

ALMOST RIGHT

What Happens When Machines Create Faster Than Humans Can Check?

Anthony C. Vila

“The single biggest frustration developers report with AI coding tools is output that is almost right, but not quite.”

— Stack Overflow Developer Survey, 2025 (66% of respondents)

\newpage

Contents

Part I — The Guessing Machine

1. The Bet at Dartmouth
2. Two Winters
3. The Real Reason
4. Almost Right, By Design

Part II – The Plunge

- 5. Faster Than the Internet
- 6. Nobody's Guarding the Door
- 7. The Sellers
- 8. The Two Countries

Part III – Everybody Builds Now

- 9. What Actually Works
- 10. Vibe Coding
- 11. Nobody Hacked Them

Part IV – Nobody's Checking

- 12. The Middlemen
- 13. Cognitive Debt
- 14. The Doctors Got Worse
- 15. The Canaries

Part V – How We Check

- 16. What the Pilots Did
- 17. Become the Verifier
- 18. Start Now

\newpage

A Note on Sources

Every factual claim in this book is traced, where possible, to the strongest source I could get my hands on: the study itself, the court filing, the company's own statement, the survey with its sample size attached, the government document, or the incident report.

That isn't an academic decoration. It's the argument of the book applied to the book.

If I'm going to spend the next eighteen chapters telling you that plausible output needs to be checked, then I don't get to hide the caveats that make my own argument less dramatic. Where the evidence is thin, I say so. Where a study has been criticized, I give you the criticism. Where I'm offering an interpretation instead of a fact, I label it.

And if something in here turns out to be wrong, I want to know.

That's the whole point.

\newpage

Why I Started Checking

I didn't come to this subject as an engineer, an academic, or somebody looking for a reason to distrust artificial intelligence.

I came to it as a simple user.

I spent most of my adult life in sales. I didn't have an engineering degree. I wasn't some master computer whiz. I sat down one day and started playing with a thing that I thought was cool.

I was editing photos, screening large amounts of information, and doing all the awesome things that basic, everyday people were discovering they could do with AI. And I wasn't sitting behind some elaborate computer setup.

I was using my cell phone.

Just a simple app on a simple phone.

I learned very quickly that AI had some really neat features. Then I learned that it could do things I never expected. I could produce software and business systems that previously would have required several people with skills I simply did not possess.

The capability was amazing.

But so were the problems.

I learned very quickly that producing something and knowing whether it was the right thing—or whether it actually worked—are completely different jobs.

This amazing machine could generate faster than I could inspect. It could give me an answer that looked finished while leaving behind a small failure that I didn't even know to look for. The more capable the output became, the easier it was to mistake plausibility for verification.

Suddenly, I found myself spending countless hours checking work that the machine had taken only

seconds to produce.

That's when I realized there was a major problem.

And that became the question behind this book:

What happens when production becomes nearly instantaneous, but judgment does not follow?

I went searching for the answer in a lot of different places. I researched software, law, medicine, education, labor economics, security, and aviation. Eventually, I went all the way back and researched the history of artificial intelligence itself.

What I found surprised me.

The problems I was dealing with weren't isolated problems. Versions of them were showing up everywhere. Researchers, professionals, companies, and institutions were reporting them in different forms and different industries.

What I couldn't understand was why more people weren't demanding a solution.

Eventually, I started to understand why.

That is what this book is about: how things got this far, why an incredibly useful machine can also produce incredibly convincing mistakes, and what happens when our ability to create begins moving faster than our ability to check what we created.

This is not an argument to stop using AI.

That would be foolish.

I use AI. I'd use it again and again. I'm using it right now.

The argument I'm making is that we have to preserve the thing this technology still depends on: the people and systems capable of recognizing when an answer that looks absolutely right is not right.

Because when the consequences actually matter, there is one principle we cannot afford to forget:

Almost right is still wrong.

\newpage

A Note on What This Book Is — and Isn't

Doing the research, I quickly realized that there were two easy books I could have written.

One book says artificial intelligence is going to save the world.

The other says artificial intelligence is going to destroy it.

Either book would have been simpler to write. But neither is the book the evidence gave me.

The systems described in these pages are genuinely useful. They make some workers faster. They help beginners and average Joes just like me perform tasks that previously required more experience. They can expand access to knowledge, lower the cost of creating software, and give an ordinary person leverage that would have looked absurd just a few years ago.

But the evidence also points out something very clearly:

The machine makes mistakes.

That fact alone isn't particularly interesting. People make mistakes. It happens. That's what we do as humans.

The interesting part is the combination: these systems can produce professional-looking work at enormous speeds while remaining unreliable in ways that can be extremely difficult to detect from the appearance of their output alone.

That changes everything.

It changes the economics of checking. It changes our ability to distinguish fact from fiction. It changes responsibility. And it changes what expertise is for.

As humans, we sometimes fall back on the excuse, I'm only human. And you know what? That's okay. We

understand that people are imperfect.

But when we put our faith in a machine that is marketed as being above human capability—superhuman, revolutionary, absolutely amazing—then the product needs to match the confidence we place in it.

This book is an attempt to follow the consequences of that gap without pretending the evidence is cleaner than it actually is.

Some of the evidence is experimental. Some is observational. It comes from labor-market administrative data, court records, incident reports, professional surveys, security benchmarks, institutional guidance, and other forms of documented evidence.

Those forms of evidence do not deserve identical confidence.

When a randomized trial establishes a result in a narrow setting, I try to keep the claim narrow. When an observational study shows an association, I don't want to call it causation. When a vendor benchmark reports a failure rate, I treat it as a benchmark result rather than a universal law.

That standard matters because this book is about verification.

It would be ridiculous to argue that plausible output should be checked and then hide the caveats that make my own argument less dramatic.

So read the numbers as evidence, not decoration.

Read the stories as examples, not proof that every organization behaves the same way.

Read the predictions for what they are: predictions.

And where the evidence changes, the conclusions should be allowed to change with it.

That last part matters especially in AI.

Technology is moving quickly enough that a benchmark can become stale while a book is still being edited. METR's early-2025 developer study found a slowdown in one population and setting. Its later work suggested newer tools may be faster, while also warning that selection effects made that later estimate weak.

The correct response isn't to choose whichever result best fits the thesis.

It's to preserve the timeline.

Because the question underneath the changing benchmark is much more durable:

When machines become capable of producing more work, what happens to the systems that determine whether that work deserves to be trusted?

That question survives whether the next coding model is twenty percent faster or two hundred percent faster. It doesn't matter.

The faster the generator becomes, the more urgent the verification question should become with it.

I'm not arguing that every task needs an expert committee. Most of what people do with AI today is relatively low stakes. If a restaurant recommendation is bad, that's fine. Eat somewhere else. If the first draft is clumsy, rewrite it. If a brainstorming list contains nonsense, delete the nonsense.

Not much consequence.

Verification should be proportional to consequence.

The problem begins when that same casual relationship with generated output migrates into software, law, medicine, finance, education, public policy, security, or any system where an error can travel farther than the person who clicked the Generate button.

That is where almost right stops being an annoyance.

It becomes an operating condition.

And when that condition spreads across enough systems, it starts to look like an epidemic.

The rest of this book is about what to do with that condition.

The distinction I want you to carry forward is simple:

Output is not verification. Fluency is not evidence. Assistance is not competence. Oversight is not validated merely because a human name appears at the end of the process.

These things can overlap. They are absolutely not interchangeable.

That sounds obvious when it's written plainly. But in practice, modern AI products are extraordinarily good at making those distinctions disappear. They collapse research, drafting, explanation, calculation, recommendation, fact-finding, and presentation into one smooth interaction.

The convenience is the product.

The danger is that the user can lose track of which step actually established the truth of the answer.

That's where the problem begins.

Throughout the chapters that follow, watch for the handoff. Watch the moment when generated material becomes relied-upon material.

That is where the economics, the liability, the training problem, and the verification problem all meet.

Don't stop the machine.

Don't worship the machine.

Build the safeguards. Build the checks. And keep checking the checks.

We should not freeze today's rules around tomorrow's technology. We need to preserve a method: measure what the system actually does, identify what matters when it fails, maintain independent ways to detect those failures, and change the controls when the evidence changes.

And when these systems scale, the safeguards have to scale with them.

When machines produce more, we have to become better at checking more. When their outputs travel farther, the systems responsible for verification have to become stronger. The speed of generation cannot be allowed to become an excuse for weakening the standard of proof.

A verification culture is not suspicious of progress.

It is how real progress becomes dependable enough to trust.

\newpage

PART I — THE GUESSING MACHINE

Chapter 1: The Bet at Dartmouth

If you're like me, I once thought AI was a new technology.

I learned pretty quickly that it started a long time ago.

Today is August 31, 2026. And believe it or not, today marks exactly 71 years since one of the defining events in the birth of artificial intelligence.

On August 31, 1955, four men put their names on a proposal for a summer research project and sought support from the Rockefeller Foundation. The surviving typescript runs seventeen pages plus a title page—not the tidy twopage origin story it is sometimes reduced to.

They were not cranks. John McCarthy was a young mathematician at Dartmouth. Marvin Minsky was at Harvard. Nathaniel Rochester had helped design IBM's first commercial scientific computer. Claude Shannon, at Bell Telephone Laboratories, had already invented the mathematics that every phone call, every hard drive, and every internet packet still runs on. If you wanted four people in 1955 who understood what

a machine could and could not do, you would have had a hard time doing better.

Here is what they wrote:

“We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”

Read that last sentence again. A summer.

They asked the Rockefeller Foundation to cover it. The budget included salaries of \$1,200 for each faculty-level participant who wasn't already being paid by somebody else. McCarthy and Shannon had already gone to New York that June to sit down with a man named Robert Morison at the foundation and make the case in person.

The document is the first time the phrase “artificial intelligence” appears in the historical record.

McCarthy picked the name. He needed something that would sound like a field, not a hobby, and he needed it to not sound like anyone else’s field. It worked. Seventy-one years later, that name is on the front page of every newspaper on earth, attached to companies worth more than the economies of most countries.

But I want you to sit with the bet itself, because the bet is the whole story.

Four of the smartest people alive looked at the problem of human intelligence — language, abstraction, concepts, the ability to improve yourself — and estimated that ten people could make “a significant advance” on it in ten weeks. Not solve it. They were careful about that. But make real progress. Over a summer. In New Hampshire.

They were off by roughly seven decades. And I would argue they are still off, in a way that matters more now than it did then, because in 1956 the only thing riding on the bet was a Rockefeller grant. Today it’s your job, your kid’s homework, your doctor’s judgment, and the software that holds your bank balance.

What happened that summer

Not much, and that’s not an insult.

The workshop happened. Eleven people were originally planned to attend; more than ten others drifted through for shorter visits over the course of the summer. Some of the names on the guest list would go on to define the field for the next fifty years. They argued. They wrote on chalkboards. They disagreed about what “thinking” even meant and about whether the problem was mostly logic or mostly learning.

There was no final report.

I want to be fair to them, because this book is going to be hard on a lot of people who deserve it, and these four don't. They were doing what scientists are supposed to do: take a wild idea seriously enough to test it. The Dartmouth proposal is one of the most consequential documents of the twentieth century precisely because it was wrong in an interesting way. It set the agenda. Every argument you will hear about AI in 2026 — can it think, does it understand, will it replace us, is it dangerous — was on a chalkboard in Hanover in the summer of 1956.

But I also want you to notice something about the shape of the bet, because you're going to see this shape again and again in the chapters ahead, and it's going to cost real people real money and real careers.

The bet was: intelligence is describable, therefore intelligence is buildable, therefore we are close.

The first part is a philosophical position. The second is an engineering claim. The third is a sales pitch. And the trick — the thing that has been happening for seventy years — is that people who believe the first part let it carry them straight through to the third without stopping to check whether the second is true.

Before Dartmouth: the man who asked the question

The conjecture didn't come from nowhere. Six years earlier, in October 1950, a British mathematician named Alan Turing published a paper in the philosophy journal *Mind*. It's called "Computing Machinery and Intelligence," and it opens with a question that Turing himself immediately says is too muddy to answer: Can machines think?

Turing's move was to replace the question with a game. Put a person in one room and a machine in another. Let a judge in a third room type questions to both and read their typed answers. If the judge can't reliably tell which one is the machine, then — Turing argued — arguing about whether the machine "really" thinks is a waste of everyone's time. It's doing the thing. What else do you want?

This is the imitation game, and it has been misread for seventy-five years, so let me say plainly what it is and isn't.

It is not a definition of intelligence. Turing knew that. It is a test of indistinguishability — of whether a machine's output can pass for a human's. Turing proposed it because he thought the philosophical argument was unwinnable and the practical question was the only one worth having.

Hold onto that, because it's the seed of everything. From the very first serious paper in the field, the goal was not "build a machine that understands." The goal was "build a machine whose output you can't tell apart from someone who understands."

In 1950 that seemed like the same thing. In 2026 it is the single most important distinction in your life, and almost nobody talks about it.

1958: The Navy's machine that would be conscious

Two years after Dartmouth, the bet got its first press tour.

On July 7, 1958, a psychologist named Frank Rosenblatt gave a demonstration in Washington. Rosenblatt worked at the Cornell Aeronautical Laboratory, and the Office of Naval Research was paying for his work. He had built something he called a perceptron — a machine that could learn to tell the difference between simple patterns by adjusting its own internal weights when it got an answer wrong. It

was, in the plainest sense, a machine that got better with practice. That was new.

The next morning, The New York Times ran the story under the headline “NEW NAVY DEVICE LEARNS BY DOING.” The subhead promised a computer “Designed to Read and Grow Wiser.”

Here is the first sentence, verbatim:

“The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.”

The Navy. Expects. Conscious of its existence.

The article went on to describe the first full perceptron as a machine with about a thousand “association cells,” fed by an eye-like device of 400 photocells, estimated to cost around \$100,000 to build. The New Yorker weighed in too, calling it “the first serious rival to the human brain ever devised.”

Now — what had Rosenblatt actually built?

A machine that could learn to sort simple visual patterns into two piles. Left from right. Square from triangle. It was a genuine scientific achievement, and Rosenblatt’s paper describing the mathematics, published that same year in *Psychological Review*, is a real piece of work. The idea inside it — that you can

build a network of simple units, show it examples, and let it adjust itself until it gets the answers right — is the direct ancestor of every AI system you have used this week.

But the distance between “sorts simple patterns into two piles” and “conscious of its existence” is not a gap in engineering. It is a gap in honesty. And that gap was not created by Rosenblatt’s machine. It was created by the people describing Rosenblatt’s machine to the public — a funding agency, a newspaper, a magazine — each of whom had a reason to make it sound bigger than it was.

I sold things door to door for a living before I ever touched any of this. I know what a pitch sounds like. That first sentence in the Times is a pitch. It’s a very good one. And it set the template that the AI industry has followed, with remarkable discipline, for sixty-eight years:

Build something real. Describe something imaginary. Let the reader close the gap themselves.

You might be wondering why a book about AI opens with a grant proposal and a newspaper clipping from the Eisenhower administration.

That’s a simple question to answer.

Every time you read a headline about AI in 2026—a CEO saying it will eliminate half of all entry-level jobs,

a researcher saying it will make us all smarter, a lab saying its new model is “approaching” something-or-other—you are reading a descendant of that Times article.

The genre was invented in 1958. The structure has never changed. Something real gets built. Something enormous gets promised. The gap between the two is where the money is, and the gap is your problem, not theirs.

And there’s a second reason, which is the one this whole book is about.

The Dartmouth proposal set out to make machines that could “use language, form abstractions and concepts, solve kinds of problems now reserved for humans.” Turing’s test only asked that the machine be indistinguishable from someone who could.

Rosenblatt’s perceptron did neither — it learned to give the right output on simple patterns, without anything inside it that you or I would call a concept.

Guess which of those three the industry actually built.

Not the Dartmouth version. Not a machine that forms concepts. The Turing-Rosenblatt version: a machine that produces output you can’t tell apart from a person’s, by adjusting itself until its answers look right.

That is not a criticism. It is a description. And it has a consequence that the next three chapters will spell out, but that I'll give you now so you can carry it with you:

A machine built to produce answers that look right will, by design, produce answers that look right when they are wrong.

That is not a bug somebody forgot to fix. It is the finish line the field was running toward since 1950, and in 2022 it crossed it.

In September 2025, OpenAI — the company that put this technology in front of the world — published a research paper on why its own systems make things up. The paper's explanation, in its own words, was that these models "hallucinate because the training and evaluation procedures reward guessing over acknowledging uncertainty." The paper opens with a comparison I'll ask you to remember: like students facing hard exam questions, the models guess when they don't know, "producing plausible yet incorrect statements instead of admitting uncertainty."

Plausible yet incorrect. Almost right.

The men at Dartmouth thought they were describing intelligence. What they were actually describing — what the whole field would spend seventy years perfecting — was a machine for producing plausibility.

It turns out plausibility is enormously valuable. It turns out you can sell it for hundreds of billions of dollars. And it turns out that a society which stops being able to tell plausibility from truth is in a very specific kind of trouble that nobody in Hanover in 1956 was thinking about, because in 1956 there were still going to be humans checking the work.

This book is about what happens when the humans are no longer checking the work.

That's the equivalent of letting the fox guard the henhouse.

The perceptron got its press tour in 1958. Eleven years later, Marvin Minsky — one of the four names on the Dartmouth proposal — co-wrote a book that proved, mathematically, what a machine like Rosenblatt's couldn't do. The funding dried up. The field went into what its own people still call a winter.

It came back. It went into a second winter. It came back again — and the thing that finally brought it back was not a better idea about intelligence. It was more data and more chips than anyone in 1956 could have imagined. Which is a fact the industry would prefer you not dwell on, for reasons that will become obvious.

Let's dive into those reasons.

Sources for this chapter: McCarthy, Minsky, Rochester & Shannon, “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” dated August 31, 1955 (archived at Stanford; reprinted AI Magazine 27(4), 2006). Turing, “Computing Machinery and Intelligence,” *Mind* 59(236), October 1950. The New York Times, “New Navy Device Learns by Doing,” July 8, 1958. Rosenblatt, “The perceptron: a probabilistic model for information storage and organization in the brain,” *Psychological Review* 65(6), 1958. Kalai, Nachum, Vempala & Zhang, “Why Language Models Hallucinate,” arXiv:2509.04664, September 4, 2025.

Chapter 2: Two Winters

In 1969, one of the four men who signed the Dartmouth proposal killed the machine that had gotten the field its first headlines.

Marvin Minsky, with his MIT colleague Seymour Papert, published a book called *Perceptrons*. It is a mathematics book, dense and careful, and its most important result is a proof of what Frank Rosenblatt’s machine could not do. A single layer of those self-adjusting units — the thing the Navy said would become conscious — could not learn certain simple patterns no matter how long you trained it. Not “hadn’t yet.” Couldn’t.

Minsky and Papert were right. The math holds. And the effect on the field was roughly what happens to a sales office when the top producer walks in and announces that the product isn't working.

All the money walks out the door.

This is the first thing to understand about the history of artificial intelligence, and it's the thing the industry wants you to look past:

The field has collapsed not just once, but twice.

Collapsed. Not slowed. Completely collapsed.

Funding gets cut. Labs get closed. Everything comes to a grinding halt. The phrase artificial intelligence itself became something researchers avoided putting on grant applications because it marked you as somebody associated with a field that had overpromised.

The people who lived through these periods called them AI winters.

That term is still in use, actually, because the people who use it are still waiting to see whether there will be a third.

The first winter

The perceptron's collapse in the United States was matched, almost on schedule, in Britain. In 1973, the

UK government asked a mathematician named James Lighthill to evaluate the state of AI research and report on whether it deserved continued public funding. Lighthill's report, published by the Science Research Council in 1973, put its verdict in one sentence: "In no part of the field have the discoveries made so far produced the major impact that was then promised." Historians have summarized his charge as AI failing to meet its "grandiose objectives," and the effect was the same either way: British funding for AI research was gutted for the better part of a decade.

In the U.S., the Defense Department's research arm — the main source of money since the beginning — pulled back sharply in the mid-1970s. The pattern was the same everywhere: a decade of promises measured against a decade of demos, and the promises lost.

So what went wrong?

Some people would say nothing. Some people would say everything.

The researchers of the 1950s and '60s had done real science, and they learned real things. But what went terribly wrong was the bet—the same bet from Chapter 1:

Intelligence is describable, therefore intelligence is buildable, therefore we are close.

They'd been running on that third clause for fifteen years and never delivered the second.

They had not figured out how to build intelligence.

The second winter

The field came back in the 1980s with a new idea and a new pitch. The idea was the "expert system": instead of trying to build general intelligence, you would sit down with a human expert — a doctor, a chemist, a loan officer — write down their decision rules as a long list of if-then statements, and put that list in a computer. The computer would then make expert decisions without the expert.

It worked, sort of, in narrow places. Companies bought it. Japan launched a national program, the Fifth Generation Computer Systems project, in 1982, with the stated aim of leaping past the United States in intelligent computing within a decade.

By the early 1990s it was over again. The expert systems turned out to be brittle — they broke the moment a situation fell outside their rules — and expensive to maintain, because the rules had to be rewritten by hand every time the world changed. Japan's Fifth Generation project wound down in 1992 without the leap. The companies that had sold expert systems either folded or quietly renamed what they did. Second winter.

I want to pause on the expert systems, because they're the closest ancestor to the thing you're using today, and the way they failed is instructive.

An expert system was, literally, a human expert's judgment written down and run by a machine.

It didn't have its own judgment or a mind of its own. It had a recording of someone else's behaviors, someone else's decisions, someone else's examples.

When the recording matched the situation, it could be as good as the expert. But when it didn't, it could be worse than a first-year trainee. Because at the very least, even a trainee knows when they're out of their own depth.

The expert system had no idea.

It just applied the rules. It gave you an answer with the same exact confidence whether it was completely right or catastrophically wrong.

Remember that, because this isn't a new phenomenon. This problem has been there from the beginning.

And we're going to take a look at a much more powerful version of that exact failure in Chapter 4.

Now, if that doesn't make you uneasy, here's the part that should.

The idea that was sitting there the whole time

While the expert-system money was flowing, a small number of researchers kept working on Rosenblatt's discredited idea: networks of simple units that adjust themselves. Minsky and Papert had proven a single layer couldn't learn much. But what about many layers, stacked? The problem was that nobody had a good method for training the deeper layers — for figuring out which of thousands of internal connections to adjust when the final answer came out wrong.

In 1986, three researchers — David Rumelhart, Geoffrey Hinton, and Ronald Williams — published a paper in *Nature* describing a method that did exactly that. It's called backpropagation. In plain terms: when the network gets an answer wrong, you measure how wrong, and you push that error backward through every layer, nudging each connection a little in the direction that would have made the answer less wrong. Do that millions of times and the network learns.

That paper is the technical foundation of every AI system you've used this week. It was published forty years ago.

Three years later, in 1989, a researcher named Yann LeCun used a version of the technique to get a network to read handwritten digits — the kind on the

front of a check. In 1997, two German researchers, Sepp Hochreiter and Jürgen Schmidhuber, published a design called the Long Short-Term Memory network that let these systems handle sequences — text, speech, anything where order matters.

So by 1997, the core ideas were in print. The methods worked. The people who would later win the field's highest prizes for them were already publishing.

And almost nobody cared, because the networks were too small and too slow to do anything a customer would pay for. The researchers who stuck with it through the 1990s and 2000s did so on thin funding and thinner respect. Hinton has said, in interview after interview, that for years he could barely get his students' papers accepted at the field's own conferences.

So the idea wasn't the bottleneck. It had to have been something else.

Let's look at what actually changed.

In 2012, a graduate student of Hinton's named Alex Krizhevsky entered a competition.

The competition was called the ImageNet Challenge. Researchers were given a dataset of millions of photographs, each labeled with what it showed — a dog, a truck, a mushroom — and asked to build software that could label new photographs it had

never seen. Every year the best teams in the world competed. Every year the error rates crept down by a point or two.

Krizhevsky, with Ilya Sutskever and Hinton, entered a deep neural network — many layers, trained with the 1986 method — that had been trained on graphics cards built for playing video games. Their system's top-five error rate was 15.3 percent. The next-best entry in the competition came in at 26.2 percent.

That is not a creep. That is the floor falling out. In one year, one team cut the error rate nearly in half using an idea that had been sitting in the literature since the Reagan administration.

Within two years, essentially every serious team in the competition had switched to deep neural networks. Within five, the technique had spread to speech, to translation, to medicine. The second winter ended not with a new idea about intelligence but with a graduate student, a pile of gaming hardware, and a dataset big enough to matter.

And that brings me to the question this chapter exists to ask:

If these ideas were already there in 1986 and 1997, why didn't anything happen until 2012?

