

Introduction

• • • •

In May 2025, the internet briefly chose a new, unfortunate main character: Marco Buscaglia, a Chicago-based freelance writer behind a syndicated summer reading list. Though the list was credited to him, he wasn't its principal author. Rather, the list was written by a chatbot.

We know this is true for two reasons. First, because ten of the fifteen books on the list were nonexistent. They were AI “hallucinations” (a term that AI boosters use because it sounds sexier than “errors”). Second, because Buscaglia confessed to using a chatbot, in an interview with the veteran tech reporter Jason Koebler, who cofounded the journalist-owned news site 404 Media.

In Koebler's article, Buscaglia admits that he used a chatbot to source the list, and blames the inclusion of nonexistent books on the list on his own oversight: “I do use AI for background at times but always check out the material first. This time, I did not and I can't believe I missed it because it's so obvious.”

Buscaglia comes across as utterly mortified, telling Koebler, “No excuses. On me 100 percent and I'm completely embarrassed. It's a complete mistake on my part ... This is just idiotic of me, really embarrassed.”

For a few days, the internet had a lot of fun with poor Buscaglia, from the baffled authors whose nonexistent books were recommended on his list to the legion of people who are sick to the back teeth of having AI jammed into their eyeballs by overcapitalized tech companies.

But Koebler’s analysis went beyond Buscaglia’s contrition. In follow-up stories and on the 404 Media podcast, Koebler dove deep into the context of Buscaglia’s unfortunate moment in the internet’s baleful glare.

Start with the list itself. The summer reading list was many lists and guides in a sixty-four-page “Best of Summer” supplement that was distributed to multiple daily newspapers by King Features Syndicate, a division of Hearst Communications, a massive and established publishing conglomerate.

Buscaglia wrote (or “wrote”) the majority of those lists. This is a remarkable feat! Koebler has experience writing lists like these, having begun his career as an intern at *Washington Monthly* magazine, where a list like this would be assigned to a team of three interns, overseen by an experienced newsroom journalist and backstopped by a large and skilled fact-checking department.

I’ve done this kind of work myself. I got my start as a journalist with assignments from *Wired* magazine that had me devoting an entire month to evaluating dozens of multitools, speaking to experts and manufacturers’ representatives, in order to pick my top three and write a brief paragraph about each. These were fun assignments, but they were also serious and involved endless back-and-forth with the fact-checkers, even down to providing proof of how many different blades and attachments each tool had.

To be frank, there’s *no way* that Buscaglia could have written all those lists and summer guides (including a guide to hammocks that features a quote from a nonexistent professor of leisure studies who is credited with writing a scholarly journal article on “hammock culture”—that paper also doesn’t exist).

In 2025, any editor who assigns a freelancer to “write” the bulk of a sixty-four-page summer guide is either tacitly or explicitly commissioning that writer to confine their efforts to prompting a chatbot to barf up a bunch of plausible verbiage and hammer it into shape.

Buscaglia was asked to do the work of dozens of writers, researchers, editors, and fact-checkers, on a short timescale that guaranteed that the resulting product would be riddled with AI-generated errors. We don’t know how much Buscaglia was paid for this job, but it’s a sure bet that it doesn’t

add up to the total salary of all the skilled professionals whose jobs he was asked to do.

In other words, Buscaglia was set up to fail. His job wasn't merely to oversee a chatbot—it was to absorb the blame for that chatbot's mistakes.

Like Buscaglia, I am a freelance writer. What's more, I'm a freelance writer who has recently used AI in the course of my work. I was trying to locate a quote from an expert I'd heard say something very smart on a podcast. The only problem was that I couldn't remember the name of the expert *or* which podcast I'd heard him on.

So I downloaded an open-source AI speech transcription model called Whisper and installed it on my laptop. Then I downloaded thirty or forty hours' worth of podcasts that I'd recently listened to, and fed them all to Whisper, asking it to produce a transcript of all of them.

While Whisper chugged away in the background on my laptop, I kept working on the article, answering emails and alt-tabbing to social media in case someone had uploaded a video of a guy putting a lemon up his nose (normal writer stuff, in other words).

A couple hours later, I opened the folder where Whisper dumps its transcripts and searched them for some phrases I remembered from the quote I was looking for, and located the file in minutes (it was Hagen Blix, coauthor of *Why We Fear AI*, discussing the plight of AI therapists on the excellent *This Machine Kills* podcast). Then I used the time code Whisper provided to open the podcast file and clean up the transcribed quote, which I pasted into my article (I was right; it was a great quote).

This was a *great* experience. It was a great *AI* experience. Thanks to AI, my computer has now acquired a permanent new feature: it can turn arbitrary amounts of recorded speech into pretty reliable transcripts, in a manner so efficient that my laptop's fan doesn't even turn on. The fact that Whisper is open-source means that it can be maintained forever, by anyone in the world who has the skills and desire to do so, no matter what happens to Whisper's manufacturer (a grossly overhyped and terrible firm called OpenAI).

Two freelance writers, using AI. One was made miserable and embarrassed by his AI usage. The other had a delightful experience, saved a bunch of time, and produced a better piece of writing. How to explain the paradox?

One possibility is that the difference lies in *how* we used AI. I used AI to transcribe some audio, whereas Buscaglia used it to generate some writing.

Though I don't brainstorm with AI or other forms of automated or random text generation, I don't have a moral objection to these practices. There are plenty of writers who find text generation to be a productive tool for producing excellent work. The Dadaists' "cut-ups" involved slicing phrases out of a manuscript and scrambling them around, looking for serendipitous juxtapositions. Writers delight in Brian Eno and Peter Schmidt's "Oblique Strategies" deck, which contains fifty-five cards printed with mysterious orders ("Be the first person to not do something that no one else has ever not done before"). The Surrealists loved their "Exquisite Corpse," in which a group of writers take it in turn to append text to the previous writer's text, but each writer can see only the last sentence of the work when they begin.

Today, I know lots of writers who use chatbots to produce fine work: they might use the chatbot as a sounding board for evaluating ideas (or variations on ideas), or to challenge them with writing prompts, or to suggest improvements. As I said, I don't use AI that way, but I know people who do, and I like the things they write.

The difference between Buscaglia's AI experience and my own—the difference between a useful tool and a technological torment—is the difference between a *centaur* and a *reverse centaur*.

