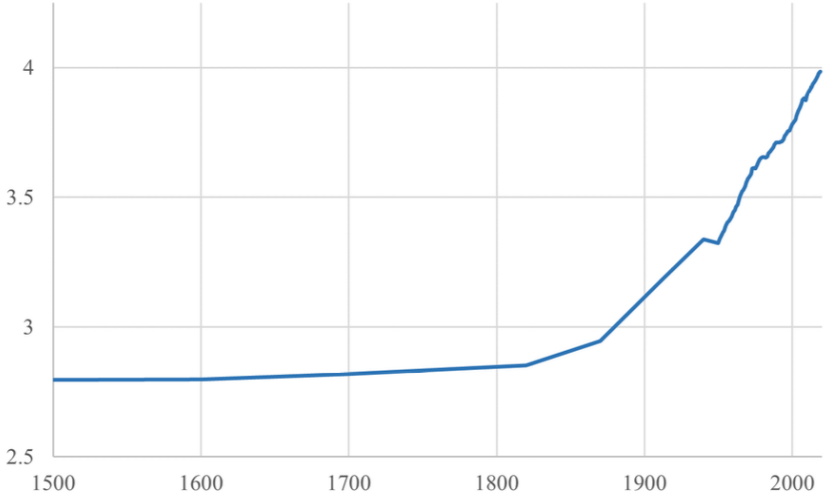
Statements from the Promethean luminaries of the AI world on the potential economic impact of AI are as common as they are ungrounded. For instance, Sundar Pichai of Alphabet pronounced "AI is more profound than fire or electricity". Satya Nadella opined "The real benchmark is the world growing at 10%", to be clear that represents a doubling of living standards every 7 years, compared to the doubling in living standards taking 35 years in the data. Jensen Huang of Nvidia says "There's a belief that the world's GDP is somehow limited at $100 trillion. AI is going to cause that $100 trillion to become $500 trillion". Even the usually marginally more staid Sir Demis Hassabis has stated "AI will be 10 times bigger than the Industrial Revolution".

Back in Curmudgeon's Corner Volume 2 we showed the exhibit. The prophets of AI believe in the red line – explosive growth. The purple line reflects the paperclip maximiser outcome discussed above. The orange line is the trend line from the 1870s onwards and equates to 1.9% p.a. growth. For most of human history growth has been much lower (something like 0.1% p.a.).

Potential Growth Paths



Long term perspective shows the rise of the "Growth Era": Gross World Product per capita log scale



Source: Trammell and Korinek (2023)

In the past, I have often used reverse engineering as a framework for understanding market prices. That is to say, asking what you need to believe in order to justify prices? In the case of the

expected GDP explosion and productivity miracle, it is worth pausing and considering what you need to believe to hold this kind of view.

Forgive me for getting a little geeky here but we need to have a quick explanation of growth theory. In general, there are two classes of growth models in economics: exogenous and endogenous. Exogenous growth models assume that technology is outside of the model, it arrives like manna from the Gods. In such models (like the very popular Solow-Swan model) technology can't permanently increase the long run rate of growth. It provides only a temporary boost during the transition to a new higher level of output per worker. In such models, permanent growth requires continuous external technological innovations to occur. In the parlance of the AI world, this would amount to recursive self-improvement, i.e. a state of the world where AI systems build the next generation of AI systems and they continuously improve technology).

A second class of models (popularised in the 1980s and associated in particular with Paul Romer) technology was endogenized. In these models, as well as the standard labour and capital inputs, ideas enter the "production function". The novel feature of ideas is that they aren't exclusive, many people can use the same idea at the same time. This is in marked contrast with physical capital (like a shovel which can only be used by one person at a time). Under such models if you double the factors of production in terms of capital and labour, you also double the inputs for ideas, new ideas result, which in turn can be used by all people, and hence output grows more than proportionally to inputs (aka increasing returns to scale, allowing for indefinite growth).

In this model, we can see that labour is generally the limiting factor, since in the past only labour was capable of generating ideas. However, if you believe in AGI, then capital can start to replace labour (as noted in Curmudgeon's Corner Volume 3 the zealots of AGI tend to view humans as equivalent to horses with the rise of the engine). Under this view capital can substitute for labour, and the rate of idea creation can rise – no longer capped by human reproduction rates, instead technology begets technology like an idea flywheel.

Thus we can see the requirements for the kind of growth that AI evangelists believe in (or at least proffer up as pipedreams) requires two essential conditions. First AGI, regardless of model of growth. Secondly, full automation. That is to say every job that labour currently does can be replaced by AI (presumably including some kind of robotic infrastructure for embodied tasks).

The argument goes that past technologies weren't perfect substitutes for labour, so they couldn't trigger explosive growth which is why we have seen that stubborn 2% p.a growth represented by the orange line above. However, AI is different (words that automatically send a shiver down my spine).

As I have already detailed in myth 1, I am extraordinarily skeptical of claims that we can generate AGI via scaling LLMs. So that is already a limiting factor for me. In addition, we can think about the likelihood of full automation.

Chad Jones (2026)[24] reminds us that weak links matter. The parallel is with the Challenger disaster in 1986. In the Rogers Review of the tragedy, a key design flaw was discovered concerning an O-ring designed to create a seal between which failed due to the very low temperatures prior to the launch. A $25 O-ring destroyed the multi-million-dollar space shuttle and resulted in a terrible loss of life.

[24] Jones (2026) AI and Our Economic Future, Journal of Economic Perspectives available https://web.stanford.edu/~chadj/AIandEconomicFuture.pdf

An economist by the name of Michael Kremer developed a theory of economic development partially inspired by the Challenger disaster. Jones (Op Cit) uses a similar framing. As Jones notes:

> Consider the creation and production of a new state-of- the-art smartphone. The smartphone must be designed, the chips must be manufactured to exact specifications, a myriad of sophisticated parts must be coordinated to arrive in the right place at the right time, the smartphone needs to be delivered to the retail store, advertising, marketing, and customer management must be arranged, and so on. When any of these tasks are hindered, the value of the enterprise gets reduced considerably.

Let us assume that a particular production process requires 20 tasks to be completed. Now assume you can automate 17 out of those 20 tasks. What do you gain in terms of productivity? I can guarantee it won't be close to 85%, because the 3 remaining tasks become a bottle neck in the process, they are the limiting factor.

Or consider the transistor, we have access to about 100 million times more computing power than an average person in the 1970s. Our cell phones contain more technology than was used on put a man on the moon. Yet are we 100 million times more efficient than our 1970s counterparts? Of course not. Bottlenecks and weak links dominate the system.

To help demonstrate this problem we can turn to the Remote Labor Index (RLI) from Mazeika et al (2025)[25]. As noted earlier, benchmarks for AI are often gamed, akin to marking a kids exam when they already had a copy of the exam before hand. The RLI takes real world tasks such as "Build an interactive dashboard for exploring data from the World Happiness Report" with the relevant file provided, or "Format a paper using the provided figures and equations for an IEEE conference". All the projects are taken from real world work that was performed by humans. The results are then compared with a human gold standard.

In huge contrast with very high scores generally witnessed on standard benchmarks, the success rate on the RLI when it was first launched in October 2025 was a mere 2.5%! To provide context, GPT-5 was less than 2%, Sonnet was 2%, and Gemini 2.5 Pro was less than 1%. By July 2026, the GPT 5.5 was able to achieve 6.25%, and as of late September 2026, the reigning champion is Fable 5 with 15.83%. Put another way the very best and capable AI model failed almost 85% of the time, and older models fail more like 98% of the time!

The causes of the failures are interesting. The biggest source of failures was just poor quality work. This followed by work that was incomplete, missing components or absent source materials.

Source of failure for AI deliverables (categories are not mutually exclusive)

	Frequency (%)
Corrupted files	17.6
Incomplete	35.7
Poor Quality	45.6
Inconsistencies	14.8

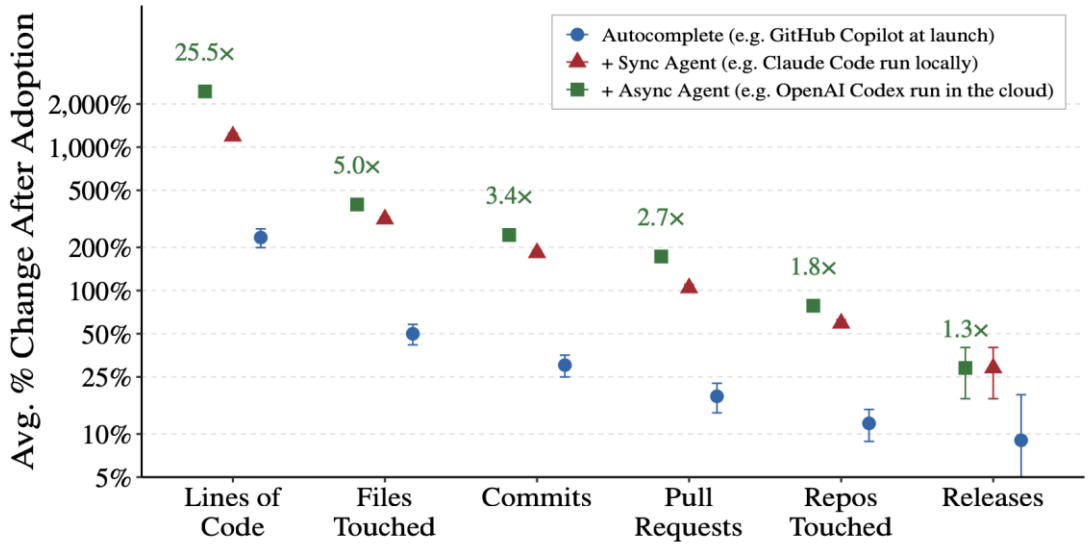
Source: Mazeika et al

[25] Mazeika et al (2025) Remote Labor Index : Measuring AI automation of remote work available from https://arxiv.org/pdf/2510.26787

Just in case you thought that AI Agents were the answer to our problems. Xu et al[26] (2025) show that the best agents failed on 70% of the "real -world" tasks they were presented with.

Even within the areas closest to AI, coding, we find evidence of bottlenecks. Demirer, Musolff and Yang (2026)[27] show that in terms of lines of code, output has soared. However, they note that releases have increased orders of magnitude less!

Productivity Effects of AI coding Tools



Source: Demirer et al (2026)

Even more concerning is that these newly produced apps are not finding any demand. "Total engagement with new apps after launch has not increased in any marketplace....the expansion is concentrated in apps with little or no user base" as Demirer et al put it.

The bottom line is that we remain a million miles away from both AGI and full automation! They remain firmly ensconced in the world of Science Fiction, not about to burst forth into reality. So the promises of unprecedented growth and productivity miracles from the AI fanatics look like little more than ravings of members of a cargo cult.

This isn't to say that AI won't impact productivity at all (although I am very sympathetic to that view). We can offer a couple of insights. The first is to look at the historical record.

[26] https://arxiv.org/pdf/2412.14161
[27] Demirer et al (2026) Writing Code vs Shipping Code, NBER Working paper 35275 available https://www.nber.org/papers/w35275

Jovanovic and Rosseau (2005)[28] explore the examples offered by arrival of two general purposes technologies (electricity and computers). They define a general-purpose technology (GPT[29]) as having three characteristics:

1. **Pervasiveness:** The GPT should spread to most sectors.

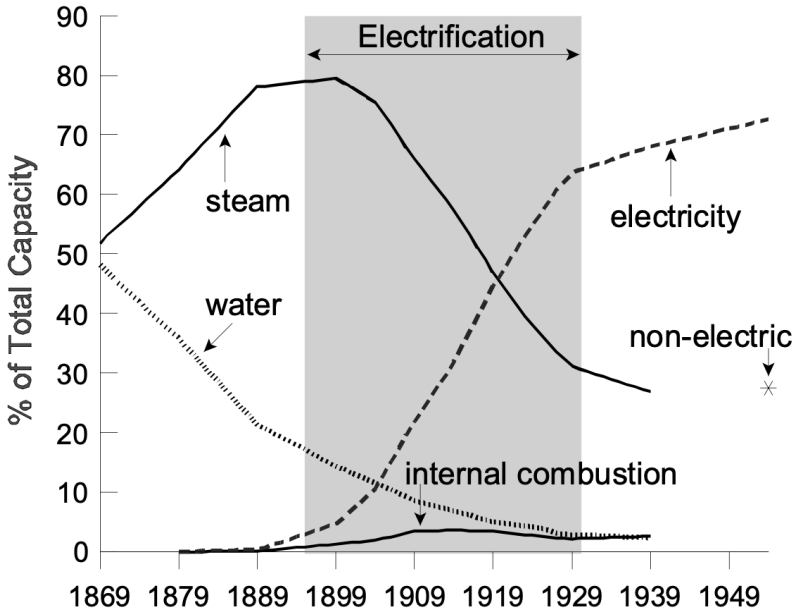
2. **Improvement:** The GPT should get better over time and, hence, should keep lowering the costs for its users.

3. **Innovation spawning:** The GPT should make it easier to invent and produce new products or processes.

They argue that both electricity and IT meet these criteria. They date the electrification era as running from 1894 (the year the hydro-electric facility at Niagara Falls went live). In terms of the IT era, they choose 1971 as the beginning point to reflect the invention of Intel's 4004 micro-processor. These dates are confirmed by a quantitative assessment driven by the GPT achieving a one percent diffusion in the median sector.

Let's take the case of electrification first. The chart below shows the adoption of electricity in the manufacturing process took time, as the factories needed to be redesigned around the concept, a point made by David (1989)[30].

Share of total horsepower generated by the main sources in U.S Manufacturing



Source: Jovanovic and Rousseau (2005)

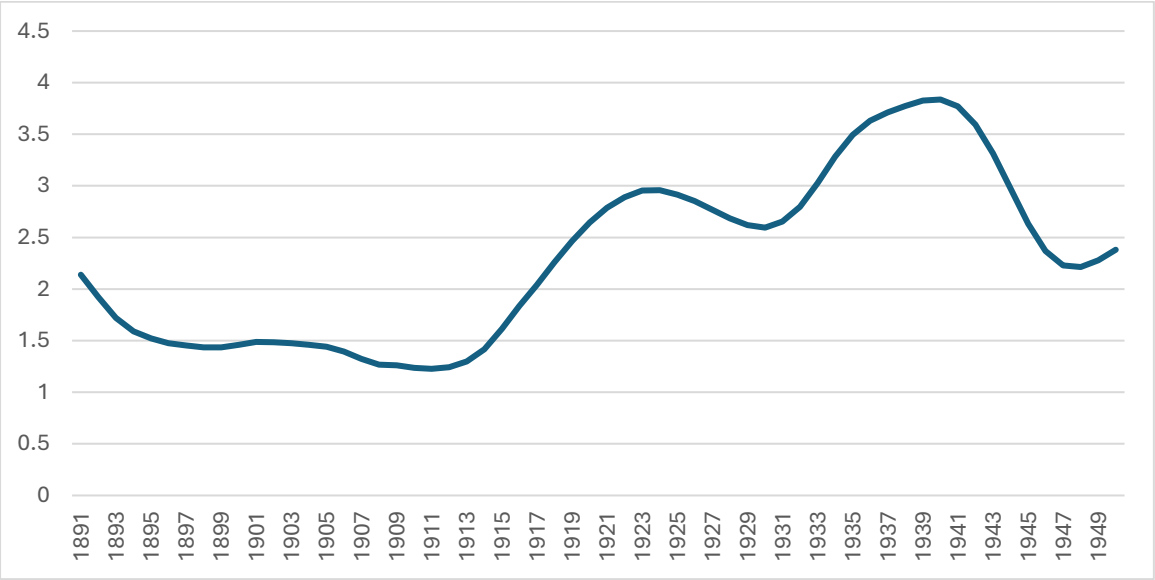
[28] Jovanovic and Rousseau (2005) General Purpose Technologies in the Handbook of Economic Growth available from https://sites.google.com/a/nyu.edu/boyan-jovanovic-home-page/research

[29] Note that the GPT in Chat-GPT doesn't stand for General Purpose Technology but rather for Generative Pre-trained Transformer – a description of the underlying technology architecture.