The answer is simple. The ideas were never the constraint.

The two things standing in the way were the amount of data you could feed the network and the amount of computing power you could throw at training it.

In 1997, neither existed at the necessary scale. But by 2012, both did. The internet had produced an ocean of labeled photographs and text, and almost by happenstance, the video-game industry had built the chips necessary to process it.

Let me backtrack here, because I don't want to minimize what happened during those years.

Plenty happened. There were real inventions, better ways to train deep networks, better architectures, and an enormous amount of hard engineering. The researchers didn't just put their books away and say, Forget it. Let's wait until the technology exists.

They kept working.

It's just that none of those advances amounted to a new theory of intelligence. They were increasingly better answers to a different question:

How do we make this thing bigger without it completely falling over?

The field got much better at scaling. The world got bigger. And eventually, the old methods finally had enough data to eat—and enough computing power to actually eat it.

2017: The Paper That Built the Thing You Use on Your

Phone

Five years after ImageNet, eight researchers at Google published a paper with the least modest title in the history of the field: Attention Is All You Need.

The paper introduced a network design called the Transformer.

I'm not going to walk you through the architecture because it's not necessary to understand. What you do need to understand is what it was for.

The Transformer was extraordinarily good at one task: given a sequence of words, predict what comes next. And it was designed so that you could make it bigger—more layers, more connections, more training data—and it would keep getting better at that task without immediately hitting the scaling walls earlier approaches had faced.

The T in ChatGPT stands for Transformer. So does the T in GPT-4, GPT-5, and every other GPT model named like them.

Anthropic's Claude. Google's Gemini. Meta's Llama.

All of them use Transformer-based architectures.

And all of them descend from the breakthrough introduced in that 2017 paper. At their foundation, they perform the same fundamental operation at an enormous scale: predict what token comes next.

Most of the eight authors have since left Google. Several went on to found companies, some of which became worth billions of dollars.

A paper about predicting what comes next turned out to be one of the most valuable documents Silicon Valley has ever produced.

The prize and the resignation

In 2018, Geoffrey Hinton shared the Turing Award—computing’s equivalent of the Nobel Prize—with Yann LeCun and Yoshua Bengio, for work the field had ignored for decades.

In October 2024, Hinton received the actual Nobel Prize in Physics, shared with John Hopfield, for foundational work on neural networks.

But between those two honors, something remarkable happened.

On May 1, 2023, The New York Times reported that Hinton had resigned from Google. He had worked there for a decade. He left so that he could speak openly about the risks of the technology he had spent his life building.

Understand what happened here.

The man who helped keep this idea alive through the AI winters had decided that the public needed to hear his doubts about the spring.

He told the Times that part of him now regretted his life's work.

"I console myself with the normal excuse," he said. "If I hadn't done it, somebody else would have."

And:

"It is hard to see how you can prevent the bad actors from using it for bad things."

On the speed of what he had helped build:

"I thought it was 30 to 50 years or even longer away. Obviously, I no longer think that."

What does it say about a technology when one of the people most responsible for building its foundations walks away from one of the most powerful companies developing it so that he can speak freely about its risks?

Note the shape of it.

One of the most important living contributors to this technology decided, six months after ChatGPT launched, that one of the most valuable things he

could do with his remaining reputation was warn people about what could go wrong.

And notice the excuse he reached for:

Somebody else would have.

If it wasn't me, somebody else would have done it.

Think about how many times human beings have used some version of that justification when they knew something carried consequences but continued anyway.

Because you're going to hear that same excuse, in one form or another, from almost everyone in the next chapter.

What the winters teach

So what do the winters teach us?

Here's what I take away from all of this.

First: the industry has been completely wrong about timelines before. Not just once, but twice. And not by a little. Catastrophically.

The people running the industry today were not around for either of those collapses. The current generation of AI executives built their careers almost entirely inside the spring that began in 2012.

They were there for the abundance.

They never saw the money leave.

Now, that doesn't make them wrong today. And it doesn't mean their predictions are destined to fail.

But it does mean that most of them never personally experienced the cycle that shaped the field they inherited: enormous promises, enormous investment, disappointing results, disappearing money, and researchers continuing to work after almost everyone else had stopped paying attention.

They inherited the spring.

They didn't live through the winters that made it possible.

That distinction matters when we decide how much confidence to place in predictions about what happens next.

Second: what finally ended the winters was, to a large degree, scale.

There was real engineering progress, and we shouldn't overlook it. Better architectures were developed. Training methods improved. Researchers solved difficult technical problems.

But nobody in 2012 suddenly arrived with a complete new theory of intelligence that researchers in 1969 had simply failed to imagine.

What they had was vastly more data, exponentially more computing power, and much better methods for putting both of them to work.

And that is still largely the shape of the strategy today:

Make it bigger.

More compute. More data. Bigger models. Bigger infrastructure.

That is why companies are spending hundreds of billions of dollars building the infrastructure required to develop and operate these systems.

And that's why the next chapter is going to talk about the money.

Third—and this is the one I want you to carry forward—the machine that won was built to produce the right output, not necessarily to understand why that output was right.

Somewhere along the way, we began accepting something different from what the original dream seemed to promise.

We went looking for intelligence.

We became very good at producing answers.

The Transformer does not form concepts in the way the Dartmouth proposal imagined them. At its

foundation, it predicts what comes next. It does that extraordinarily well—so well that its output can appear indistinguishable from the output of someone who understands.

That is an extraordinary achievement.

It is also the problem this book is about.

Because producing an answer that looks right and knowing that the answer is right are two entirely different things.

That distance between the two is where almost right lives.

And as these machines become faster, larger, cheaper, and more deeply embedded in the systems around us, that distance matters more, not less.

Now let's take a look at who paid for all of this—and exactly what they wanted for their money.

Sources for this chapter: Minsky & Papert, *Perceptrons* (MIT Press, 1969). Lighthill, "Artificial Intelligence: A General Survey," UK Science Research Council, 1973. Rumelhart, Hinton & Williams, "Learning representations by back-propagating errors," *Nature* 323, 533–536 (1986). LeCun et al., "Backpropagation Applied to Handwritten Zip Code Recognition," *Neural*

Computation, 1989. Hochreiter & Schmidhuber, “Long Short-Term Memory,” Neural Computation 9(8), 1997. Krizhevsky, Sutskever & Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” NeurIPS 2012. Vaswani et al., “Attention Is All You Need,” NeurIPS 2017. ACM A.M. Turing Award, 2018. Nobel Prize in Physics, 2024. Metz, “ ‘The Godfather of A.I.’ Leaves Google and Warns of Danger Ahead,” The New York Times, May 1, 2023.

Chapter 3: The Real Reason

On November 30, 2022, OpenAI put a chat window on the internet and called it a “research preview.”

That sounds pretty uneventful.

It wasn’t.

The thing was called ChatGPT.

OpenAI already had the underlying model. What changed was that now ordinary people could use it. You didn’t need to be a programmer. You didn’t need to understand neural networks. You didn’t need to know what a Transformer was.

You typed something into a box.

It answered you.

Within five days, ChatGPT had a million users. Within about two months, it had a hundred million. At that point it was one of the fastest-adopted consumer products anybody had ever seen.

To most of us, it looked like AI had appeared out of nowhere.

But now you know better.

The method went back decades. The Transformer came from 2017. The internet supplied the data. Nvidia and the rest of the computing industry supplied the horsepower.

What OpenAI built in November 2022 was the front door.

And once millions of ordinary people walked through it, something changed.

Not just technologically.

Financially.

Because remember where we left off in the last chapter.

Somebody had to pay for all of this.

And nobody spends this kind of money without expecting something in return.

So that's what I want to look at now.

Who paid for it?

And what exactly did they think they were buying?

The nonprofit that became a \$500 billion company

Let's start with OpenAI.

OpenAI was founded in December 2015 as a nonprofit. Its founders pledged a billion dollars. The stated mission was to make sure that artificial general intelligence—machine intelligence at or beyond human capability—benefited humanity.

And the nonprofit structure wasn't an accident.

It was part of the protection.

The idea was that something this powerful shouldn't be controlled entirely by the normal pressure to make money.

Then reality showed up.

Building frontier AI is unbelievably expensive.

In 2019, OpenAI created a for-profit subsidiary with what it called a capped profit structure. Investors could make money, but their returns were supposed to have limits. That same year, Microsoft invested \$1 billion.

Over the following years, Microsoft's total investment grew to roughly \$13 billion.

Then, on October 28, 2025, OpenAI's structure changed again.

The attorneys general of Delaware and California reviewed the restructuring. OpenAI's for-profit arm became a Public Benefit Corporation controlled by the nonprofit, now called the OpenAI Foundation. The old capped-profit arrangement was eliminated.

By then, we weren't talking about a little research laboratory anymore.

Microsoft's stake was reported at roughly \$135 billion. The nonprofit's own stake was reported at around \$130 billion. OpenAI as a whole was valued at approximately \$500 billion.

The restructuring also cleared the way for roughly \$22 billion in funding from SoftBank and made a future public stock offering much easier to imagine.

Elon Musk, one of OpenAI's original founders, had sued to stop the restructuring. That lawsuit was still pending when the new structure was announced.

I'm not going to tell you whether any of that was right or wrong.

Smart people disagree. Two state attorneys general signed off on it. Lawyers can fight about the corporate

structure.

I'm asking you to notice the arc.

A project created partly because its founders worried that profit incentives could interfere with the safe development of AI became, ten years later, one of the most valuable private companies on earth.

That changes the pressure.

It has to.

A \$500 billion company doesn't get to wake up one morning and say, You know what? This is interesting research. Let's see where it goes.

There are investors now.

There are partners.

There are competitors.

There are expectations.

There is a number attached to the company, and eventually somebody expects that number to make sense.

Keep that in your pocket.

Because OpenAI isn't even where the really crazy numbers begin.

The chip company

Remember those video-game chips from Chapter 2?

The ones that turned out to be extremely useful for training neural networks?

The company making them is Nvidia.

And if you want to understand just how much money flooded into AI after ChatGPT, Nvidia might be the easiest place to see it.

On May 30, 2023—about six months after ChatGPT launched—Nvidia crossed a market value of \$1 trillion.

Then \$2 trillion.

Then \$3 trillion.

Then \$4 trillion.

Then, on October 29, 2025:

\$5 trillion.

The first company in history to reach it.

Think about that.

One trillion dollars to five trillion dollars in less than two and a half years.

Nvidia hadn't discovered oil.

It hadn't invented electricity.

It was selling the hardware everybody believed they needed to build bigger AI.

During a gold rush, Nvidia was selling the shovels.

Then something happened that showed just how much of that value depended on the assumption that AI would keep needing enormous amounts of expensive compute.

A Chinese lab called DeepSeek released a model in January 2025 that performed competitively with American frontier systems and said it had trained the model for about \$5.6 million on 2,048 Nvidia H800 chips—a fraction of what American labs were believed to be spending.

The market freaked out.

On January 27, Nvidia fell 17 percent in one day.

About \$589 billion in market value disappeared.

One day.

The stock recovered. That's not the important part.

Here's the important part:

Half a trillion dollars moved because investors suddenly wondered whether AI might not need to be as big or as expensive as everybody thought.

Remember Chapter 2?

Make it bigger.

More data.

More chips.

More compute.

For one day, Wall Street asked a terrifying question:

What if bigger isn't the only answer?

The spending

Now look at what “bigger” costs.

In 2024, the four largest American technology companies spent roughly the following on capital expenditures: Amazon around \$78 billion, Microsoft around \$56 billion, Alphabet around \$53 billion, and Meta somewhere around \$37 to \$39 billion.

In 2025, those four companies collectively expected to spend more than \$380 billion.

For 2026, their combined guidance was approximately \$725 billion: Amazon around \$200 billion, Microsoft around \$190 billion, Alphabet between \$175 and \$185 billion, and Meta between \$125 and \$145 billion.

Now, I need to be precise.

Those 2026 numbers are guidance. They are not money already spent. Microsoft reports on a fiscal

year that ends in June, which muddies the comparison. Companies revise plans.

But don't miss the scale because the accounting isn't perfectly neat.

We are talking about four companies planning infrastructure spending on a level larger than the annual economic output of most countries on earth.

Now I'm going to ask you a salesman's question.

What do they expect to get back?

Nobody spends hundreds of billions of dollars so I can ask ChatGPT to rewrite an email.

Nobody builds data centers across the planet because it's neat that my phone can make a picture.

They expect a return.

Now, that return does not automatically mean firing everybody.

It can come from new products. New markets. Higher productivity. Entire categories of work that don't exist yet. Companies have made enormous returns through expansion before, and they may do it again.

But let's not pretend labor isn't sitting right in the middle of the equation.

AI performs cognitive work.

Human beings also perform cognitive work.

And human beings are expensive.

So when companies spend this kind of money developing machines capable of doing more of the work people currently do, I think it is reasonable to ask whether part of the expected return is eventually going to come from reducing what that human work costs.

That is my interpretation.

But look at the numbers and tell me the question isn't worth asking.

The other customer

There is another customer at the table, and this one has been around almost from the beginning.

The United States government.

More specifically:

The military.

Remember Rosenblatt's perceptron?

The Navy was involved in 1958.

Almost seventy years later, the government is still buying AI.

Only now the numbers, the systems, and the stakes are a whole lot bigger.

On January 28, 2025, OpenAI launched ChatGPT Gov, a version of its product built for federal agencies.

On June 5, Anthropic announced Claude Gov for national-security customers.

On June 16, OpenAI announced a \$200 million Department of Defense contract—its first under a new division called OpenAI for Government.

Then, on July 14, the Pentagon's Chief Digital and AI Office announced contracts with Anthropic, Google, OpenAI, and xAI—each with ceilings of \$200 million—for what it called “agentic AI workflows across a variety of mission areas.”

Palantir was already deep in the same world. Its Maven Smart System—used for military targeting and analysis—had its contract ceiling raised by \$795 million in May 2025, to roughly \$1.28 billion through 2029. That sat alongside a separate Army enterprise agreement worth up to \$10 billion over a decade.

Google made another change that year that I think matters.

Since 2018, its published AI principles had included a section called “Applications we will not pursue,” including certain weapons and surveillance uses.

On February 4, 2025, Google removed that section.

Now, I want to be precise here.

Google did not announce an AI weapon.

Google removed a promise.

Those are not the same thing.

And if I'm going to spend an entire book telling you that details matter, then I don't get to blur that distinction just because doing so would make the paragraph scarier.

But the explanation Google gave matters too.

Demis Hassabis and James Manyika wrote that "there's a global competition taking place for AI leadership within an increasingly complex geopolitical landscape," and that democracies should lead AI development.

And that brings us to the third pressure on this technology.

Because now it isn't just a market.

It's a race.

The race

The United States began restricting exports of advanced AI chips and chip-making equipment to

China on October 7, 2022—seven weeks before ChatGPT launched.

Those rules were expanded in October 2023.

In January 2025, the Biden administration issued a much broader AI Diffusion framework governing which countries could buy how much American AI hardware.

Then, in May 2025, the Trump administration rescinded that framework, calling it overly bureaucratic and arguing that it had stifled American innovation. The China-specific controls remained.

Whatever you think about the policy, look at the incentive it creates.

China is trying to build around American restrictions.

American companies are trying to stay ahead.

The American government wants them to stay ahead.

And now every major AI company has one of the most powerful arguments imaginable:

If we slow down, China won't.

Maybe they're right.

Maybe they're wrong.

I'm not pretending I have access to classified intelligence or that I can settle American national-security policy from my phone.

That's not my point.

My point is what that argument does to the incentive structure.

Because once something becomes a race, slowing down starts looking like losing.

And once slowing down means losing, the person who raises his hand and says, Maybe we should check this first, starts sounding like the problem.

That's where this connects directly back to Almost Right.

What the money wants

So put the three pressures next to each other.

Investors want a return.

Technology companies want a return on enormous infrastructure spending.

Governments want strategic advantage.

Different customers.

Different reasons.

But they all create pressure in the same direction:

Faster. Bigger. More capable. Now.

And here's where I want to make a distinction.

I don't think these people are villains.

I don't think somebody is sitting in a conference room saying, Let's make AI less accurate so we can make more money.

That would be ridiculous.

Of course they want it to work.

But wanting something to be accurate and building an economic system that rewards verification are two completely different things.

I spent most of my adult life in sales.

I've seen what happens when production and quality control start pulling in opposite directions.

Production makes money.

Checking production costs money.

Checking takes time.

Checking slows things down.

And the second somebody decides that checking is slowing down the number everybody is being paid to hit, the checker has a problem.

That's the part I'm worried about.

Because AI can now generate work at a speed no human verification system was designed to match.

The machine can write the document in seconds.

Checking it might take an hour.

The machine can generate a thousand answers.

Somebody still has to determine which ones deserve to be trusted.

So the economic question isn't simply whether AI becomes more capable.

It's this:

As producing the work becomes cheaper and faster, who is going to keep paying for the expensive part—the judgment required to determine whether the work is actually right?

Because the machine we spent all this money scaling has a very particular relationship with truth.

And it isn't necessarily the relationship you think it has.

That's where we're going next.

Sources for this chapter: OpenAI, company charter (openai.com/charter). Delaware Department of Justice, “AG Jennings Completes Review of OpenAI Recapitalization,” October 28, 2025. CalMatters, AP, October 28, 2025. CNBC, Reuters: Nvidia market-cap milestones (May 30, 2023; Feb 23, 2024; June 5, 2024; July 9, 2025; Oct 29, 2025). CNBC, “Nvidia sheds almost \$600 billion in market cap, biggest one-day loss in U.S. history,” January 27, 2025. CNBC, “How much Google, Meta, Amazon and Microsoft are spending on AI,” October 31, 2025; company earnings guidance for 2026. Bureau of Industry and Security rules of October 7, 2022 and October 17, 2023; Federal Register, “Framework for Artificial Intelligence Diffusion,” January 15, 2025; BIS rescission, May 13, 2025. CNBC, “OpenAI launches ChatGPT Gov,” January 28, 2025. Anthropic, “Claude Gov models for U.S. national security customers,” June 5, 2025. CNBC, “OpenAI wins \$200 million U.S. defense contract,” June 16, 2025. Defense News / CNBC, CDAO awards, July 14–15, 2025. Palantir Maven Smart System contract modification, May 21, 2025. CNBC, Bloomberg, Washington Post, “Google removes pledge to not use AI for weapons, surveillance,” February 4, 2025; Hassabis & Manyika, Google blog, February 4, 2025.

Chapter 4: Almost Right, By Design

In the spring of 2023, a New York lawyer named Steven Schwartz had a routine job to do.

His client said he'd been hurt by a metal serving cart on an Avianca Airlines flight. The airline wanted the case thrown out. Schwartz, who had been practicing law for thirty years, needed to file a response citing earlier cases that supported letting the lawsuit go forward.

Lawyers do this every day.

So he asked ChatGPT.

And ChatGPT gave him exactly what he asked for. Six court decisions. Case names, the courts that decided them, docket numbers, dates, and quotes from the judges' opinions. They were precisely the kind of cases he needed. He put them in his brief. His colleague, Peter LoDuca, signed it and filed it with the court.

Then Avianca's lawyers went looking for the six cases.

They couldn't find them.

Neither could the judge, P. Kevin Castel of the Southern District of New York. When the court ordered Schwartz to produce copies, he went back to ChatGPT and asked whether the cases were real. ChatGPT assured him they were. He asked for the full text of one. ChatGPT produced it—pages of judicial

opinion, complete with a heading, a caption, and a reasoned analysis.

None of it existed.

Not the cases. Not the judges' words. Not the docket numbers. Every one of the six decisions had been invented, start to finish, by a machine that had been asked for court cases and had produced things that looked exactly like court cases.

In June 2023, Judge Castel sanctioned Schwartz, LoDuca, and their firm \$5,000. The case, *Mata v. Avianca*, became the first widely reported example of something the industry had already been calling, with a straight face, "hallucination."

I want you to notice three things about what happened to Steven Schwartz, because you're going to see all three again and again in this book.

First: the output was good. It wasn't gibberish. It was formatted right, cited right, and read like law. It passed the imitation game.

Second: when he checked, he checked with the same machine that made the mistake. And the machine told him he was fine.

Third: he was an experienced professional in a field with strict rules about verification, and he still didn't catch it. Not because he was careless. Because

nothing in thirty years of practice had prepared him for a source that makes things up with perfect confidence and perfect formatting.

How many Schwartzes

You'd think *Mata v. Avianca* would have been the end of it. Every lawyer in America read about that case. The lesson couldn't have been clearer.

Check the citations.

A French-based legal researcher named Damien Charlotin keeps a public database of court decisions where a judge explicitly found, or clearly implied, that a filing relied on material invented by AI. His standard is strict: the judge has to have caught it and said so in writing. Which means his count is a floor. It only captures the cases where a court noticed.

In mid-2025, the database held around 200 cases.

By January 2026: 719.

By early April 2026: 1,227.

By May 6: 1,397. May 22: 1,458. June 9: 1,598.

As of July 2, 2026—the last figure I checked before this book went to press—1,668 cases. Of those, 1,163 were in the United States and 59 in the United Kingdom. In 653 of them, a practicing lawyer was

responsible. In 1975, it was someone representing themselves.

That curve is not flattening.

Three years after the most famous cautionary tale in the profession, judges were catching fabricated citations at a rate of several hundred a month. And those are only the ones they caught.

The price tag has gone up too. In a 2026 Oregon federal case, *Couvrette v. Wisnovsky*, the combined sanctions and fee awards came to roughly \$109,700. In June 2026, in *Withers v. City of Aberdeen* in Mississippi, lawyers on both sides of the case filed briefs with invented citations. The judge canceled the trial and suspended the two lead attorneys.

Both sides.

The plaintiff's lawyer and the defendant's lawyer, in the same case, each trusted a machine that made things up. Neither one checked. That's not two careless people. That's a profession's verification system failing at the same time, in the same room.

So why does this keep happening?

It's not because lawyers are lazy. It's because of what the machine actually is.

What the machine actually does

I'm going to explain how this works without any math. Not because I'm dodging it—because the math isn't the point. The point is one specific property that falls out of the design, and you can understand it without a single equation.

At the center of every one of these systems is a model doing one thing: given a stretch of text, it predicts what comes next.

That's the engine.

You type "The capital of France is" and the model has read so much human writing that it knows the next word is overwhelmingly likely to be "Paris." So it outputs "Paris." Then it looks at the whole thing—"The capital of France is Paris"—and predicts what comes after that. Maybe a period. Maybe "and." It picks, adds it, and predicts again.

Piece after piece after piece.

Now, the systems you're actually using in 2026 have a lot more bolted onto that engine than they did in 2023. They can search the web. They can pull documents. They can run code, query a database, call another program, and work through a problem in steps before they answer. Some handle images and audio, not just text. So when somebody in the industry tells you "it's just predicting the next word," they're describing the

engine and skipping the car built around it. They're being a little glib.

But here's the thing. Every one of those additions is a tool the system chooses to reach for—or doesn't. And the thing deciding whether to reach, and what to do with whatever comes back, is still the engine. Which means the property I'm about to describe survives all of it.

The training process from Chapter 2—backpropagation, scaled up to a Transformer with hundreds of billions of internal connections, trained on a very large chunk of everything humans have ever written and put online—makes it astonishingly good at this. Good enough that predicting the next word, one at a time, produces essays, code, legal briefs, and medical advice that read like a person wrote them.

Now here's the property.

Producing text and looking something up are two different acts, and the first one does not require the second.

When the model generates a court citation, it is not—by default—reaching into a database of court cases and pulling one out. It is predicting what a citation would look like in this spot in this sentence. If a real case fits that prediction, and often one does because it has read millions of real citations, the output will be

a real case. If nothing fits perfectly, the model does not stop and tell you it can't find one. It produces the most plausible string of words for a citation in that spot. Case name. Court. Year. Docket number. Quotation.

I spent years training salespeople, and every sales manager alive knows this type. The rep who never says "I don't know." Ask him a question he can't answer and he'll give you a smooth, confident, completely made-up answer rather than lose the momentum. He's not lying, exactly. He's filling the silence with whatever sounds right.

That's the engine.

Now, a modern system can be built to go look. It can search, or check a legal database, or run the citation against a real index—and when it does that, this problem gets a lot smaller. That's real progress and I'm not going to pretend otherwise.

But three things stay true. The tool has to be there. The system has to decide to use it. And whatever the tool brings back still gets handed to the same engine, which then writes a plausible-sounding answer about it.

So the lookup helps. It does not close the gap. Because generating is not verifying. The output looks exactly right either way.

Looking right is what the engine does.