In automation theory (the academic study of automation), a *centaur* is a person who is assisted by a machine. Think of a clever human's head, arms, and hands atop a horse's strong body. Riding a bicycle or driving a car makes you a centaur; so does using the mute button on your TV remote when an ad comes on. Wearing a hearing aid makes you a centaur, and so does using a calculator to multiply large numbers. Being a centaur can be glorious. My whole writing career has been a series of centaur moves, from the used IBM Selectric my parents bought for me to play with when I was six or seven to the Apple II Plus we got when I was nine, which kicked off an unbroken string of better and better writing tools, with spellcheckers, version control, collaboration features, autosave, and more. Best of all, I get to choose *exactly* which of these features I use. I work for myself, after all, so if I don't want to use the grammar checker (hell no), that's my business. No one expects me to write more pages when I get a new tool. A couple of

years ago, a former student of mine asked me to try out his LLM tool that would help me write dialogue and flesh out characters. I played with it for ten minutes and then never went back to it. No one told me I was being uncooperative or spoiling a grand plan to increase efficiency and realize cost savings by refusing to use AI. Sometimes, when I'm really stuck, I write with a pen in a notebook. No one cares, except me.

A *reverse centaur* is a human who is conscripted into acting as an assistant to a machine. There's a classic *I Love Lucy* episode where Lucy and Ethel are working on an assembly line at a chocolate factory, taking bonbons off the belt and wrapping them in paper. As the belt goes faster and faster, Lucy and Ethel have to work at superhuman speed. They're reverse centaurs: the machine can move the chocolates from one place to another, but it needs a human to pack them into the box, and the humans who act as its assistant are made to work at a pace that exceeds all human capacity until calamity ensues.

I Love Lucy played that bit for laughs. Amazon warehouse workers get the horror-movie version. They are observed by AI-equipped cameras that time their movements and monitor their "time off-task," penalizing them if they fail to make quota. Amazon warehouse workers experience the highest level of on-the-job injuries in the U.S. warehouse sector, and many of them have to resort to urinating in bottles because a visit to the toilet would blow their quotas. An Amazon warehouse is full of machines, but there are jobs the machines can't do, and that's where the Amazon warehouse workers come in—they assist the machines. They are reverse centaurs, and they have been conscripted to serve as peripherals for the warehouse's automation systems. They aren't merely used by those machines—they are used *up*.

It's not hard to imagine how a warehouse worker might *choose* to use AI in their daily work—for example, a computer vision system in a pair of smart glasses highlights an item they're looking for in a bin (there've been many instances where I've stared directly at something without seeing it in which *I* would have loved this). Automation isn't *necessarily* the enemy of warehouse work: there's nothing wrong with a forklift! The difference between automation that helps a warehouse worker and automation that torments that worker is whether the worker gets to choose where, when, and

how to use that automation. It's the difference between a centaur and a reverse centaur.

I love the automation system in my car that warns me if I'm drifting out of my lane, and I recently discovered (the hard way!) that if the person ahead of me brakes suddenly, my car will let out a sphincter-puckering series of beeps and activate its own brake. I was pretty happy about that, even if it did come as a hell of a surprise.

Compare that with the drivers in those Amazon vans rolling around your neighborhood. They have to sign into at least *nine* separate apps, and they are continuously scored based on their driving performance, as assessed by various AI tools. Drivers lose points for braking or swerving (even if that's the only way to avoid a sudden road hazard) or for deviating from the proscribed route set by the AI (even if there are obstructions or hazards). Drivers have impossible-to-meet quotas and the per-parcel compensation rate drops if they fail to meet it. Drivers are forbidden from peeing in bottles, but also given no time to urinate. The driver is just a peripheral for the van, present only because the van can't drive itself or get your parcel onto your porch. They are a reverse centaur.

This centaur/reverse centaur distinction is the heart of the paradox at the heart of the debate about the usefulness of AI tools. When you find yourself surrounded by people swearing that a given tool is worse than useless and others swearing that it has made their lives easier and better, you can bet that the former group is made up of reverse centaurs who've had AI imposed upon them, the latter group is all centaurs who've gotten to make up their own minds about where, when, and how to use AI tools. The solution to the paradox is to stop thinking about what the gadget does, and pay attention to who the gadget does it *to* and who the gadget does it *for*. The important part isn't the technical characteristics of the device, it's the power relationships of the people who use the device.

It's been a long time since I last held a factory job. I'm a science fiction writer by trade, which means that I get asked about AI about a million times per day. Some people think science fiction writers are seers or prophets (regrettably, some extremely delusional science fiction writers share this view). They imagine that science fiction is a literature that predicts the future.

This is nonsense, of course. The future isn't predictable, which is a damned good thing, because if the future were predictable, it would be foreordained, which would mean that the actions taken by people don't matter.

But science fiction writers *do* have an intimate relationship with the future. By creating futuristic parables—made-up stories about futures that may never come to pass—we SF writers remind everyone that the future is up for grabs, full of possibility.

That's in stark contrast to tech bosses, who claim total authority over our possibilities, both present and future. There's Mark Zuckerberg: *Yes, well, obviously you'd like to talk to your friends without my spying on you all from asshole to appetite, but let's be realistic here. That's like asking for water that isn't wet.*

Or Apple CEO Tim Cook: *Yes, certainly, it would be nice if we could provide you with a reliable mobile device without locking you into our App Store, where we rake in a 30 percent commission on everything you spend, and where we get to decide which apps you can use, and which ones you can't (like the privacy apps that we banned in China or the ICE-tracking app we banned in America). But come on, it's just not realistic to want a computer that a) works and b) takes orders from you, not me.*

This is a cheap bully's trick: insisting that their abusive behavior is out of their hands, that they are merely acting in accord with some kind of iron law or great force of history. Margaret Thatcher was the world champion of this: she insisted that her cruel program of austerity and privation was inevitable, repeating the phrase "There is no alternative" so often that wags started referring to her as "Tina" ("There Is No Alternative").

Thatcher's contemporary successors are to be found in the C-suites of the tech industry.

Call this philosophy *inevitabilism*, the insistence that there is only one conceivable way to do things, and any problems you're experiencing aren't anyone's fault, they're just inescapable *reality*. The fact that the world is filled with sneaky Bluetooth sniffers that track your location as you walk through shopping centers or even just down the street isn't something you can blame a person or a company for; it's just a fact.

Science fiction is an anti-inevitalist literature. Science fiction is only incidentally about thinking up new gadgets and explaining what they do.

Far more important than what the gadget *does* is who it does it *to* and who it does it *for*. As the legendary science fiction editor Gardner Dozois once wrote, “Most SF can predict the car, some SF can predict the drive-in theater, but SF that can predict the changes in teen-age sexual behavior as a result of the drive-in is vanishingly rare.” To this I would add, “Truly visionary science fiction might think about how a world where you need a government-issued ID like a driver’s license in order to engage in sexual experimentation might plausibly lead to a ‘database nation’ of ubiquitous surveillance.”

The mere existence of a literature in which many different writers have imagined many different futures, each of which can *feel* possible and even desirable, is a rebuke to inevitabilism. The fact that social relations between people and their technology can be different means that the current arrangement is a choice, and—crucially—it means that we can choose something else.