[30] Paul David (1989) Computer and the Dynamo https://warwick.ac.uk/fac/soc/economics/research/workingpapers/1989-1994/twerp339.pdf

As the chart below shows, productivity was falling and continued to fall prior to and after the 'introduction' of electricity. It was only around 1913 that productivity began to increase (somewhere over 30% of the total horsepower was then supplied by electricity). This pattern is often called the productivity J curve. By the end of the era (around 1929, productivity had risen to around 3%). So, if we are prepared to attribute all the increase to the process of electrification (probably an overstatement but let's be generous) then we can say electrification raised productivity by around 1.5% p.a.

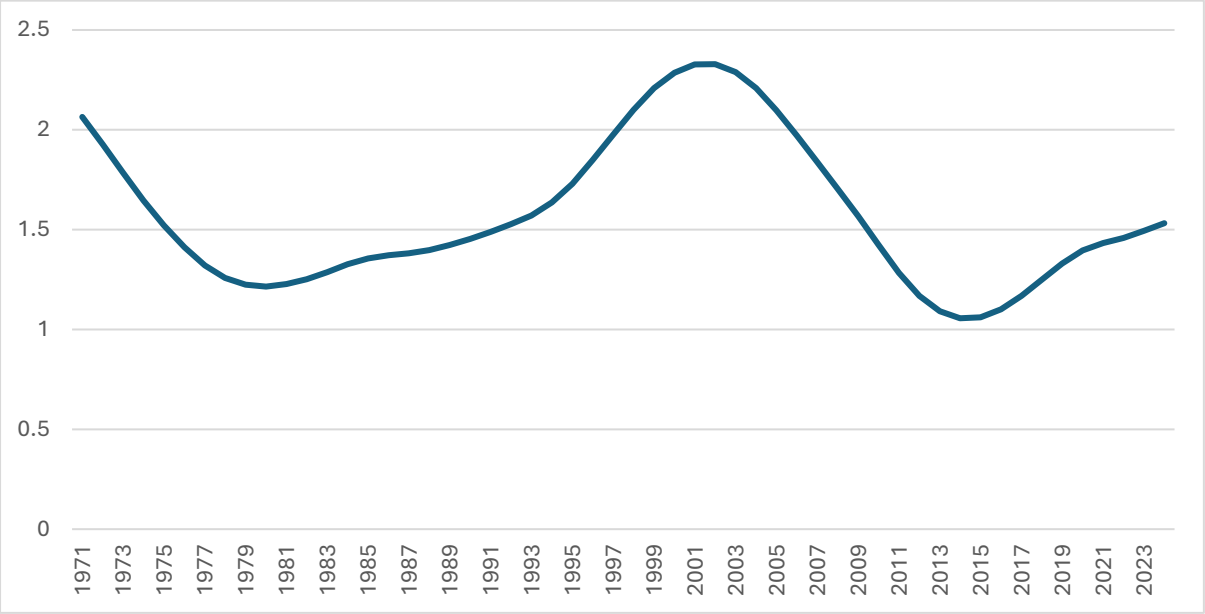
Labour Productivity % p.a. (HP filter) during the electrification era.



Source: labour productivity from Bergeaud, A., Cette, G. and Lecat, R. (2016): "Productivity Trends in Advanced Countries between 1890 and 2012," Review of Income and Wealth, vol. 62(3), pages 420–444 updated, and HP filter applied

The second example that Javanovic and Rousseau examine is the computer. As noted above, they trace the introduction to the arrival of the 4004 chip from Intel. As with electrification, a productivity J curve can be seen, with productivity initially declining before then increasing. Productivity fell from around 2% p.a. to 1.2% in the decade after the arrival of the 4004 chip. Again, if we are willing to attribute all of the increase to computers then by the peak productivity was running at 2.3%, so roughly a 100bp increase in productivity.

Labour productivity, %p.a. during the computer era



Source: labour productivity from Bergeaud, A., Cette, G. and Lecat, R. (2016): "Productivity Trends in Advanced Countries between 1890 and 2012," Review of Income and Wealth, vol. 62(3), pages 420–444 updated, and HP filter applied

So, if AI proves to be a "normal" technology to borrow the term from Narayanan and Kapoor (2005)[31], as opposed to the herald of AGI, then history would suggest a productivity growth improvement of around 100-150bps taking a prolonged period to actually show up in the data.

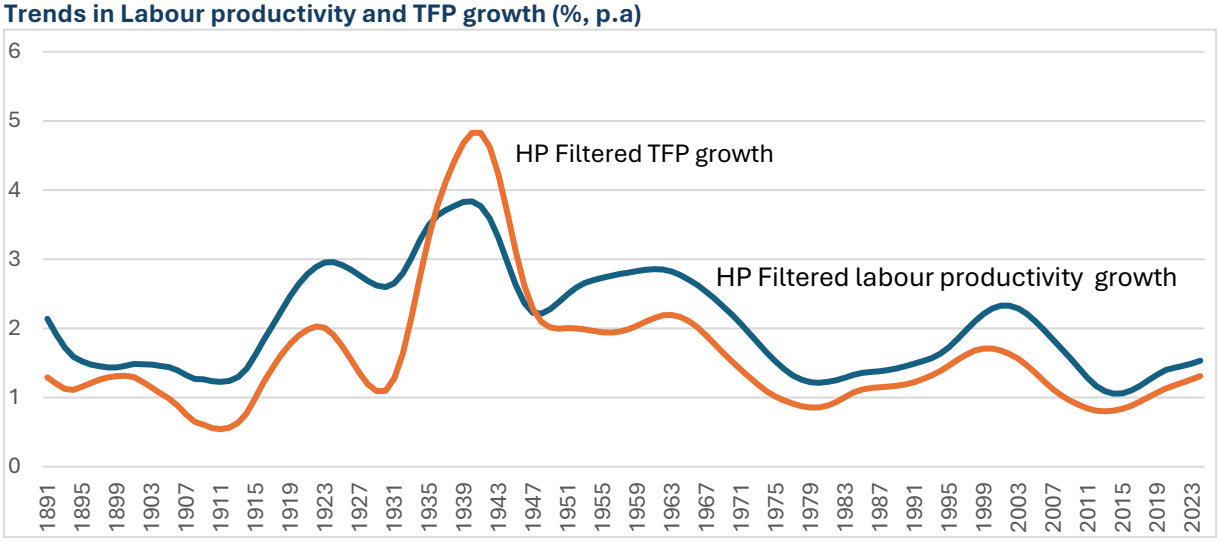
The second approach to thinking about the possible productivity gains is to follow the framework provided by Acemoglu (2024)[32]. He lays out a simple model for thinking about the improvement in total factor productivity.

Before we go any further, I ought to briefly explore the different types of productivity we are talking about here. Labour productivity is the output per worker (or per hour worked) – a relatively straight forward concept. Although I have noted elsewhere that it is a better (more meaningful) concept at the micro level than the macro level where it is essentially just the result of a calculation.

Mainstream economists (never a term I use as a sign of praise) like to think of labour productivity growth as a combination of so called total factor productivity growth (TFP) and the changes in the capital stock (weighted by the capital share). Such economists tend to view TFP growth as a measure of innovation or technological progress. I take a rather different view and can show that TFP growth is just a weighted average of wage growth and profit growth, but that is for another day. For now, we will humour the mainstream and follow their (incorrect) view. Across the longest run sample we have (1890 onwards), US labour productivity growth was averaged 2.1% p.a, whilst TFP growth has averaged 1.6% p.a.

[31] Narayanan and Kapoor (2005) AI as a Normal Technology available https://kfai-documents.s3.amazonaws.com/documents/c3cac5a2a7/AI-as-Normal-Technology---Narayanan---Kapoor.pdf
[32] Daron Acemoglu (2024) The Simple Macroeconomics of AI, available https://economics.mit.edu/sites/default/files/2024-04/The%20Simple%20Macroeconomics%20of%20AI.pdf

Trends in Labour productivity and TFP growth (%, p.a)



Acemoglu sets up a simple framework for thinking about how TFP growth will change under AI. In essence, the growth of TFP is a function of the labour share (how much labour gets paid out of total GDP) multiplied by the share of total labour tasks exposed to and economically viable for AI to replace multiplied by the microeconomic cost savings achieved by this AI replacement.

Acemoglu constructs a deliberately conservative estimate. He starts with a labour share of 50%, assumes that 4.6% of tasks can be replaced with AI, and that the micro cost saving will be of around 27%. Given this he estimates that TFP growth will rise by just 0.06% p.a. over a decade!

Using the same framework, but different assumptions, Aghion and Bunel (2024)[33] arrive at a higher estimate. They use a labour share of 64%, assume that 24.3% of tasks can be replaced by AI, and that the cost saving will be 43.5%. With these assumptions, TFP growth would be 0.68% p.a. higher.

For what it is worth, I created my own run, using the actual labour share (52.8%), assuming as per the RLI above that 15% of tasks could be replaced by AI, and took the middle ground between Acemoglu and Aghion and Bunel assumed a 35% cost saving. This combination generates TFP growth of 0.28%.

For fun, I ran a scenario I labelled Wide Eyed Optimist (WEO) to try and capture the essence of the beliefs of the AI leaders quoted above. Again, I took the actual labour share (53%) but now I assumed that 100% of tasks could be done by AI (in keeping with full automation) and cost saving would be 100%. This gave a boost to TFP growth of 5.3% p.a. So at least if you believed in their view of the world, then their growth statements are not completely irrational. Of course, I would argue that their beliefs are completely irrational given the evidence!

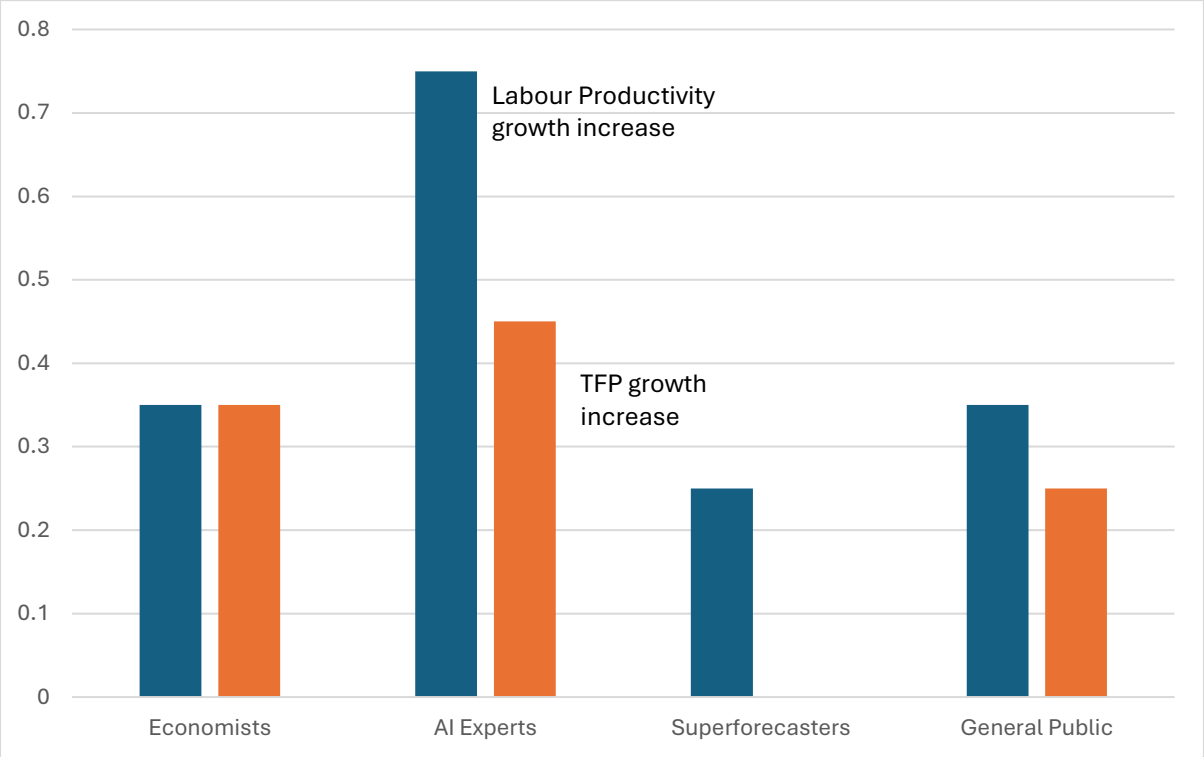
We can compare these evaluations to a set of forecasts collected by Karger et al[34]. They create a panel of different groups of forecasters:- economists, AI experts, superforecasters(a group who have exceled at forecasting in past events), and the general public. The chart below shows the increase in productivity these groups anticipate. It transpires that all expect relatively small increases in productivity (regardless of which measure we are using). In terms of labour

[33] Aghion and Bunel (2024) AI and Growth available from https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf
[34] Karger et al (2026) Forecasting the economic effects of AI available fromn https://forecastingresearch.org/pdf/economic-effects-of-ai.pdf

productivity, the estimates range from 0.25% increase (taking overall labour productivity to 2.4% p.a) to 0.75% increase (taking overall labour productivity to 2.65%p.a.)

For total factor productivity growth, the range is from 0% increase (from the Superforecasters, very close to Acemoglu's estimate) to 0.45% increase. This would take TFP growth overall to between1-1.45% p.a. The range is hardly the stuff of economic booms of unprecedented parallel.