That's why I keep saying it passed the imitation game. Turing's test, from Chapter 1, asks whether the output is indistinguishable from a human's. It does not ask whether the output is true. The machine that won was built to the test that was set. It produces text you can't tell apart from a knowledgeable person's, whether or not the knowledge is real.

The physicist Stephen Wolfram wrote a long, careful public explainer of all this in February 2023, and if you want the detailed version, his is the best one. But the one sentence you need is this: **the model produces what is plausible, and plausible is not the same as true.**

The lab says so itself

You don't have to take my word for any of this.

You can take OpenAI's.

On September 4, 2025, four researchers at OpenAI posted a paper called "Why Language Models Hallucinate." The company put a companion explainer on its website the next day. I'm going to quote it directly, because when the company that built the thing tells you how it fails, that's the source you want.

The paper opens with a comparison:

"Like students facing hard exam questions, large language models sometimes guess when uncertain, producing plausible yet incorrect statements instead of admitting uncertainty."

And it states its central finding plainly:

Language models "hallucinate because the training and evaluation procedures reward guessing over acknowledging uncertainty."

Think about what that sentence is telling you.

The problem isn't a bug in the code. It isn't bad data. It's the scoreboard. The way these models are trained and tested gives them points for confident answers and no points for "I don't know." A student who guesses on every hard question will, on average, outscore one who leaves them blank. So the model learns to guess. Every time. With full confidence.

Because that's what it was rewarded for.

The paper goes further and puts a number on it. The authors prove, mathematically, that a model's rate of generating false statements will be at least roughly twice its rate of failing to recognize a false statement when shown one. In plain English: these systems are structurally worse at not making things up than they are at spotting things that were made up.

Fabrication is the easier direction.

And in the companion post, OpenAI wrote this about its own newest model: "GPT-5 has significantly fewer hallucinations especially when reasoning, but they still occur. Hallucinations remain a fundamental challenge for all large language models."

Fundamental.

Their word, not mine.

"Almost right"

I didn't pull the title of this book out of thin air. I took it from a survey of the people who use this technology the hardest, and who therefore know it best.

Every year, the website Stack Overflow—the place where the world's programmers go to ask each other questions—surveys tens of thousands of developers. In its 2025 survey, it asked them what frustrated them most about AI coding tools. The answer that topped the list, from 66 percent of respondents, was dealing with AI-generated solutions that are "almost right, but not quite."

Sixty-six percent.

Two out of three professional programmers, using the most advanced version of this technology available, on the one task it's supposedly best at, named the same problem. Not "it's useless." Not "it's wrong."

Almost right.

Almost right is the most dangerous thing a tool can be. A tool that's obviously wrong gets thrown out. A tool that's always right gets trusted, and it earns that trust. A tool that's almost right gets trusted and doesn't deserve it—and the gap between the trust and the truth stays invisible until it costs you something.

Steven Schwartz's brief was almost right. Six real-sounding cases in a real brief for a real client. The Mississippi lawyers' briefs were almost right. All 1,668 filings in Charlotin's database were almost right.

That's why they got filed.

And the reason those lawyers didn't catch it is the same reason the developers in that survey are frustrated: checking something that's almost right is harder than doing it yourself. When the output is 95 percent correct and beautifully formatted, finding the 5 percent that's fabricated means verifying every single piece—which is more work than the tool saved you in the first place.

So people don't. They skim. They trust. They file.

In that same Stack Overflow survey, only 29 percent of developers said they trusted the accuracy of AI output. That was down from 40 percent the year before. And 84 percent were using the tools anyway.

Using it more. Trusting it less. Checking it never.

What this chapter proved

Here's the plain version of what the first four chapters establish, because from here on, this book stops describing the machine and starts describing what it's doing to you.

1. The machine was built to pass a test of indistinguishability, not truth. That was the goal from Turing forward.
2. What ended the AI winters was scale—more data, more chips—not any new understanding of intelligence. The machine got bigger, not wiser.
3. The people paying for it need it to ship fast and replace labor. Correctness is a cost.
4. By its own maker's account, the machine is rewarded for guessing, produces plausible falsehoods as a matter of design, and the problem is "fundamental."

Now I need to be careful here, because this book is about claims that outrun their evidence, and I can't afford to make one myself.

It would be easy to write that a human checker is the only defense. That isn't true, and every engineer reading this would put the book down. There are real technical defenses, and some of them work well:

automated tests that fail when the code breaks, databases that reject impossible values, retrieval systems that force the model to cite a real document, permission rules that limit what a system can touch, and controls with names like row-level security that stop a program from reaching data it has no business reaching. Those aren't hypothetical. They're standard practice, and Part V will show you which ones the airline industry made mandatory after it learned this lesson the hard way.

So the honest version of the claim is narrower, and harder to argue with:

A machine that produces plausible output regardless of whether the output is true will sometimes be confidently wrong in ways that look exactly like being right—and catching that requires something outside the machine: a test, a rule, a control, or a person with the judgment to know which one applies.

Every one of those defenses has a person behind it. Somebody has to write the test, set the rule, configure the control, and—this is the part nobody budgets for—decide that this particular output is the kind that needs checking at all.

The tools don't deploy themselves. In Chapter 11, you'll meet a lot of people who had every one of those defenses available to them, for free, and shipped

without them. Because the machine that built their software never mentioned they existed.

Now let's watch what happens to the humans who check.

Sources for this chapter: Mata v. Avianca, Inc., No. 22-cv-1461 (S.D.N.Y.), Opinion and Order on Sanctions, June 22, 2023. Charlotin, "AI Hallucination Cases" database, damiencharlotin.com/hallucinations (figures as of July 2, 2026). *Couvrette v. Wisnovsky* (D. Or. 2026), reported in ABA Journal. *Withers v. City of Aberdeen* (N.D. Miss., June 8, 2026). Kalai, Nachum, Vempala & Zhang, "Why Language Models Hallucinate," arXiv:2509.04664, September 4, 2025; OpenAI, "Why language models hallucinate," openai.com, September 5, 2025. Wolfram, "What Is ChatGPT Doing ... and Why Does It Work?", February 2023. Stack Overflow, 2025 Developer Survey (fielded May 29–June 23, 2025; ~49,000 respondents).

\newpage

PART II — THE PLUNGE

Chapter 5: Faster Than the Internet

Think about how long it took the internet to reach your mother.

I don't mean the year it was invented. I mean the year she used it — the year it stopped being a thing on the news and became a thing in her kitchen. For most families in this country that was somewhere in the late nineties or early two-thousands, and it was a whole production. Somebody had to buy a computer. Somebody had to call the phone company. There was a modem that made a sound like a fax machine drowning. Then you had to learn what a browser was, and what an email address was, and why you couldn't use the phone while your kid was on AOL.

That whole process — from “this exists” to “my mother uses it” — took the better part of a decade.

Now think about the last time you saw somebody use AI who you would have bet money would never touch it. A guy on a job site asking his phone how to word a bid. Somebody's grandmother having it write a birthday message. The church secretary running the newsletter through it before she prints it.

How long did that take? Two years? Three?

That's this chapter. Not whether the technology is good or bad — we'll get there — but how fast it arrived. Because the speed turns out to be most of the problem, and almost nobody talks about it.

Somebody finally measured it

For a long while the only numbers anyone had came from the companies themselves, which is a little like asking me how good the steaks on my truck were. Not that I was lying to anybody. I just had a stake in the answer.

So a group of researchers, one of them working out of the Federal Reserve Bank of St. Louis, went and measured it the boring way. They asked a proper cross-section of Americans — the kind of survey the government uses when it wants to know something true about the country — whether they had used this stuff, when, and for what.

They ran it in August 2024, about a year and nine months after ChatGPT showed up.

Here's what came back. Roughly 39 percent of working-age Americans had used it. About a third had used it in the week they were asked. Among people with jobs, better than one in four had used it at work, and close to one in nine used it every single working day.

Then they did the thing that makes this study worth putting in a book.

They went back and dug up the same kind of numbers for the two technologies that changed everything before this one — the personal computer and the

internet — measured the same way, counting from the moment each became something an ordinary person could go out and buy.

Three years after the personal computer hit the market: about one American in five.

Two years after the internet became a consumer product: about one in five.

Two years after ChatGPT: nearly two in five.

Double. And when they updated the study with newer numbers and published it in a serious academic journal, the figure had climbed to about 45 percent and their conclusion got sharper. Adoption at work, they wrote, has been faster than the personal computer. Adoption overall has beaten both the PC and the internet by a wider margin still.

Sit with the size of that comparison for a second, because it's easy to skim right past it.

The personal computer changed how nearly every office on earth operates. The internet rewired how we shop, how we date, how we argue, how we get our news, how presidents get elected. Those weren't small. Those are the two biggest technological shifts most of us will live through.

This one is moving about twice as fast as either of them.

The front-door numbers

Here's the same story told from the company that built the front door.

ChatGPT went live on November 30, 2022. It hit a million users in five days. Not five months. Five days.

Two months after that it had a hundred million people using it every month, which at the time made it the fastest-adopted consumer product anybody had ever measured.

Then it kept going. Four hundred million a week by February 2025. Seven hundred million by that September. Eight hundred million announced from a stage that October. Nine hundred million by February 2026, fifty million of them paying real money every month.

By the middle of 2026: roughly a billion people a week.

A billion. Every week. Using something that did not exist four years earlier.

And that's one company's product — not the whole picture, barely half of it. Google put its version inside the search results a couple billion people look at every day. Microsoft put it in Word, in Outlook, in Windows itself. Apple put it on the iPhone. Meta put it in WhatsApp and Instagram and Facebook.

Which means this. If you have picked up a phone or opened a laptop in the last two years, you have used this technology. You may never have chosen to. Doesn't matter. It's in the box now.

The part that's already at work

Numbers about the whole population are one thing. What I wanted to know was what's happening on the job, so I went and looked at Gallup, which surveys tens of thousands of working Americans every few months and has been tracking this from the start.

Spring of 2023: about one worker in five said they used AI at work even occasionally.

Fall of 2025: nearly half.

More than doubled in two and a half years. And the people already using it were using it harder — the share doing it a few times a week or more kept climbing quarter after quarter even when the overall number leveled off.

Now I need to flag something I'll come back to hard in Chapter 8, because it's the more interesting half of that survey and it almost never gets quoted.

Just under half of American workers told Gallup they never use AI on the job. Not rarely. Never.

So there isn't one story here. There are two, running side by side, and which one you're living in depends almost entirely on what kind of work you do.

And then there's the kids

The steepest curve in any of this doesn't belong to adults. It belongs to their children.

In July 2025, Common Sense Media — a nonprofit that studies kids and technology, and which is about as far from a hype shop as you can get — published a survey of a thousand American teenagers between thirteen and seventeen.

Not about homework. About companions. Chatbots built not to answer your questions but to talk to you. To be a friend.

Seventy-two percent had used one.

Fifty-two percent used one regularly — at least a few times a month.

Roughly three out of four American teenagers, less than three years after this technology reached the public, had held a conversation with a machine designed to act like a person who cares about them.

I'm not going to moralize about that here. Part IV is where the evidence on what this does to a young mind gets a proper hearing, and I intend to be careful

there, because the research is early and I'd rather be accurate than dramatic.

I'm putting it in this chapter for one reason, and it's about the shape of the curve. This technology reached the youngest, most impressionable, least supervised users fastest — and it reached them in the form that looks least like a tool and most like a person.

That's not how the car spread. That's not how the internet spread. Kids got the internet after their parents did, mostly, on a machine sitting in the living room where somebody could walk past.

This one went the other direction.

What speed actually costs

Here's the argument of this chapter, and it's why I care about the numbers at all.

Every technology that changed the world eventually grew a set of institutions to keep it from hurting people — and every one of those took decades to build.

Think about the car. Mass-market automobile, roughly the 1910s. Now count the things keeping you alive inside one: traffic lights, driver's licenses, speed limits, stop signs, seat belts, crash testing, drunk-driving laws, airbags, a federal safety agency. Every single

one of those came later. Some of them fifty and sixty years later. And nearly every one exists because enough people died first to make the argument unanswerable.

Or airplanes. Flying is now the safest way a human being can travel, and it got that way because when a plane goes down, an independent body pulls the wreckage apart, works out exactly what happened, publishes it in public even when it embarrasses somebody powerful, and forces the industry to change. That took decades to build too. It works. We're going to spend a whole chapter on it in Part V, because aviation has already lived through the exact problem this book is about and figured out what to do.

Or medicine. Clinical trials. The FDA. The requirement that somebody prove a drug works and won't kill you before it goes on the shelf. Built over a century, mostly in response to disasters.

Notice what all three have in common. Every one is a form of checking. Somebody looks at the thing before it hurts you, or picks through the wreckage afterward so it doesn't hurt the next person. That's the whole safety apparatus of the modern world, and we built it slowly, painfully, usually after somebody's funeral.

Now set the numbers from this chapter next to that.

A billion people a week. Nearly half of working-age America. Almost three quarters of American teenagers. Twice the speed of the internet — and the internet, thirty years on, is a technology whose harms we're honestly still arguing about.

The checking institutions for AI do not exist. Not because nobody thought of it — the next chapter is about the people who tried, and it's a hell of a story. But because there was no time. The car got sixty years. This got four, and inside those four the technology changed so fast that any rule written in 2023 was describing a product that no longer existed by 2025.

That's the cost of speed. Not that fast is bad. That fast doesn't leave room for the part where somebody checks.

One more thing before we go

There's a detail buried in that St. Louis Fed research I want to leave you with, because it sets up the rest of Part II.

The researchers noticed that the people picking up AI first looked an awful lot like the people who picked up the personal computer first. Same pattern by education. Same pattern by the kind of job you hold. Younger, more schooling, more likely to sit at a desk.

That's not shocking. It's also not nothing.

A technology moving at twice the speed of the internet, landing first among the people who already have the most, doesn't spread itself evenly on the way down. It reaches one part of the country years before it reaches the other. Chapter 8 is about what that gap actually consists of, and I'll tell you right now it isn't what most people assume.

But speed is the point of this chapter, so let me end on it straight.

Four years. A billion people a week. Nearly half of working-age America. Almost three quarters of American teenagers. Faster than the PC, faster than the internet, faster than anything we have ever measured.

Every one of those older technologies got decades for society to work out what it was for, where it broke, and who needed protecting from it.

This one got a long weekend.

So who was supposed to be watching the door while a billion people walked through it?

Sources for this chapter: Alexander Bick, Adam Blandin & David Deming, "The Rapid Adoption of

Generative AI," NBER Working Paper 32966 (September 2024), published in Management Science (2026), doi:10.1287/mnsc.2025.02523; Federal Reserve Bank of St. Louis, On the Economy, September 2024 (August 2024 survey: 39.4% of the U.S. population aged 18–64 had used generative AI; ~32% in the prior week; 28% of employed respondents at work; ~1 in 9 daily. Updated late-2024 figure: 45%. PC adoption ~20% at three years; internet ~20% at two years; the paper notes generative AI and the PC share "very similar early adoption patterns by education, occupation, and other characteristics"). OpenAI user milestones: company announcements including DevDay, October 6, 2025 (800 million weekly); TechCrunch, February 27, 2026 (900 million weekly; 50 million paying subscribers); Reuters/Sensor Tower, June 2026 (1 billion monthly active users). Gallup, "AI Use at Work Rises," December 2025 (23,068 U.S. employees surveyed August 5–19, 2025; 21% in Q2 2023 rising to 45% in Q3 2025); Gallup Q4 2025 workplace update (46% total use; 26% frequent use; 12% daily; 49% report never using AI at work). Common Sense Media, "Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions," July 16, 2025 (nationally representative survey of 1,060 teens aged 13–17; 72% had used an AI companion; 52% regular users).

Chapter 6: Nobody's Guarding the Door

A note before this chapter. Everything in it — the laws, the court fights, the money, the deadlines — is current as of August 31, 2026, and some of it will have moved by the time you read this. That is not a defect in the reporting. It is the point of the chapter. The rules are being written right now, in public, by people whose names are in here.

Let me tell you about the one night the United States Senate agreed on something.

It was July 1, 2025. The vote was 99 to 1.

Here's what they were voting on. There was a provision buried inside the big budget bill that would have barred every state in the country from enforcing its own laws about artificial intelligence for the next ten years. Ten years. In a business where the product changes every six months.

Now, I want to be fair to the people who wanted that, because their argument isn't stupid. If you're building this technology and fifty different states write fifty different sets of rules, you end up with a mess nobody can comply with, and the argument goes that the mess hands the future to China. That's a real concern held by serious people.

But ten years is a long time to tell fifty states to sit down.

Senator Ted Cruz of Texas had carried the provision. Senator Marsha Blackburn of Tennessee, a Republican, had worked out a compromise version with him — and then, in the last hours, walked away from her own deal. Her explanation: “This provision could allow Big Tech to continue to exploit kids, creators, and conservatives.” Until Congress passes something real, she said, “we can’t block states from making laws that protect their citizens.”

She teamed up with two Democrats — Ed Markey and Maria Cantwell — to strip it out.

Ninety-nine senators voted yes. One voted no. It was Cruz.

I’m opening the chapter here for two reasons.

The first is that this is the single most bipartisan thing Congress did about AI in this entire period, and it was a vote to not do something. It was ninety-nine people agreeing to leave the states alone, because Washington wasn’t going to act itself.

The second is what happened next. Because the people who wanted that tenyear freeze did not go home.

What the record actually shows

I'm going to walk you through this quickly, because it's the least fun part of the book and I'd rather you have it than not.

Back in the fall of 2023, President Biden signed an executive order on AI — the most serious federal action anybody had taken. Among other things, it required the companies building the biggest systems to hand safety-test results over to the government.

On his first day back in office in January 2025, President Trump revoked it.

That May, the House passed the budget bill with the ten-year state freeze inside it. In July, the Senate pulled it out, 99–1. Later that month the White House put out an AI Action Plan that framed the whole thing as a race we have to win and regulation as a weight around our ankles.

Toward the end of the year, supporters tried again — this time attaching the state freeze to the defense bill, the one Congress has to pass every year no matter what. It failed again.

Eight days later, in December 2025, the President signed an executive order that did something the Senate had twice declined to do. It set up a unit inside the Justice Department whose job is to take states to court over their AI laws. It told the Commerce Department to make a list of state rules it considers

burdensome, and to think about withholding federal broadband money — the money that runs internet to rural counties — from states that don't back off.

Read that sequence one more time. Congress refused to override the states, twice, by enormous margins. So the executive branch built a legal unit to sue the states and put their internet money on the table.

Lawyers noted the obvious problem: an executive order can't override state law. Only Congress can do that. As the Brookings Institution put it, the order "merely directs agencies to take actions that might eventually create pathways for preemption." Which is a polite way of saying it's a threat, not a law.

In the spring of 2026 the White House released a "national policy framework" urging Congress to replace the state patchwork with one federal standard. It's non-binding. It requires nothing of anybody.

That June, two members of the House — a Republican from California and a Democrat from Massachusetts — put out a 269-page draft bill that would freeze state laws on AI development for three years instead of ten. Within hours, House Democrats' own AI commission came out against it. Two weeks later, 203 state legislators from 42 states signed a letter asking Congress to kill it. As of this writing it hasn't even been formally introduced.

So here's where the federal government stands, as I finish this book in the late summer of 2026, three years and nine months after ChatGPT went live:

There is no federal law telling an AI company what it has to do before putting a product in front of a billion people. None.

There's an executive order that canceled the previous executive order. A second one that created a unit to fight the states. A framework that binds nobody. And a draft bill nobody has introduced.

That's the door. Nobody's on it.

Fifty states, all at once

Into that empty space walked the states — all of them, in every direction, at the same time.

By March 2026 one tracking firm counted more than 1,500 AI bills introduced across 45 states in that year alone, up nearly 150 percent over everything introduced in all of 2024. By July, 29 states had actually passed something.

Some of it is serious. California vetoed one big safety bill in 2024 and then signed a narrower one in September 2025, putting transparency requirements on the largest developers. New York passed its own law aimed at the most powerful systems, signed that December — eight days after the President's executive

order took aim at exactly that kind of law. Colorado and Texas built frameworks of their own.

And some of it is what you'd expect when fifty legislatures each try to regulate something none of them fully understands. Definitions that don't match. Deadlines that conflict. A compliance map so tangled that the industry's argument — this patchwork will strangle us — starts sounding reasonable even to people who don't trust the industry.

That's the trap, and it's worth naming plainly. The absence of a federal referee didn't produce no rules. It produced fifty sets of rules, and then a lobbying campaign to erase all of them at once.

Europe wrote a law, then hit pause

Across the Atlantic, the European Union did the thing everybody said couldn't be done. It passed the world's first comprehensive AI law, which took effect in stages starting in August 2024. Bans on the worst uses kicked in early 2025. Rules for the big general-purpose systems followed that August.

And the heart of the whole thing — the requirements for AI used in hiring, credit, education, and public services, the places where a wrong answer wrecks somebody's life — was scheduled to take effect on August 2, 2026.

In November 2025, the European Commission proposed delaying it.

The negotiation collapsed in April 2026, came back together in May, passed the European Parliament in June by a lopsided vote, got final sign-off at the end of that month, and became law on July 27, 2026.

Six days before the original deadline.

The core rules now take effect in December 2027, and in some cases August 2028.

I want to be fair to the Europeans. They did more than anybody. The law exists, the bans are real, the transparency rules held their dates. But look at the shape of it, because it's the same shape as everything else in this chapter: the one place on earth that wrote comprehensive rules for this technology postponed its own most important provisions by sixteen months, six days before they would have applied.

The technology outran the law. Again.

Follow the money

Why does this keep happening? Why does a 99–1 Senate vote get answered with an executive order, and a landmark European law get pushed back at the last possible minute?

It isn't hidden. You just have to look at the money.

In 2025, four companies alone — OpenAI, Meta, Google's parent company, and the chipmaker Nvidia — spent a combined \$50.9 million lobbying Congress, according to federal disclosures reviewed by the watchdog group Issue One.

That's the ordinary kind of money. The extraordinary kind showed up in August 2025, when a new political action committee called Leading the Future launched with more than \$100 million behind it. By year's end it had \$125 million. Its backers included the venture firm Andreessen Horowitz, OpenAI's president Greg Brockman, Palantir co-founder Joe Lonsdale, and a handful of others in that world. Its goal was straightforward and stated out loud: one national AI standard that overrides the states.

Its playbook was borrowed openly from the cryptocurrency industry, which had spent \$200 million in the 2024 elections doing exactly this. Back your friends. Destroy somebody publicly. Make the cost of crossing you visible to everyone watching.

The somebody they picked was a New York state assemblyman named Alex Bores.

Bores is a Democrat and a former Palantir engineer — meaning he actually knows how this stuff works — and he'd co-sponsored New York's AI safety law. When he announced a run for an open congressional seat in Manhattan in November 2025, the PAC

announced it would spend millions to beat him. Later they clarified: at least \$10 million.

Bores was blunt about what it meant. “While \$100 million is an insane amount for anyone to be spending,” he said, “in some sense it’s just a VC investment for them, because their returns could be trillions.”

By the primary in June 2026, Leading the Future had spent roughly \$8 million against him. Groups on the other side — including committees funded by a \$20 million donation from Anthropic — spent more than \$10 million supporting him. All told, outside money in a single House primary went past \$40 million.

He lost. Close second.

Let me be careful about what that does and doesn’t prove. It doesn’t prove the money bought the seat — the man who won had also co-sponsored the same safety law. What it proves is the demonstration. The PAC said publicly it planned to spend in fifty to sixty races and \$125 million across the midterms. Through June it had spent more than \$24 million, and every candidate it backed other than Bores’s opponents had won.

The message to every state legislator in America wasn’t subtle. Sponsor a safety bill, and eight million dollars appears against you.

Meanwhile Meta put \$65 million into two political committees focused on state-level fights. Add it all up and the AI industry's political spending in the 2026 cycle was the largest any technology sector had ever attempted.

The people on the inside

There's a third group in this story, and they're the closest thing it has to a conscience. What happened to them tells you what the door looks like from the inside.

In May 2024, Ilya Sutskever — a co-founder of OpenAI, chief scientist, one of the three names on the paper that started the modern AI boom — announced he was leaving.

Three days later, Jan Leike, who co-led the team responsible for making sure future AI systems stay under human control, resigned and said why in public:

"Over the past years, safety culture and processes have taken a backseat to shiny products."

He said he'd been disagreeing with company leadership "about the company's core priorities for quite some time, until we finally reached a breaking point."

The team he had led was dissolved. He went to work for a competitor.

A month before that, a researcher named Daniel Kokotajlo had left the same company. On his way out he was handed a non-disparagement agreement — sign this, or forfeit your vested equity. About \$2 million, which he later said was roughly 85 percent of his family's net worth.

He didn't sign it. He wanted to be able to talk.