AI hucksters want you to believe that all the things they call “AI”—an incoherent grab bag of many technologies, some of them not especially related to the rest—are coming, and that when they arrive, there is only one conceivable way that we could use them. This is just high-tech Thatcherism, the inevitabilist move of a bully who insists that they’re only doing what implacable reality demands of them.

This is a book about what AI can and *cannot* do, but even more important, it’s about the possible social arrangements of AI, from not using some AI technology at all, to using it in ways that let some of us choose to be centaurs, while saving our friends and neighbors from being conscripted into reverse-centaury.

The current wave of AI is full of software performing impressive feats in both parsing and generating language and images and sounds. There are lots of interesting, fun, and productive ways to use this technology. There is nothing about the *technology* of AI that determines how it *must* be used. We can choose to use it sometimes, or never, or all the time, depending on our needs and proclivities. We don’t have to let billionaires tell us how it must be used.

Margaret Thatcher’s “There is no alternative” has a fine rejoinder in the science fiction writer William Gibson’s famous maxim, “The street finds its own use for things.”

SF writers make lousy prophets, but we can be pretty good technology critics.

OceanofPDF.com

BEFORE



OceanofPDF.com

How We Got Here

• • • •

The AI hype machine *runs* on badly considered criticism, and it's not unique in this regard.

Lee Vinsel, a Science, Technology and Society scholar at Virginia Tech, described this process in a widely cited essay entitled “You’re Doing It Wrong: Notes on Criticism and Technology Hype,” which defined a critical misstep that Vinsel dubbed *criti-hype*, which he defines as “criticism that both feeds and feeds on hype.”

When a tech company tells a fanciful story about the terrible things its products can do, they do so in anticipation of some advantage from convincing the public of their fearsome capabilities. For example, the advertising industry has long touted its ability to bypass our rational faculties in order to sell us things, aided by some kind of mind-control technique.

In the 1950s, ad companies provoked a moral panic by claiming that “subliminal advertising” could implant ideas directly into consumers’ minds. They claimed that subtle images and words hidden in the shadows of illustrations, eyeblink single-frame insertions in films, and backward audio or audio that was too high or too low to be perceived by the human ear would act directly upon our desires, causing us to rush out and buy whatever they were selling.

It’s easy to see why the advertising industry would want to tell this story: because it helps them sell ads. Just like your boss is primed to believe a story about how someday he’ll be able to fire you and replace you with a

docile chatbot, so, too, are merchandisers primed to believe that there is a tool that will cause you to march into their store, wallet in hand, and buy whatever they're offering at whatever price they're charging.

Many of the people who promoted the story of subliminal advertising doubtless believed it. Others may have been cynics who gave the clients what they wanted to hear.

Today, many of the practitioners of surveillance advertising techniques also doubtless believe their own stories (and Rasputin may have believed that he could hypnotize people with his mystical stare).

★ ★ ★

That doesn't make it true. The evidence for mind control via targeted advertising is awfully thin,¹ and Big Tech critics who breathlessly repeat these claims are unwitting accomplices to the sales departments of ad-tech companies: *Why should you pay a 40 percent premium to advertise on Facebook? Just ask my critics: I've invented a mind-control ray.*

Contrast this with an anti-criti-hype approach: *Ad tech claims it has a mind-control ray, and it's using that outlandish claim to bilk advertisers out of hundreds of billions while amassing galactic-scale surveillance dossiers on every living person, which it both leaks and hands out to any cop or authoritarian state that asks for it. What they're calling "mind control" is, at best, just fine-grained targeting. They're not convincing people they're thirsty and then selling them a drink; they're listening in on billions of online conversations to find the people complaining about how parched they are and then showing them an ad for Coke.*

★ ★ ★

This is an approach that strikes directly at the source of ad tech's power, which comes from its profits, which come in turn from its ability to convince rubes that there's such a thing as a Big Data mind-control ray.

Moreover, this is an approach that (correctly) places the ad-tech industry on one side of a struggle, with advertisers (who are being sold a defective product) and advertisees (that's us, being spied upon and exploited) on the

other side. Obviously, there's a certain irreducible degree of enmity between advertisers and the public, but we all share a common enemy in the ad-tech industry.

There's a certain kind of policymaker or governmental enforcer whose priority is keeping businesses from being bilked by ad-tech hucksters. There's a different group of powerful people who are charged with protecting the public from surveillance. Framing ad tech as a scam—rather than as a miracle of persuasive technology—puts these two groups on the same side as well. It's a recipe for victory.

The companies that made billions selling (and mis-selling) ad tech are now on the vanguard of the AI bubble. It's the same scammers, pulling the same scam. The last time around, a bunch of tech's loudest critics got sucked into criti-hype, breathlessly warning about how our “dopamine loops” were being hacked by these evil sorcerers. Far from puncturing the ad-tech bubble, this criti-hype inflated it even further, making billions for the people currently driving AI psychosis among the world's deepest-pocketed, most easily gulled investors and policymakers. It will be downright embarrassing (not to mention terribly destructive) if we let these guys trick us into helping them sell their snake oil in back-to-back scams.

To win the AI fight, we need to enlist allies. If you're a worker whose job is in some AI pitchman's crosshairs, then every time you repeat claims about AI's current—or future—ability to do your job, you help that hustler convince your boss to fire you and replace you with an AI.

But it's worse than that, because repeating the AI barker's patter *also* alienates the people who benefit from your work. Some of those people will be grateful to you for the espresso shots you've pulled for them, or the contracts you've drafted for them, or the education you've provided to their kids, but a sizable fraction of this cohort will shed no tears for your technological obsolescence, not if it means that they're going to be able to get an unlimited supply of whatever you produce at prices so low they don't even register.

Here's an example that makes this distinction super clear: for decades, the largest, most profitable companies in the world have been outsourcing their customer service to overseas call centers. This move coincided with a trend toward steeply declining quality and thus a concomitant increase in how often you need to talk to a customer service rep. To make things even

worse, these outsourced customer service reps don't work for the company and have severely curtailed agency and flexibility, meaning they generally can't solve your problem.

Long before AI came along, these people were being used by giant corporations as catchall accountability sinks. Their job was to get yelled at by you, at minimal expense to the company, which is why they earn terrible money and also why they can't give you a refund, change your plane ticket, or compensate you for the fact that your hotel room was covered in meth and dog hair.

No matter how astute you are, it's hard to see these people as class allies in a war against the company that indirectly employs them and directly shafts you.

But that becomes a lot easier when these human reps are replaced by chatbots, who are so much more obviously put in place to simply absorb your frustrations and abuse without helping you.

In 2022, a Canadian named Jake Moffatt contacted Air Canada customer service to find out about bereavement fares so that he could travel to his grandmother's funeral. This was shortly after Air Canada had wired up some chatbots to its customer service desk, and one of these bots helpfully advised Moffatt that all he needed to do to make use of a bereavement ticket was to buy a regular, full-fare ticket, attend the funeral, and then submit a letter with proof of his grandmother's death, whereupon Air Canada would provide him with a refund for the difference between the full-fare ticket and the discounted bereavement fare.