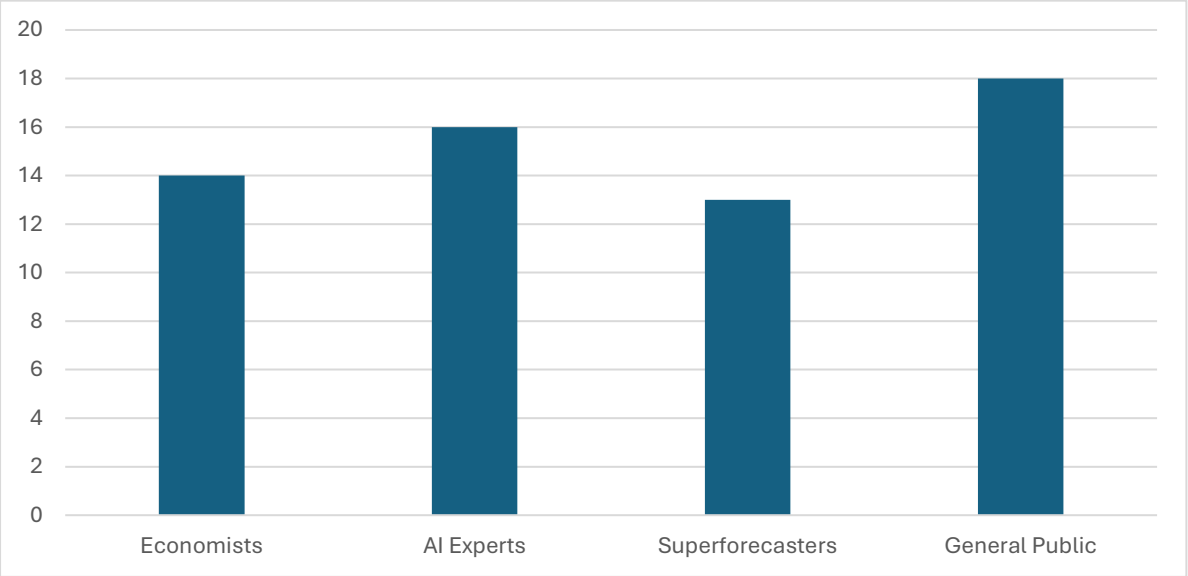
Increase in productivity (% p.a. over approximately 10 years)



Source: Data from Karger et al. Note the base level of growth was 1.9% for labour productivity growth, and 1% for TFP growth.

These forecasts represent the base case assumptions. Karger et al asked the participants for their forecasts of both the probability of a "rapid improvement" scenario and their productivity views under such a scenario. The rapid improvement scenario was defined as AI that "outperforms top humans in research, coding and leadership, capable of award-winning creative works, and nearly all physical tasks" that is effectively, full automation and AGI. The range of probabilities was reasonably tightly clustered ranging from 13% from the superforecasters to 18% from the general public.

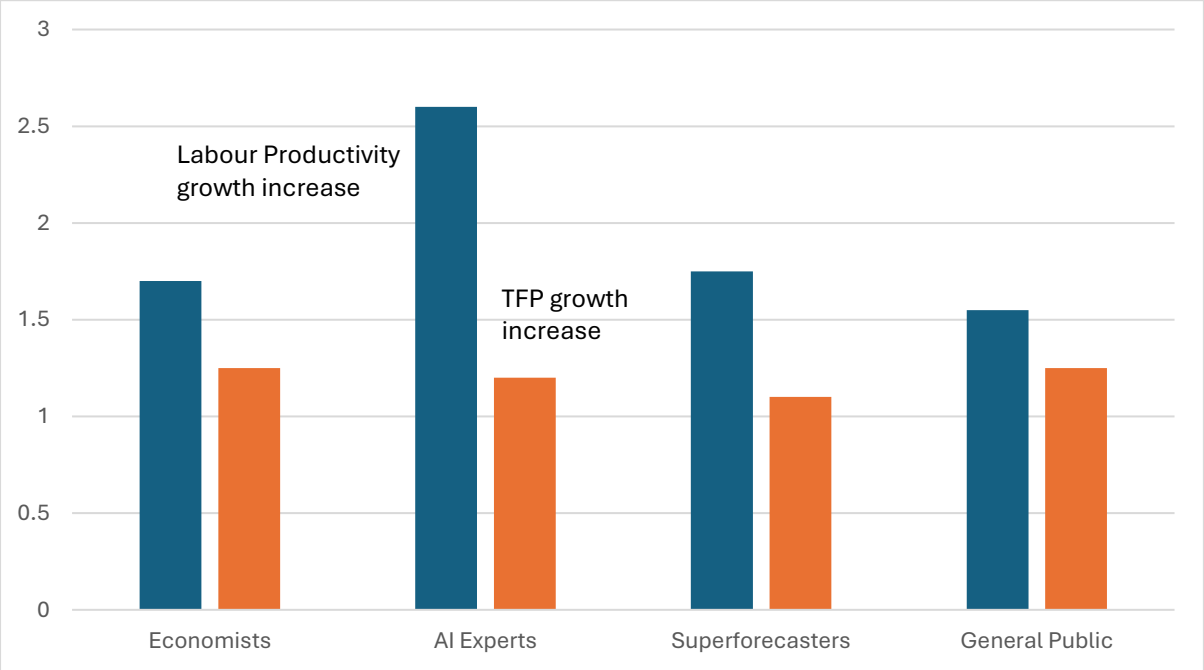
Probability of full automation and AGI over ten years



Source: Karger et al (2026)

Under this "tail" scenario, total labour productivity growth is expected to rise from 1.9% to anywhere between 3.5% - 4.5% p.a. TFP growth is expected to rise from 1% p.a. to anywhere between 2.1-2.3% p.a.

Increase in productivity under "Rapid improvement" (% p.a. over approximately 10 years)



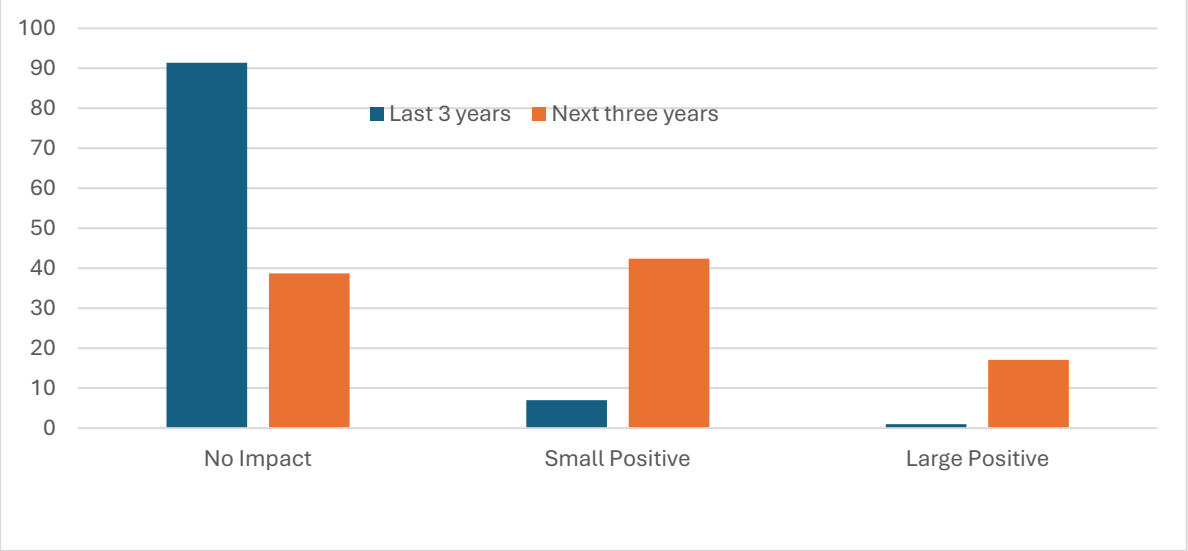
Source: Data from Karger et al. Note the base level of growth was 1.9% for labour productivity growth , and 1% for TFP growth.

Recall the exhibit of trends in productivity growth shown in the exhibit above to provide some context for these estimates. Labour productivity growth reached a trend peak of just short of 4% in 1939, so under this rapid improvement scenario participants anticipate a return to such boom times. Their TFP growth estimates are below the trend peak of 4.8% seen in 1941 but a significant improvement over the long-term average of 1.6% p.a.

We can also turn to the evidence collected by Yotzov et al (2026)[35]from nearly 1000 US firms. Over the last 3 years, 91% of respondents said AI had resulted in no impact on productivity at all! 7% said it had increased productivity by a small amount (less than 5%). Given this it should be no surprise that the average cumulative impact on firm level productivity was just 0.24%.

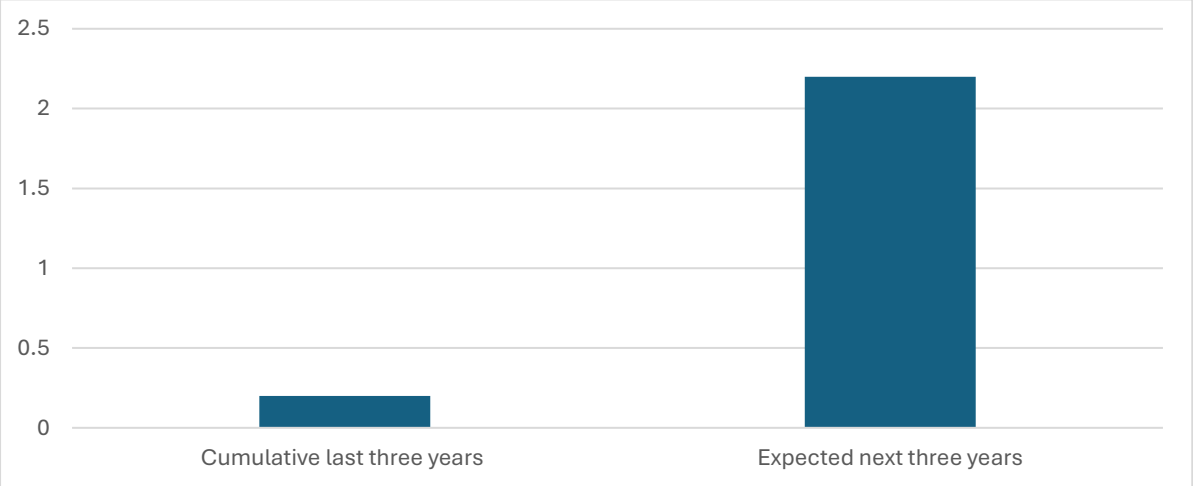
When asked about the next three years, 39% of respondents still expected no productivity increase! 42% expected a small positive impact, and 17% expected a large positive impact. The cumulative increase expected by US firms over the next three years was 2.2%. Firms from elsewhere where more muted in their expectations (for instance, Australian firms only expected a cumulative productivity increase of 0.9% over three years).

% of US firms reporting impact on productivity



Source: Yotzov et al

Scale of Productivity Increase (US firms, 3 year cumulative)



Source: Yotzov et al

So, AI has certainly not proved revolutionary yet and isn't really expected to be by the majority of those running US firms.

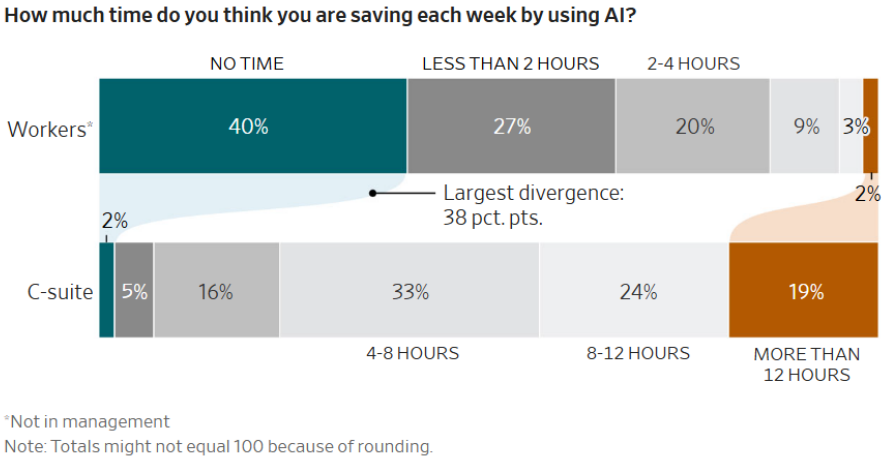
[35] Yotzov et al (2026) Firm data on AI available from https://www.nber.org/papers/w34836

No discussion of productivity improvements could be complete without my favourite study. METR (an AI evaluation center) conducted an experiment in 2025 using developer productivity improvements using AI[36]. The initial study found that developers forecast that using AI would allow them to reduce completion time on projects by around 24%. The data actually showed that the developers took on average of nearly 20% longer to complete the task with AI than without it! Admittedly the interval was huge, ranging from 2% longer to nearly 40%...but in all cases longer!

In a subsequent follow up experiment[37] using more tasks and a smaller set of developers, speed improvement was at least positive, 4%. Again, with a wide range from -15% (once again working with AI taking longer) to +9% (a productivity enhancement).

Recently new experimental evidence casts some more doubts on the likely productivity benefits of AI. Yu et al (2026)[38] explored how much people used AI and how much it improved their time to complete a task compared to their expectations. The first finding was people used AI much more than they thought they would. Across their sample people predicted they would use AI 33% of the time, and in fact used AI 47% of time. On the easiest tasks, people thought they would use AI only 18% of the time, they actually used it nearly twice as often (38% of the time). Across all the tasks, people thought that using AI would save them 58% of their time. It actually saved them 7%! In the case of the easiest tasks, using AI was actually slower than not using AI!

It is also worth asking who sees the benefits from AI, as the exhibit below shows, it appears that Yotzov et al not only asked managers, but also asked workers what they thought about the likely impact of AI on productivity. Workers were much more pessimistic about the productivity gains from AI, expecting a cumulative increase of just 0.9% over the next three years, compared to the management's expectations of 2.2%. This is consistent with the data shown in the chart below from the Wall Street Journal. A stunning 40% of workers found they saved no time with AI! 19% of those in management reported saving more than 12 hours a week! Guess we know who we can replace with AI!



Source: WSJ

[36] Becker et al (2025) Measuring the Impact of early 2025 on Experienced open-source developer productivity available from https://arxiv.org/pdf/2507.09089

[37] https://metr.org/blog/2026-02-24-uplift-update/#wider-adoption-of-ai-has-made-it-more-difficult-to-measure-task-level-productivity

[38] Yu et al (2026) The efficiency-gain illusion available from https://arxiv.org/abs/2605.22687

At best, evidence on productivity improvement is minimal. At worst, AI may make tasks longer than just human work! Either way, unless you believe in the imminent arrival of AGI and full automation, business as usual seems like the best viewpoint to hold as a base case.

One other aspect with respect to productivity and benchmarks should be noted. The benchmarks that are commonly used generally measure technical feasibility or capability. They don't measure the cost effectiveness of deployment. As any number of CFOs have discovered this can't be ignored. Recall the Uber management who reported that the company had burnt through their entire 2026 AI budget in just four months and then went on to question the connection between the use of AI and more useful consumer features!

Artificial Analysis (artificialanalysis.ai) provides us with some insights into the relationship between the capability and the cost. They measure "intelligence" score of various LLMs and also compute the costs (based on things like reasoning and output token prices, and number of tokens used etc).

For instance, running Claude Fable (with fallback) results in an "intelligence" score of 62, but comes with a heft cost tag of $5455. Compare this to say Gemini 3.7 Flash (high) with an intelligence score of 56, but a cost of $485.