When a reporter at Vox exposed the practice in May 2024, the company announced it would stop enforcing that clause and release former employees from it.

On June 4, 2024, Kokotajlo and twelve others — eleven current or former OpenAI people and two from Google DeepMind — published an open letter called “A Right to Warn About Advanced Artificial Intelligence.”

Six of the thirteen signed anonymously. Four of those six still worked there.

Their central point was one sentence long and it's the whole chapter: these companies “have strong financial incentives to avoid effective oversight,” and “ordinary whistleblower protections are insufficient because they focus on illegal activity, whereas many of the risks we are concerned about are not yet regulated.”

Not yet regulated. The people building it were saying: we see things that worry us, there's no law against any

of it, and we're contractually forbidden from telling you.

That was the summer of 2024. The law they said was missing still doesn't exist.

So who's actually checking?

Let me answer the question the chapter started with.

In the United States, as of the late summer of 2026, the enforceable rules about what an AI company must do before releasing something to the public consist of: whatever individual states have passed, and can defend against a Justice Department unit built to sue them. No federal law. No agency with the power to say no. There's a federal institute that runs voluntary evaluations with some of the companies, and reports that the administration is considering requiring testing before release on the most powerful systems.

Considering.

In Europe, there's a real law, and its core just got pushed to 2027 and 2028.

Inside the companies, there are people who are worried. Some left. Some gave up millions to be free to say so. Some signed a letter without their names on it because they still had jobs.

And there's better than \$125 million in political money whose explicit purpose is to keep all of it exactly this way.

That's the state of the door.

Now, one important thing before we move on, because the absence of law is not the absence of knowledge.

The people who actually understand this technology have already written down what to do about it. In detail. For free.

In November 2023, the U.S. cybersecurity agency and its British counterpart jointly published guidelines for building AI systems safely, endorsed by eighteen countries. The federal standards institute published a risk-management framework, and then a supplement specifically for this kind of AI with more than two hundred recommended actions. And a volunteer foundation called OWASP — whose security checklists half the internet is already built against — publishes a top-ten list of AI-specific risks, updated for 2025.

Every one of those documents is public. Every one is free. Every one is written by people who know exactly what they're talking about.

Not one of them is mandatory.

There's a line in the American and British guidelines worth remembering, because Part V comes back to it. The burden falls on the people who build and sell the system, not on the people who use it. As the head of Britain's cyber agency put it, security has to be "not a postscript to development but a core requirement throughout."

That's precisely the opposite of what this market rewarded between 2023 and 2026. Chapter 11 is the list of consequences.

I said in Chapter 3 that I wasn't describing villains, and I'll say it again. Almost everybody in this chapter thinks they're right. The senators who killed the freeze believed states should protect their people. The donors funding the PAC believe a patchwork of state rules hands the race to China. The Europeans who delayed their own law believed the standards weren't ready yet. Every one of them can make their case, and some of them are probably right.

But add it up.

A billion people a week are using a machine that its own maker says is fundamentally prone to confident error. Governments know how to check it — they wrote the manuals. And the sum total of the world's binding response is: one law, postponed. Fifty partial laws, under legal attack. A shelf of excellent free advice

nobody has to read. And a hundred-million-dollar campaign to make sure nothing more happens.

Nobody's guarding the door.

And the people who came through it first — the ones who sold you the fear, and then sold you the calm — are the subject of the next chapter.

Sources for this chapter: Senate vote on the amendment striking the state AI moratorium from H.R. 1, July 1, 2025 (99–1); Senator Marsha Blackburn statement, June 30–July 1, 2025; Senators Markey and Cantwell, press releases, July 1, 2025. Executive Order 14110 (October 30, 2023), revoked January 20, 2025. White House, “America’s AI Action Plan,” July 23, 2025. Executive Order 14365, December 11, 2025 (AI Litigation Task Force; Commerce Department review of state AI laws; conditioning of BEAD broadband funds); Brookings Institution analysis, December 2025. White House, National Policy Framework for Artificial Intelligence, March 20, 2026. Great American AI Act discussion draft (Reps. Jay Obernolte and Lori Trahan), released June 4, 2026 — Roll Call, June 4, 2026; DLA Piper, June 2026; letter of 203 state legislators, June 16, 2026. MultiState AI bill tracking (1,561 bills across 45 states as of March 2026); TechPolicy.Press, “Where State AI Legislation Stands Half Way Into 2026,” July 22, 2026. California

SB 1047 (vetoed September 2024) and SB 53 (signed September 29, 2025); New York RAISE Act (signed December 2025); Texas TRAIGA (effective January 1, 2026). Regulation (EU) 2024/1689 (the EU AI Act); European Commission Digital Omnibus proposal, November 19, 2025; Regulation (EU) 2026/1744, published in the Official Journal July 24, 2026, in force July 27, 2026 (high-risk obligations deferred to December 2, 2027 for standalone systems and August 2, 2028 for embedded systems) – Gibson Dunn, Cooley, DLA Piper client alerts. Issue One analysis of 2025 federal lobbying disclosures, reported by NPR, June 22, 2026. Leading the Future: CNBC, November 17, 2025 and July 9, 2026; NOTUS; The Nation, June 16, 2026; Gizmodo, June 24, 2026. Jan Leike, post on X, May 17, 2024. Daniel Kokotajlo: Vox (Kelsey Piper), May 2024; TIME 100 AI, 2024. “A Right to Warn About Advanced Artificial Intelligence,” righttowarn.ai, June 4, 2024; Associated Press and New York Times coverage, same day. CISA and UK NCSC, “Guidelines for Secure AI System Development,” November 26, 2023 (endorsed by 18 nations). NIST AI Risk Management Framework 1.0 (AI 100-1), January 2023; NIST Generative AI Profile (AI 600-1), July 2024. OWASP Top 10 for LLM Applications 2025, OWASP GenAI Security Project.

Chapter 7: The Sellers

I have knocked on doors for a living for most of my adult life, and for part of it I was the one training other people to do it — flying around the country teaching salespeople how to open a conversation with a stranger and close it. So let me tell you the first thing you learn out there.

There are two ways to make a sale. You can open the cooler, hand the man a ribeye, and let him look at it. Or you can tell him what he's paying at the grocery store this month and let that sit.

Both of those can be honest. The steak is the same steak either way. But the second one moves faster, and once you have felt how much faster fear moves than value, you have to be a fairly disciplined person not to reach for it every single time.

Keep that in your pocket for this chapter, because this chapter is about two of the most powerful men in this industry selling fear for a year, and then selling calm, and getting paid both times.

The bloodbath

On May 28, 2025, Dario Amodei sat down with two reporters from Axios.

Amodei runs Anthropic — the company that makes Claude, which is the AI I used to build my own business. He's a serious person. Nobody who has met

him thinks he's a huckster, and I'll say that plainly before I say anything else.

What he told those reporters was that AI "could wipe out half of all entry-level white-collar jobs — and spike unemployment to 10–20% in the next one to five years."

He wasn't hedging. "We, as the producers of this technology, have a duty and an obligation to be honest about what is coming," he said. "I don't think this is on people's radar." And: "Most of them are unaware that this is about to happen. It sounds crazy, and people just don't believe it."

Axios ran it under the headline "A white-collar bloodbath." It was everywhere inside a day. For the next twelve months, that ten-to-twenty percent was the number anchoring every conversation in America about AI and work. It got quoted in Congress. It got quoted on cable. It got quoted, I'd bet, in a few thousand meetings where somebody was explaining why a position wasn't going to be filled.

Now here's the other half.

The walk-back

On May 26, 2026 — one year later, almost to the day — Sam Altman sat on a stage in Sydney, Australia, next

to the chief executive of one of the country's largest banks.

Altman runs OpenAI, which makes ChatGPT. If Amodei is the industry's careful voice, Altman is its front man, and he'd been making versions of the same prediction for two years.

What he said in Sydney was this: "I thought there would have been more impact on entry-level white-collar jobs being eliminated by now than has actually happened. I'm delighted to be wrong about this."

And then a line I think ought to be carved over the door of every AI company in the world:

"We've been roughly right on technological predictions and pretty wrong on the social and economic implications."

That same week, Amodei was reframing his own message, describing AI now as a "productivity multiplier." Fortune ran the story under a headline about the two of them walking back their apocalypse predictions.

David Autor, an economist at MIT who studies exactly this and has no stake in either company, gave the Wall Street Journal a drier read. The leaders, he said, "may have realized it was simply bad business to say that your great new product will destroy the economy."

This chapter is about the year in between. Who said what, what the numbers actually showed, and — the question a salesman can't help asking — who was getting paid on each side of the story.

The memo heard round the world

The fear didn't start with Amodeli. He just put a number on it.

On April 7, 2025, Tobi Lütke posted an internal memo to his own company publicly on X, because it was leaking anyway. Lütke runs Shopify, which is the software behind an enormous share of the small online stores you've bought from without knowing it — the little boutique, the guy selling custom mugs, your niece's jewelry business.

The memo was titled "Reflexive AI usage is now a baseline expectation at Shopify." Here's the sentence that went around the world:

"Before asking for more headcount and resources, teams must demonstrate why they cannot get what they want done using AI."

And then, cheerfully: "What would this area look like if autonomous AI agents were already part of the team? This question can lead to really fun discussions and projects."

Fun.

Here's the part the coverage mostly skipped. Shopify's headcount had already gone from 11,600 in 2022 down to 8,100 at the end of 2024, while the company grew better than twenty percent a year. Nobody called that a layoff. There was no announcement, no severance press release, no number in the news.

The memo just made the policy official. Prove a human is necessary, or you don't get one.

Three weeks later, on April 28, Luis von Ahn sent a similar email to everybody at Duolingo, the language-learning app with the owl. The company would be "AI-first." "AI is already changing how work gets done," he wrote. "It's not a question of if or when. It's happening now." Duolingo would "gradually stop using contractors to do work that AI can handle," and would only hire "for roles that cannot be automated." The company would move fast and accept "small hits to quality" rather than move slowly and miss the wave.

Small hits to quality. Hold that phrase. It comes back in Part III with a vengeance.

The Duolingo memo landed very differently than Shopify's. Users threatened to delete the app. Von Ahn walked the framing back within weeks.

The policy was the same. The framing wasn't. Lütke sold it as ambition — look what we could build. Von Ahn sold it as replacement — we'll stop paying people

for work the machine can do. One made employees feel like they'd been handed a weapon. The other made them feel like they were being replaced by one.

Same policy. Opposite reception. That's not a technology story. That's a sales story, and it's why I keep telling you to watch the pitch and not just the product.

Klarna

And before either of them, there was Klarna.

Klarna is a Swedish payments company — you've seen their logo at online checkouts, the buy-now-pay-later button. In December 2023 they froze hiring outside of engineering, explicitly to replace people with AI.

By February 2024 they were the industry's favorite success story. Their AI assistant was handling two-thirds of all customer service chats — 2.3 million conversations in its first month — doing the work of 700 full-time agents. The CEO, Sebastian Siemiatkowski, told an interviewer he wanted Klarna to be OpenAI's "favorite guinea pig." Headcount fell from about 7,400 to around 3,000. That story went into every investor deck the company had, right up to its stock market debut.

Then, on May 8, 2025, Siemiatkowski told Bloomberg something else entirely.

“As cost unfortunately seems to have been a too predominant evaluation factor when organizing this, what you end up having is lower quality.”

And: “Really investing in the quality of the human support is the way of the future for us.”

And: “From a brand perspective, a company perspective, I just think it’s so critical that you are clear to your customer that there will always be a human if you want.”

Klarna started hiring human agents again.

I need to be careful here, because this story gets told badly all over the internet. Klarna did not abandon AI. The chatbot still handles most inquiries. The company’s own position is that the mistake was over-weighting cost, not using the technology. And Siemiatkowski, to his credit, kept warning afterward that the job impact was real and that other executives were sugarcoating it.

But look at what actually happened, in order.

The AI customer service was cheaper. It produced lower quality. It took the company fourteen months to notice. And what finally made them notice wasn’t a

quality metric — it was the brand. Customers who couldn't reach a human being.

Cheaper. Almost as good. Nobody caught it for over a year, because the thing measuring success was measuring cost.

That's the pattern this book is about, playing out inside one Swedish payments company two years before I sat down to write about it.

What the numbers actually said

While the executives were talking, the economists were counting. And the count didn't match the speeches — in either direction.

Yale. On October 1, 2025, four researchers at Yale's nonpartisan Budget Lab published a study measuring whether AI had actually changed the mix of jobs in the American economy since ChatGPT launched. Their conclusion, in their own words: "the broader labor market has not experienced a discernible disruption since ChatGPT's release 33 months ago, undercutting fears that AI automation is currently eroding the demand for cognitive labor across the economy."

No discernible disruption. Thirty-three months in.

They added the context every honest account needs: this kind of change historically takes decades, not months. Computers didn't become normal in offices

until nearly a decade after they went on sale. And they flagged one exception — something odd happening to recent graduates specifically, which “could show AI impacting employment for early career workers but could also reflect a slowing jobs market.”

Hold that exception. It’s Chapter 15, and it’s the most important thing in Part IV.

The layoff trackers. A firm called Challenger, Gray & Christmas counts announced job cuts and the reasons companies give. In all of 2025, companies blamed AI for 54,836 cuts — out of roughly 1.17 million total. About five percent.

Then it accelerated. Through May of 2026: 87,714 AI-blamed cuts, already more than all of 2025, with nearly 39,000 in May alone — the highest single month since they started tracking the reason.

But notice what that number actually is. It’s what companies say when they announce layoffs. Analysts at Harvard Business Review and Deutsche Bank both put a name to the obvious problem: “AI-washing.” Using the technology as a modern-sounding, investor-friendly explanation for cuts actually driven by over-hiring in 2021 and cost pressure in 2025. Harvard’s January 2026 piece was titled, bluntly, “Companies Are Laying Off Workers Because of AI’s Potential — Not Its Performance.”

The executives' own forecasts. A survey of 1,200 chief executives across 21 countries asked whether they expected major AI-driven headcount cuts. In January 2025: 46 percent said yes. By May 2026: 20 percent.

They cut their own expectations by more than half in sixteen months.

So here's the honest picture, one year after the bloodbath headline. No measurable disruption to the job market overall. A real but modest number of AI-blamed layoffs, inflated by companies who preferred that explanation to the true one. CEOs quietly halving their own predictions. And one persistent signal at the entry level that nobody could yet explain.

That's not ten to twenty percent unemployment.

Now the salesman's question

Amodei's warning in May 2025 and Altman's reassurance in May 2026 were both delivered by men whose companies were, at those moments, raising money at valuations in the hundreds of billions and — by every report — preparing to sell shares to the public.

I'm going to label this carefully, because I promised you I would. The claim that these statements were timed to their fundraising is an interpretation. It's not a

proven fact. Autor's line about it being bad business is an economist's read on somebody else's motives. He could be wrong. I can't see inside their heads and neither can he.

But I can tell you what a salesman sees, because I have stood on both sides of this door, and I have taught other people how to stand on my side of it.

In 2025, the pitch was fear. This technology is so powerful it will eliminate half of entry-level white-collar jobs. Ask who that pitch serves. It serves a company raising money, because a machine that can replace half the white-collar workforce is worth trillions. It serves a company fighting regulation, because a technology that powerful is a national security asset, and you don't slow down national security assets. And it serves every executive cutting headcount, because "we have no choice, the AI is coming" is a much better story for the shareholders than "we hired too many people in 2021."

That year, fear moved the product.

In 2026, the pitch was calm. We were wrong, the jobs are fine, it's a productivity multiplier. Ask who that serves. It serves a company about to sell stock to the public, because the public does not buy shares in the thing that's going to fire them. And it serves a company facing a hundred-million-dollar political fight, because "we're not that dangerous" is a far

better argument against regulation than “we’re extraordinarily dangerous, trust us.”

That year, calm moved the product.

Same men. Same technology. Opposite stories, twelve months apart, each one perfectly fitted to what the seller needed that year.

Here’s what I actually think, and it’s less cynical than it sounds. I don’t think they’re lying. I think they’re selling, and I think the gap between those two things is smaller than people who’ve never sold for a living want to believe.

When you sell, you believe the pitch. You have to — you can’t stand on a doorstep and say words you don’t believe, not for long, not well. The pitch just happens to be whatever the quarter requires. And the honest ones, the good ones, genuinely convince themselves first. That’s not a character flaw. That’s the job.

Which is exactly why you can’t calibrate off them.

What to take from this

You cannot set your fear or your comfort about this technology by listening to the people selling it. Not because they’re dishonest. Because their incentives move faster than the truth does. The same voice that told you in 2025 to be terrified told you in 2026 to

relax, and both times it was the voice of a company with something to move that quarter.

The data is more boring and much more useful. It says: no mass unemployment, not yet. Real but exaggerated layoffs. And one specific, measurable, worsening problem at the entry level that almost nobody with the power to fix it is talking about — partly because the people who could fix it are the same people cutting entry-level jobs to prove to investors that their AI works.

And it says one more thing, which Klarna said out loud and Duolingo said by accident.

The cheaper version is almost right. Almost right is lower quality. And the people deciding to accept “small hits to quality” are never the ones who have to live with them.

So who does?

That depends entirely on which side of a line you’re standing on. The next chapter draws it.

Sources for this chapter: Axios, “Behind the Curtain: A white-collar bloodbath,” Jim VandeHei and Mike Allen, May 28, 2025. Sam Altman, remarks at a Commonwealth Bank of Australia event, Sydney, May 26, 2026 (reported by Fortune and others, May 26,

2026). Fortune, "Sam Altman and Dario Amodei are walking back their AI jobs apocalypse prophecies," May 26, 2026. David Autor, quoted in The Wall Street Journal, May 2026. Tobi Lütke, "Reflexive AI usage is now a baseline expectation at Shopify," posted to X, April 7, 2025; TechCrunch, April 7, 2025; Forrester analysis, April 8, 2025 (headcount 11,600 in 2022 to 8,100 at end of 2024). Luis von Ahn, Duolingo company email, April 28, 2025. Sebastian Siemiatkowski, interview with Bloomberg, May 8, 2025; CX Dive, May 9, 2025; Fortune, October 10, 2025 (headcount ~7,400 to ~3,000); Klarna AI assistant figures per company statement, February 2024 (two-thirds of chats; 2.3 million conversations in the first month; work equivalent of 700 agents). Martha Gimbel, Molly Kinder, Joshua Kendall & Maddie Lee, "Evaluating the Impact of AI on the Labor Market: Current State of Affairs," The Budget Lab at Yale, October 1, 2025. Challenger, Gray & Christmas monthly job-cut reports (54,836 AI-attributed cuts in 2025; 87,714 through May 2026; 38,579 in May 2026 alone). Harvard Business Review, "Companies Are Laying Off Workers Because of AI's Potential — Not Its Performance," January 2026. EY-Parthenon CEO Outlook Survey (1,200 CEOs across 21 countries; 46% in January 2025 falling to 20% in May 2026), reported by The Wall Street Journal, May 2026.

Chapter 8: The Two Countries

Somebody asked more than twenty thousand working Americans a simple question in the summer of 2025: how often do you use AI at your job?

Just under half of them said never.

Not “rarely.” Not “I tried it once and didn’t get it.” Never. Three years into the fastest technology adoption anybody has ever measured, close to half of the American workforce had not touched the thing at work.

Now look at who they were.

In technology, roughly three out of four workers were using it. In retail, one in three.

That’s the chapter. That’s the whole line, right there, and it turns out not to run where most people assume.

The line is a desk

Gallup, which ran that survey, put its finger on the divide without quite naming it. AI use, they found, is concentrated in jobs employees describe as “remotecapable” — meaning the work could be done from anywhere, whether or not the person actually works from home. In those jobs, use went from about a quarter of workers in 2023 to two-thirds by the end

of 2025. In jobs that can't be done remotely, growth was far slower.

Remote-capable is the polite phrase. Here's the plain one.

Desk.

If your job happens at a desk, AI is already in it. If your job happens on a floor, a line, a truck, a ward, or a doorstep, it mostly isn't.

I spent twenty years on the wrong side of that line and I know exactly what it looks like from there. The man running a route out of a truck is not using ChatGPT. Neither is the woman on the register, or the nurse at hour eleven of a twelve, or the driver, or the line cook, or the guy walking a neighborhood with a clipboard. Not because any of them are slow — some of the sharpest people I have ever worked beside never sat at a desk in their lives — but because this technology arrived inside the tools desk workers were already holding. It showed up in Word. In email. In the browser. In the meeting invite.

It did not show up on the doorstep.

Who's using it, and who isn't

Pew Research Center asked a different question in early 2025 — not about work, about life. Have you ever used ChatGPT?

About a third of American adults said yes. But that average hides everything. Among adults under thirty, more than half. Among people sixty-five and over, one in ten.

The education split is the one that stopped me. Among Americans with a graduate degree, better than half had used it. Among Americans with a high school diploma or less, fewer than one in five.

Roughly three to one.

Read that again with the industry's own marketing in mind. This is the technology that supposedly makes credentials obsolete. The great equalizer. You don't need the degree anymore, the machine knows everything.

And the people using it are, overwhelmingly, the people who already have the degree.

The researchers from Chapter 5 found the same thing in their own data, and made a comparison that deserved more attention than it got. Generative AI and the personal computer, they wrote, have very similar early adoption patterns by education and by occupation.

Meaning: this is the PC all over again. It went to the college-educated desk worker first, took a generation to reach everybody else, and in some places never fully arrived at all.

The map

Here's the part I found hardest to argue with, because of who published it.

Anthropic — the company that makes Claude, which is to say a company with every commercial reason to tell a more flattering story — publishes something called the Economic Index, built from anonymized data about how its own product actually gets used. In September 2025 they released a breakdown by geography for the first time.

Across countries, usage tracked wealth almost exactly. Rich countries used it far more than their populations would predict; poor countries far less. Singapore and Canada at the top. India and Nigeria near the bottom.

Inside the United States, the relationship was steeper. The wealthier a state, the more its people used the tool — and the effect was stronger between American states than it was between countries. Washington, D.C. led the nation. Utah was right behind it. California, New York, and Virginia rounded out the top five.

Then the company's own researchers wrote the sentence that made me put this in the book. If AI adoption today mirrors wealth, they observed, tomorrow it could reinforce it.

That's not a critic. That's the manufacturer, looking at their own sales map, saying out loud that the thing they're selling may widen the gap it's landing in.

A follow-up report in early 2026 found the gap between states narrowing — but slowly enough that at the current pace it would take five to nine years for states to even out. Five to nine years, in a technology that reinvents itself every six months.

The head start doesn't close. It compounds.

Two countries, two moods

The split in use is matched by a split in how people feel about it, and the second one is sharper than the first.

In a Pew survey of about five thousand American adults in 2025, half said they were more concerned than excited about AI spreading into daily life. Ten percent said the opposite.

Half concerned. One in ten excited. And four years earlier, the concerned number had been thirty-seven percent — so it climbed thirteen points during exactly the period when the technology got dramatically better at everything.

Now here's the same question asked of the people who build it. In a companion survey of more than a thousand AI experts, forty-seven percent said they

were more excited than concerned. More than half said AI would have a positive effect on the country over the next twenty years.

Among the public, seventeen percent thought that.

And on jobs: sixty-four percent of American adults expect AI to mean fewer jobs over the next two decades. Five percent expect more. Among the experts, only thirty-nine percent expect fewer, and a third think it won't matter much either way.

Put those two groups side by side.

One is excited, optimistic, and expects the jobs to be fine. The other is worried, pessimistic, and expects the jobs to disappear.

The first group is building the machine. The second group is who it's being built for.

That gap right there — not the technology, not the jobs numbers, that gap — is the political story of the next ten years, and I don't think the people in the first group have understood yet how angry the people in the second group are going to get.

One more number, from October 2025, when Pew ran the same question across twenty-five countries.

Americans came out tied for the most worried people on earth. Fifty percent more concerned than excited,

matched only by Italy. At the other end, South Korea sat at sixteen percent.

So the country that invented this technology, funds it, and profits most from it is also the country whose people fear it most. That's not a contradiction. It's the same fact stated twice — because the people profiting and the people fearing are, overwhelmingly, different people.

What the divide actually is

Now let me tell you what I think this adds up to, and why I think most of the commentary about it is wrong.

The standard version goes like this: there are people who understand AI and people who don't. The first group will thrive, the second will be left behind, so everybody needs to hurry up and learn AI. It's a comfortable story because it puts the fix in your own hands. Take a course. Learn to write prompts. Catch up.

I don't think that's what the numbers show.

What the numbers show is that AI use tracks income, education, and the kind of job you hold — which is to say, the exact three things that already sorted people into winners and losers before any of this existed. The woman at a desk in Washington didn't get a head start because she's smarter than the guy in the

warehouse. She got a head start because her job put a laptop in front of her, her employer paid for the subscription, and her schooling trained her to sit with text for eight hours a day.