Which is exactly what he did, only to be told by a (human) Air Canada customer service rep that the chatbot had "hallucinated" its advice about bereavement fares, and that company policy required fliers to provide proof of death *before* buying their tickets, and so no refund would be forthcoming.

Moffatt pursued this through every level of Air Canada customer service support, laboriously escalating his claim until he had exhausted all appeals. Finally, he filed a case with the British Columbia Civil Resolution Tribunal, and more than two years later, he was refunded CAD812.02.

There's never just one ant. The likelihood that the only lie an Air Canada chatbot told to one of its customers was this one is vanishingly small. Far more likely is that most of the Air Canada customers who lost

money because one of the company's chatbots lied to them were simply unwilling to spend two years playing Air Canada's Kafka LARP just to get their money back. Their loss is Air Canada's gain.

Now, I happen to be Canadian (all the best Americans are, you know), and that means that I have spent hundreds of hours on hold with Air Canada, only to be connected to overseas call-center employees who were systematically curtailed from solving my problems. I am the *last* person to defend the old system of resolving Air Canada complaints.

But.

Those customer service reps—low-paid and ineffectual as they are—have repeatedly and sincerely tried to help me. When they made mistakes, it was due to misunderstandings. No one ever told me that the errors these people made were between me and them, and had nothing to do with Air Canada.

Someone sold Air Canada a chatbot by promising to put these people out of work. That same person knew—or should have known—that the chatbot would “hallucinate” a stream of incorrect advice that would cost Air Canada fliers like me money. In the war against Air Canada's remorseless enshittification, I am on the same side as those poor call-center folks.

OceanofPDF.com

Against Inevitabilism

• • • •

The social arrangements of technology are a choice, not an inevitability. Consider the various iterations of the cash register: prior to the cash register, a grocery store clerk did skilled labor that made them hard to replace. While basic arithmetic isn't that hard to master, reliably summing up orders consisting of many items and making change for eight hours straight, without making major mistakes, requires a rare combination of talent and temperament.

From the grocery clerk's skill came power. The harder it is to replace a grocery clerk, the more the clerk can demand of their boss, including higher wages and better working conditions.

After a grocery store invests in a cash register, the clerk becomes much easier to replace, because the pool of people who can reliably punch buttons on an automatic tabulator and make change according to its calculations is much larger than the pool of people who can manage the same trick with pencil and paper.

This means that the boss can now lower the clerk's wages, add more duties to their workday, or fire half their clerks and give their work to the remaining half.

This isn't inevitable! The boss's freedom to fire a worker or change their job description or wages does not exist in a vacuum. It is socially determined: if the workers are in a union (or even if they aren't but they live in a country where the union movement has successfully lobbied for strong

labor protections), the benefits of the cash register might be more equitably divided between labor and capital.

Depending on these social arrangements, a grocery clerk with access to a cash register might be freed from boring labor and given more time to chat with their customers. If the store attracts more customers, the clerk can serve those customers thanks to the labor-saving properties of the cash register, and demand a share of the higher profits that come from serving more customers without hiring extra staff. The boss who tells you that the only way to use a cash register is to fire your coworkers and make you do their jobs, too, is practicing vulgar Thatcherism. They're trying to bamboozle you with inevitabilism.

After all, there are many users of the cash register for whom it is an unalloyed good, like the grocery across the street that's organized as a worker's co-op, where the worker/owners enjoy the register for its labor-saving properties without having to fight with a boss over their share of the productivity gains of the new technology.

Successor technologies to the cash register—like mobile phone attachments that let people process credit card transactions—allow creative workers to sell directly to passersby at art fairs, comic-cons, and flea markets. They enhance the welfare and improve the material circumstances of workers. The most important thing about the gadget isn't what it does, it's who it does it *for* and who it does it *to*.

Inevitabilists will tell you a version of this story whose moral is, *See, you take the good with the bad*. They're wrong. You don't have to take the good with the bad. You can get the good *without* the bad. The difference isn't to be found in what a technology does; it's in what your boss uses the technology to inflict upon you. If you want to get the good without the bad, you need to switch from fighting *technologies* to fighting *bosses*.

Really Weird Math

• • • •

Your boss almost certainly loves AI. Bosses will not shut up about AI. Adopting AI, integrating AI, using AI—these subjects seem to weigh more heavily on our bosses’ minds than objectively more important issues, like their quarterly profits. When bosses get so excited about a new tech trend that they forget about profits, it’s a good bet that you’re living through a tech bubble.

Even by the standards of tech bubbles, the AI bubble is a *whopper*. Hundreds of billions of actual dollars have been spent on AI hardware and data centers since OpenAI shipped ChatGPT 3 in 2020.

But even more hundreds of billions of *imaginary dollars* have been spent on AI. For example, Microsoft has a long-standing partnership with OpenAI. As part of this partnership, Microsoft gives OpenAI “tokens” that OpenAI can spend to access the computers in Microsoft’s data centers.

OpenAI books these tokens as investment revenue, at face value. This is some very funny accounting. The tokens Microsoft invests in OpenAI can only be redeemed for Microsoft data-center access. When Microsoft “invests” \$10 billion worth of tokens in OpenAI, that “\$10 billion” figure assumes that if you or I showed up at Microsoft and bought \$1 worth of computing, and then Sam Altman showed up and ordered \$10 billion worth of computing, Microsoft would simply multiply the amount of computing it sold to us by ten billion.

Think of it this way: say an ice cream cone costs \$1 and contains one cup of ice cream. If you order ten billion cups of ice cream, that isn’t \$10

billion worth of ice cream. Anyone who orders ten billion cups of ice cream would expect a massive discount relative to the customer who buys just *one* cup—though of course it remains highly unclear what you’re planning to do with all that ice cream.

So OpenAI is booking Microsoft’s tokens as an investment that is grossly inflated—but that’s just the warm-up.

The really weird math comes after OpenAI redeems its tokens with Microsoft to power ChatGPT and its other products: Microsoft books that transaction as \$10 billion worth of AI-related *revenue* to its cloud-computing division.

Imagine that the ice cream parlor has some video-game machines on the back wall, and to play these, you need tokens that you can get from the cashier in exchange for dollars. But the cashier likes you, so they give you \$10 worth of video-game tokens to play, which you promptly pump into the Galaga machine.

Using Microsoft and OpenAI’s funny accounting, the ice cream parlor has “invested \$10” in you (by giving you \$10 worth of funny-money tokens), and then the ice cream parlor “earned \$10 in revenue” (when you pumped the tokens into the machine).

Microsoft and OpenAI aren’t alone in using these accounting gimmicks to create the impression that more investment and revenue are flowing into AI companies. Every AI company is using all kinds of cheap tricks to juke the stats to make AI seem like a bigger phenomenon than it actually is.