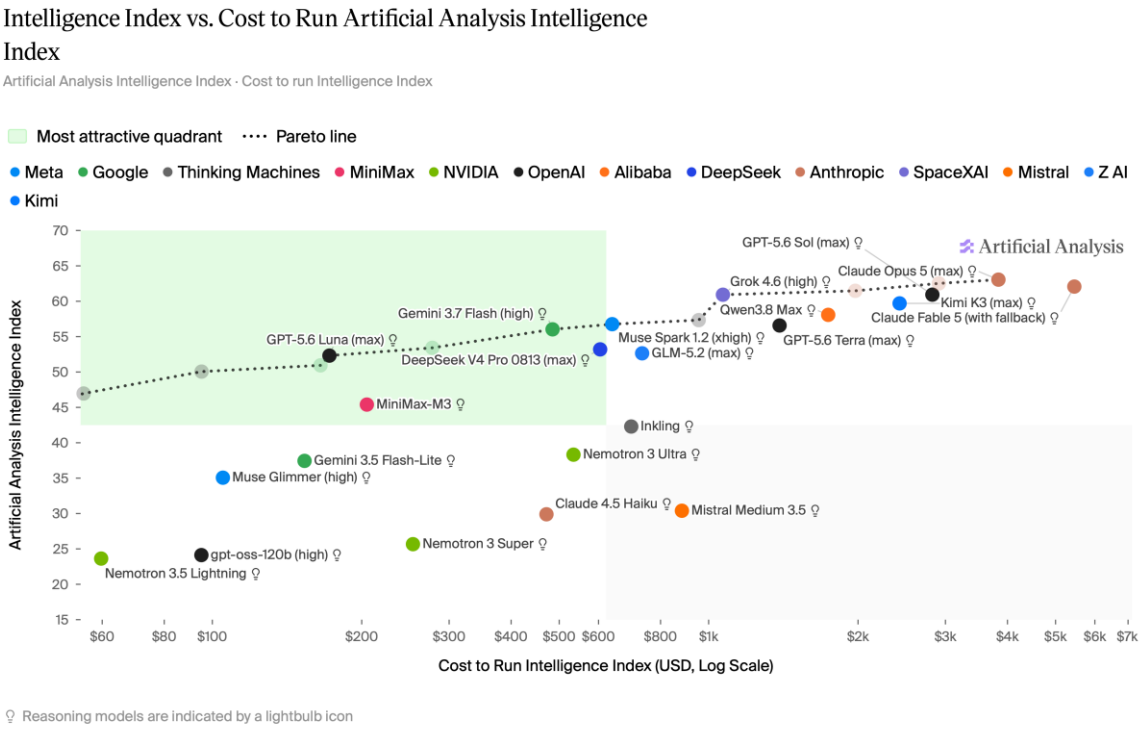
We can also look at the expense incurred for the very cutting edge of AI. OpenAI reported that it had "solved" the Navier Stokes problem (a problem related to general applicability of the equations used in fluid dynamics, and one of the Millenium Prize Problems as prescribed by the Clay Mathematics Institute). Leaving aside the very real controversy over possibility of plagiarism/front running, it is far from clear that we have actually learnt anything from AI solving this problem. The proposed proof runs to some 166 pages but is largely impossible to understand even for specialist mathematicians. When a mathematician normally solves such a problem, they are forced to explain their workings, sharing their insights and potential new methods.

OpenAI provided us with an essentially impenetrable "proof". In many way this is akin to Deep Thought, a computer built to answer the most important question....the question of life, the universe and everything (in Douglas Adam's wonderful Hitchhiker's Guide to the Galaxy) which after 7.5 million years of computation reached the conclusion 42. This of course, raises the question, what was the question? If AI can produce proofs that no one understands, what has it really achieved. The cost of looking for a solution according to the company was around 300 billion tokens, based on their own pricing scheme this equates to around $15-20million.

In fact, it transpires that the natural language proof provided by OpenAI (see Bastounis, Circello and Hansen[39]) doesn't match the Lean code (the verification program). So although Lean verified the proof, the Lean proof isn't the same the proof written in the paper.

We will have more to say about the cost side of running AI when we cover the next myth.

[39] Bastounis, Circello and Hansen (2026) Navier-Stokes lost in translation https://arxiv.org/abs/2610.08144



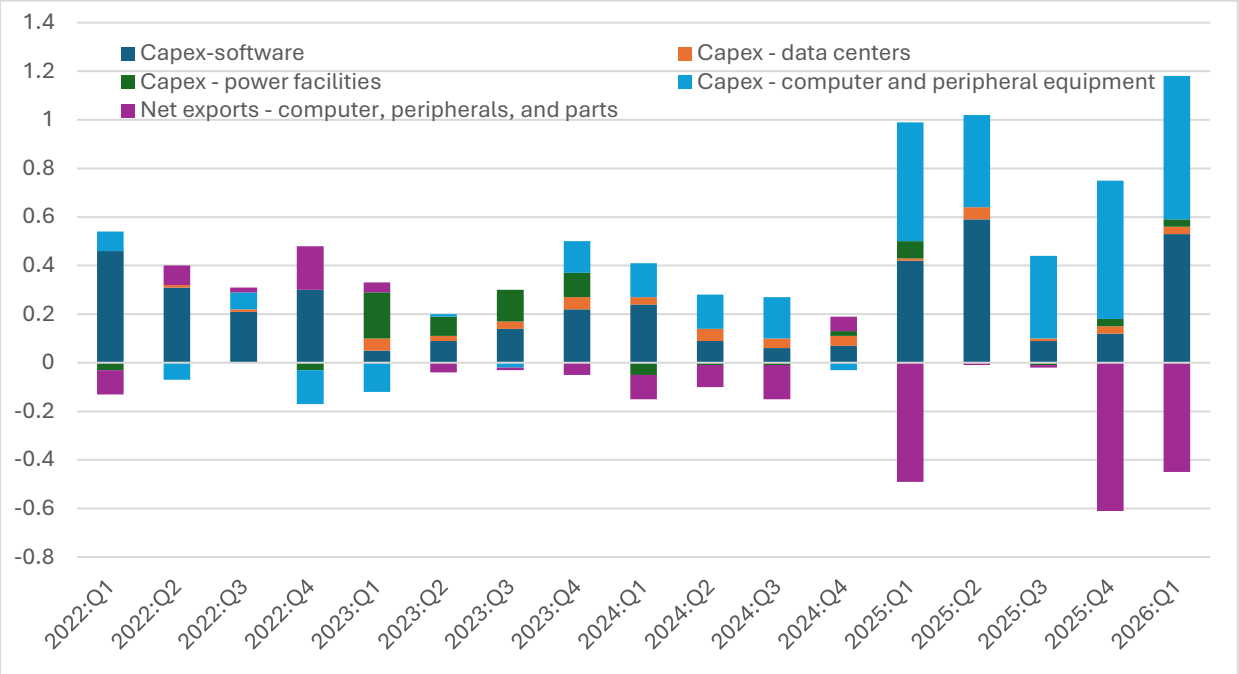
Source: Artificial Analysis

Given all of the above, it shouldn't be a surprise when we see report after report showing the AI is struggling to make a difference in the real world. For instance, RAND (2024)[40] found that 80% of AI projects fail (twice the rate of failure for IT projects). MIT Nanda (2025)[41]that 95% of organizations were seeing no measurable return to their AI projects. S&P Global Market Intelligence found that 42% of businesses were scrapping their AI initiatives.

While the evidence of any productivity improvements is meagre, to say the least, the impact on GDP growth is not as hard to come by. Soto, Thieu and Allen (2026)[42] from the Federal Reserve Board of Governors produce the breakdown shown below. As we have documented many times in previous Curmudgeon's Corners (and will again in the next section) we have witnessed an investment boom (particularly amongst the hyperscalers). However, from a GDP perspective we have to subtract off the net exports relating to the purchase of many technology components (think GPU's made in Taiwan, Korean memory chips).

[40] https://www.rand.org/pubs/research_reports/RRA2680-1.html
[41] https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf
[42] https://www.federalreserve.gov/econres/notes/feds-notes/the-ai-buildout-and-the-economy-publicly-available-data-to-assess-ais-impact-20260717.html

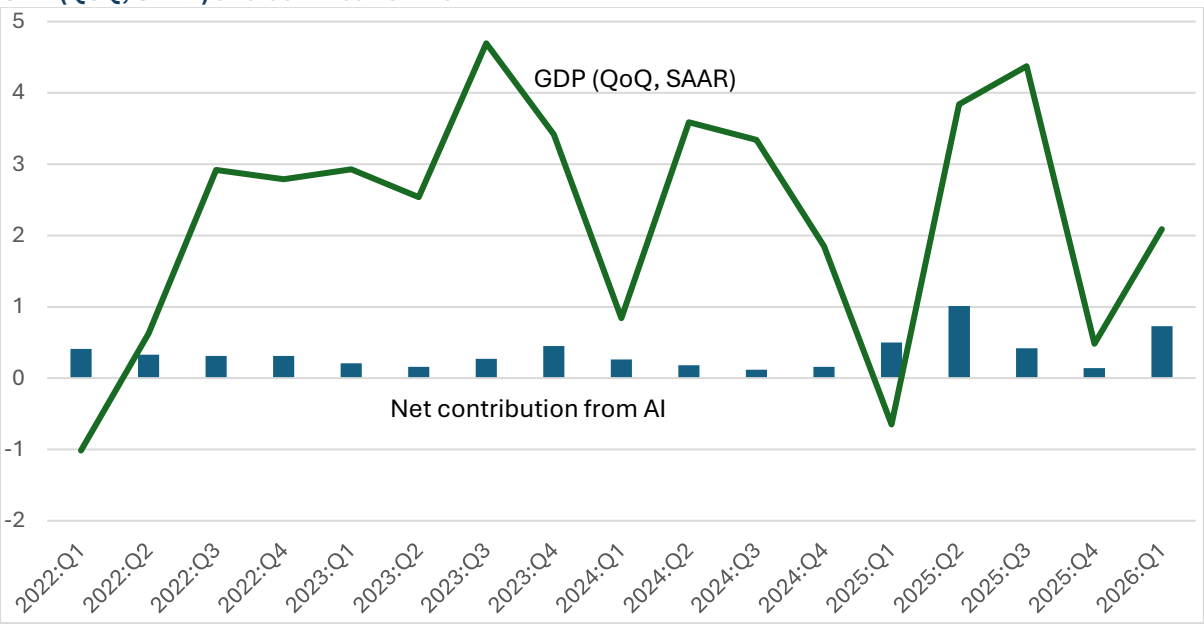
Contribution to US GDP (QoQ)



Source: Soto, Thieu and Allen

As the chart below shows the net contribution to GDP growth is an average 40bps since 2022, but closer to 60bps in the last year and half. So AI clearly matters for GDP growth, but hasn't yet become the engine of the US.

GDP (QoQ, SAAR) and contribution from AI

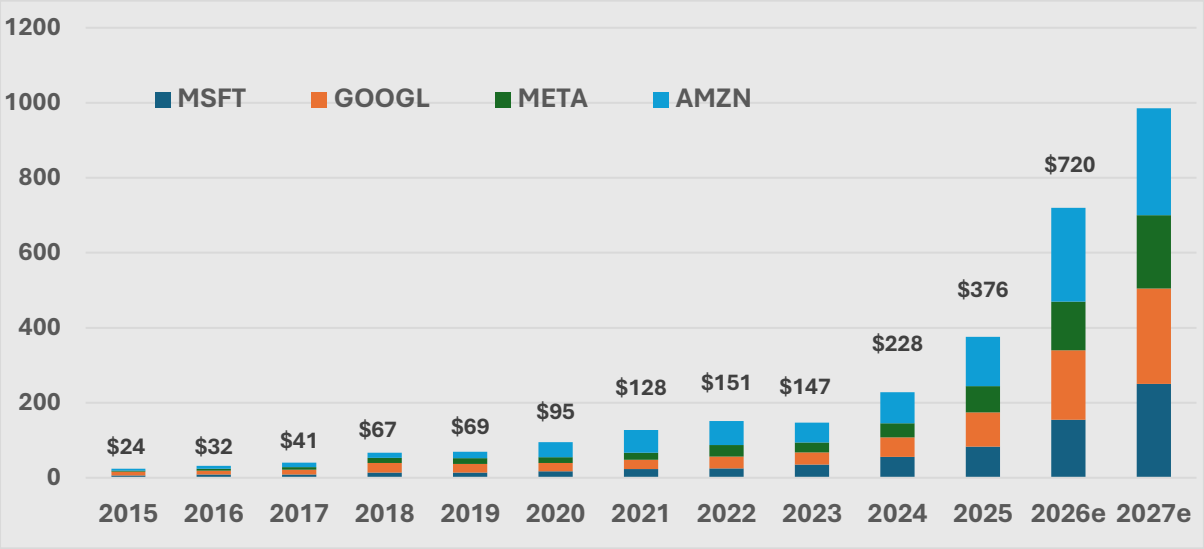


Source: Soto, Thieu and Allen and NIPA

Myth 3: This is a boom, not a bubble

As we have pointed out in Curmudgeon's Corner Volume 3, the scale of the capex boom we are witnessing is breathtaking. Secondly only in magnitude to the UK railroads during the 1840s. If the estimates from Wall Street are correct for 2027, the Hyper-scalers will be spending around 3% of GDP! Three times the peak of capex boom seen during the TMT bubble!

Hyper-scalers capex (US$bn)



Source: S&P Cap IQ

As we noted in Volume 3, this expenditure seems to be driven in large part by the fear of missing out (FOMO), or as my colleague Willem put it, perhaps more accurately, the fear of losing out (FOLO). Indeed, the late great economist James Crotty used the term coerced investment to describe such situations, "competitive pressure is so intense that firms are induced and often "coerced" into investing all the financial resources – here profits – available to them.... Invest or die"[43] becomes the firms mantra.

The hyper-scalers are playing a dangerous game of prisoner's dilemma. For instance, Mark Zuckerberg of Meta has said "We are going to spend aggressively. Even if we lose a couple of hundred billion. It would suck, but it's better than being behind the race of super intelligence", Sundar Pichai of Alphabet "The biggest risk could be missing out", or Sam Altman of OpenAI "Whether we burn 500 million a year, or 5 billion or 50 billion I don't care. I genuinely don't...we're making AGI, it's going to be expensive, it's totally worth it"!

Ultimately, the hyperscalers are involved in an auction like set up, bidding to see who can produce the most compute. Like many auctions, they may end up suffering the winner's curse (the auction is won by the most optimistic of the bidders).

The power of simple models

One of the lessons from Curmudgeon's Corner Volume 3 was that simple models can often be used to expose the extreme nature of beliefs that underlie market pricing. So, Steve Fazzari, ACR's Academic Advisor of Macroeconomic Research, and I designed an incredibly simple little

[43] James Crotty (1993) Rethinking Marxian Investment Theory: Keynes-Minsky Instability, Competitive Regime Shifts and Coerced Investment , Review of Radical Political Economics

model to explore what you need to believe in order to justify the scale of capex that we are witnessing.