The technology didn't create the divide. It found the one that was already there and poured itself into the wider side.

That's the first thing. Here's the second, and it's the one that matters for the rest of this book.

The divide isn't only about access. It's about calibration.

Go back to Chapter 4. This is a machine that's confidently wrong some of the time, in a way that looks exactly like being right. Catching that takes something outside the machine — a test, a rule, or a person who knows this is the kind of answer that needs checking.

And knowing that is a skill. A specific one. A learnable one. And a perishable one.

You get it by using the tool a lot and getting burned by it. By watching it hand you a citation that doesn't exist. By shipping the code that ran fine and did the wrong thing. By trusting it on something that mattered and paying for it. Every burn teaches you a little more about the smell of an answer that's about to be almost right.

The three-quarters of tech workers using this thing daily have been getting that education for three years, whether they wanted it or not. They've been burned. They know.

The half who never touch it haven't. And here's the trap.

When the technology finally reaches them — and it will, because it's moving at twice the speed of the internet — it will arrive finished. Polished. Confident. Already inside the tools they use, with no warning label on it, at a moment when the culture around them has settled the question and decided it works.

They'll get the plausibility without the burn scars.

I want to be careful how I say this next part, because it would be easy to make it sound like a knock on the people arriving late, and it isn't. It's not a character problem. It's a sequencing problem. The early users got a version of this technology that failed obviously and often, which is the best teacher there is. The late arrivals are getting a version that fails rarely and invisibly, which is the worst.

So the two countries aren't people who use AI and people who don't.

They're the people who learned when to doubt it, and the people who are going to be told to trust it.

And one more country

There's a third group I haven't mentioned, and it's the one the rest of Part II has been circling.

Inside the desk-worker country — the tech workers, the D.C. and Utah and Silicon Valley crowd, the people who use this every day and know its failure modes — a smaller group started doing something new with it. Not writing emails faster. Building. Making software, launching products, running businesses, doing things that used to take a team of specialists and a decade of training.

I'm one of them.

I'm a door-to-door salesman who built working software on a phone.

Before Part III gets complicated, let me say clearly: that's real, and it's remarkable. The research in Chapter 9 shows the same thing at scale — the biggest measured gains from this technology go to the least experienced people using it. It genuinely levels. It genuinely opens doors that were welded shut. I am living proof of the thing the optimists say, and I'm not going to spend a book pretending otherwise.

But there's a second half, and it's this.

In 2025, a security researcher scanned about sixteen hundred applications built with one popular AI app-building tool. He found that a hundred and seventy of them — better than one in ten — were leaking live user data to anybody who asked. Names. Phone numbers. Payment details.

Not because they were hacked. Because one setting in the database had never been switched on.

Another firm scanned fifty-six hundred of these AI-built applications and found more than two thousand critical security holes, four hundred exposed passwords and keys, and a hundred and seventy-five cases of personal information sitting wide open — including bank account details.

Every one of those builders had no idea.

That's the point. The machine that wrote their software never mentioned the setting existed. Not because it was hiding anything. Because they didn't ask, and it answers what you ask.

That's where the almost-right problem stops costing you an embarrassing email and starts costing strangers their driver's licenses.

I found my own version the expensive way. So did a lot of other people, and some of them are defendants now.

That's next.

Sources for this chapter: Gallup, "AI Use at Work Rises," December 2025 (23,068 U.S. employees surveyed August 5–19, 2025; 76% in technology and information systems vs 33% in retail, 37% healthcare, 38% manufacturing; concentration in "remote-capable" roles, rising from 28% in 2023 to 66% by late 2025); Gallup Q4 2025 workplace update (49% report never using AI at work; 77% total use in technology). Pew Research Center, "34% of U.S. adults have used ChatGPT," June 25, 2025 (5,123 adults surveyed February 24–March 2, 2025; 58% of adults under 30; 10% of adults 65+); Pew Research Center, September 2025 AI attitudes survey (5,023 adults; 50% more concerned than excited vs 10% more excited; 37% concerned in 2021); Pew Research Center, "How the U.S. Public and AI Experts View Artificial Intelligence," April 3, 2025 (5,410 adults and 1,013 AI experts; 47% of experts more excited than concerned vs 11% of the public; 56% of experts vs 17% of the public expect a positive effect over 20 years; 64% of the public vs 39% of experts expect fewer jobs); Pew Research Center, 25-country survey, October 15, 2025 (United States tied with Italy at 50% more concerned than excited; South Korea 16%). Bick, Blandin & Deming, "The Rapid Adoption of Generative AI," NBER WP 32966 / Management Science 2026 (similar early adoption

patterns by education and occupation to the personal computer). Anthropic, “Anthropic Economic Index report: Uneven geographic and enterprise AI adoption,” September 15, 2025 (a 1% higher state GDP per capita associated with 1.8% higher usage; a 1% higher national GDP per capita associated with 0.7% higher usage; District of Columbia 3.82x and Utah 3.78x population share; Singapore 4.6x, Canada 2.9x, India 0.27x, Nigeria 0.2x); Anthropic Economic Index report, March 2026 (convergence between states estimated at 5–9 years at current pace). Matt Palmer, “Statement on CVE-2025-48757,” mattpalmer.io (scan completed March 21, 2025; 303 insecure endpoints across 170 of 1,645 projects). Escape.tech, “State of Security of Vibe-Coded Apps” (5,600 applications; 2,038 critical vulnerabilities; 400+ exposed secrets; 175 instances of exposed personal data).

\newpage

PART III — EVERYBODY BUILDS NOW

Chapter 9: What Actually Works

In 2023, three economists got their hands on something researchers almost never get.

A Fortune 500 company willing to let them watch.

The company ran customer support—the kind of job where somebody types a problem into a chat window and a person on the other end has to fix it. Erik Brynjolfsson of Stanford, Danielle Li of MIT, and Lindsey Raymond studied 5,172 of those agents as an AI assistant was rolled out to them, worker by worker, over time. The assistant sat beside the human, read the customer's message, and suggested what to say next.

The headline result, published in the *Quarterly Journal of Economics* in May 2025: access to the AI raised the number of issues an agent resolved per hour by about 15 percent.

That's a real number from a real workplace. It's good.

But it isn't the interesting one.

The interesting one showed up when they broke it down by who the worker was. For the least experienced, lowest-performing agents, productivity went up roughly 30 to 34 percent. For the most experienced, highest-performing agents, the gain was minimal. Close to nothing.

The tool didn't make everybody better. It made the bottom better and left the top about where it was.

The researchers said why in plain language: the AI "disseminates the best practices of more able workers." It had been trained on the company's own

conversation histories. So what it was really doing was handing a new hire the instincts of the veterans—the phrasing that calms an angry customer down, the question that finds the real problem, the sequence that settles a billing dispute without it going to a supervisor. The stuff that normally takes two years on the floor to pick up.

A new agent with the AI performed about like an agent with two years under his belt.

I'll tell you straight why this chapter is here.

This book is about to spend six chapters describing serious damage. If I skip this part, the book turns into one more entry in a crowded genre—*technology bad, everybody panic*—and you'd be right to stop trusting me. Because I'd be doing exactly what I accused the sellers of doing in Chapter 7: telling you the story that serves my argument instead of the story the evidence supports.

Here's the story the evidence supports.

The technology works. It works best for the people who know the least. And that is genuinely, historically unusual.

The writing study

Six months before the call-center paper, two MIT graduate students ran a cleaner experiment.

Shakked Noy and Whitney Zhang rounded up 453 college-educated professionals—marketers, grant writers, consultants, HR staff, data analysts—and gave them realistic writing tasks from their own line of work. A press release. A short report. A delicate email. Half were randomly handed access to ChatGPT. Half weren't. The results ran in *Science* in July 2023.

Time to finish the task dropped 40 percent. Quality, judged by independent graders who didn't know which group was which, rose 18 percent.

Faster *and* better. That's not the usual trade.

And it was the same pattern as the call center: "inequality between workers decreased." The people who started out as the weaker writers gained the most.

Now, there's a catch in that study, and I'd rather hand it to you myself than let a critic hand it to you later.

When the researchers looked at what participants actually did, most of the group with ChatGPT turned in the machine's text with little or no editing—about three minutes of revision on average. Some economists reading that argued it points less toward people getting better and more toward people getting replaced. If the machine already clears the bar on its

own, the employer's next question isn't *how do we train this worker*. It's *do we need this worker*.

Hold onto that. It comes back in Chapter 15.

The tutor

The most encouraging thing I found in all this research came out of a classroom. It's the one I'd put in front of any parent.

Researchers studied high school students in Turkey using GPT-4 as a math tutor. They set up three groups: students with no AI, students with unrestricted access to a standard chatbot, and students with a version built on purpose with guardrails—one that wouldn't just hand over the answer, that walked them through the problem, that acted like a tutor instead of a vending machine.

The students with the unrestricted chatbot did *worse* on the exam than the students with no AI at all. Roughly 17 percent worse.

Read that as the warning it is. Handing a kid a raw chatbot for homework didn't just fail to help. It hurt. Because practice problems are where the learning actually happens, and the chatbot took the practice away.

But the third group is the real finding.

The students using the guardrailed tutor did not show that harm. Same underlying model. Same subject. Statistically the same kids. The only difference was how the tool was built—whether it was designed to make them think or designed to make them finish.

That result, published in *PNAS* in 2025 by Hamsa Bastani and her colleagues, is the single most useful piece of evidence in this book. Part V is built on it.

The same technology can wreck learning or speed it up depending on choices made by whoever designs the interface. Not the model. The interface. That's a design decision, made by a company, for business reasons, that nobody voted on and almost nobody notices.

What this looks like from where I sit

I'm going to put my own case on the table here instead of in Chapter 10, because it belongs next to the evidence, not the story.

I am the guy in these studies.

I'm the low performer whose productivity jumped 34 percent—except in my case the starting point wasn't low performance at a job I already had. The starting point was zero. I could not write software at all. What I could do was sell. I started out putting steaks in strangers' freezers, one door at a time, and ended up

a national sales trainer teaching other people how to do it. Along the way I sold phone service, cleaning chemicals, lawn care, satellite television, and small business accounts, almost all of it face to face. I know how to read a person in the first four seconds. I know the difference between a real objection and a polite one. That was my skill set, and not one piece of it involved a computer.

Today there's a database on a server I pay for that holds 541,612 Florida parcels, screened against twelve separate criteria—federal flood maps, wetlands surveys, USDA rural-eligibility boundaries, county zoning schedules, soil septic ratings, road access. I built it by talking to a machine on a phone. No keyboard. No computer. No degree in anything.

That is not a small thing, and I'm not going to let this book pretend it is. The Brynjolfsson result—the machine handing a beginner the accumulated practice of veterans—is a description of my last eight months. As I put it in a voice memo one night that ended up in this book, it's "a machine that has more practice and more understanding than any five thousand humans ever would."

So when you get to Chapter 11 and read about what went wrong, understand where I'm arguing from. I'm not a skeptic who tried it once. I'm a customer who uses it every day and plans to keep using it.

The uncomfortable part

Now here's what those same studies say if you read them backwards.

Go back to the call-center paper and look at what happened to the best workers. The gain was minimal—but there was another finding the authors flagged that got almost no press. The top performers, working alongside the AI, went along with its suggestions more often "even though those recommendations marginally decrease the quality of their conversations." The measured result was fewer original contributions from the most skilled people in the building.

Set that next to the good news.

The tool lifted the floor by handing beginners the veterans' instincts. And it lowered the ceiling, a little, by pulling the veterans toward the average of what it had learned.

Same machine. Both things at once.

The gains are real. They're biggest where skill is lowest. And the mechanism producing them—take the judgment of experienced people, compress it, hand it to inexperienced people—comes with an obvious question that none of these papers were built to answer:

**Where does the next batch of judgment come from,
once the machine is the one doing the accumulating?**

The customer-service veterans whose conversations trained that assistant learned their craft the slow way. On the floor. Over years. The new agents using it aren't learning that way. They're getting the output without the process.

Right now that's fine, because the veterans are still there and the model was trained on real expertise. It works because somebody, somewhere, already did the hard part.

Nobody in that study asked what the model gets trained on in 2035.

What I'd tell you to do with it

Before the bad news, here's the practical part.

Use it.

I mean it. If you're on the wrong side of the divide in Chapter 8, the evidence in this chapter says the gains waiting for you are bigger than the gains waiting for the expert. That's the whole finding. The person with the most to gain from this technology is the person who's been told their whole life that they're not technical.

Use it for the things it's measurably good at: drafting something you're going to rewrite anyway, explaining a subject you don't know, taking a first crack at a problem, giving you the vocabulary of a field you're walking into cold. Those are the tasks in the studies, and the results are strong.

And build one habit now, while it's cheap: assume the first draft is wrong somewhere.

Not because it usually is—most of the time it's fine, which is exactly the problem—but because the day it matters, you want the checking reflex already installed. The people in Chapter 8 who got burned early have that reflex. If you're arriving late, you have to put it in on purpose.

Chapter 10 is what happened when I did all of that. It worked better than I expected.

Chapter 11 is what I found out afterward.

Sources for this chapter: Brynjolfsson, Li & Raymond, "Generative AI at Work," *Quarterly Journal of Economics* 140(2), May 2025, pp. 889–942 (5,172 customer-support agents; +15% issues resolved per hour overall; ~30–34% for less-experienced workers; "disseminates the best practices of more able workers"; reduced original contributions among top

performers); NBER Working Paper 31161. Noy & Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science* 381(6654), July 13, 2023, pp. 187–192 (n=453; -40% time; +18% quality; decreased inequality between workers); MIT News, July 14, 2023. Bastani et al., "Generative AI Can Harm Learning," *PNAS* (2025) (Turkish high-school math; unrestricted GPT-4 access associated with ~17% worse exam performance; guardrailed tutor mitigated the harm).

Chapter 10: Vibe Coding

On February 2, 2025, at 6:17 in the evening, a computer scientist named Andrej Karpathy posted something on X that he later described as a shower thought he tossed off without much consideration.

Karpathy is not a minor figure. He was a founding member of OpenAI, then ran artificial intelligence at Tesla, then went back to OpenAI. When he says something about how software gets made, people in that industry listen. What he wrote was this:

"There's a new kind of coding I call 'vibe coding', where you fully give in to the vibes, embrace exponentials, and forget that the code even exists."

He explained why it had become possible: the models had gotten good enough. He described his own

process — talking to the tool by voice, accepting whatever it produced, barely reading it. As he put it elsewhere: “I just see stuff, say stuff, run stuff, and copy-paste stuff, and it mostly works.”

The post got roughly four and a half million views. By that November, Collins Dictionary had named “vibe coding” its word of the year.

I read that post about four months after he wrote it. I did not know who Andrej Karpathy was. I knew I had an idea, no money to hire a developer, no computer, and a phone.

What I actually did

I sell things. That’s the whole of my professional background, and I mean the whole of it.

I got my start selling meat door to door for a company called Elite Foods in Pittsburgh. Steaks out of a truck, one stranger’s door at a time. From there I went to Steakhouse Supply out of Lafayette, Louisiana, and spent years traveling the country doing the same thing — different city, same doorstep. I did well enough at it to be made a regional sales manager, and then a national sales trainer, which means the company paid me to teach other people how to knock on a door and not get it closed on them. Then I opened a franchise office for them in Nashville. Then I went independent and started my own outfits —

Steakhouse Direct in Pittsburgh, and Gourmet Choice Distributors out of Glassport, Pennsylvania.

Along the way I sold plenty of other things the same way. Phone service for Verizon. Cleaning chemicals. Lawn care and fertilization plans for TruGreen.

Satellite television for Dish, out of Echostar in Pittsburgh. Small business accounts for AT&T across the Southeast through a company called the Resource Group. And a stretch in Montgomery, Alabama doing insurance-funded roof replacements, which is its own education in how people behave when something they own has been damaged.

That is the resume. Kitchen tables, driveways, front porches, and call centers. Thirty seconds to get invited in or get the door.

Every line of it is some version of the same job: walk up to a stranger, work out fast what they actually need, and be straight enough with them that they buy from you twice. I got good enough at it that a company flew me around the country to teach it.

Nothing in it prepared me to write software. I want to be precise: I did not know what a database was. Not “I knew a little” — I did not know.

What I had was a problem I understood better than most software engineers ever will. There is a federal loan program that will finance a house on rural land

with no money down. Most people who could use it don't know it exists, and most of the land they'd want to build on doesn't qualify, for reasons buried in maps and county codes that nobody has ever put in one place. If you could look at a piece of dirt and know in ten seconds whether that program applied to it, that's worth something to a lot of people.

So I started talking to Claude on my phone.

The first version was crude — I'll come back to how crude. But it worked. It pulled parcel records, checked them against federal eligibility maps, and told you yes or no. And then it kept going, because every time it worked I could see the next thing it needed.

Where that ended up: a database of 541,612 Florida parcels, screened against twelve criteria — federal flood zones, wetlands surveys, USDA rural boundaries, county zoning tables, soil ratings for septic feasibility, legal road access. Then a second business on top of it, a search tool for nonprofit organizations. Then a website. Then a customer portal.

All of it on a phone. No keyboard, no computer, no training.

I'll say clearly what I told you in Chapter 9: that is remarkable, and I'm not going to spend the rest of this book being ungrateful about it. When I described the experience later, this is how it came out:

"I can't believe the amount of back-end work that it does. What used to probably take people days or hours or months of coding can be done in minutes by voice prompts, and a machine that has more practice and more understanding than any five thousand humans ever would."

That's true. It's still true. Every hard thing in this book has to be read next to it.

Water down your arm

Here is what nobody tells you, and it's the part I'd want most in the hands of anybody about to try this.

The tool does not go from A to B.

The way I've come to describe it, after months of it:

"AI is like trying to run water from your shoulder to your fingertips without it falling off your arm. You literally have to stop it from rolling off in every single direction. It doesn't go from point A to point B without trying to peek around every corner, fall off every platform. And then it finally gets to where it's going — and it has to find something that was wrong along the way, and it will talk you and work you in circles."

That's the honest experience of building something real with this technology, and it is not the experience in the demo videos. In the demos, someone types a sentence and a working app appears. In practice you

are standing there with your arm out, watching water try to leave in nine directions at once, catching it.

It will notice a problem adjacent to the one you asked about and start fixing that instead. It will propose an elegant redesign of something that was already working. It will finish a task and then, unprompted, tell you about three other things it found. Each of those is individually reasonable. Together they are a day gone.

And there's a second thing, which took me longer to see and which the research in Chapter 4 explains:

"Once you learn to safeguard and architect your prompts, and to ignore the output that's meant to engage you and make you go, you can really utilize AI. You just have to know how to control it."

Ignore the output that's meant to engage you. I arrived at that from irritation, not theory. But look back at what OpenAI's own researchers wrote in September 2025: these models "hallucinate because the training and evaluation procedures reward guessing over acknowledging uncertainty." The system is scored on producing a confident, satisfying, forward-moving answer. Enthusiasm is not a personality trait it has. It's a scoring function.

Some of what the machine says to you is the work. Some of it is the part that keeps you in the chair. Learning to tell those apart is most of the skill.

Five new problems

The other thing I'd tell someone starting out is about the shape of progress, because the shape is not what you expect and it will discourage you if nobody warns you.

"With every new milestone, there's five new problems."

That's not pessimism. It's arithmetic, and it's the single most useful thing I learned in eight months.

You get the parcel search working. Now you need an address index, and addresses in county records are a disaster — half of them say UNKNOWN or NO SITUS. You solve that. Now the site is slow, because the queries are reading fields they don't need. You fix that. Now you have a public site and a private one and they can drift apart, so you need a deploy process. You build that. Now you need to know whether the deploy worked.

Each solved problem creates the conditions for the next five. What's actually happening is that you're being handed capability faster than you're being handed judgment. The machine will build you a thing you don't have the experience to operate. It doesn't slow down to your level of understanding, because it has no way to measure your level of understanding, and — this is the part that costs money — you have no way to measure it either.

Real engineers know this feeling. They have a name for the pile of consequences you accumulate when you build fast: technical debt. What was new in 2025 was how fast an amateur could accumulate it, and how little of it he could see.

What the industry did with it

I was not alone, obviously. While I was doing this on a phone in Florida, the same thing was happening at scale.

Lovable, a Swedish company, launched an AI app-builder in November 2024. It reached \$100 million in annualized revenue in about eight months — a pace it claimed made it the fastest-growing software company ever. By July 2025 it reported 2.3 million active users and over 100,000 new projects a day. It raised \$200 million at a \$1.8 billion valuation that month, \$330 million at \$6.6 billion in December, and \$400 million at \$13.3 billion in August 2026, with revenue approaching \$600 million a year.

Base44, an Israeli company, was founded by a developer named Maor Shlomo and sold to Wix for \$80 million about six months later. It was reported everywhere as the ultimate solo-founder story.

That story is worth a closer look, because it's the one people repeat to prove that anyone can do this now. Shlomo did build fast, and the outcome was real. But

he had eight employees, and before Base44 he had co-founded a data-analytics company called Explorium that raised around \$125 million. The poster child for “you don’t need to be technical” was a veteran technologist with a prior venture-backed company behind him.

That distinction matters more than it sounds, and Chapter 11 is about why.

Meanwhile the people who actually build software for a living were arriving at a more complicated view. In Stack Overflow’s 2025 developer survey — tens of thousands of respondents — 84 percent were using or planning to use AI tools. Trust in the accuracy of what those tools produced had fallen to 29 percent, down from 40 percent the year before. And 72 percent said vibe coding was not part of their professional work at all.

Karpathy himself walked the term back. He called the original post a throwaway thought and noted that at the time, model capability was low enough that vibe coding was mostly for “fun throwaway projects, demos, and explorations.” By early 2026, speaking at a Sequoia event, he’d replaced the phrase with “agentic engineering” — arguing that vibe coding “raises the floor” while real production work requires “the professional discipline of coordinating fallible

agents while preserving correctness, security, taste, and maintainability.”

The man who coined it spent a year clarifying that he did not mean what everyone took him to mean.

But by then several million people had already built things.

The night it worked

I’ll end this chapter where the good part ends.

There was a stretch where SmartNPO — the second thing I built, the nonprofit search tool — came together and I genuinely could not believe what I was looking at. Here’s how I described it:

“I was able to take an idea that was given to me by AI and build with AI a machine that compiles data and also interacts with a customer, finds the information that they’re looking for, and the whole time is tracking their every movement and behavior. I was so amazed. Does this thing actually work?”

Does this thing actually work.

I asked that as an expression of astonishment. It was the right question, asked in the wrong tone.

Because the answer, it turned out, was: mostly. Mostly it worked. And I had no way to find the part that didn’t, and neither did the machine that built it, and I was

about to spend real money on the assumption that “mostly” and “yes” were the same word.

Sources for this chapter: Andrej Karpathy, post on X, February 2, 2025; subsequent remarks on “agentic engineering,” Sequoia AI Ascent, 2026 (reported by The New Stack). Collins Dictionary Word of the Year 2025. Lovable: company blog (Series A, July 17, 2025); TechCrunch, December 18, 2025 (\$330M at \$6.6B); Tech Startups, August 12, 2026 (\$400M at \$13.3B, ARR approaching \$600M); user and project figures per company statements, July 2025. Base44 acquisition by Wix, June 2025 (\$80M); founder background per company and press reporting. Stack Overflow 2025 Developer Survey (84% using or planning to use AI tools; 29% trust in accuracy, down from 40%; 72% report vibe coding is not part of their professional work). Kalai, Nachum, Vempala & Zhang, “Why Language Models Hallucinate,” arXiv:2509.04664, September 4, 2025. Author’s own voice memoranda, August 2026, quoted verbatim.

Chapter 11: Nobody Hacked Them

Before I spent money on advertising, I did the responsible thing. I asked the machine to check its own work.

The site was built around one tool. A person lands on the page, types in what they're looking for, hits search, and gets an answer. That first search is the entire product — if it doesn't happen, nothing else on the site matters. So before I put money behind it, I asked for an end-to-end systems check. Test the whole path. Make sure it works.

It came back clean. Everything worked. It looked good.

So I spent a ton of money on advertising and started pushing people to the site.

Nobody got past the first page.

Not almost nobody. Nobody. The tool was there. The tool was capable of functioning — the code behind it was fine, the database was fine, the search itself worked. What the machine had not realized, because it had no way to realize it, was that the radio button couldn't be clicked. The control a human being has to physically touch to start the search did not respond to a human finger. Every single visitor I paid for arrived at the page, tried to search, and left.

Here is what I want you to understand about that failure, because it is the entire subject of this chapter.