For example, you will often see established tech companies like Meta and Google trumpeting the amount of “engagement” their AI products received during the previous quarter. A naive person—or a credulous investor—might reasonably assume that this means that AI is popular among users.

But that is far from the truth. In a paper published in June 2025 in *The Journal of Hospitality Marketing & Management* with the catchy title “Adverse impacts of revealing the presence of ‘Artificial Intelligence (AI)’ technology in product and service descriptions on purchase intentions: the mediating role of emotional trust and the moderating role of perceived risk,” two business-school professors reported on their work surveying everyday people about their attitudes toward AI. They concluded that more

than 90 percent of us are *less* likely to use a product that’s advertised as AI-based or AI-enabled.

If people don’t like AI, what accounts for those glorious quarterly numbers about “increasing AI adoption” and “gains in AI interactions”?

As with centaurs and reverse centaurs, the seeming paradox disappears once you apply the lens of incentive and power to the system. To unpack those power relationships and incentives, we need to step out of AI for a moment and expand our focus to look at the tech industry that gave us the AI bubble.

Where do AI products come from?

Tech companies are among the most *instrumented* firms in world history. Tech companies pioneered the use of “productivity monitoring” tools for their tech staff, counting keystrokes per hour or lines of code per day.

Savvy tech bosses have long known that these measurements were at best imperfect and even counterproductive. You don’t want your coders to type as many lines as possible—you want them to type exactly the right number of *correct* lines. Famously, a young Bill Gates mocked IBM’s practice of awarding bonuses to programmers based on the number of lines they generated as “the race to build the world’s heaviest airplane.”

But while many tech bosses moved past measuring productivity by counting keystrokes, they continued to market those counterproductive productivity tools to their customers. Bill Gates may sneer at “the race to build the world’s heaviest airplane,” but a key selling point for Office365, Microsoft’s flagship enterprise cloud productivity suite, is a set of monitoring tools that provides bosses with a dashboard that ranks their juniors’ activities—how often they type, click, and backspace; how many times they open, close, or create a new document.

Managers can find out how individual employees stack up against one another, see how their departments compare, and even compare their company’s “productivity” with productivity at rival firms in their sector (apparently, no one who uses this feature to see proprietary information about their direct competitors ever worries that this means that those direct competitors can see *their* data).

Meanwhile, tech workers may have largely escaped the kind of granular monitoring that their companies inflict on less-technical workers, but their

bosses haven't given up on measuring their performance. After all, the reason that tech *workers* are so easy to monitor is that tech *products* are so good at monitoring their users.

The digital tools you use can gather endless statistics on your usage. The Instagram app, for example, tracks everything from how quickly you scroll, what's on the screen when you stop scrolling (or just slow down), even readings from your phone's accelerometer that tell Meta how you're holding and moving your phone while you interact with their app. All this can be correlated with your location, the specifications of your device, and your other activities—the places you've visited recently, the places you subsequently visit, the things you buy, and the people you converse with.¹



The fact that your technology tools spy on *you* means that the tech workers who *make* those tools can be rewarded or punished based on your usage.

Perhaps you're old enough to remember Google+, which was Google's 2011 full-court-press effort to launch a social media service to compete with Facebook. Google's product teams were ordered to find ways to integrate G+ into their products (for example, by shifting YouTube comments to Google+ message boards). To impress a sense of urgency upon their employees, Google management made Googlers' bonuses and stock options contingent on how often the users of their products interacted with Google+ (it's common for techies' bonuses to account for much, even most, of their annual pay).

In tech management jargon, G+ interactions were made into the company's overriding "Key Performance Indicator," the number that everyone had to hit in order to end the year in the black. KPIs are used in many sectors, but tech is uniquely amenable to them because everything a user does with a product can be monitored and recorded, creating many categories of usage statistics that can be turned into KPIs.

KPIs definitely have a focusing effect on the working lives of tech workers, but they are among the most perverse of all incentives. A piece of common wisdom goes, "You treasure what you measure." Or, more

pointedly, there's the version known as Goodhart's Law: "When a measure becomes a target, it ceases to be a good measure."

If we're being generous, we can imagine that the Google bosses who turned G+ interactions into a KPI were thinking, "If users are interacting with G+ more, that means that they're enjoying it." Even if they don't get that far, they were almost certainly thinking something like, "The fact that so many users are interacting with G+ will surely convince some investors that Google could be a significant player in social media, in addition to dominating search."

There's a big reason corporate leaders at large tech firms are so concerned with what investors think of their future prospects. That's because as long as a large tech firm can claim that it is growing, it will enjoy a substantially higher stock valuation than a comparable firm that is "mature."

Let's unpack that: publicly traded corporations all have something called a *price-to-earnings ratio* (or P/E). If a company has a P/E of ten, then for every dollar the company brings in, its corporate valuation is \$10 (a company with a P/E of ten that makes \$1,000,000 per year will be valued at \$10,000,000).

P/Es are not created equal. Consider two companies in similar lines of business, say, a taxi company and an app-based ridehail company. The taxi business has been around for fifty years, is handsomely profitable, and has never had a bad quarter. Meanwhile, the ridehail company is only two years old and it *loses* money every year.

But there's another key difference: their rate of growth. The boring taxi company made a million bucks this year. Last year, it also made a million bucks. As inflation rises, the boring taxi company's annual revenue rises with it, but the earning power of the total revenue of the taxi company doesn't change. Ten years ago, you could buy a million loaves of bread with the taxi company's total revenue. This year, you can also buy a million loaves of bread with the take.

This is called a "mature" company. There are lots of them, and there are good reasons to invest in them—say, if you want a safe bet on a small but steady return. They're the boring, safe stock-market equivalent of a mom-and-pop landlord with a duplex that provides the owners with a home

downstairs and a rental flat on top that covers their expenses, year in and year out.

It's a safe bet—unlike the ridehail business, which is operating in a risky field that has seen a lot of failures and just a few successes: Uber and Lyft.

Uber's never been profitable, at least not by the bedrock principles by which investors measure profitability for normal companies. It resorts to the weirdest accounting gimmicks imaginable to make the case that it's actually in the black: for example, many of Uber's foreign operations were such disasters that Uber eventually "sold" them to overseas rivals, who are themselves failing, and who bought out the local Uber operation with illiquid private stock (stock that isn't traded on a stock market and can't be readily converted to cash). Periodically, Uber will declare that their holdings in some dog-shit failing overseas ridehail company have massively increased in value, and, on that basis, declare that it had a profitable year. This is the corporate version of your weird uncle who strains your Thanksgiving dinner conversation by bragging that his holdings of precious Beanie Babies and NFTs have skyrocketed in value and therefore he is now a millionaire.

Lyft, meanwhile, took years to turn even a modest profit, and its actual rate of profit is very low, even with the majority of its expenses being borne by drivers who are misclassified as contractors and who supply the vehicles, fuel, and maintenance to keep Lyft's operation running.