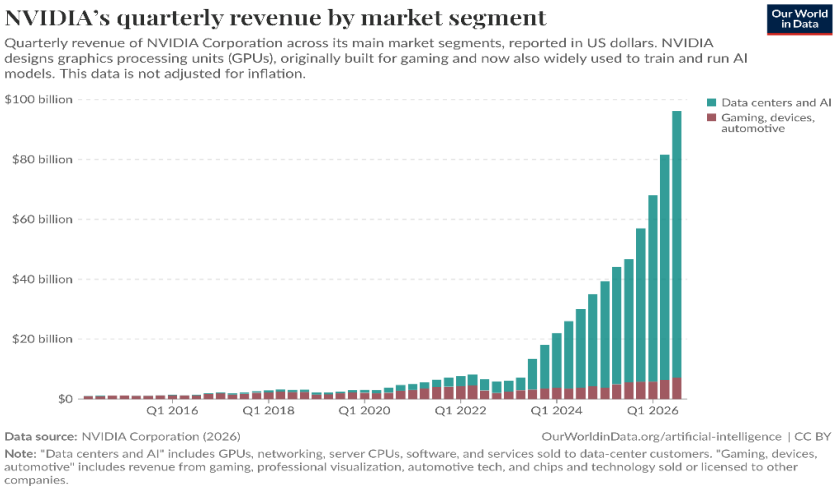
How much productivity growth do you need to see for today's level of hyper-scalers capex to make sense? We assume a few things, an economy wide profit share of 9.5%, a discount rate of 12%, a depreciation rate of 16.7% (more on this later), an original rate of productivity growth of 2%. We also assume that all of the increase in profits is captured by the hyper-scalers (an exceedingly generous and unrealistic assumption)! Given these assumptions, in order to generate $2.6trn (the cumulative capex spend seen by the hyper-scalers since 2022, including 2027 estimates), productivity would need to permanently rise 4 percentage points[44], taking overall productivity growth to 6% p.a. This is two percentage points higher than has ever been seen in the trend data, and you have to assume it is permanent!

Assuming a shorter (more realistic?) depreciation life (of 4 years rather than 6) results in the required productivity to be even more extreme. The required productivity increase would be 6.3 percentage points, taking overall productivity growth to 8.3% - more than twice the highest rate of trend productivity growth we have seen!

Put another way, the observed behaviour (what economist's call revealed preference) of hyper-scalers capex spend requires you to believe that AI is going to be truly transformative, effectively that AGI and full automation are here. Contrast with the modest ranges of productivity growth forecast in the previous section! Although they may be acting out of FOLO, the actions of the hyper-scalers are consistent with those of a true believer. If you don't share their belief, then we are witnessing one of the biggest capital misallocations ever!

Intriguingly, in September 2023, David Cahn of Sequoia Capital (one of the leading lights of venture capital in the technology world) published a blog post entitled AI's $200bn Question[45] in which he explored the required AI revenue to justify the infrastructure build out.

The intuition and the maths are simple. Take Nvidia's run-rate data center revenue, "multiply it by 2 to reflect total data center Capex, including non-chip expenses)" then multiply by 2 again "to reflect a 50% margin across the hyperscalers and the AI product company". The end number is "an end customer revenue requirement for a single year of CapEx." Cahn has updated his model a few times since the original post[46]. The table below shows the maths. To justify the kind of expenditure we have been witnessing, we would need to see around $1.5trn of customer spend on AI! These numbers are also cumulative, so to justify the build out we have witnessed since 2022 would require a cumulative expenditure of $3trn!



[44] Illustrative ACR analysis based on the assumptions stated. Actual results could differ materially if the assumptions change.
[45] https://sequoiacap.com/article/follow-the-gpus-perspective
[46] Most recently https://dcahn.substack.com/p/ais-15t-question

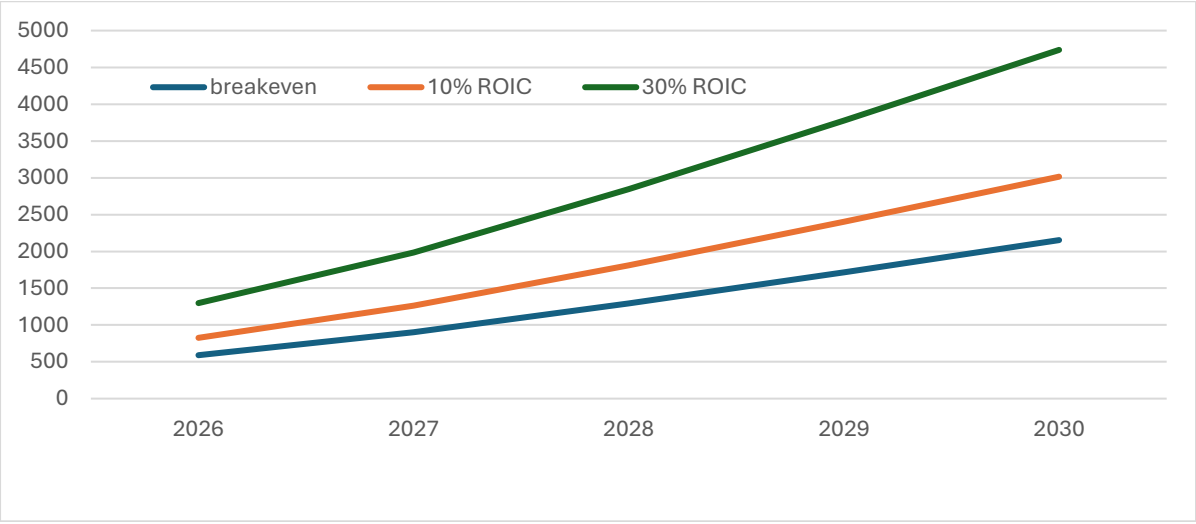
$bn	2022	2023	2024	2025	2026e
NVDA Data center run-rate revenue ($)	15	47.5	115	193	375
Data center Facility Build and cost to operate %	50	50	50	50	50
Implied Data Center AI Spend	30	95	230	386	750
Software Margin %	50	50	50	50	50
AI Revenue required for Payback	60	190	460	772	1500

You can arrive at roughly the same number from the perspective of the hyper-scalers as well. If you double the hyper-scalers capex (to reflect the margins required as above) then you also get a figure close to $1.5trn, helping to triangulate the evidence.

We can also use a similar simple model to project forward, given estimates of the hyper-scalers capex. Consensus estimates suggest that capex will rise from roughly $1trn in 2027 to $1.5-1.6trn in 2030. We can reverse engineer given assumptions about depreciation rates, EBITDA margins and required investment returns (20%, 55%, and 0-30% see below) we can estimate the annual revenue required from the end users to make this add up.

The chart below shows the required revenues to justify the consensus view of hyper-scalers capex. I've varied the required return across 3 levels. Breakeven (i.e. no return on invested capital), a 10% return, and a 30% return. The last of these is close to what the hyper-scalers used to earning in the years before the AI boom!

Required revenue from AI end users ($ bn pa)



Source: ACR Alpine Capital Research

Assuming a 10% return, then AI spending would need to be grow the current required $852bn to $3trn in 2030! Truly breathtaking sums. Also note I have assumed a 55% EBITDA margin, compared to margins that were currently negative, so this growth alone is not enough, we need to see a massive switch in profitability as well as the revenue growth.

As a check on this model, I compared it with a model from a new paper by Van Nieuwerburgh[47]. His work is based from the data center side of the AI eco-system, and prices the build out using a 10% return, a 6 year depreciation rate, a 50% margin. He covers out to 2032. Given the 188.3GW of capacity expected to be built, he estimates that AI revenue must "correspond to mature annual revenue of approximately $3.549 trillion by 2032" equivalent to "8.8% of 2032 GDP"! A very similar to my model above.

Table 6. Capital and Required Revenue by Completion Cohort

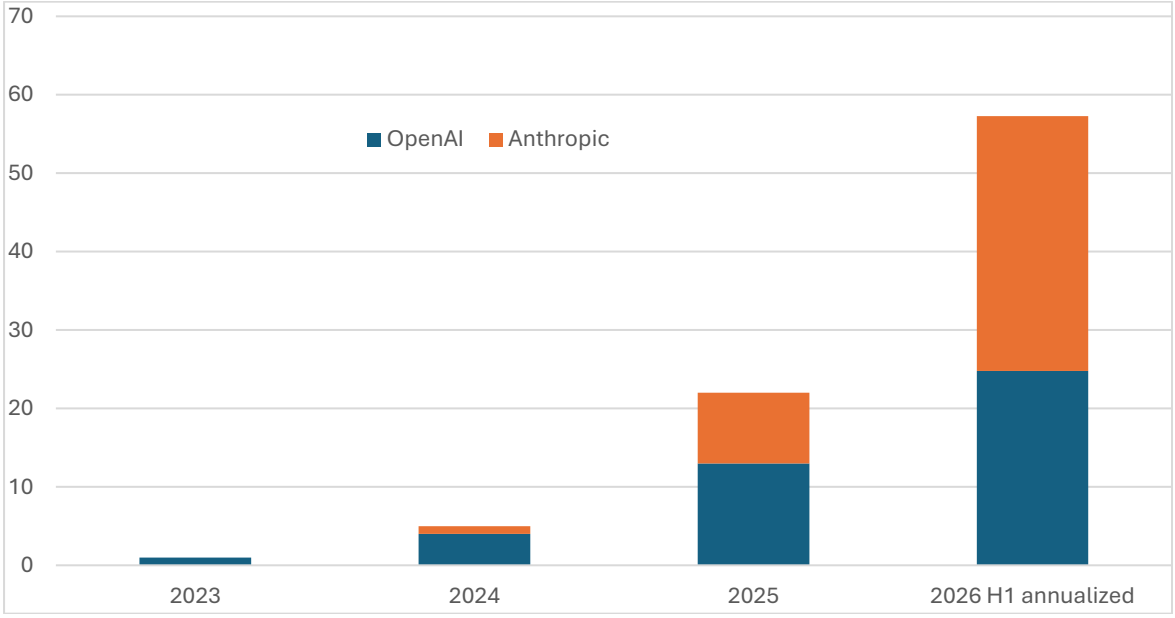
Year	GW	Cost	Capital	Req. CF	Req. Rev.
			($ billions)		
2025	13.8	563.9	589.4	114.2	228.4
2026	26.6	1,133.6	1,184.9	229.5	459.1
2027	35.8	1,588.8	1,660.7	321.7	643.4
2028	43.5	2,005.5	2,096.3	406.1	812.2
2029	21.5	1,029.6	1,076.2	208.5	417.0
2030	17.7	881.9	921.9	178.6	357.2
2031	16.1	835.5	873.3	169.2	338.3
2032	13.4	723.4	756.2	146.5	293.0
Total	188.3	8,762.2	9,158.8	1,774.3	3,548.5

Notes: Capacity is based on the outcomes in Table 4. Reported completion years are retained, and the 54.0 GW of undated capacity assumed to be completed by 2032 is allocated equally across 2028–2032. Cost is $Q_t = 0c_t \times K_t$, using $41 million per MW in 2025 and 4 percent annual cost growth. Capital is $B_t = \phi \times Q_t$. Required operating cash flow assumes a $r = 10\%$ unlevered return, a six-year life for the 68 percent IT share, and a 20-year life for the remaining assets. These values imply $\phi = 1.04527$. Required revenue assumes a 50 percent operating cash-flow margin. Figures may not sum exactly because of rounding.

Source: Van Nieuwerburgh (2026)

To put these figures in some context, the exhibit below shows the revenues achieved by OpenAI and Anthropic. If we annualize the H1 2026, total revenue is $57bn, a far cry from $1.5trn required! Cumulatively since 2023, these two companies have produced just over $85bn, a far cry from the required $3trn.

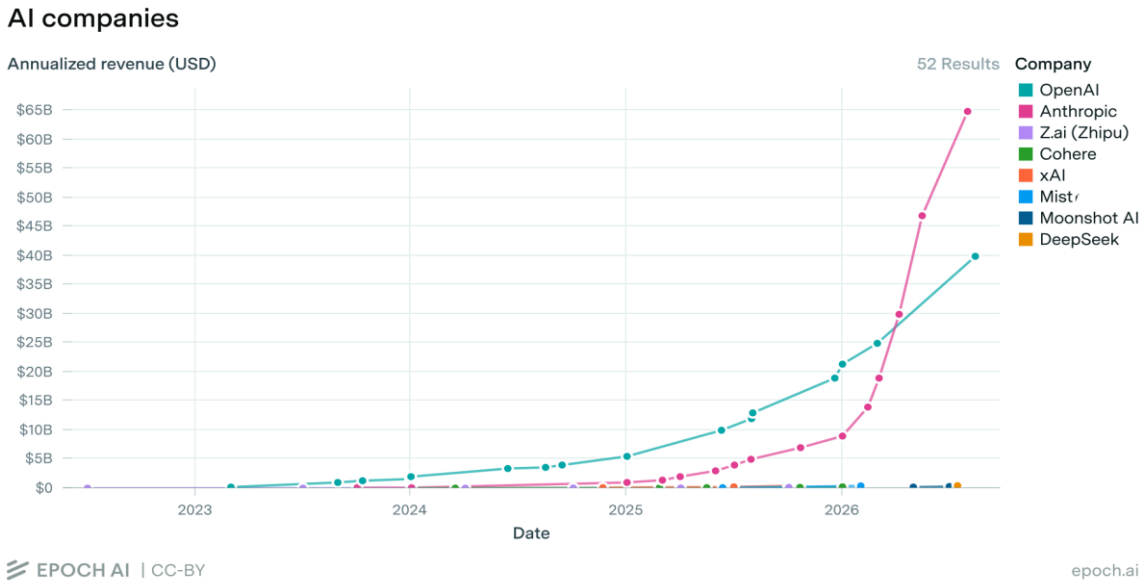
OpenAI and Anthropic Annual Revenues ($bn)



These numbers may seem smaller than the ones you often hear quoted. For instance, Anthropic has recently been cited as telling investors that it hit $65bn in revenue in July of this year. However, this is a monthly annualized number, rather than a real recorded annual revenue number.

[47] Stijn Van Nieuwerburgh (2026) Financing the AI buildout
https://www.nber.org/system/files/working_papers/w35865/w35865.pdf

It is also worth noting that OpenAI and Anthropic report their revenues on a different basis. OpenAI reports its revenue net of the share they need to pass to Microsoft. However, Anthropic recognizes the full value of any partner-based token sales as revenue. From what we can gather from the leaks of Anthropic's S-1, it pays 16c of every dollar of revenue to the services providers. To put their revenue on the same basis as OpenAI would result in the last data point below being around $54bn, rather than $65bn.