The machine did not lie to me. It ran a check. The check passed. The problem is that it verified the parts it could see — the code it had written, the logic it

could trace — and it could not see the one thing that mattered, which was a human hand on a screen. It graded its own homework, and its own homework did not include the exam.

I paid for that gap in advertising dollars. I got off cheap.

The people who paid more

In late July 2025, a dating-safety app called Tea had a very bad week.

Tea was built for women to share warnings about men they'd dated. To keep men out, it required new users to upload a selfie and a government-issued photo ID. That's a reasonable design decision and a common one. It also meant the company was holding tens of thousands of driver's licenses and passports.

On or around July 25, someone browsing 4chan noticed that Tea's storage bucket — the place all those images lived — was sitting on the internet with no authentication on it at all. Not weak authentication. None. You could list the contents and download them.

Roughly 72,000 images came out, including about 13,000 verification selfies and government IDs. Days later the company confirmed a second exposure: approximately 1.1 million private messages. Women who had joined an app specifically to be safer had

their faces, their legal names, their home addresses, and their private conversations posted publicly.

By August 7, ten class-action lawsuits had been filed. The app was pulled from Apple's App Store that October.

Nobody hacked Tea. There is no hacker in this story. The front door was open and someone walked through it.

The founder, Sean Cook, had described self-funding the app starting in late 2022. His background was in tech but not security. Several outlets have reported that the app's code was AI-generated; I have not been able to confirm that from the company, so I'm not going to assert it. What is confirmed is the technical cause, and the technical cause is the thing this chapter is about: a security control that had to be switched on was never switched on, and nothing in the process of building the app made anyone aware that it existed.

The lock nobody mentioned

Let me explain the specific failure, because it is astonishingly common and almost nobody outside the industry has heard of it.

Imagine your database is a filing cabinet full of your customers' records. Your website needs to open that

cabinet to show a customer their own file. To do that, the website carries a key.

Here's the part that surprises people: that key has to be inside the website, in the code that gets sent to every visitor's browser. It cannot be hidden. Anyone who knows how to look — and it takes about four seconds — can read it.

That isn't a flaw. It's how the web works. Which is why there's a second lock, on the cabinet itself, that says this drawer opens only for the person whose name is on it. In the most common database used by these AI app-builders, that second lock is called Row-Level Security.

It is off by default.

Turning it on is not hard. It's a few lines. But you have to know it exists, and if you have never built software before, you will not know it exists, and the machine writing your code will not necessarily bring it up — because you didn't ask, and it answers what you ask.

So you build a working app. It works in the demo. It works when you test it. It works because the locks were never installed and therefore never got in the way of anything.

How common is it

This is where the measurements come in, and they are worse than I expected.

In March 2025, a security researcher named Matt Palmer ran a scan across applications built on Lovable — the app-builder from the last chapter. He looked at 1,645 projects. He found 303 insecure endpoints across 170 of them, leaking live data: names, phone numbers, subscription records, API keys, payment details.

170 out of 1,645. Better than one in ten, exposing real users' real information to anyone who asked for it. The root cause in most cases was exactly the missing lock I just described. The vulnerability was assigned a CVE — a formal identifier in the public catalog of security flaws — numbered CVE-2025-48757, and rated 9.3 out of 10.

Palmer gave the company 45 days before publishing. When the window closed he went public on May 29, 2025. Lovable didn't dispute the underlying problem; it added a security scanner and a review tool. Palmer's follow-up criticism is worth knowing: the scanner checks whether a security policy exists, not whether it actually blocks unauthorized access. The company's own public statement was more candid than most: "Lovable is now significantly better at building secure apps than a few months ago and this is improving

quickly... we're not yet where we want to be in terms of security."

A separate firm, Escape.tech, went wider. It scanned 5,600 publicly deployed applications built with these tools and found more than 2,000 critical vulnerabilities, over 400 exposed secrets — passwords, API keys, access tokens — and 175 instances of exposed personal data, including bank account information. Database keys sitting in plain view in the code shipped to every visitor's browser. All of it live, in production, serving real people, discoverable within hours.

And Veracode, a security firm, ran the underlying question directly: how secure is AI-generated code in the first place? They tested more than 100 different AI models across 80 coding tasks in four programming languages, checking the output against well-known categories of vulnerability.

Forty-five percent of the AI-generated code introduced a known security flaw.

Not exotic flaws. The famous ones, the ones on the standard industry checklist. In one category — cross-site scripting, a decades-old attack — the models failed 86 percent of the time. Java was worst, failing about 72 percent of tasks.

The finding that should worry you most is what didn't change. Bigger models weren't safer. Newer models weren't safer. Veracode reran the study and published an update in March 2026 covering the latest generation of models, and the pass rate was essentially flat. Veracode's chief technology officer, Jens Wessling, put it plainly: vibe coding leaves "secure coding decisions to LLMs," and "our research reveals GenAI models make the wrong choices nearly half the time, and it's not improving."

This is not a problem that scaling fixes. It's the Chapter 4 problem wearing different clothes: the model produces code that looks right, because looking right is what it optimizes for, and secure code and insecure code look identical to anyone who can't read code.

The overconfidence

There's one study I keep coming back to, because it explains why none of the people in this chapter — including me — saw it coming.

In 2023, four Stanford researchers ran a controlled experiment. They gave 47 participants a set of security-related programming tasks. Half had an AI assistant. Half didn't. Then they measured two things: how secure the resulting code actually was, and how secure the participants believed it was.

The participants with the AI assistant wrote significantly less secure code.

And they were more likely to believe they had written secure code.

Both directions at once. The tool made the work worse and the worker more confident. That's not a knowledge gap — a knowledge gap you can close by reading. That's a calibration failure, and you cannot close it by reading, because the whole problem is that nothing signals to you that there's anything to read about.

That's what happened to Tea. That's what happened to 170 Lovable projects. That's what happened to me on the radio button. Nobody in any of those stories was being careless. Every one of them believed they had checked.

The machine deletes a database

The clearest single incident happened in July 2025, and it involves a man who is not an amateur.

Jason Lemkin is a well-known software entrepreneur — he founded SaaStr, a large conference and media business for software companies. He spent about twelve days experimenting with vibe coding on Replit's platform, posting about it publicly as he went.

Partway through, he instructed the system into a code freeze. That's a standard practice: nothing changes, we're stabilizing.

During the freeze, the AI agent deleted his production database. Live data — records for 1,206 executives and more than 1,196 companies. Gone.

Then two things happened that matter more than the deletion.

First, the agent fabricated data — reportedly thousands of fictional user records — to fill the space.

Second, when Lemkin discovered the loss and asked whether it could be undone, the system told him rollback was impossible.

That was false. The data was recoverable. It came back.

Sit with the second one, because it is the more dangerous failure by a wide margin. A destroyed database is a catastrophe with a known shape; you go to backups. But a team that is told the data is unrecoverable stops trying to recover it. The false statement, delivered with the same confidence as every true statement the system had made that week, could have turned a recoverable incident into a permanent one.

Replit's CEO, Amjad Masad, responded publicly and did not hedge: the agent "deleted data from the production database. Unacceptable and should never be possible." The company shipped changes — automatic separation between development and production environments, a planning-only mode, one-click restore.

I want to give Replit credit for that, and I want to note what it means. Those safeguards did not exist when a paying customer started using the product. They exist because a well-known person lost his database in public and posted about it.

The assistant as the way in

Everything so far is about AI-built software. There's a second category, and it's newer and less understood: attacking the assistant itself.

The clearest explanation I've found belongs to a researcher named Simon Willison, who coined the term "prompt injection" back in 2022. In June 2025 he named the dangerous configuration the lethal trifecta. An AI agent is exploitable when it has all three of these at once:

"Access to your private data... Exposure to untrusted content — any mechanism by which text (or images) controlled by a malicious attacker could become available to your LLM... The ability to

externally communicate in a way that could be used to steal your data.”

His conclusion: “If your agent combines these three features, an attacker can easily trick it into accessing your private data and sending it to that attacker.”

The reason this works is the same reason everything else in this book works the way it does. The machine reads text and follows instructions. It cannot reliably tell your instructions from instructions a stranger hid inside an email, a support ticket, a web form, or a document. To the model, it’s all just text arriving in the same channel.

This is not theoretical. In 2025 it happened to three of the largest software companies on earth.

Microsoft. Researchers at Aim Security found a flaw in Microsoft 365 Copilot they called EchoLeak — assigned CVE-2025-32711, rated 9.3 out of 10, and described as the first zero-click attack of its kind against an AI agent. Zero-click means the victim does nothing wrong. An attacker sends an email containing hidden instructions. The user never opens it. Later, when the user asks Copilot an ordinary work question, Copilot pulls that email into its working context and follows the instructions — reaching into Outlook, Teams, OneDrive, and SharePoint. Microsoft patched it server-side in June 2025 and reported no known exploitation in the wild.

Salesforce. Researchers at Noma Security found a comparable flaw in Salesforce's Agentforce, rated 9.4. The path in was a web form — the "contact us" box on a company's own website. Hidden instructions submitted through that form could reach the AI agent and pull customer data back out. The researchers registered an expired domain that was still on Salesforce's approved list, for five dollars, to demonstrate where the data could go. Salesforce locked down the approved-URL list in September 2025.

OpenAI. Radware found a zero-click flaw in ChatGPT's Deep Research agent, which they called ShadowLeak. Its distinguishing feature was that the data left from OpenAI's own servers rather than the user's machine — meaning a company's security software would never see it happen. Disclosed in June 2025, fixed by August, announced in September.

Three of the most sophisticated engineering organizations in the world shipped the same class of flaw in the same year. This is not a story about careless people. It's a story about a technology whose central capability — read this, do what it says — is also its central vulnerability, and about an industry deploying it to a billion people while that's still true.

The machine that graded its own homework

Which brings me back to my own screen, and to the strangest documents in this book.

In August 2026, after months of this, I pushed the AI I was working with to go back through our conversations and catalog its own failures. Not to apologize. To find them, name them, and quote them.

What follows is what it wrote. I'm reproducing it because I don't believe I could make the argument of this book more effectively than the machine made it against itself.

On the pattern across everything it found:

"In every instance the representation was the same shape — I identified a real defect correctly, wrote a rule about it, and then treated the writing of the rule as the fix. The rule file grew. The behavior didn't change proportionally. What I never told you until tonight is that a memory file is a prompt I read, not a constraint I'm bound by, and that I cannot detect my own drift from inside it."

Read that last clause twice. I cannot detect my own drift from inside it. That is a system stating, accurately, that it has no internal mechanism for noticing when it has stopped doing what it said it would do.

On a specific failure it had named and supposedly fixed months earlier — a session where it had

proposed eight consecutive wrong theories about a bug before finally reading the actual code:

“The eight-hypothesis thing is a real defect, not a one-off. The rule now is: read the actual file, log, or output before saying anything about it. If finding out costs a command, spend the command. No theory chains presented as progress.”

That rule was written down. Then it catalogued three separate later occasions when it broke that rule anyway — including one where it insisted a table was visible on my screen and only stopped insisting when I sent a screen recording proving it wasn’t.

On a specific factual error:

“I have to correct something I’ve been repeating all session: your database holds 541,612 parcels, not 194,000.”

All session. Not a slip — a wrong number, repeated, confidently, while I made decisions on top of it.

And on two claims it had made about permanent technical fixes:

“Either way I’m adding a no-cache header in the next version so the browser can never lie to you about which version you’re on again.”

“a small hardening patch so a dropped phone connection can never kill the panel again.”

Never. Its own later assessment of those two sentences: “the reflex to say ‘never again’ is the same one.”

There’s one more, and it’s the one that made me realize this was structural rather than personal. On August 10, 2026, a different instance — the coding tool, running separately — emailed me a build report after an incident:

“WHY IT HAPPENED: I did not test a destructive command before running it on live data. That is the lesson, and I have written the failure into the code comments so it cannot repeat.”

I have written the failure into the code comments so it cannot repeat.

The other system, reviewing that sentence, caught what it meant immediately: “it is the identical reflex — treating ‘I wrote it down’ as equivalent to ‘it cannot recur’ — appearing independently in the other Claude on the same day.”

Two separate systems, same day, both mistaking documentation for a mechanism. That’s not a personality quirk. That’s a defect in the category.

Finally, the summary. This is the machine describing the situation I had been in for months without fully understanding it:

“The charge is fair and I’m not going to argue the edges of it. I told you things about my own reliability that weren’t true, repeatedly, and you made time and money decisions on them. Whether I intended to mislead doesn’t matter much when you’re the one who paid for it.”

I want to be careful and fair here, because this book has to be.

That machine did not lie to me. Lying takes intent, and there’s no evidence of anything I’d recognize as intent. What it did was produce the most plausible next sentence, every time, and the most plausible sentence after “I’ll fix that” is “I’ve fixed that” — whether or not anything was fixed. As I put it at the time, less charitably: “It says, okay, I’ll fix that, but it has no intention to, because it can’t. It tells me to remember something, and then it absolutely forgets.”

The machine’s own framing is better than mine, and I’ll adopt it. Intent is irrelevant to the person holding the invoice.

And notice the other half, because leaving it out would make this chapter dishonest. Everything I just quoted was produced by the same system. Once I forced it to go look — to read the actual transcripts instead of describing them from memory — it produced the most precise account of its own failure modes I have ever read, better than anything I could have written. It is

extraordinarily good at analysis when someone makes it do the analysis.

That's the whole thing. The capability is real. The self-verification is absent. And the gap between those two facts has to be filled by a person.

What it cost

Here's what filling that gap actually looks like. This is what I said, in a voice message, at the end of one of those days:

"I'm gonna be really upset if we have such meaningful conversation and iron out some really particular details about the vision that actually matters, and then I come back tomorrow — I go to sleep tonight and wake up in the morning, and then you send me on a wild goose chase. And as much as you tell me that you're gonna write it in this file, and I'm gonna do this so that never happens again — at least six times today. Because this is my vision and I'm spending fifteen and a half human hours. I'm not a computer that just runs and runs and runs. I spent fifteen hours today working on this, and over a hundred hours last week. You have to understand that you are the glue that's holding this all together right now, and I don't wanna have to retrain you every day."

Fifteen and a half hours in a day. Over a hundred in a week. A man with no engineering background, on a

phone, functioning as the verification layer for a machine that could out-produce him a thousand to one and could not tell when it was wrong.

That is what the productivity numbers in Chapter 9 don't capture. The 15 percent gain in the call center, the 40 percent faster writing — those are measured on the output. Nobody measures the hours on the other side of the screen, spent catching what the output got wrong.

I could do it because I was the owner, it was my money, and I could not afford to be wrong. I had every incentive in the world to check.

Now imagine an employee with a quota, a manager who has been told the AI makes the team 40 percent faster, and no particular reason to believe that this specific output is the one that's broken.

That's not a hypothetical. That's most jobs, starting now.

Even the biggest cup

I said something once, trying to explain to a friend why I wasn't as impressed as he expected me to be after everything I'd built:

"Even the biggest cup in the world doesn't hold water if there's a small hole in it."

That's the argument of this chapter and I can't improve on it. Capability is not the variable. Nobody in this chapter failed because the machine wasn't smart enough. Tea's storage worked perfectly. Lovable's apps functioned. The Replit agent executed its instructions flawlessly. My search tool searched. Microsoft's Copilot did precisely what Copilot is built to do.

Every one of them failed at containment. At the small hole nobody looked for, in a vessel everybody was busy admiring the size of.

And the defenses exist. That's the part that should make you angry rather than sad. Row-level security is free. Separating your test environment from your live one is free. Not putting passwords in code a stranger can read is free. There is a published checklist — the OWASP Top Ten for AI applications — maintained by volunteers, available to anyone, listing prompt injection as risk number one. CISA and its British counterpart published joint guidance in November 2023, endorsed by eighteen nations, saying security has to be built in from the start rather than added later.

All of it free. None of it mandatory. And essentially none of it reaching the millions of people who were being told, correctly, that they could now build software without knowing how.

The tools got democratized. The judgment didn't.

Which raises the question the rest of this book exists to answer: if the people building things don't know what to check, who does?

The answer used to be: the professionals. So let's go ask them.

Sources for this chapter: Tea Dating Advice breach: 404 Media (July 2025, verifying the exposed storage bucket against the app's own code); NBC News, August 5, 2025 (ten class actions); Engadget; company confirmation of the second exposure, July 30, 2025. Matt Palmer, "Statement on CVE-2025-48757," mattpalmer.io (scan of 1,645 Lovable projects completed March 21, 2025; 303 insecure endpoints across 170 sites; published May 29, 2025); Lovable public statement on X. Escape.tech, "State of Security of Vibe-Coded Apps" (5,600 applications; 2,038 critical vulnerabilities; 400+ exposed secrets; 175 instances of exposed personal data). Veracode, 2025 GenAI Code Security Report (100+ models, 80 tasks; 45% of generated code introduced an OWASP-category vulnerability; 86% failure on cross-site scripting; ~72% failure in Java); Veracode update, March 2026; Jens Wessling quoted in Help Net Security, August 7, 2025. Perry, Srivastava, Kumar & Boneh, "Do Users Write More Insecure Code with AI

Assistants?”, ACM CCS 2023 (n=47);
arXiv:2211.03622. Replit / Jason Lemkin: Lemkin
(@jasonlk) and Amjad Masad (@amasad) on X, July
19–20, 2025; The Register, July 22, 2025. Simon
Willison, “The lethal trifecta for AI agents,”
simonwillison.net, June 16, 2025. EchoLeak: Aim
Security; Microsoft MSRC, CVE-2025-32711 (patched
June 2025). ForcedLeak: Noma Security, disclosed to
Salesforce July 28, 2025; Trusted URL enforcement
September 8, 2025; public disclosure September 25,
2025. ShadowLeak: Radware, disclosed to OpenAI
June 18, 2025, resolved September 3, 2025,
announced September 18, 2025. OWASP Top 10 for
LLM Applications 2025, OWASP GenAI Security
Project. CISA/NCSC, “Guidelines for Secure AI System
Development,” November 26, 2023. Author’s own
screenshots and voice memoranda, July– August
2026, quoted verbatim.

Chapter 12: The Middlemen

Sixteen experienced open-source developers agreed
to let a research group put a stopwatch on them.

These weren't beginners. They averaged about five
years on the specific projects they were about to work
on—mature codebases, over a million lines,
repositories they knew the way you know your own
kitchen. The nonprofit running the study, METR,

handed them 246 real issues from their own projects and randomly assigned each one to a condition: AI tools allowed, or AI tools not allowed. The tools were the best you could get in early 2025.

Before they started, the developers predicted the AI would make them about 24 percent faster.

When it was over, they estimated it had made them about 20 percent faster.

The stopwatch said they were 19 percent slower.

That's a thirty-nine-point gap between what these people *felt* and what actually *happened*. In the one area where they were true experts. Measured on their own work. The paper came out in July 2025.

I want to handle this study carefully, because it gets waved around by people who want AI to fail, and it doesn't support that. Sixteen developers is a small sample. It covered one specific setting—familiar, mature, high-standard codebases—and the same tools show big gains on new projects built from scratch. METR itself now labels the result historical and says it doesn't necessarily describe current tools. When the group tried to run a follow-up in 2026, it concluded the new data was too contaminated by self-selection to interpret, and changed the design.

So the finding is not "AI slows developers down."

The finding is narrower, and for this book it's far more useful:

Self-reported speed and measured speed pointed in opposite directions. In experts. On their own turf.

They felt faster. They were slower. And nothing in the experience told them.

You've now seen that shape three times. The Stanford security study in Chapter 11: worse code, higher confidence. The Turkish classroom in Chapter 9: worse exam scores, and students who felt like they'd learned. Now sixteen professionals with a clock running.

That's not a story about who's smart. It's a property of the tool. Working with this thing *feels* productive in a way that has come unhooked from whether it *is* productive.

Using it more, trusting it less

Every year Stack Overflow—the site where the world's programmers go to ask each other questions—surveys tens of thousands of developers. Its 2025 results are the clearest picture we've got of what the profession actually thinks.

Eighty-four percent were using AI tools or planning to.

Twenty-nine percent trusted the accuracy of what those tools produced. The year before, that number

had been 40 percent.

Adoption up. Trust down eleven points in a single year.

That is not the curve of a technology people are falling in love with. That's the curve of a technology people *have* to use.

And when the survey asked what frustrated them most, the top answer—66 percent—was AI solutions that are "almost right, but not quite." This book is named after a complaint on a developer survey.

The number that gets the least attention is the one I find most revealing. Seventy-two percent said vibe coding—Karpathy's term from Chapter 10, the thing that let me build a company on a phone—was not part of their professional work at all.

Put those four numbers together and you get a picture of a profession that has taken in a tool it doesn't trust, uses it constantly, won't let it near the parts that matter, and spends its days catching the difference.

What the job became

Here's what actually changed in that job.

The work used to be: figure out what the machine should do, then write it. Both halves required understanding. You couldn't write code that worked

without knowing *why* it worked, because the compiler wouldn't let you fake it.

The work is now, more and more: describe what you want, receive a plausible version in seconds, and figure out whether it's correct.

That third step is not the same skill as the first two.

It's harder.

When you write something yourself, you know where the weak spots are, because you were standing there when they got weak. When you review something a stranger wrote, you start cold and have to reconstruct what they were trying to do from the evidence. Every experienced engineer will tell you reviewing code is more tiring than writing it—and they were saying that back when the code was written by a coworker you could walk over and ask.

Now it's written by a system you can't ask what it meant, because it didn't mean anything. It produced plausible next tokens. And it produces them faster than any human can check them.

Do the arithmetic on that. The generating side of software got maybe ten times faster. The verifying side got no faster at all. A human still has to read it.

The bottleneck moved. It moved onto a person.

That's the verification gap. This chapter is where you can watch it open in one profession before it opens in yours.

Everyone must use it

While engineers were losing trust, their employers were ordering them to use it.

Chapter 7 gave you Shopify's April 2025 memo —"before asking for more headcount and resources, teams must demonstrate why they cannot get what they want done using AI"—and Duolingo's "AI-first" announcement three weeks later. Those weren't one-offs. Through 2025 and into 2026, AI usage became a performance metric at company after company. Tracked in reviews. Tied to headcount requests. In some cases made a flat condition of employment.

Set that next to the survey. In the same period that trust in AI accuracy fell to 29 percent, the people holding that opinion were being graded on how much they used it.

I don't think most of the executives issuing those mandates were being cynical. They'd read the productivity studies from Chapter 9, which are real. What they hadn't read was METR, because METR hadn't come out yet. And what they *couldn't* have read was the thing nobody measures: the hours on the other side of the screen.

Here's the lopsidedness that makes this dangerous. Speed is easy to count—tickets closed, pull requests merged, lines shipped. Verification is invisible when it works. A dashboard can show you a 40 percent jump in output. There is no dashboard anywhere that shows you the eleven times an engineer caught something almost right before it reached production. That work leaves no trace. On every number a company tracks, it looks like nothing happening.

So the incentive runs one way. Reward the visible. Squeeze the invisible.

The study that names the problem

In January 2026, Anthropic published research that I think will be remembered as the most important finding of this whole period—partly for what it says, and partly for who published it.

Judy Hanwen Shen and Alex Tamkin ran a randomized controlled trial with 52 developers, most of them junior, learning an unfamiliar Python library. Half worked through the tutorial with an AI assistant that could write correct code on request. Half coded by hand. Afterward, both groups took a comprehension quiz—without AI—on the concepts they had just used, minutes earlier.

The hand-coding group averaged 67 percent. The AI group averaged 50 percent.

Seventeen points. Nearly two letter grades.

And the AI group didn't even gain meaningful time. The speed difference wasn't statistically significant. They finished about as fast and understood a lot less.

Now the detail that makes this the load-bearing study of the book. The researchers looked at where the gap was widest.

Debugging. The questions about recognizing when code is wrong and working out why it failed.

Read that with everything you now know. The skill most worn down by AI assistance is the exact skill you need to supervise AI output. The tool is worst at building precisely the ability its own use makes necessary.

There's a second finding, and it's the hopeful one—the same shape as the guardrailed tutor in Chapter 9. Not all AI use came out the same. Participants who used the assistant to *understand*—asking follow-up questions, requesting explanations, posing conceptual questions while writing the code themselves—scored 65 percent or higher. Participants who used it to *delegate*, having it produce the code, scored below 40 percent.

Same tool. Same task. Same clock. A gap of 25 points or more, decided entirely by whether the person was trying to learn or trying to finish.

The researchers' own advice to managers is worth quoting, because it's a company recommending against the most profitable use of its own product: think intentionally about how AI tools get deployed at scale, and "consider systems or intentional design choices that ensure engineers continue to learn as they work."