So our ridehailing business—only two years old, yet to turn a profit—looks like a worse bet than the taxi company. But there's one way in which the ridehail service and the taxi company differ that makes all the difference in the world: *growth*.

The ridehail service has increased its ridership—and its revenue—tenfold for every year it's been in operation. Like the taxi company, it is booking a million dollars in revenue per year. Unlike the taxi company—whose profits are \$200,000 a year—the ridehail company is *losing* a million dollars a year.

But two years ago, the ridehail company made \$1,000. Last year, it made \$100,000. This year, it made \$1,000,000. It is a *growth* company. And the thing is, it's impossible to predict when that growth will end. Take Uber: in the company's S-1 (the document given to potential investors before a

company debuts on the stock market, enumerating its strengths, weaknesses, and plans for the future) from 2019, Uber told investors that it planned to displace *every* ride in a motorized land transport in the *world*: every bus, taxi, streetcar, and subway ride *on the planet*.

If Uber succeeds, then buying a share in the business isn't merely staking a claim on the revenue it's making today—it's a bet that pays out with a share of the incredible, world-swallowing revenue the company will generate later on, as it grows. A share of Google in 1998 was a claim on the nonexistent revenue of a small but well-regarded search engine that lacked any kind of business model. By 2025, it's turned into a share of a company that controls *90 percent* of the world's search traffic, bringing in \$350,000,000,000 the previous year.

When investors believe a company is a "growth stock," they are prepared to pay *much* more than they are for a share in a "mature" company. Back in 2020, Tesla hit a P/E of more than *1,000:1*. Around that time, Ford's P/E was less than 10:1.

Why does this matter? After all, valuation isn't the same thing as *revenue*. The company doesn't own any shares in itself—those shares are all disbursed among the original investors, the company's executives, the institutional investors like insurance companies and pension funds, and retail investors like you and me, who might buy a few shares through a Schwab account for our retirement. When the company's share price goes up, the company doesn't get richer, its *shareholders* do.

But here's the thing: when a company has a solid growth stock, other people *want* that stock. The company can use stock to buy stuff, like other companies. Key personnel can be hired at whopping compensation rates, with the majority of that compensation coming in stock.

Where does the company get the stock to buy these companies and hire these genius coders?

By typing zeros into a spreadsheet.

A company's own stock is an *endogenous* substance—it emanates from within the body corporate. There are some limits on how much stock a company can issue without getting into trouble, but the fact remains that *the company issues its own stock*. It doesn't have to rely on any external force to carry off this trick.

Contrast that with a mature company, one that is bidding on the same resources: that mature firm is expected to pay for things with *money*, which is *exogenous* to the company. Money comes from your country's central bank and its fiscal agents (such as chartered banks). Money does *not* come from within a company. If you doubt it, by all means, put this book down, head into the office, and use the company laser printer to run off a stack of U.S. \$100 bills (drop me a postcard from federal prison and let me know how the experiment went).

Say our taxi company wants to compete with the ridehail company. The taxi company's got some cash in the bank, of course, so it can make a job offer to some smart app coders and a project manager to oversee the production of a taxi-summoning app. If they're feeling really ambitious, they can go to their bank manager and ask for a loan so that they can buy a little app startup, or they can even approach an outside investor to buy a share of the company with funds that can be used to bid on those key personnel and acquisitions.

But the ridehail company *also* wants to hire key app programmers and buy promising app startups. Unlike the taxi company, they don't have to dip into their (nonexistent) savings to make these purchases: they can offer stock instead, which they generate by typing zeros into a spreadsheet. Not only that, but they can approach their own bank manager and take out a loan, and *stake their stock as collateral*, securing a preferential interest rate that's better than anything the taxi company can get (the billions Elon Musk uses to buy companies like Twitter and institutions like the U.S. presidency are all loans, collateralized with Tesla stock, which is why he's so vulnerable to fluctuations in Tesla's share price). They can also offer a stake in their company to a new investor—who will accept far *less* stock for a far *higher* price.

In other words, growth companies find it far easier to *grow*. They can hire key personnel and buy key firms at a price that's *far* lower than the price paid by their mature rivals. Being a growth company is awfully nice.

But it's also awfully precarious. In finance, they speak of Stein's Law, which holds that "If something cannot go on forever, it will stop." Companies that grow at a fierce clip make for attractive investments—but they also demand close attention on the part of investors, because the faster a company grows, the faster it will reach some factor that *ends* its growth.

Which brings us back to Google: the company commands a 90 percent market share in search. That means that virtually *everyone* uses Google for search. The 10 percent who don't use Google are almost certainly people who've made a conscious, firm choice to use something else. Remember: in 2024, Google was convicted of operating an illegal monopoly because of the tens of billions of dollars the company spent every year, buying up default search status on every operating system, device, and browser. It's very hard to even *find* a search box that isn't wired into Google's servers, and even if you do somehow discover a search engine you prefer to Google, you will have to change your defaults in every browser and device you use, and remember to do so every time you get a new phone or a new laptop, or just install a new browser. Ninety percent of the market and true systemic ubiquity is—or should be—the very definition of a “mature business.”

But what it really means is that Google's growth is unlikely to come from signing up new searchers. I mean, sure, they could try raising another billion human beings to adulthood while convincing them to be Google users (you may be familiar with this project, which Google calls “Google Classroom”). This might just work, but it's gonna take more than a decade, and markets are fickle and impatient.

Google *can* try to juice their search revenue without adding more search users. The company lost *another* monopolization case in 2024, over one tactic they deployed to extract more revenue per searcher: according to records in *US and Plaintiff States v. Google LLC*, Google executives deliberately made their search results *worse*, in the expectation that users would have to repeatedly search Google to get the information they were seeking. Every time you search anew, that's another chance for Google to show you more ads.

But this is a gimmick. Google might be able to juice its search revenue by making you search two or three times to get to your answer. They may be able to double the number of ads on every page. But they can't make you search *ten* times more to get the information, and they can't make you wade through *ten* times more ads. Eventually, you'll get so fucked off with Google's enshittification that you'll find some other search engine to try.

For Google to maintain a credible story about how it will continue to grow, it has to find new lands to conquer. That's why Google was so desperate to make Google+ happen: if they could convince Wall Street that

they were a credible competitor to Facebook, then the market would treat Google as though it might conceivably double or triple in size, and value its stock accordingly.

What's more, that sky-high share price will let Google buy the companies and hire the personnel to make that growth possible, poaching Facebook's best product managers and coders, buying up buzzy new social media startups (just as Facebook bought Instagram and WhatsApp) and folding them into Google's social media offerings.

Which is to say that if the market believes you can grow, the resulting P/E ratio provides the resources you need to turn that belief into reality.