It is worth noting that other layers of the AI stack also produce revenue. For instance, Databricks produced $4.8bn, Snowflake $4.68bn, and MongoDB $2.46bn in 2025 from the data layer. The application layer also generates some revenues (entities such as Cursor (taken over by SpaceX as I was writing this, around $500mn revenue in 2025), Perplexity $200mn in 2025 and Harvey to name but a few). But as is abundantly clear, these are not earthshattering amounts. Even adding say an exceedingly generous $20-30bn to the revenue of OpenAI and Anthropic doesn't dent the deficit outlined above. Hell, even doubling the level of OpenAI and Anthropic revenues doesn't come close to touching the level of revenue needed to justify the scale of current (let alone future) capex.

Now of course, OpenAI and Anthropic are both projecting enormous revenue growth ($350bn and $300bn in 2030 respectively). This represents an incredible 85% CAGR for OpenAI and 101% CAGR for Anthropic. Anything short of this is likely to prove ruinous for the AI labs. As I was finishing this essay, I came across the data on the next page from Ramp which shows that at least in the recent past Token spend may have actually hit a ceiling – if this persists there will be trouble ahead for the AI labs.

Total token spend by business (indexed)



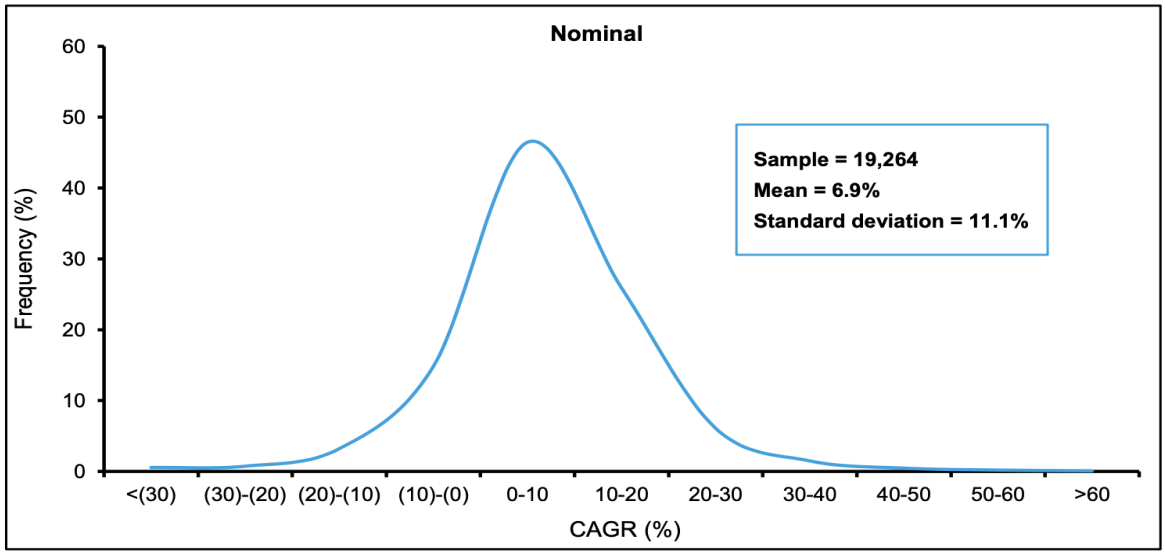
Source: Ramp

My old friend, Michael Mauboussin[48] and his colleague Dan Callahan wrote a report in May[49] showing the following distribution. It charts the frequency of five year sales growth for some 19,300 firm-period observations. The mean nominal CAGR is 6.9% with standard deviation of 11.1%. Thus OpenAI and Anthropic are anticipating essentially an 8-10 standard deviation event! To put this into context, this should be observed once every 4.43 quintillion years!

As Mauboussin and Callahan note America Online (AOL) did record a five year CAGR of 103% between 1997-2001. However, this was driven by the merger with Time Warner in early 2001. "Time Warner's sales were more than five times of AOL at the time of the merger agreement".

Nvidia is expected (by analysts) to produce some $610bn in revenue by 2030, a CAGR of 49%. Not as extreme as the AI labs, but still right at the very edge of the distribution.

Empirical distribution of nominal 5-year sales growth for firms with $2-5bn in sales 1950-2025



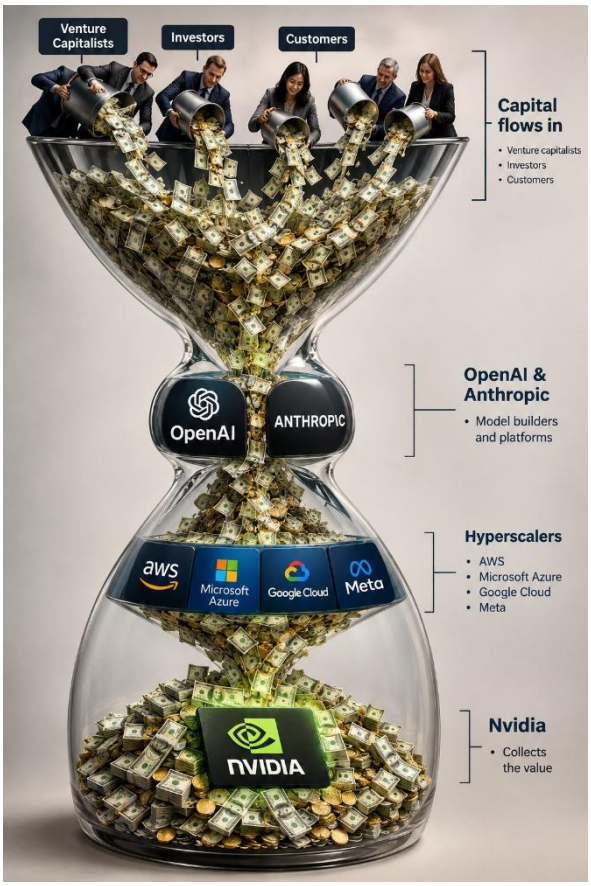
[48] Maybe it is just me, but I have often found Michael and I agreeing, in fact, I often think of Michael as me if I could only hold it together and wear a suit! Neither of which I do!

[49] Mauboussin and Callahan (2026) Bayes and Base Rates 2.0 available from https://www.morganstanley.com/im/en-us/individual-investor/insights/series/consilient-observer.html

In essence, it doesn't matter which of the frames you adopt, Steve and my simple productivity framework, the Sequoia Capital model, or a look at the expectations for growth from the AI Labs, the conclusion is the same....to believe this isn't a bubble requires you to believe that this time is truly different!

Financing of the AI bubble

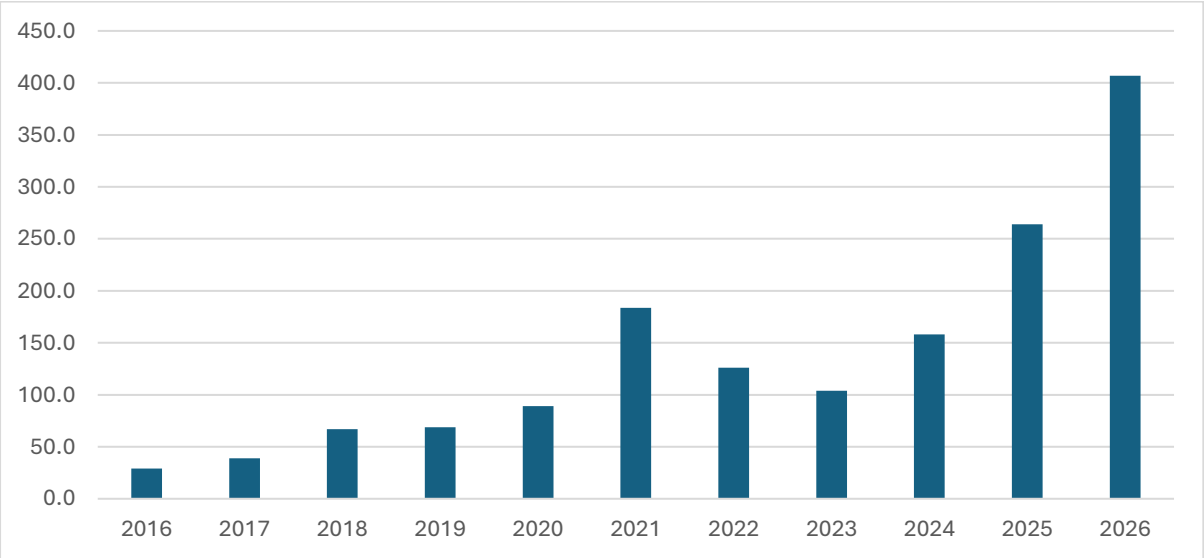
All of the above talk about revenue makes me uneasy. It is profits that actually matter at the end of the day, and the focus on revenue is typical of growth driven mania – don't look at my losses just focus on the sales growth. As I have pointed out many times before; I would have amazing growth selling $20 bills for $10, but I guarantee that I will never turn a profit.



I confess I used AI to generate the funnel image which probably accounts for the face of one of the Venture Capitalists appears to have a bin for his head! (The only time I used AI in the production of any of the Curmudgeon's Corners). I think it captures the concepts I am going to articulate well....at least much better than anything I could have drawn myself.

Let's start with the top of funnel – what we might call exogenous finance, source of capital from outside the AI ecosystem. Venture Caps have obviously been the principal source of cash for the AI labs since 2022, an impressive $1.1 trillion has been poured into AI eco-system, of which $400bn has been in 2026 alone (of which over 50% went to OpenAI and Anthropic).

Venture Cap inflows into AI ($BN)



Source: Pitchbook

Equity investors are also now being dragged into the funding morass as Anthropic and OpenAI prepare to IPO (maybe), following on from SpaceX (which raised $85bn earlier this year). As noted earlier Anthropic is proceeding to IPO (having filed its S-1 this week), and OpenAI has delayed its IPO until 2027.

Customers are another source of exogenous capital, (as detailed above in discussing the revenue of OpenAI and Anthropic). They have provided cash inflows of, let's be generous, $150bn since around 2022.

The final group are the hyper-scalers who as we have repeatedly seen have injected vast sums via their capex programs, approximately $1.6trillion since 2022. Adding all of this together, gives us an exogenous money inflow of around $2.9trn since 2022.

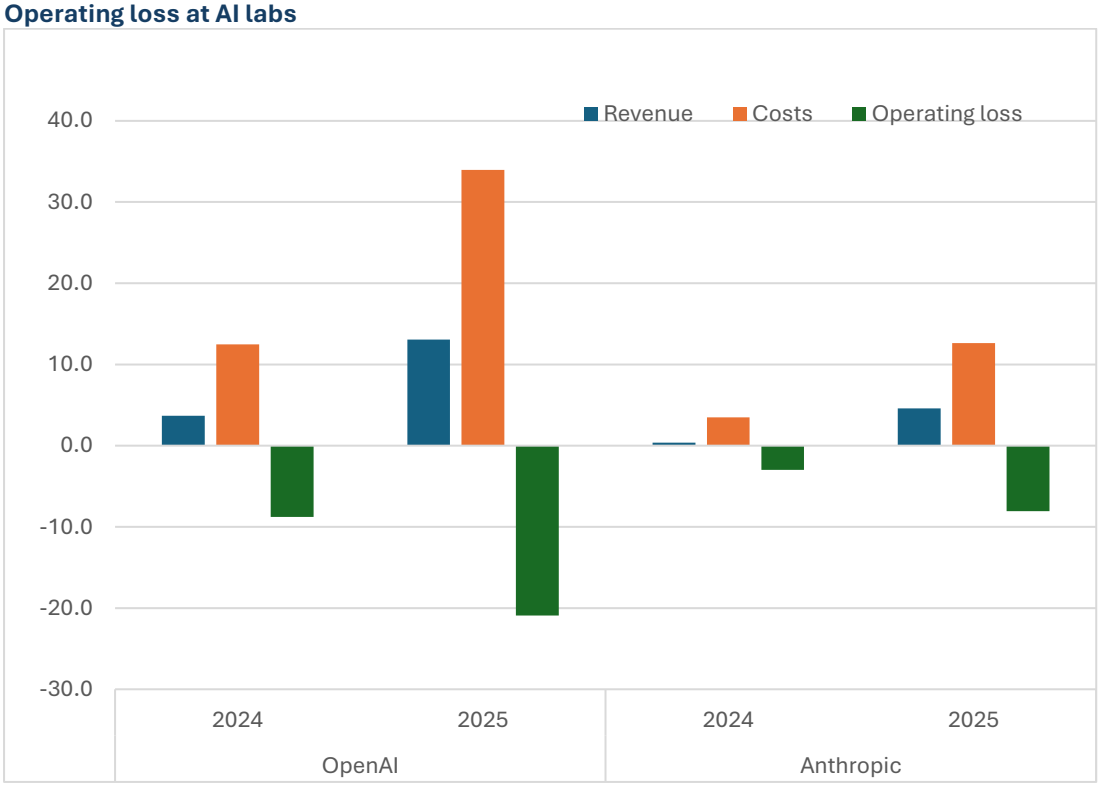
The flows from the VCs and customers have gone into keeping the loss-making entities of OpenAI and Anthropic afloat. The sheer scale of the losses is breathtaking. I've ignored the financing related losses that both OpenAI and Anthropic incurred in 2025, and the chart below focuses on the operating losses.

For every $1 of revenue that OpenAI achieved in 2025, it spent $1.6. For every $1 of revenue that Anthropic delivered it spent $1.76. No wonder they want you to focus on their revenue growth!

The recently leaked elements on Anthropic's S-1 show a commitment of $518bn over the next decade in terms of compute. Some 80% of which is uncancellable. Assuming (an admittedly unlikely) straight line scenario, Anthropic is already committed to spending at least $41bn per year on compute.

OpenAI shows compute commitments of at least $710bn over the next ten years. Again assuming a straight-line scenario, OpenAI is on the hook for $71bn per year for compute!

Operating loss at AI labs



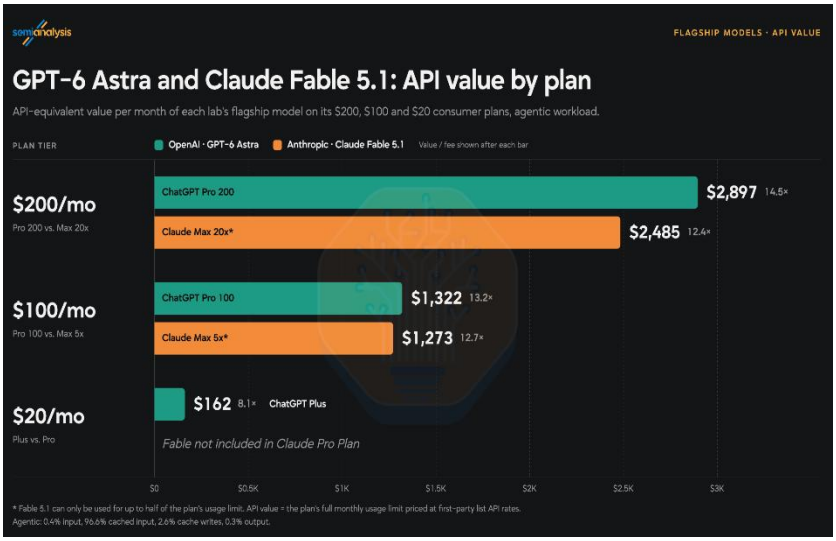
Source: Media Reports

However, one of the key issues that needs to be resolved is the way in which compute costs scale with use. In the world of software, the cost of a marginal user is very low. The big costs are the upfront development costs and the infrastructure costs, but each new user doesn't increase either of these significantly. So as the business scales the growth drops to the bottom line.