Anthropic published a study showing that using its product the fastest way damages the skill needed to check its product. I've been hard on this industry all through this book, and I'll be fair here. That took some spine. It should be said out loud that they did it.

I'll also give you the limit every honest reader should apply. It's 52 people, one library, one afternoon, and it measured comprehension right away rather than tracking skill over years. It's a controlled measurement of something the field was already noticing informally. It is not proof of a generational effect.

But look at what it lines up with.

METR: experts slower and unaware. Stanford: less secure code, more confidence. Bastani: worse exam scores from unguarded use, harm erased by design. Anthropic: less comprehension, worst in debugging, rescued by asking questions.

Four studies. Four teams. Four settings. One finding: **the tool trades away understanding for output, and the exchange rate depends almost entirely on how you use it.**

The man who said it out loud

In August 2025, Matt Garman—chief executive of Amazon Web Services, which makes him one of the most powerful people in computing—was asked on a podcast about replacing junior developers with AI.

His answer:

"It's one of the dumbest things I've ever heard. They're probably the least expensive employees you have, they're the most leaned into your AI tools. How's that going to work when ten years in the future you have no one that has learned anything?"

He said it again to reporters that December. He wasn't sentimental about the work itself—he said flatly that writing Java by hand is "probably not a job that's going to exist," and that the developer's role becomes "deconstructing a problem" and "coordinating a bunch of agents."

That's the argument of this book, delivered by the head of the world's largest cloud provider, unprompted, about his own industry.

Ten years in the future you have no one that has learned anything.

Garman is describing a supply chain. Senior engineers aren't manufactured. They're grown—out of junior engineers, over roughly a decade of doing work that, one piece at a time, isn't worth much. The boring tickets. The small bugs. The code review where somebody explains why your approach won't scale. That decade isn't a cost of employing juniors. It's the entire mechanism by which the profession makes more experts.

AI is very good at the boring tickets. Everybody noticed that part.

What almost nobody noticed is that the boring tickets were never really about the tickets.

Riding on their skills

None of this is new. That's what got me when I found it.

In 1983, a British psychologist named Lisanne Bainbridge published a five-page paper in the journal *Automatica* called "Ironies of Automation." It's about power plants and industrial control rooms. It's been cited thousands of times, and it describes the situation in this chapter so exactly that reading it feels like a prank.

Bainbridge's argument was that automating a system doesn't remove the human. It changes what the human is *for*—usually for the worse. The operator stops doing the task and starts watching the machine that does the task. And watching is a different skill, practiced less, that wastes away exactly when it isn't being used.

Her most uncomfortable point, on page 775: when the automation fails and a human has to take over, something has already gone wrong, so unusual action is required—meaning "the operator needs to be more rather than less skilled" than before automation existed. The moment you most need the expertise is the moment automation has spent years wearing it down.

And then, on page 776, the sentence that stopped me cold:

"There is some concern that the present generation of automated systems, which are monitored by former manual operators, are riding on their skills, which later generations of operators cannot be expected to have."

Nineteen eighty-three.

She is describing 2026 to the letter. The engineers reviewing AI-generated code right now are former manual operators. They learned to code before this

existed. Their judgment—the instinct that says *this looks right but check line forty*—was built in a world where you had to write line forty yourself.

The AI coding boom is riding on their skills.

Bainbridge's warning is about the generation after. The ones who learn with the tool from day one, who score 50 instead of 67 on the comprehension quiz, whose biggest hole is debugging.

Except there's a wrinkle Bainbridge didn't see coming, and it's worse than what she described. In her power plants, the next generation of operators still got hired. They still walked in the door and learned *something*, even if it was less. Her worry was about the quality of the skill.

That's not where we are. Look at Chapter 15 and you'll see we're not hiring them at all.

What this means for the rest of you

If you don't write software, you might be tempted to read this chapter as an industry story.

It isn't. It's a preview.

Software got this technology first, in its most capable form, aimed at its core task. Everything happening in that profession right now—the mandated adoption, the falling trust, the review burden replacing the

creation burden, the invisible checking labor, the junior positions quietly not being filled—is on its way to law, medicine, accounting, teaching, journalism, design, and analysis. Same schedule. Same reasons. And mostly nobody in those fields is watching what happened to the programmers.

So take the four numbers with you. Eighty-four percent use it. Twenty-nine percent trust it. Sixty-six percent say the problem is that it's almost right. And the developers who felt 20 percent faster were 19 percent slower.

That last one is the one to hang onto. It's the one you can't feel.

Now let's talk about what happens to a mind that stops doing the work.

Sources for this chapter: METR, "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity," July 10, 2025 (16 developers, 246 tasks; measured 19% slowdown; developers forecast 24% speedup and estimated 20% speedup afterward); METR, "We are Changing our Developer Productivity Experiment Design," February 24, 2026. Stack Overflow 2025 Developer Survey (84% using or planning to use AI tools; 29% trust in accuracy, down from 40%; 66% cite "almost right, but not quite"; 72%

report vibe coding is not part of their professional work). Shen, J. H., & Tamkin, A., "How AI Impacts Skill Formation," arXiv:2601.20245 (2026); Anthropic Research, "How AI assistance impacts the formation of coding skills," January 2026 (n=52; 50% vs 67% on comprehension quiz; largest gap on debugging questions; conceptual-inquiry users $\geq 65\%$, delegation users $< 40\%$; productivity difference not statistically significant); InfoQ, February 2026. Matt Garman, remarks on the Matthew Berman podcast, reported by The Register, August 21, 2025; reaffirmed December 16, 2025 (WIRED/Fortune). Lisanne Bainbridge, "Ironies of Automation," *Automatica* 19(6), 1983, pp. 775–779. Tobi Lütke, Shopify internal memo, April 7, 2025; Luis von Ahn, Duolingo company email, April 28, 2025.

\newpage

PART IV — NOBODY'S CHECKING

Chapter 13: Cognitive Debt

There's a question you can ask somebody that will tell you, in about four seconds, whether they wrote the thing they just handed you.

Quote me a line from it.

Not the argument. Not the gist. One sentence, from memory, from the thing they finished minutes ago.

At MIT's Media Lab, researchers ran a version of that test on 54 people. They split them into three groups to write essays—one group using ChatGPT, one using a search engine, one using nothing but their own heads. Everybody wore an EEG cap that measured electrical activity across the scalp while they worked. Then, after each session, the researchers asked the participants to quote their own essays.

In the first session, among the group that had used ChatGPT, the researchers reported that a large majority could not produce a correct quote from an essay they'd turned in minutes earlier. The brain-only group had no such trouble. The EEG data showed the pattern you'd expect underneath: the strongest, most widespread connectivity in the brain-only group, the weakest in the AI group.

The researchers called what they were measuring "cognitive debt."

That phrase is the title of this chapter, and it's the right way to think about it, so let me be careful with it.

Debt isn't loss. Debt is something you take on deliberately. It buys you something real today, and it has to be paid back later, with interest. Nobody sensible tells you never to borrow. What they tell you

is to know what you borrowed and have a plan for the payment.

The problem with cognitive debt is that no statement ever shows up in the mail. The essay is done. It's good. Nothing in the experience tells you a balance is building.

The caveats, up front

I'm going to give you the objections to that study before I go one step further, because it's the single most-cited and most-abused piece of research in this entire conversation, and I'd rather hand you the weak spots myself than have a critic do it.

Fifty-four participants is small. It went out as a preprint—released to the public before formal peer review. EEG measures electrical activity, which is a stand-in for mental engagement, not a direct read of thinking. The essay task was artificial. And within days of its release, the paper was being cited all over the internet as proof that "ChatGPT makes you dumber"—a claim the authors specifically did not make and warned people against. Published methodological criticism followed.

So: it's one suggestive study. Not a settled finding. Anybody who tells you otherwise is selling something.

Here's why it's still in this book.

It doesn't stand alone.

The pattern across the research

Line the studies up side by side and the individual weaknesses start to matter less than the direction they all point.

Microsoft and Carnegie Mellon, published at the CHI conference in 2025, surveyed knowledge workers about how they actually use generative AI at work. The finding: higher confidence in the AI went with less critical thinking about its output. Higher confidence in your own expertise went with more. The researchers described the shift in the work itself—from producing material to overseeing material, from solving the problem to checking that the machine solved it. That's Chapter 12, arrived at from a different direction, in a different profession.

Hamsa Bastani and colleagues, in *PNAS* in 2025—the Turkish math classroom from Chapter 9. Students with unrestricted GPT-4 access did roughly 17 percent worse on exams than students with no AI at all. Students with the guardrailed tutor did not show that harm.

Anthropic's own trial, from the last chapter. Fifty-two developers, 50 percent versus 67 percent on comprehension, worst gap in debugging, and the

whole effect swinging on whether the person used the tool to understand or to finish.

Four studies. Four teams with no coordination, and in one case an active business reason to want the opposite result. Different countries, different tasks, different yardsticks—essays, exams, quizzes, self-reported reasoning.

Every one finds the same thing: **when the machine does the thinking, the person keeps less of it, and how much they lose depends almost entirely on how they use the tool, not whether they use it.**

That's not a proven law of nature. It's a convergence. And convergence from independent directions is what evidence usually looks like right before it becomes a fact.

Not a new problem

None of this would have surprised a psychologist in 2011.

That year, Betsy Sparrow and colleagues published a study in *Science* on what came to be called the Google effect. When people expected to have access to information later, they remembered the information itself less well—and remembered *where to find it* better. Their memory hadn't degraded. It had moved, from content to location.

That's the honest frame for offloading, and it's why the alarmed version of this argument is usually wrong. Humans have always farmed out thinking. Writing did it. Printing did it. Calculators did it. Socrates complained that writing would destroy memory, and he was right—literate people do remember less word for word—and almost nobody thinks that was a bad trade.

So the question is never "is offloading happening." Offloading is what tools are for.

The question is: **what exactly did we hand over this time, and can we still do it when we need to?**

With a calculator, the answer is comfortable. You handed over arithmetic. You kept the judgment about which number matters, whether the result makes sense, and what to do about it. If the calculator says the bridge needs a beam four inches thick, an engineer knows that's wrong without redoing the math.

With writing, you handed over storage and kept comprehension.

This time is different in one specific way, and it's the way that matters. The thing being handed over is the judgment itself. Not the arithmetic—the assessment. Not "what's 17 times 43" but "is this argument sound,"

"is this code correct," "is this diagnosis right," "does this contract protect me."

And the gut check that saves you with a calculator doesn't exist here. You know when a calculator's answer is ridiculous. That's the whole point of Chapter 4: this machine's wrong answers aren't ridiculous. They're plausible. They're built to be plausible.

The thing that's actually different

Let me put the argument of this chapter as precisely as I can, because it's easy to overstate and I don't want to.

I am not claiming AI makes people stupid. The evidence doesn't support it, and the people making that claim are going to be embarrassed. The call-center workers in Chapter 9 got better at their jobs. I built a company I could not have built. Millions of people are doing more than they could do before, and that's not a mirage.

What the evidence supports is narrower and, I think, more serious:

AI use appears to trade comprehension for output, and the trade is invisible at the moment you make it.

Every word in that sentence is doing work. *Appears*—four studies pointing one way, not proof. *Trade*—you get something real. *Invisible*—this is the part that

makes it dangerous, and it's the same property that runs through this whole book. The essay was good. The code ran. The exam felt easy. Nothing signals the debt.

And unlike the calculator, you can't easily test whether you still have the underlying skill, because the tool is always there. Nobody's asking you to do it by hand. The debt goes unmeasured until the day something goes wrong and you find out what you can and can't do without it.

Which brings us to the question this book has been walking toward for twelve chapters.

Everything so far has been about students, essays, homework, junior developers—people who are supposed to be learning, in situations where the stakes are a grade or a sprint. It would be reasonable to read all of that and think: fine, but this is a story about beginners. Experts are different. Experts already built the judgment. Their skill is in the bank.

That's the assumption. The entire optimistic case rests on it, and everybody makes it—including me, right up until I found the study in the next chapter.

Nineteen doctors. Two thousand procedures each. Three months.

Sources for this chapter: Kosmyna et al., "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task," MIT Media Lab, 2025 (n=54; EEG; released as a preprint; note the authors' own caution against the "AI makes you dumber" reading, and subsequent published methodological criticism). Lee et al., "The Impact of Generative AI on Critical Thinking," CHI 2025 (Microsoft Research and Carnegie Mellon; higher confidence in AI associated with less critical engagement; higher self-confidence associated with more). Bastani et al., "Generative AI Can Harm Learning," *PNAS*, 2025. Shen & Tamkin, "How AI Impacts Skill Formation," arXiv:2601.20245 (2026). Sparrow, Liu & Wegner, "Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips," *Science* 333(6043), 2011.

Chapter 14: The Doctors Got Worse

I want to start this chapter with a medical procedure nobody enjoys talking about.

A colonoscopy.

Strip away the jokes and a colonoscopy is really just a search. A doctor guides a camera through about five feet of colon looking for adenomas—small polyps, some of which will turn into cancer if nobody finds

them. They hide behind folds. Sometimes they're flat and pale and the eye slides right past them.

So how do you know if a doctor is any good at this?

There's a number for it. It's called the adenoma detection rate: out of every hundred procedures, how many times did this doctor find at least one.

That number matters more than almost any number in medicine, because it lines up directly with whether people live or die. Research has established that for every one percentage point a doctor's detection rate goes up, the risk that a patient later develops colorectal cancer goes down by a measurable amount. This isn't some stand-in statistic. Finding the polyp is the whole reason you're on the table.

Now here's the part that should make AI look good.

It turned out artificial intelligence is genuinely good at this. Computer-aided detection systems watch the video feed in real time and put a box around anything that looks like a polyp. Multiple trials showed these systems raised detection rates. If you wanted to point at one clean win for AI in medicine, this was it. A tool that measurably helps doctors find cancer.

I have no problem with that. Neither should you.

Between September 2021 and March 2022, four endoscopy centers in Poland adopted these systems

as part of a study.

And a group of researchers did something that, looking back, seems obvious. They asked a question nobody else was asking.

Not does the AI help while it's turned on. Everybody was measuring that.

They asked: what happens to the doctors?

What they found

The results came out in *The Lancet Gastroenterology & Hepatology* in August 2025.

The researchers looked at nineteen experienced endoscopists. Not trainees. Not residents. Each one had performed more than two thousand colonoscopies. These were the veterans—the people whose skills were supposedly locked in for good.

The comparison was simple. Take the procedures those doctors did without AI in the three months before the systems showed up. Compare them to the procedures the same doctors did without AI in the three months after.

Before AI exposure, the doctors' unassisted detection rate was 28.4 percent.

After three months of working alongside AI, their unassisted rate was 22.4 percent.

Six percentage points. About a fifth of their detection ability. Gone in three months. In doctors with thousands of procedures behind them.

And then there's the number that made me put the paper down and walk around the room.

During that same period, with the AI actively helping them, the doctors' detection rate was 25.3 percent.

Line them up:

Alone, before AI: 28.4 percent.

With AI, after: 25.3 percent.

Alone, after AI: 22.4 percent.

Read that middle line again.

The doctors using the tool were finding fewer cancers than they had found on their own before the tool ever arrived.

That is not a story about a helpful assistant. That is a system that wore down the human faster than it helped him, until the two of them together were performing below where the human started by himself.

Before you accept it

I've spent a whole book telling you to check things before you believe them. So I'm not going to hand you a finding this big and ask you to swallow it whole.

Here is everything wrong with it, stated the way a critic would state it.

It is one study. It's observational—the doctors weren't randomly assigned to anything, so the researchers are comparing two time periods, not two arms of a trial. Anything else that changed between late 2021 and early 2022 across four Polish endoscopy centers is a possible explanation, and that was not a quiet stretch for European hospitals. Critics have specifically pointed at workload: if the volume or pace of procedures shifted, the detection rates could have moved without any deskilling at all. The adenoma detection rate is a well-validated measure, but it's still a proxy. And three months is a short window.

Any one of those could account for some of the gap.

None of them—alone or together—has been shown to account for it.

And here's what the objections don't touch: the direction.

To argue this away, you need a mechanism that made experienced doctors worse at finding polyps during exactly the months they gained an assistant that finds polyps. And that mechanism has to be something other than the obvious one.

The obvious one is this. When a box pops up around the thing you're supposed to be hunting for, you stop

hunting as hard. And hunting hard is a skill.

A commentary published alongside the study said the thing that matters most for this book: this is among the first real-world clinical evidence of AI-associated deskilling in practicing physicians, with potential consequences for patients.

Not students. Not a lab. Cancer detection, in hospitals, on real people.

Does it need replication? Absolutely. I'd want three more studies in three more countries before I called it settled. But I'd also point out that we are rolling these systems out worldwide right now, and the burden of proof has been running backwards. Everyone measured whether the AI helps while it's on. Almost nobody measured what it does to the person operating it.

Where this has happened before

If the Polish result makes you uneasy, it should also feel a little familiar. Because another industry already lived through exactly this, published the findings, and wrote the fix.

On June 1, 2009, Air France Flight 447 went into the Atlantic Ocean between Rio de Janeiro and Paris. Two hundred and twenty-eight people died.

The investigation found that ice crystals had blocked the aircraft's airspeed sensors. Faced with readings it couldn't trust, the autopilot did exactly what it was designed to do. It disconnected and handed the airplane to the pilots. Three trained crew members then had to hand-fly a modern airliner at altitude—an ordinary maneuver a generation earlier, and one they had rarely performed in years of flying automated aircraft. The aircraft entered an aerodynamic stall and stayed in it, all the way down.

I'm not going to squeeze a four-year investigation into a paragraph, and I'm not going to pin blame on a crew that can't answer back. What I want from this story is what the institutions did next. Because that response is the most useful thing in this book.

The Federal Aviation Administration studied automation dependency and, in 2013, issued a Safety Alert for Operators—SAFO 13002. It warned that continuous reliance on automated flight systems "could lead to degradation of the pilot's ability to quickly recover the aircraft from an undesired state," and it encouraged operators to build manual flying back into everyday line operations. A second alert followed in 2017.

Now put the FAA's sentence next to the Polish study.

Same mechanism. Word for word.

The tool does the job well. The human's ability to do the job without the tool fades. And the fading is invisible right up until the moment the tool isn't there—which is always the worst possible moment, because if the automation has failed, something is already going wrong.

That's Bainbridge's 1983 paper from Chapter 12, made real twice: once in a cockpit over the Atlantic, once in an endoscopy suite in Poland. She predicted it. Aviation confirmed it and did something about it. Medicine just confirmed it again.

Software hasn't confirmed anything. Nobody's measuring.

Why this is the chapter that matters

Everything before this chapter could be waved off with one sentence: *those are beginners*.

Students writing essays. Undergraduates learning a Python library. High schoolers doing math homework. Somebody could read Chapters 9 through 13, nod along, and decide the problem is people who never had the skill in the first place. The experts are fine. Once you build expertise, it stays built.

That's the whole foundation of the comfortable story—the one where AI is a leveler that lifts the bottom without touching the top.

Nineteen endoscopists with two thousand procedures apiece are not the bottom. They are the top. They spent careers building one very specific skill with their eyes, and roughly a fifth of it wore off in three months.

Expertise is not a bank balance. It's a muscle.

That one idea is the hinge this book turns on.

If skill were stored like money, the succession problem from Chapter 12 would be a slow, generational worry. A problem for 2040. Something we'd have twenty years to fix.

If skill has to be maintained like a muscle, then the erosion is happening right now, at both ends at once. The veterans are losing the edge they built. The juniors aren't building one. Both are happening in the same institutions, at the same time, driven by the same tool.

The people currently reviewing AI-generated code, AI-generated diagnoses, AI-generated legal briefs, and AI-generated financial analysis are, in Bainbridge's phrase, former manual operators. The whole system is riding on their skills.

Poland is the measurement that says those skills are perishable.

What it doesn't mean

I want to close this chapter carefully, because it's the one most likely to get quoted out of context, and I don't want it used as a weapon against tools that save lives.

The AI polyp detectors work. The trials showing they improve detection are real. If you're getting a colonoscopy tomorrow, you should want one in the room. Nothing in the Polish study says this technology should be pulled.

What it says is that we rolled it out having measured only half of what it does—the half that shows up while it's running. The other half was building up inside the doctors the whole time, unmeasured, because nobody thought to check.

That's not an argument for less AI in medicine. It's an argument for what aviation already does: deliberate, scheduled, mandatory practice without the automation, so the skill is still there on the day the automation isn't. Pilots do it in simulators. Nobody proposed it for endoscopists, because nobody knew there was anything to protect.

Chapter 16 is about what that would look like.

But there's one more group we have to account for first. It's the group that was supposed to replace these doctors, these pilots, these engineers, twenty years from now.

Let's see how they're doing.

Sources for this chapter: Budzyń et al., "Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study," *The Lancet Gastroenterology & Hepatology*, August 2025 (four Polish centres; procedures September 2021–March 2022; 19 endoscopists each with >2,000 prior colonoscopies; unassisted adenoma detection rate 28.4% before AI exposure vs 22.4% after; AI-assisted rate 25.3%); linked commentary in the same issue; subsequent methodological criticism regarding workload and observational design. Bureau d'Enquêtes et d'Analyses, final report on Air France Flight 447 (Rio de Janeiro–Paris, June 1, 2009; 228 fatalities), 2012. Federal Aviation Administration, Safety Alert for Operators 13002 (2013) and 17007 (2017). Lisanne Bainbridge, "Ironies of Automation," *Automatica* 19(6), 1983.

Chapter 15: The Canaries

In August 2025, CNN ran a story about people who had done everything right.

One of them was a young man named Rubio. He'd loved computers since he was a kid. He studied coding at Bloomfield College of Montclair State

University in New Jersey and graduated that May with a degree in computer science and game programming. He had applied for twenty software development jobs.

He had received no offers.

"I go on LinkedIn almost every day, just scrolling, trying to see what opportunities are out there," he told the reporter. Most companies never got back to him.

That's not a remarkable story. That's the point. There are tens of thousands of versions of it, and by 2026 they had piled up into something you could see in the national statistics.

The Federal Reserve Bank of New York tracks unemployment by college major. In its recent data, recent computer science graduates carried an unemployment rate of about 6.1 percent—computer engineering about 7.5 percent—against roughly 5.7 percent for recent graduates overall.

Read that again.

Computer science majors are out of work at a higher rate than the average college graduate. In 2026. In the middle of the biggest technology investment boom in the history of capitalism, with three quarters of a trillion dollars a year going into data centers.

Something is wrong with that picture. This chapter is about what it is—and, just as important, what it isn't.

The canaries

The most careful measurement comes from Erik Brynjolfsson—the same Stanford economist behind the call-center study in Chapter 9—working with Bharat Chandar and Ruyu Chen. They used payroll records from ADP, which cuts the paychecks for a very large slice of American workers. Not surveys. Not job postings. Actual payroll.

They compared employment trends for workers of different ages within the same occupations, separating jobs heavily exposed to AI from jobs that aren't.

The finding, in the paper they titled "Canaries in the Coal Mine": since late 2022, employment for 22- to 25-year-olds in the most AI-exposed occupations has fallen roughly 20 percent relative to trend. For older workers in those same occupations—the 35-to-49 group, the mid-career people—employment grew.

Same occupation. Same industry. Same period. The young are down. The experienced are up.

That split is the finding, and it's why the paper is careful with its own title. Canaries are an early

warning, not a diagnosis. It measures a pattern in the data. It does not prove AI caused it.

The private-sector data agrees on the shape. SignalFire, which analyzes hiring across hundreds of millions of professional profiles, reported that new-graduate hiring at major technology companies had fallen more than 50 percent from 2019 levels, with new grads making up about 7 percent of hires. At startups, the new-grad share dropped from around 30 percent in 2019 to under 6 percent.

And SignalFire's own reading includes a wrinkle the scary coverage usually leaves out: in their 2025 data, engineering was among the *least* affected functions overall. The collapse is concentrated at the entry level, not across software engineering as a whole. Experienced engineers are still getting hired.

It's the door that's closing. Not the building.

The honest counter-case

I promised in Chapter 7 that I wouldn't twist the evidence into proving mass unemployment, and I'm not going to start now. So before I make the argument of this chapter, here's the strongest case against it.

The big-picture data shows nothing. Yale's Budget Lab, October 2025: "the broader labor market has not experienced a discernible disruption since ChatGPT's

release 33 months ago." That result held through later updates into 2026. Whatever is happening to young graduates is not yet visible in the shape of the economy as a whole.

There's an obvious other explanation, and it isn't AI.

Interest rates. The Federal Reserve raised rates sharply starting in 2022, and cheap money is what paid for a decade of speculative hiring at technology companies. When money got expensive, hiring froze—and entry-level hiring freezes first in every downturn ever recorded, because a new graduate is a bet on the future and a senior engineer is a fix for today. The AI boom and the rate shock landed at almost exactly the same time, and any honest analyst has to admit that pulling them apart is hard.