But this is a double-edged, razor-sharp sword. The corollary of the idea that "a growing company is worth several times more than a similar, mature company" is "once a company stops growing, it becomes vastly overvalued, because it is now a mature company."

If you're holding a lot of stock in a growth company, you can certainly enjoy the ride up, but you need to sleep with one eye open and one fist poised over the sell button. The minute the market decides that the company's growth has petered out, there will be a mass sell-off, and not just because investors have lost confidence in the company's growth prospects, but because investors believe *other investors* have lost confidence in the company's growth prospects.

It doesn't matter if you think the company might keep growing: if no one else believes that, then the price of the shares you're holding are going to plunge, and as they do, the company will lose one of the key factors that will help it grow—namely, that high P/E ratio.

This is why extremely profitable Big Tech companies experience flash crashes in their share price whenever there's a hint of bad news. A recent, spectacular example came in January 2022, when Meta warned investors that it had experienced *less growth* among U.S. users than projected. Within twenty-four hours, investors had staged a panicked mass sell-off totaling \$230 *billion*. At the time, this was the largest one-day devaluation ever experienced by any corporation in human history.

But Meta quickly lost its all-time-loser title. In 2025, there was the release of DeepSeek, a chatbot backed by the Chinese hedge fund High-Flyer. The chatbot, which runs on low-end chips, attained performance

benchmarks heretofore associated with AI running on massive server arrays in power- and water-hungry data centers.

Within one day, Nvidia—the leading graphics-card company that makes the specialized processors that the biggest AI companies rely on for their flagship “foundation models”—had lost \$600,000,000,000 (two! thirds! of! one! trillion! dollars!) in market capitalization.

These mass sell-offs are panicky, but they’re not irrational. Though many companies’ share prices recover from these flash crashes, each one is a crapshoot. The higher a company flies, the more likely it is that if it stalls out, it will never recover from its tailspin, because the more a company is growing, the higher its P/E ratio will be and the more its stock will drop once it is perceived as “mature.”

As painful as the growth-to-maturity share price crash is for investors, it’s even harder on the people who run the company itself. Since the 1970s, stock options and stock grants have played an ever-larger role in executive compensation. These are tax-advantaged: you don’t pay tax on them until you sell them, and even then, it’s taxed as a capital gain at about half the rate of an executive’s wages. And, of course, companies love paying key staff in stock because, as noted, they make the stock right there on the premises with a keyboard and a spreadsheet, while dollars have to come from an investor, a customer, or a lender.

That means that corporate execs at fast-growing companies have a disproportionate amount of their *personal* net worth tied up in their company’s shares. From the CEO on down, the management of a growing firm is committed to growth not just because they seek to be prudent overseers of a going concern, but also because they don’t want to see their own fortunes cut in half or worse.

This is why tech companies are so fantastically, monumentally fixated on growth. It’s not the “growth for the sake of growth” that Edward Abbey called “the ideology of the cancer cell.” It’s because investors want to make sure that their pension savings and personal riches (and the savings and riches they manage on behalf of others) will grow, and that means that they bet very big on growing companies, and bail out as soon as those companies’ growth starts to slow.

This is why we have tech bubbles. It’s why Google jammed Google+ into every conceivable corner of every product it offered. It’s why Facebook

changed its name to “Meta” and spent \$60 billion on a doomed metaverse. Sure, Zuckerberg would have preferred that you actually *use* the metaverse, but that was really secondary. The main point was to be sure that investors continued to treat Facebook/Meta’s stock as a growth stock, to buy some time while Zuck hunted around for another growth story to tell (which is how I came to be writing a book about AI, and how you came to be reading it).

KPIs are a very powerful tool for selling this narrative. As conservative ideologues never tire of repeating: “incentives matter.” An executive who makes the company’s employee bonuses contingent on getting you to “interact” with a new technology can turn the entire firm’s prodigious brainpower and imagination to the task of producing an impressive and growing number that can be shown to investors in order to convince them that the company is successfully expanding into a new line of business.

The problem, of course, is the curse of Goodhart’s Law: when management converts the metric (“how many users are interacting with the hot new technology?”) to a target (“you will get paid based on how many users interact with the hot new technology”), then it stops being a useful metric.

That’s because employees cheat. Tell a product team that they can double their salary by making a given number go up, and they’ll find lots of ways to make that number rise that are far easier and more reliable than getting a bunch of users to try out a product or a feature.

Take Facebook’s infamous “pivot to video.” In 2015, Mark Zuckerberg and company decided that Facebook needed to grow into YouTube’s territory. They wanted to convince video creators to make a ton of video for Facebook, and not just video that overlapped with traditional YouTube fare: Facebook was chasing timely, news-oriented videos that could go viral based on their currency and relevance (and not just because they featured someone putting a lemon up their nose or eating an anvil).

So Facebook juiced the stats. They told the media that the popularity of video had exploded on their platform. Publishers who sued them alleged that they tweaked their algorithms so that videos stood a *much* higher chance of being “recommended” to Facebook users, and then they fudged the stats so that a video racked up a new “view” if it stayed on your screen for just a few seconds before you angrily swiped it away.

This “worked.” The media industry borrowed billions of dollars and raised billions more in investment capital, mass-fired veteran print journalists, built video production studios, and pumped out an endless scroll of video. Some of it was pretty good, too.

When we tell the story of Facebook’s pivot to video, we usually focus on these media companies, which makes sense. After years of Facebook users clicking away from the videos Facebook tried to cram into their eye sockets, eventually Facebook gave it up as a bad job and stopped pushing video. Views and ad revenue for media companies’ videos fell off a cliff. Investors pulled the plug on the new “video-first” media firms, and bankers called in their loans. A second wave of video journalists joined their former print colleagues on the breadlines. A wave of bankruptcies swept the media industry.

It’s only natural that this very visible outcome takes center stage in accounts of the pivot to video, but they were the *effect* of the pivot to video scam, not the *cause*.

Why would Mark Zuckerberg run this colossal sleight of hand? What was it about dominating video streaming that attracted all this nefarious energy?

In short: a growth story. The audience for the pivot to video wasn’t the media execs who fired their reporters and retooled around short-form video. It wasn’t the users who tuned in and watched those videos (or didn’t, as it happened).

The audience was *investors*. The story wasn’t “video is popular on Facebook.” The story was, “Facebook—which is losing young users and has saturated its market—has discovered a new field to grow into: becoming the dominant video service. This means that Facebook’s revenues can grow by at *least* the amount that YouTube brings in every year (assuming Facebook supplants YouTube), but there’s more, since Facebook is targeting a new online video niche, and no one can say just how big that might get.”

It might even have worked. If Zuckerberg had been right about the public’s appetite for newsy, current-eventsy videos, then the trick he pulled on the media industry might have made all those media companies and their investors very rich, and made Zuck richer still.