It doesn't appear that the same can be said of AI. The more you use AI, the greater the required compute. It appears to me that AI has the worst of all possible worlds, high fixed costs (development and infrastructure) and high variable costs! Yes, the fixed cost element benefits from scale, but the variable cost element doesn't.



As an illustration of the potential problem of costs scaling with revenue, consider the chart below of Token usage. Agents (which, remember, are really to do lists or loops plus an LLM) are now using vastly more to tokens than humans (by nearly a factor of 8.5 times more).



Source: Semianalysis

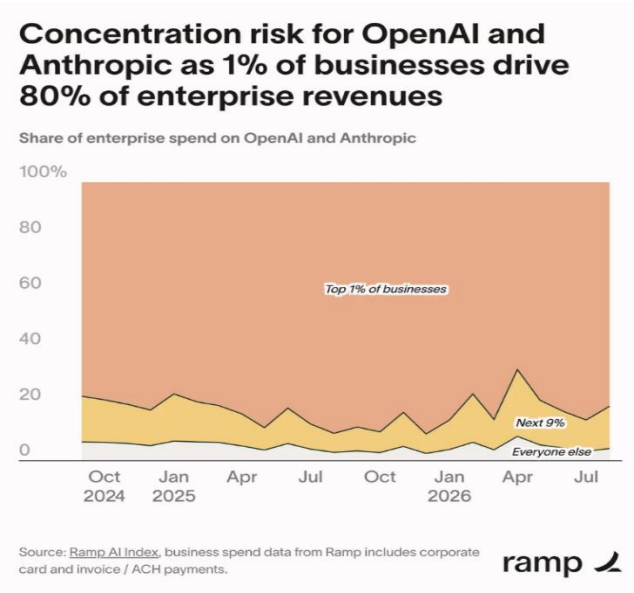
Now consider the two business models that AI labs are pursuing. One is the subscription approach where you pay a flat rate and can use a certain number of tokens (often dynamically capped). But in all cases the subscription price is way below the cost of the tokens. The other route is a pay per use basis, where you can truly see how much the service is costing you – although you don't know in advance.

Neither of these is without issues, on the subscription model token usage easily surpasses the amount paid (especially amongst the highest users, who are probably the ones using complex agent based requests, rather than simple questions).

On the pay what you use model, the end client is faced with uncertainty about their bill until the end of the project, and having to pay the true cost, both of which could well inhibit usage. Neither of these approaches look like easy wins in terms of being able to generate profits.

It is also worth noting the concentration risk that OpenAI and Anthropic both face. According to Ramp data, 1% of businesses drive 80% of the revenues from enterprise across the two companies.

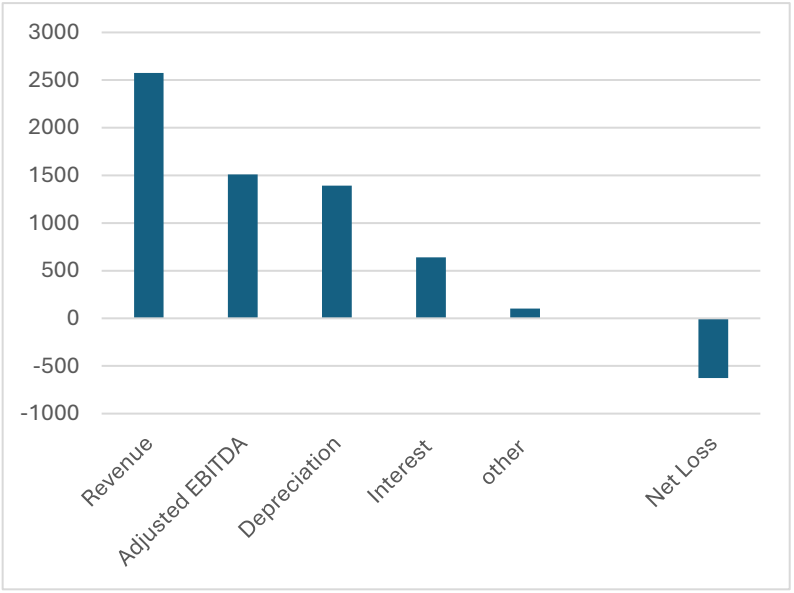
Concentration Risk



Anthropic discloses that two customers, Cursor and GitHub Copilot, alone make up 25% of their 2025 revenue. This strongly shows the dominance of coding in Anthropic's revenue.

In turn these AI labs (and the application layer – who also pay the labs) pay the hyper-scalers and the neo-clouds for compute (time and access to machinery needed to train, test and run LLMs). Microsoft states their dedicated AI business is running at a $37 bn (ARR). We know from filings that $24.1bn is the result of flows from OpenAI. Amazon says that the AI business within Amazon Web Services (AWS) is running at $25bn. Google cloud is expected to produce $80bn in revenue this year (of which it is estimated 27% comes from AI labs). In all cases, the hyper-scalers capex dwarfs their revenues achieved from AI.

CoreWeave: Q2 2026 ($mn)



Source: CoreWeave earnings presentation

The hyper-scalers report "remaining performance obligations" which are unearned revenue that they expect to be invoiced in the future. As of the middle of 2026, Microsoft, Google, Amazon and Oracle are showing RPOs of just over $2tn! Roughly 50% of those RPOs are from OpenAI and Anthropic! Just recall that as of now these two companies are producing revenue of around $57bn. That's some incredible growth expected in the next 2-4 years! And as noted above a huge liability for OpenAI and Anthropic. While the market trumpets the RPO as signs of future demand growth, a more considered investor would frame these as claims against two already loss making companies and perhaps give pause for thought.

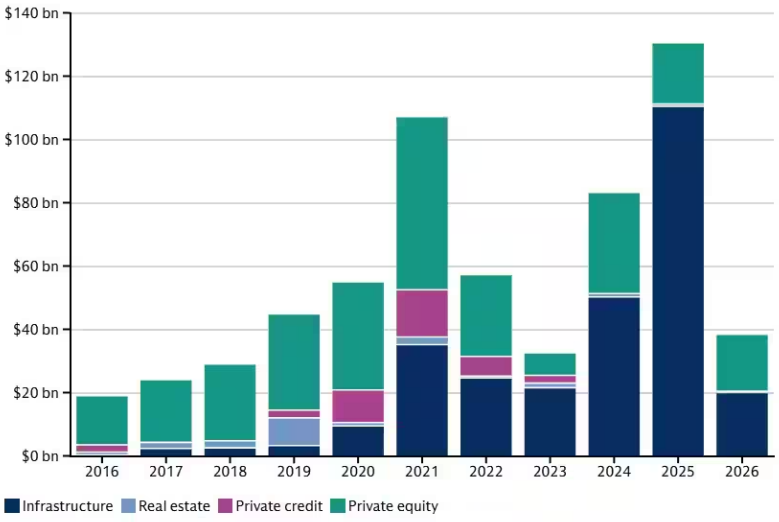
The neo-clouds, such as CoreWeave, show us the challenges of supplying compute. CoreWeave reported a $1.5bn adjusted EBITDA for Q2 2026. Now as soon as you hear the term adjusted EBITDA you should be nervous. This is back to the old days of the internet bubble when we used to call such figures EBBS (earnings before bad stuff). You simply can't go around ignoring the depreciation charges for a company who is providing infrastructure such as data centers, nor can you ignore the interest payments required to service the debt that such companies use to finance themselves. Depreciation charges alone wipe out 92% of the adjusted EBITDA! Making these changes eventually reveals that CoreWeave made a net loss of $0.6bn.

The extent of the interest charges reminds us that alongside the exogenous capital flowing into the AI system detailed above, when it comes to data centers there is another inflow, coming from the private markets (credit and equity).

So, what is all this capex spending by the likes the hyper-scalers and neo-clouds going on? According to Epoch AI, a one gigawatt AI data center costs around $38bn in front capex. The vast majority of this is taken up the technology within the center (GPU etc.) accounting for over 50% of the cost! Recently Zuckerberg of Meta stated that the upfront capex was more like & 50-60bn today, with annual running costs of around $15bn!

A range of alternative assets have helped finance AI data centers

Global data center deal value, by year and alternative asset class (in USD or equivalent).

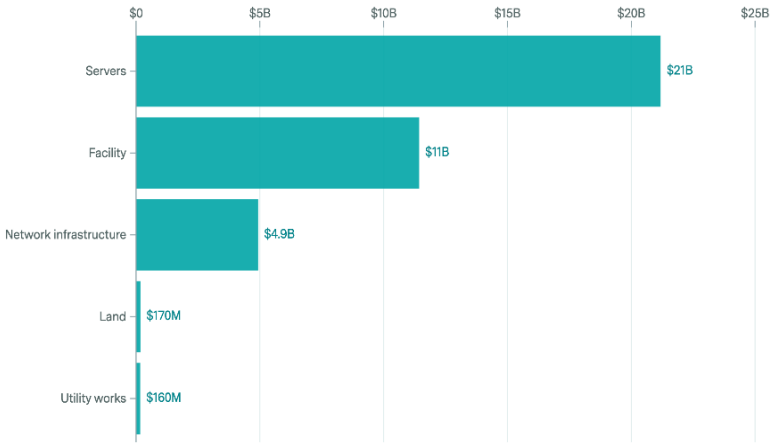


Source: Preqin, Goldman Sachs Research
2026 data is year to date; as of May 29, 2026

Goldman Sachs

According to Currence, some 26GW are supposed to be constructed and operational this year. With an average capex cost $37bn, this amounts to $962bn. And this is expected to be doubled in 2027! Now as the chart below shows only around half of this amount is actually under construction currently. One can't help but wonder if this will the hyper-scalers saving grace.

One-gigawatt AI data center up-front CapEx



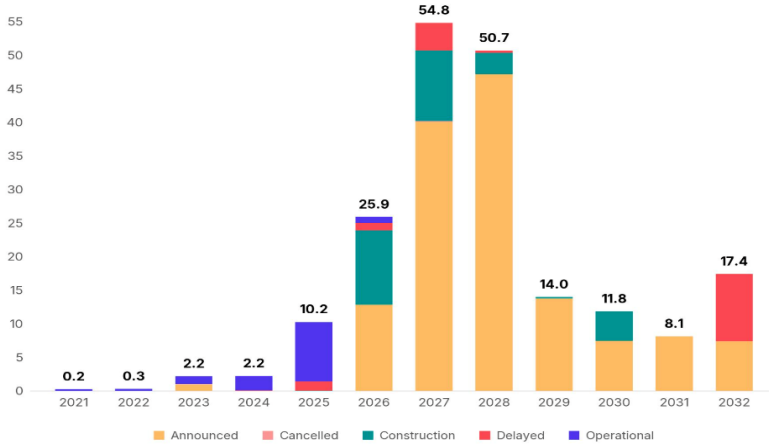
Based on a stylized model of a US hyperscaler AI data center that has 1 GW of nameplate capacity for the IT equipment. "Facility" includes the building shell, mechanical systems, and electrical equipment.

EPOCH AI | CC-BY epoch.ai

Let's assume that all $962bn of data centers get built. Furthermore, let's assume a depreciation life of 8 years (composed on facilities with an

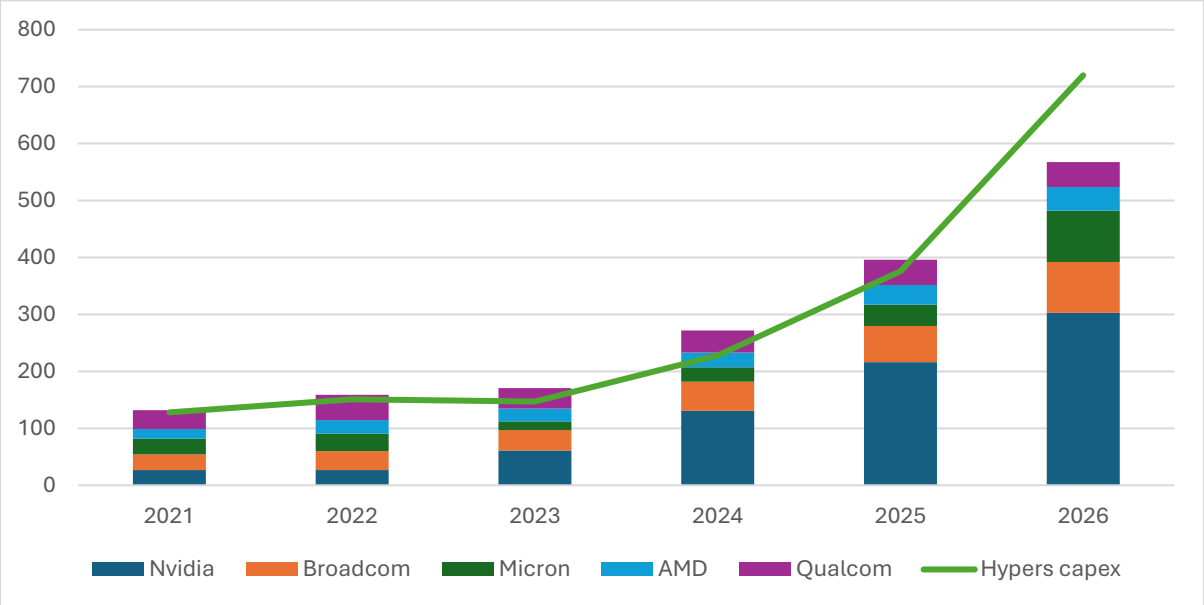
average lifespan of 15 years, and servers and networks with an average lifespan of 5 years). The depreciation charge on this years build alone would be $120bn! Collectively, data center owners need to make this much each and every year just to break even – assuming they aren't debt financed which we know is not a good assumption. By the time we add in a decent return (and the interest costs), the economics aren't looking good.

Project pipeline, by operational date (GW)



As is obvious from the quick look at data centers, the major expense is servers and chips, so we know where the money is flowing. Nvidia, and to lesser extent Broadcom, Micron, AMD and Qualcomm. It is no coincidence that sales of these companies track the expenditure of the hyper-scalers capex.

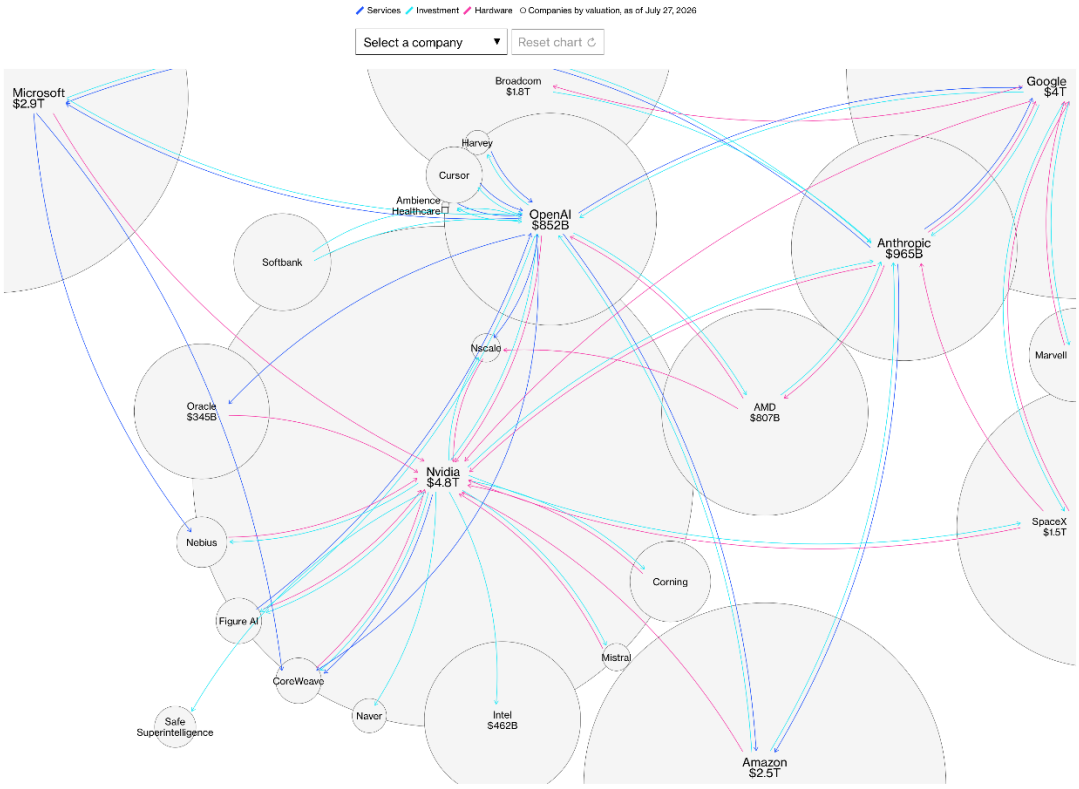
Chip producers revenue tracking Hyper-scalers capex (US$bn)



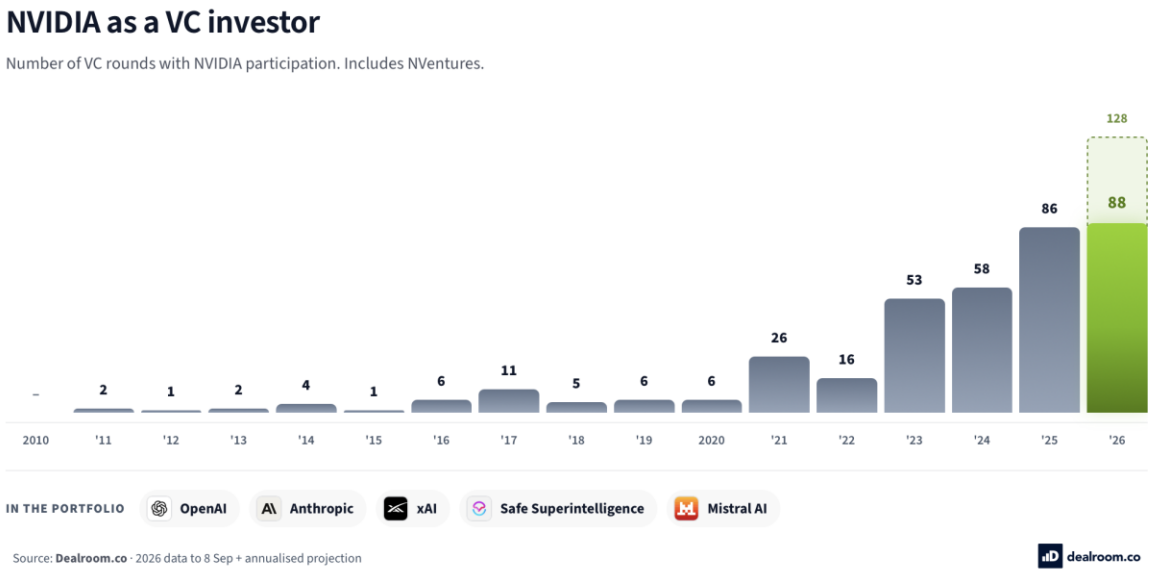
Source:S&P Cap IQ

Nvidia looks priced for belief in this new era. I have often articulated the power of a reverse engineered residual income model. Running such a model on Nvidia reveals that to justify its current market price (August 2026), Nvidia will need to deliver some 30% p.a. earnings growth each and every year for the next 10 years! This despite the fact that pretty much everyone from OpenAI to Microsoft is in the process of trying to reduce their reliance on Nvidia chips by producing their own. Yet another example of the egregious faith that market participants are placing on this time truly being different!

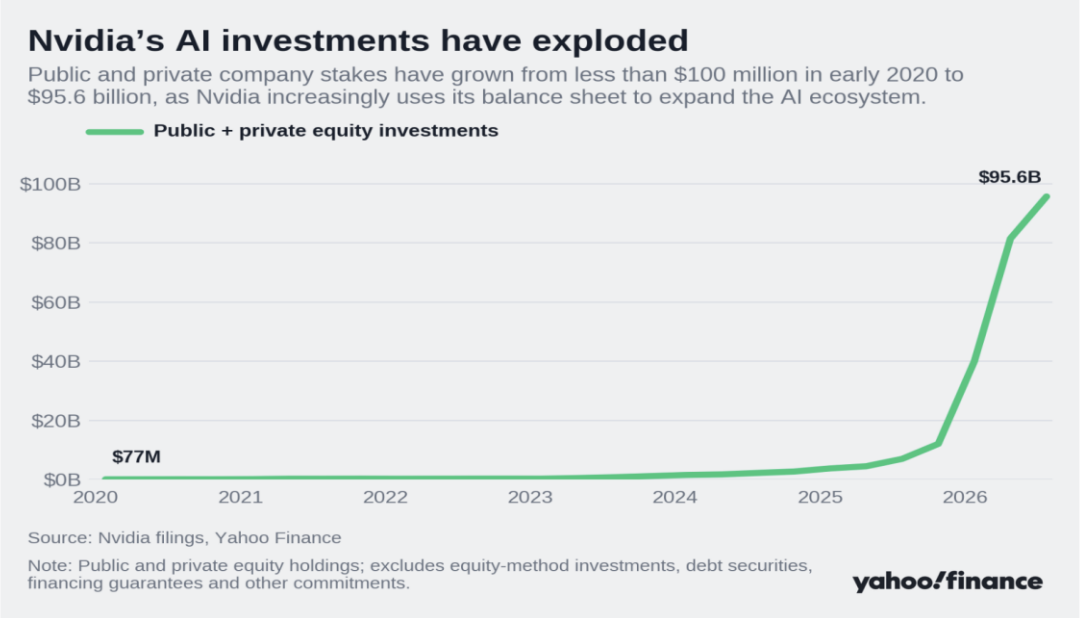
As we pointed out in Curmudgeon's Corner 3, Nvidia sits, like a funnel web spider, at the very heart of the circular financing web that is the hallmark of the AI ecosystem.



Indeed, Nvidia is a massive VC participant. According to Dealroom in the first eight months of 2026, Nvidia was an investor in 53 venture rounds of $100mn or more. Counting rounds of any size, this rises to 88 rounds.



Combine this with the Nvidia's stakes in public companies and it is increasingly looking like the central bank of AI.



Not content with this, Huang along with Apollo, BlackRock, and Goldman Sachs inter alia is attempting to turn GPUs into an asset class. "This is really the first time that technology chips have become an investable asset class" as Huang put it. He even went as far as to claim "These are revenue-generating assets now. They're productive, they're long lived, they're fungible, they're flexible…you have to think about it like it's infrastructure".

Larry Fink, BlackRock's CEO, described this as "next future of financial engineering".
It is noteworthy that every major financial bubble seems to have been accompanied by the creation of some form of financial engineering. All I can hear is J.K Galbraith's wonderful warning in A Short History of Financial Euphoria "The world of finance hails the invention of the wheel over and over again, often in a slightly more unstable version. All financial innovation involves in one form or another, the creation of debt secured in greater or lesser adequacy by real assets… All crises have involved debt, that in one fashion or another, has become dangerously out of scale in relation to the underlying means of payment".

The Rise of an Earnings Bubble - The impact on the S&P 500

The final thing I want to touch on in this already overly long essay, is the impact of all the above on earnings of the S&P 500. In the past, I have delineated amongst different types of bubbles. I will spare the reader the full taxonomy here (I feel I have tortured you enough if you have made it this far, dear reader). Suffice it to say, that pantheon of bubble types contains a class known as intrinsic bubbles – perhaps better explained in investment parlance as fundamental or earnings bubbles. Perhaps the best example of this kind of bubble was the financials during the run up to the GFC. Their extraordinary earnings growth was driven by the unsustainable housing bubble. However, because the earnings were high, the valuations didn't look stretched, as the expense was masked by overinflated earnings.

As we noted in Curmudgeon's Corner Volume 3 and shown above, we have a situation where one company's expenditure is another company's income (sometimes with that expenditure financed by the vendor). When that expenditure is treated as capex then a mismatch occurs, as the supplier (Nvidia for example) recognizes that revenue straight away, as they should since they have sold goods (GPUs). However, the purchasers will depreciate the assets (most of the hyper-

scalers are using 5-6 year). Thus, only twenty percent of the cost is taken in the first year after the capex. This is all perfectly normal, of course, but the sheer scale of the undertaking with respect to AI makes the situation unusual. The arrangement obviously leads to profits being boosted by the mismatch.

Of course, in the current situation, listed sector profits are further inflated. The hyper-scalers put their GPUs to work (maybe). The AI labs are not currently listed, which means their losses (see above) are outside of the scope of the S&P 500.

I built a very simple top-down model to see how much of the S&P 500 earnings were the result of the hyper-scalers capex. I modelled two rounds of flow through from the hyper-scalers capex. The first one representing the immediate beneficiaries of the capex such as Nvidia, (i.e. high flow through and high margins) and the second one representing the construction companies, power suppliers etc. (lower flow through and lower margins). Using this simple model, I estimate that without the hyper-scalers capex, S&P500 earnings would be about $40 per share lower (i.e. rather than the consensus expectation of around $330 per share, earnings would be $290 per share). This represents almost half of the total earnings growth expected in 2026!

Of course, this isn't the whole story. As noted in Curmudgeon's Corner Volume 3, the web of entanglement also includes a lot of shareholding. Nvidia as noted above has become the central bank of AI investing in just about every AI lab and neo-cloud company. The hyper-scalers are also invested in the AI labs. As the paper worth of these investments has risen, Nvidia and the hyper-scalers are able to claim these gains as profits! Stripping out the impact of these holdings from 2026 earnings removes another $27.5 per share from earnings.

The total impact of these two illustrative assumptions described above, I estimate the S&P 500 EPS at around $260 per share, and 2026 growth around 6%, rather than the 33% that is expected. Caveat Emptor.

Conclusions

Keynes once urged "Economists must leave to Adam Smith alone the glory of the Quarto, must pluck the day, fling pamphlets into the wind, write always sub specie temporis". I have singularly failed to heed his advice with this missive. So in a desperate attempt to make amends, and in the futile hope of appealing to the Gen Z worshippers of TL; DR let me offer a brief summary of this work.

Myth 1: LLMs can be scaled to AGI

Myth 2: There will be an ensuing productivity and growth miracle

Myth 3: This is a boom not a bubble

All three myths are connected. If LLMs could be scaled to AGI, then they might produce a productivity and growth miracle, and if this was the case, then perhaps we might be witnessing a boom not a bubble.

You can think of them as compound probabilities if you like. Even if you assign a 50% chance to each of the myths the chances that all three are true are only 12.5%.

However, to me, all three myths are exactly that, fictions fabricated to drive a narrative. Each one has a vanishingly small probability of being true, let alone the compound.

LLMs are distorted mirrors of their training sets, stochastic (sycophantic) parrots, and compressed correlation matrices. Nothing more, nothing less. They offer the illusion of thought but no real cognitive ability. No amount of scaling will lead to them becoming artificial general intelligence.

The impact of LLMs on economic growth is likely to be paltry. At best it may result in a very small rise in productivity. We aren't on the cusp of explosive growth.

We are witnessing possibly the biggest capital misallocation the world has ever seen. It appears as if everyone is going all in on LLMs and their incumbent hardware. Unlike many investment booms of yesteryear, the legacy of LLM boom is short duration (5-6 years at best, not the 50 years for railroads or the 30-40 years for fiber optic).

We have a terrible confluence of a stock market valuation bubble, an earnings bubble (making the scale of the valuation bubble understated) and debt bubble (in the arena of the financing of data centers).

The AI Emperor is hallucinating clothes!

Important Disclosures & Disclaimers

ACR Alpine Capital Research LLC is an SEC-registered investment adviser. For more information, please refer to Form ADV on file with the SEC at www.adviserinfo.sec.gov. Registration with the SEC does not imply any particular level of skill or training.

The views expressed are those of the author and reflect an interpretation of currently available academic research. Other researchers and practitioners may reach different conclusions.

Data and statistics are derived from the sources identified in the text or accompanying exhibits. Calculations labelled as ACR analysis reflect ACR's application of the assumptions described and have not necessarily been independently verified.

It should not be assumed that recommendations made in the future will be profitable or will equal the performance of the examples discussed. You should consider any strategy's investment objectives, risks, charges, and expenses carefully before you invest.

This information should not be used as a general guide to investing or as a source of any specific investment recommendations and makes no implied or expressed recommendations concerning the manner in which an account should or would be handled, as appropriate investment strategies depend upon specific investment guidelines and objectives. This is not an offer to sell or a solicitation to invest.

This research commentary presents the author's views concerning artificial intelligence, productivity, corporate investment and financial-market developments. It is provided for informational purposes and does not constitute individualized investment advice or a recommendation to buy or sell any security.

Opinions and estimates offered constitute our judgment as of the date set forth above and are subject to change without notice, as are statements of financial market trends, which are based on current market conditions. There are risks associated with purchasing and selling securities and options thereon, including the risk that you could lose money. All material presented is compiled from sources believed to be reliable, but no guarantee is given as to its accuracy.

The investment outlook represents ACR's views on the economic factors that may affect the global capital markets. There can be no guarantee that these factors will necessarily occur as ACR anticipates, nor that if they do, they will lead to positive performance returns. There can be no assurance that any objective will be achieved.

Forward-Looking Statements: This commentary contains opinions and observations regarding potential future economic, market, and technological developments. Such views are inherently uncertain and are based on current information and assumptions. Actual outcomes may differ materially from those discussed, and there can be no assurance that any forecast, expectation, or opinion expressed herein will prove correct. Most charts contain source references, but some do not clearly identify whether data is: third-party sourced, adapted by ACR, or proprietary calculations.