But even if it didn't work, it still *worked*. Facebook's P/E ratio stayed high, even as its core business stagnated. Facebook retained the growth company's competitive advantage and was able to snap up dozens of companies and hire more of Silicon Valley's most gifted technologists.

Magicians make a big deal out of never revealing their secrets, though of course, many magicians do, and they make YouTube videos and write books showing exactly how their tricks are done. If you watch the videos and read the books, you'll quickly realize that the reason to keep the method a secret is that most magic tricks have a really simple mechanic, and the "magic" comes from a magician's dexterity, showmanship, and patter.

Magicians have lots of ways to distract and entertain you while they trick you into thinking that you have "freely chosen" a card, and then they lead you around by the nose, convincing you that if the card isn't in their hat, up their sleeves, or in their mouth, it must have disappeared. Learn the method, though, and you'll discover that the magician forced you to choose that card, and the whole point of that long list of places the card *isn't* is to distract you from the place where the card *is*.

The pivot to video was a cheap trick. Zuckerberg chose the victory condition (that is, *more video views prove we are going to grow by hundreds of billions*), defined the metric that proved the victory condition has been met (a "video view" is a two-second, fleeting window that scrolls down the user's screen), designed the system that determined the ability to attain the metric (*the Facebook algorithm will push videos in preference to all other kinds of posts*), and then pulled the world's least convincing rabbit out of his hat.

Now, as George W. Bush said, "Fool me once, shame on—shame on you. Fool me—you can't get fooled again." The simple gimmicks used with Google+ and Facebook's pivot to video eventually wore thin as market analysts wised up to them, so tech upped its game.

Have you ever noticed that when you load a video using any of the major streaming apps (Disney+, Netflix, HBO Max, etc.), you suddenly have to handle your phone like a photonegative, touching only the bezel at the edge of the device? Just brushing against *any* part of the screen instantly switches you to an unrelated video, while getting back to the video you were just watching takes an agonizing eternity.

This isn't an accident: it's just a KPI, rearing its ugly head.

The major expense for video streaming services is, well, *video*. Streaming services have to either create or license video content, and either way, this costs money. These companies make their money through recurring subscription fees, and viewers don't want to pay every month to watch the same videos over and over again.

To retain users, the streamers have to convince them that the video on the service is worth another month's subscription fees. The easiest way to do this is to commission or license a steady stream of blockbuster programs, the way HBO did in the 2000s, when everyone was talking about the "golden age of prestige television."

This is "easy" in the sense that Hollywood is full of extremely talented people: writers endlessly shopping amazing ideas for great TV shows, as well as actors, directors, costumers, editors, postproduction staff, and everyone else needed to make these TV shows into reality.

But it's *hard* in the sense that all these people expect to get paid, and the very best of them command rather handsome salaries.

For a streaming service hoping to limit costs and retain customers, making or buying an endless stream of new, high-quality TV shows is a losing gambit. Instead, these platforms try to convince their subscribers to watch the shows the service already owns—the vast troves of material that you haven't watched yet because it's old, or obscure, or not to your taste, or because you just missed it the first time around.

To convince you to watch this back catalog, the services build recommendation systems, and—crucially—tout these recommendation systems to investors as the way that they will be able to retain subscribers without breaking the bank on new programming. If investors buy this story—if they treat the streaming services as growth companies—they will keep their precious P/E ratios, and with them, the ability to acquire key companies and personnel, and secure loans at preferential rates. In other words, they'll be able to continue growing.

Obviously, the most salient number in this tale is "How many users stopped paying for a subscription this month?" But the streamers insist that there's another number that's nearly as important: "How often did we successfully recommend a back-catalog program to a subscriber?"

After all, even if subscribers are subscribing *this* month, maybe we'll retain *next* month's subscribers if we can convince them that we have a bottomless well of content the existence of which they never suspected, and which they *love*.

So the word goes down (in the form of a KPI): *Your bonus depends on getting a user to follow a recommendation, and don't think you can fool us by just cycling users to some crappy show—the user has to spend at least ten seconds watching your recommended show for it to count.*

And *that* is why merely grazing any part of your screen while watching a show on a streaming app causes it to switch to some other show, *and* it's why it takes ten seconds to get rid of this new show and find your way back to the show you *wanted* to watch.

You are only the secondary audience for streaming video recommendations. The primary audience is investors.

Which brings me, at long last, to AI.

The tech platforms are desperate to convince Wall Street that you love AI, which is very different from convincing *you* that you love AI. Obviously, it would be nice if you loved AI (just like it would have been nice if you'd watched all those Facebook videos), but for the individuals who stand to make titanic amounts of money in salaries and bonuses from keeping the growth stock story going, that's just a sideshow.

Even if you eventually reject AI so comprehensively that the investment bubble collapses, these individuals who are setting the agenda for their companies will either be long gone (with so much money in hand that they can found a dynastic fortune that stretches down through the ages), or will have found another bubble to inflate.

Avoiding AI is even harder than avoiding video recommendations on a streaming service. Every button you used to click to do something useful has been moved to a different part of your screen and replaced with a button that looks nearly identical, and which summons an AI genie that refuses to be dispelled until you've wrestled with the user interface for endless, agonizing seconds.

I use a stock Android phone, a Google Pixel, their flagship device. In the time since I started writing this book, I've accidentally conjured an AI demon by:

- taking a picture;
- switching apps;
- sending a text;
- replying to a text;
- dictating an email; and
- searching the settings for a way to turn off AI.


Every time I make an AI appear on my screen, a team at Google gets a little KPI score boost, and a process at Google gleefully notes the fact that I have “engaged” with AI. I wouldn’t be surprised if these logs were granular enough to note that I have “engaged” with AI six times in a week. At some point near the end of this quarter, a team at Google will comb through these statistics, pick the most impressive-seeming, and make a beautiful series of charts illustrating them: “Users who interact with Google AI returned to it an average of six times in the first week, and their engagements doubled every week.”

Google needs to do this, because the alternative to growth isn’t stasis, it’s collapse. Without growth, Google will see its share price decline to that of a “mature” company, and then it will lose key employees, miss out on key acquisitions, and its cost of capital will go through the roof. Combine all that with the fact that a failure to grow will devastate the personal finances of every decision-maker at Google, and it’s easy to see why they would run things this way.

This is how AI has colonized our economy. As Ed Zitron writes in his “The Hater’s Guide To The AI Bubble” from July 2025, seven giant AI firms—Nvidia, Microsoft, Alphabet (Google), Apple, Meta, Tesla, and Amazon—account for 35 percent of the value of *all* U.S. stocks. Nvidia accounts for 19 percent of the value of these seven stocks. Nvidia’s entire valuation is based on the fact that the other six companies are spending *hundreds of billions of dollars* on Nvidia’s GPU chips, which power their AI data center.

Keeping the growth story alive isn’t about one company, or one sector. The entire U.S. economy hangs in the balance.

BUBBLE



OceanofPDF.com