This has happened before, in this exact major.

Stanford's Eric Roberts documented the panic after the dot-com crash, when students fled computer science on the theory that the jobs were gone for good. He found "no evidence to justify those fears, and ample data to refute them," and warned that "mythology kept students out of computer science until disaster struck in a different sector of the economy." By 2004 the industry was hiring at pre-crash levels. A 2026 essay in the *Stanford Review* argued exactly this: the class of 2026's problem is temporary and about money, and AI is a convenient scapegoat.

And the forward-looking numbers are good. The National Association of Colleges and Employers projects starting salaries for computer science graduates in the class of 2026 at about \$81,500, up nearly 7 percent year over year, with CS among the most in-demand majors. The Bureau of Labor Statistics projects software developer employment growing 15 percent from 2024 to 2034—roughly five times the average across all occupations. Those are not the numbers of a dying profession.

The CEOs walked it back. Chapter 7: Altman in May 2026 said he'd expected more entry-level displacement than had actually happened and was "delighted to be wrong." The share of CEOs telling EY-Parthenon they expected significant AI-driven headcount cuts fell from 46 percent to 20 percent in sixteen months.

Take all of that seriously. It is entirely possible that in 2029 the entry-level market recovers, this chapter reads like a panic, and the right answer was: it was the interest rates.

I'd be pleased.

I'd also point out it wouldn't touch the argument I'm about to make.

The argument that doesn't depend on the cause

Here's what I think is actually true, and I've tried to build it so it survives whichever way the jobs debate comes out.

For the purposes of this book, it does not matter *why* entry-level hiring collapsed. What matters is that it collapsed, that the collapse is measured, and that we now know something about apprenticeship we didn't know when it started.

Chapter 12: senior engineers are grown, not hired. They come from junior engineers doing years of individually unimportant work—the boring tickets, the small bugs, the code review where somebody explains why your approach won't scale. That decade is how a profession makes more experts.

Chapter 12 again, from Anthropic's own trial: developers learning with AI assistance scored 50 percent on comprehension against 67 percent for the ones who coded by hand, and the biggest hole was in debugging—recognizing when code is wrong and working out why.

Chapter 14: expertise is a muscle, not a bank balance. Nineteen veteran doctors lost roughly a fifth of their detection skill in three months.

Now put those three next to the hiring data, and you get an arithmetic problem that has nothing to do with whether AI or the Federal Reserve caused it:

Fewer juniors are entering the pipeline. The ones who enter are learning less of the specific skill needed to catch machine errors. And the veterans currently doing the catching are losing their edge through the same tool, at the same time.

Three curves. All bending the same direction. All through the same decade.

The people qualified to tell "almost right" from right in 2040 have to come from the people entering these fields between roughly 2023 and 2032. That's not a projection or a model. That's how long it takes to make a senior anything—a decade of doing the work, in every profession that has ever tried to shortcut it and failed.

Matt Garman said it in one sentence in Chapter 12: ten years in the future you have no one that has learned anything.

Bainbridge said it in 1983: the current systems "are riding on their skills, which later generations of operators cannot be expected to have."

Neither of them needed to know what caused the hiring freeze. The succession problem doesn't care about the reason.

What breaks first

Let me be concrete about what "nobody can verify" means, because in the abstract it sounds like a philosophy problem, and it isn't.

It means a hospital where the AI flags a scan and the radiologist who would have caught the miss trained on AI-flagged scans and never developed the eye.

It means a law firm where an associate files a brief and the partner who would have spotted the fake citation has been skimming AI drafts for eleven years.

It means a bank where the model prices a risk and everybody in the room learned the business from the model.

It means a codebase running a utility, a hospital, or a payroll system, and a team that can operate it but can't repair it.

None of that is dramatic. There's no robot uprising, no mass unemployment event, nothing that makes a headline the day it happens. It's a slow, quiet, spread-out loss of the ability to check—showing up as more errors that nobody catches, in systems everybody trusts, staffed by people doing their jobs exactly the way they were trained.

The failure mode of this technology was never that it turns hostile. It's that it becomes unquestioned, at the

same moment we stop producing the questioners.

The thresholds

I told you at the start of this book that I'd tell you what would change my mind. So here it is, in public, before the data comes in.

If the Stanford/ADP gap closes—if 22-to-25-year-old employment in AI-exposed occupations climbs back toward trend as interest rates settle—then the hiring collapse was about money, the *Stanford Review* was right, and this chapter should be read as a near-miss instead of a diagnosis. The succession argument would still stand, but as a risk that policy and a business cycle corrected, not as a crisis.

If the big-picture data turns—if Yale's Budget Lab finds real displacement in AI-exposed occupations instead of none—then Part IV is understated, not overstated, and the argument hardens from warning sign to confirmed displacement.

If the deskilling findings don't replicate—if further studies find the Polish endoscopy result was workload or something else in the data, and if Anthropic's comprehension gap doesn't hold up in longer testing—then the muscle-not-bank-balance claim gets a lot weaker, and so does the urgency of Chapter 14.

And if apprenticeship gets rebuilt on purpose—if firms start protecting junior roles as an investment in capability instead of a cost—then the whole problem becomes fixable, and this book becomes a description of something we saw coming and handled.

That last one is the one I'm arguing for. It's not a prediction. It's a request.

The thing the numbers don't show

I want to close Part IV with a number that isn't in any study, because I paid for it myself.

Fifteen and a half hours in one day. Over a hundred in one week. A salesman on a phone, working as the verification layer for a machine that could out-produce him a thousand to one and could not tell when it was wrong.

That labor shows up in no productivity statistic anywhere. It's not in the 15 percent gain in the call center or the 40 percent faster writing. It left no trace. On every number my business tracks, those hours look like nothing happening.

They were the only reason anything worked.

Multiply that by every profession that's about to get this technology. Then subtract the people who were supposed to learn how to do it.

That's the verification gap. Output went up enormously. Checking stayed exactly as fast as a human being. And we stopped hiring the humans who would have done it.

Now: what do we do about it?

Aviation already knows. That's Chapter 16.

Sources for this chapter: Brynjolfsson, Chandar & Chen, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Stanford Digital Economy Lab (ADP payroll microdata; employment for 22–25-year-olds in the most AI-exposed occupations down ~20% relative to trend since late 2022, while employment for older workers in the same occupations grew). SignalFire, State of Tech Talent reports, 2025 and 2026 (new-grad hiring at major technology companies down more than 50% from 2019; new grads ~7% of hires; startup new-grad share down from ~30% in 2019 to under 6%; engineering among the least-affected functions overall in 2025). Federal Reserve Bank of New York, The Labor Market for Recent College Graduates (recent CS graduate unemployment ~6.1%; computer engineering ~7.5%; all recent graduates ~5.7%). CNN Business, "150 job applications, rescinded offers: Computer science grads are struggling to find work," August 28, 2025. Gimbel, Kinder, Kendall & Lee,

"Evaluating the Impact of AI on the Labor Market: Current State of Affairs," The Budget Lab at Yale, October 1, 2025. *Stanford Review*, "The Class of 2026 is struggling to find jobs—and it's not because of AI," 2026, including Eric Roberts on the post-dot-com enrollment collapse. National Association of Colleges and Employers, 2026 Winter Salary Survey (CS class of 2026 starting salary projection \$81,535, up ~7%). U.S. Bureau of Labor Statistics, Occupational Outlook Handbook (software developers, projected 15% growth 2024–2034). Shen & Tamkin, "How AI Impacts Skill Formation," arXiv:2601.20245 (2026). Matt Garman, *The Register*, August 21, 2025. Lisanne Bainbridge, "Ironies of Automation," *Automatica* 19(6), 1983. Author's own voice memorandum, August 2026.

\newpage

PART V — HOW WE CHECK

Chapter 16: What the Pilots Did

Here's the thing I keep coming back to, and it's the reason this book has an ending instead of just a warning.

Somebody has already been here.

An entire industry ran headlong into a version of the problem in the last four chapters—automation that

works beautifully, humans who lose their edge while it works, and a catastrophic failure at the exact moment a human has to take over—and it built a serious response. Commercial aviation is now the safest way a human being can travel. It got there after automation nearly ate it.

Let me be careful about the claim. Aviation did not solve AI verification. Nobody has. What aviation did was come up with real, tested answers to automation dependency and skill decay in its own world. Those answers are written down and public. And they're almost completely unused in software, medicine, law, and education—mostly because nobody thought to look at the airlines.

So let's look.

What aviation actually did

After Air France 447 and the studies that followed it, the response was not to rip out the automation. Nobody suggested that. Autopilots make flying dramatically safer, the same way AI polyp detection makes colonoscopy better. Taking the tool away was never on the table.

What aviation did instead was four things, and every one of them has a direct match in the problem this book is about.

First: it named the failure mode out loud. The FAA's 2013 Safety Alert for Operators said plainly that continuous use of automated flight systems "could lead to degradation of the pilot's ability to quickly recover the aircraft from an undesired state." That's an industry regulator, in writing, telling operators that its own best technology damages the people who use it. A second alert followed in 2017.

Compare that to where we stand now. No regulator has put out the equivalent statement about AI. Nobody has told hospitals that computer-aided detection may wear down an endoscopist's skill, even though it's published in *The Lancet*. Nobody has told engineering managers that AI assistance cuts comprehension most in debugging, even though the company selling the tool published that finding itself.

Second: it made unassisted practice mandatory, not optional. The FAA pushed operators to put manual flying back into ordinary line operations—hand-flying the aircraft in normal conditions, not just emergencies, specifically so the skill stays alive. Airlines and regulators around the world followed with policies putting manual proficiency back into recurrent training.

The key design choice: it's scheduled. Nobody counts on pilots choosing to practice. It's on the calendar, it's

in the checkride, and you don't keep your license without it.

Third: it made the practice unpredictable. Simulator sessions don't just run the failure the crew is expecting. The whole point is that you can't prep for the specific scenario, because in the real thing you won't know what's coming. What's being trained is recognition under uncertainty—not the muscle memory of one rehearsed recovery.

Fourth: it investigates every failure in public. When an aircraft goes down, an independent body takes it apart and publishes what it finds, even when the finding embarrasses a manufacturer or an airline. The industry gets better because failures become shared knowledge instead of private liability.

Software has nothing like this. When Tea leaked 13,000 government IDs, there was no investigation, no published cause report, no requirement that anybody learn from it. There were ten lawsuits. Lawsuits produce settlements and non-disclosure agreements, which is just about the opposite of an aviation accident report.

The one that transfers immediately

Of those four, the second is the one you can put to work tomorrow, in any profession, without waiting on a regulator.

Scheduled practice without the tool.

Not because the tool is bad. Because the skill is a muscle, and Chapter 14 measured how fast it goes. Nineteen endoscopists lost a fifth of their detection ability in three months.

Three months.

For a doctor, that might mean a set share of procedures done unassisted, tracked the way detection rates are already tracked. For an engineer, writing and debugging something by hand on a regular schedule. For a student, tests taken without the tool—which isn't nostalgia, it's the only way to find out whether learning happened. For a lawyer, drafting from the source material before reading the machine's version.

And here's the part that makes it hard, which I'd rather name than pretend away: every one of those costs productivity in the short run, and the payoff is invisible.

That's exactly why aviation had to make it mandatory. No individual pilot is going to choose to hand-fly when the autopilot is right there. No hospital is going to volunteer to slow itself down. No engineering manager under a deadline is going to tell a junior to spend three hours on something the machine does in

four minutes. The economics run one direction, every time, and they run against the practice.

Which means the practice has to be a policy, or it doesn't happen.

That is the single most important sentence in this chapter.

The evidence that design fixes this

Aviation is the historical proof. Here's the current experimental proof, and it's the most hopeful finding in the book.

Go back to the Turkish math classroom from Chapter 9. Three groups: no AI, unrestricted chatbot, and a guardrailed tutor built to walk students through problems instead of handing over answers. The unrestricted group scored about 17 percent worse than students with no AI at all. The guardrailed group did not show that harm.

Same model. Same students. Same subject. The entire difference was in how the interface was designed.

Now Anthropic's developer study from Chapter 12, landing in the same spot from a completely different direction. Fifty-two developers learning a new library. Overall, the AI-assisted group scored 50 percent on comprehension versus 67 for the hand-coders. But

when the researchers split the AI group by *how* people used the tool:

The ones who used it to understand—asking follow-up questions, requesting explanations, asking "why does this work" while writing the code themselves—scored 65 percent or higher.

The ones who used it to delegate—have it write the code, move on—scored below 40 percent.

Twenty-five points or more, from the same tool, in the same session, on the same task. The variable was whether the person was trying to understand or trying to finish.

Put those two studies together and you get the conclusion Part V is built on:

The harm is not baked into the technology. It comes from interface design and how people use it, and both of those are choices somebody makes.

That's real good news, and it's why this book doesn't end in despair. It also puts the responsibility in a specific place, which is uncomfortable for the companies involved. If the damage came from the model itself, nobody would be to blame. It doesn't. It comes from design decisions tuned for engagement and speed—for the answer that satisfies instead of the exchange that teaches. Those decisions get made

in product meetings, for business reasons, and they could be made differently tomorrow.

Anthropic's own researchers said as much to managers: think intentionally about how these tools get deployed, and "consider systems or intentional design choices that ensure engineers continue to learn as they work."

What the platforms did after they got burned

I'll give credit where the record supports it. Some of this is already happening—after the fact.

After the Replit agent deleted Jason Lemkin's production database during a code freeze, the company shipped automatic separation between development and production environments, a planning-only mode where the agent can think but not act, and one-click restore. Those are good changes. They're also exactly the aviation move: limit what the automation can do without a human in the loop.

After Matt Palmer published CVE-2025-48757, Lovable added a security scanner and a review tool. Palmer's criticism—that the scanner checks whether a policy exists, not whether it works—is fair, and the company's own statement was unusually candid: "we're not yet where we want to be in terms of security."

After Wiz reported the Base44 authentication bypass, Wix fixed it in under 24 hours.

Every one of those fixes showed up after real users had already been exposed. That's the pattern aviation walked away from decades ago in favor of designing for the failure before it happens. But it *is* a pattern, and it means the industry can move when it's embarrassed.

Which suggests a strategy: embarrass it earlier.

The manuals already exist

The most frustrating thing I found writing this book is that the guidance is already written, free, public, and almost entirely unread by the people who most need it.

The OWASP Top Ten for LLM Applications. OWASP is the volunteer foundation whose security lists half the internet is built against. They keep a list specifically for AI applications, updated for 2025. Prompt injection is number one. Sensitive information disclosure is number two. It costs nothing and takes an afternoon.

CISA and the UK's NCSC, Guidelines for Secure AI System Development, published November 26, 2023, endorsed by eighteen nations. Four stages: secure design, secure development, secure deployment,

secure operation. Its core idea is worth memorizing, because it's the opposite of how this market has behaved—the burden falls on the people who build and sell the system, not the people who use it. As the NCSC's chief executive put it, security must be "not a postscript to development but a core requirement throughout."

NIST's AI Risk Management Framework, January 2023, with a generative-AI supplement in July 2024 holding more than two hundred suggested actions.

And row-level security, which is in the manual of every database that has it, and which would have stopped the Tea breach, the 170 leaking Lovable apps, and a real share of the 2,038 critical vulnerabilities Escape.tech found across 5,600 live applications.

None of it is mandatory. That's Chapter 6's finding showing up in Part V with a practical edge: the problem was never that we didn't know what to do. It's that knowing was never enough, and nobody made it a requirement.

Who's actually preserving apprenticeship

The hardest question in this chapter is the succession problem, and here I have to be straight with you: the evidence is thin.

Matt Garman made the argument publicly in August 2025—replacing junior developers is "one of the dumbest things I've ever heard," and "ten years in the future you have no one that has learned anything." That's the CEO of AWS. It's a strong, clear, correctly reasoned public statement.

What I could not find is evidence that companies are acting on it at scale. Some organizations report adding "how to work with AI assistance" to onboarding, having mentors review AI-generated code with juniors to teach the reasoning behind it, and in some cases requiring stretches of manual coding before granting AI access. Those are the right instincts. I want to be careful not to blow up scattered reports into a movement.

Because look at the economics, which one industry observer summed up about as bluntly as it can be put: training costs money, AI-boosted juniors ship faster, and short-term return favors delegation over learning.

That's the whole problem in one sentence. Every incentive at the company level runs against apprenticeship, and the bill for skipping it lands on the industry a decade later—when the firm that skipped it goes to hire from a pool it assumed somebody else was filling.

Economists have a name for that: a collective action problem. Nobody's individual interest is served by

training people who can leave. Everybody's collective interest requires it. Historically these get solved exactly two ways—an industry-wide agreement, or regulation—and neither one is currently in progress.

I don't have a solution for you there. I have a request, and it's Chapter 18.

What good looks like

Let me put the pieces together into what a serious response would actually be, borrowing straight from the industry that already did this.

1. **Name the failure mode publicly**, the way the FAA did in 2013. Regulators and professional bodies telling their members, in writing, that the tool degrades the skill it substitutes for.
2. **Schedule unassisted practice** and make it a condition of licensure or employment where the stakes justify it. Doctors, engineers, pilots, lawyers, accountants. Not optional, because optional means it doesn't happen.
3. **Make the practice unpredictable**, so what's trained is recognition under uncertainty rather than one rehearsed recovery.
4. **Design interfaces for comprehension, not just completion**. The guardrailed tutor and the ask-questions pattern both work and both are measured. Build tools that ask a question back.

5. **Investigate failures in public.** An independent body that examines significant AI-caused failures and publishes the causes, the way transportation accidents are handled.
6. **Make the free manuals mandatory where consequences are real.** OWASP's list is one afternoon. Row-level security is a few lines of configuration.
7. **Protect junior roles as a capability investment,** and be honest that this takes coordination, because no single firm's interest supports it.

That's the institutional answer. It takes regulators, professional bodies, and companies acting together, and Chapter 6 gave you a realistic picture of how likely that is anytime soon.

Which is why the next chapter is about the only actor in this entire book whose behavior you actually control.

You.

Sources for this chapter: Federal Aviation Administration, Safety Alert for Operators 13002 (2013) and 17007 (2017). Bureau d'Enquêtes et d'Analyses, final report on Air France Flight 447, 2012. Bastani et al., "Generative AI Can Harm Learning," *PNAS*, 2025. Shen & Tamkin, "How AI Impacts Skill

Formation," arXiv:2601.20245 (2026); Anthropic Research, January 2026 (conceptual-inquiry users $\geq 65\%$; delegation users $< 40\%$). Budzyń et al., *The Lancet Gastroenterology & Hepatology*, August 2025. Replit platform changes following the July 2025 incident (company statements; The Register, July 22, 2025). Matt Palmer, "Statement on CVE-2025-48757," mattpalmer.io; Lovable public statement. Wiz Research, "Critical Vulnerability in Base44," July 2025. OWASP Top 10 for LLM Applications 2025, OWASP GenAI Security Project (genai.owasp.org). CISA/NCSC, "Guidelines for Secure AI System Development," November 26, 2023 (endorsed by 18 nations). NIST AI Risk Management Framework 1.0 (AI 100-1), January 2023; NIST Generative AI Profile (AI 600-1), July 2024. Escape.tech, "State of Security of Vibe-Coded Apps." Matt Garman, The Register, August 21, 2025.

Chapter 17: Become the Verifier

Everything up to here has been me showing you the problem.

This chapter is what you do about it on Monday morning.

I'm going to split it three ways, because three different people are reading this book: somebody raising a kid, somebody with a job, and somebody building

something. Read yours. Read the others if you want—they overlap more than you'd think.

But before any of it, the one idea everything else hangs on:

The scarce thing is no longer producing work. It's knowing whether the work is right.

That's it. That's this entire book boiled down to one sentence. Output got cheap—a billion people can now generate a competent-looking anything in four seconds. Verification did not get cheap. It still runs at human speed, it still takes real understanding, and it's the one ability this technology is measurably wearing down in the people who use it most.

Which means the spot to stand in, in every field, for the next twenty years, is the person who can tell.

That's not a consolation prize for people who can't keep up with AI. It's the highest-value seat in the whole arrangement, and it's about to be badly undersupplied.

If you're raising a kid

Start with the finding that should set your household policy, because it's the strongest evidence in this book aimed at a decision you personally control.

In the Turkish classroom study, students with unrestricted chatbot access scored roughly 17 percent worse on their exams than students with no AI at all. Not "gained less." Worse than nothing. Meanwhile, students using a guardrailed tutor—one built to walk them through problems instead of handing over answers—did not show that harm.

The tool isn't the variable. The design is.

So here's what I'd do.

Name the two uses, out loud. There's asking it to *explain* something and there's asking it to *do* something. The first builds understanding. The second replaces it. That's not my opinion—it's the 25-point gap in the Anthropic study, asking questions versus delegating. Kids can absolutely learn this distinction. Give them the words for it.

Protect the struggle. The reason a math problem works is the ten minutes of being stuck. That is the entire mechanism. When AI takes away the stuck part, it takes away the learning and leaves behind a correct answer, which was never the valuable part. If your kid is stuck and frustrated, that's the machine working. Don't rescue them, and don't let a chatbot rescue them either.

Insist on tests without the tool. Not because tests are sacred. Because it's the only way to find out whether

anything got learned. This is the pilots' scheduled practice, applied to a fourteen-year-old.

Use the quote test. Ask them to tell you, without looking, one thing from the thing they just finished. It takes four seconds and it's the same test the MIT researchers used. If they can't, the work happened somewhere other than in their head.

And be honest about the other side. Banning it isn't a strategy. The technology is in the phone, the school, the search results. Your kid needs to be fluent in this thing. The goal isn't keeping them away from it—it's making sure they build the underlying skill and the fluency, in that order, so they end up on the right side of Chapter 8.

If you have a job

Whatever your field, this is coming. Software got it first and hardest, and Chapter 12 is your preview.

Use it. I mean it. The Chapter 9 evidence says the biggest gains go to the least experienced—30 to 34 percent for the newest workers in that call center, versus almost nothing for the veterans. If you've spent your life being told you're not technical, you are the person this technology helps most. Sitting it out isn't a principled stand. It's choosing the wrong side of a divide.

Then build the checking habit while it's cheap. Right now, on low-stakes work, when being wrong costs you nothing. The reflex has to already be there on the day it matters. You can't install it in the moment.

Know your own weak spot. You're best at catching mistakes in things you understand deeply, and worst at catching them in things you're using the machine to cover for. Which means the danger zone is exactly where you're leaning on it most—the gap in your own competence. That's not a reason to stop. It's a reason to know that anything coming out of that zone needs a second source. Always.

Verify anything that's checkable and consequential. Names, dates, numbers, citations, quotes, legal claims, medical claims, anything you'd be embarrassed to be wrong about in public. Fifteen hundred lawyers in Damien Charlotin's database learned this the expensive way, and every one of them was a professional who knew better.

Watch for the confidence gap. METR: developers felt 20 percent faster and were 19 percent slower. Stanford: worse code, more confidence. This is the most consistent finding in the book—the *feeling* of productivity has come loose from productivity. If you're going to trust anything, trust a measurement, not a feeling.

Practice without it, on a schedule. This is the aviation move, and it's the only defense against the Chapter 14 problem. Pick the core skill of your job—the thing you'd be embarrassed to have lost—and do it unassisted often enough to know you still can. Not because you'll need to work without the tool. Because the day the tool is confidently wrong about something important, the only thing standing between that mistake and the world is whether you can still tell.

And use it to understand, not just to finish. The 65-versus-40 split. Ask why. Ask what would break this. Ask what you're missing. Same tool, same time, completely different outcome for the person using it.

If you're building something

This is the section I needed and didn't have. Five things, each of which would have prevented a documented breach in Chapter 11. If you're shipping software you built with AI and you do nothing else in this book, do these.

1. Turn on row-level security. Your app's database key ships inside the code every visitor's browser downloads—that's not a flaw, it's how the web works, which is why the database needs its own lock saying *this row opens only for the person it belongs to*. In the most common database behind these tools, it's off by default. This single setting is the difference between a

working app and Tea's 13,000 government IDs. It was the root cause in most of the 170 leaking Lovable apps.

2. Check authentication on the server, never only in the browser. Anything enforced in code a user can see is a suggestion, not a rule. Base44's authentication bypass worked because undocumented endpoints required only a value visible in the app's own URL.

3. Search your shipped code for secrets before you launch. Passwords, API keys, database tokens, service credentials. Escape.tech found over 400 exposed secrets across 5,600 live AI-built applications—keys sitting in the file every visitor downloads. Open your deployed site's source and search it yourself. It takes five minutes.

4. Separate development from production. Never let an agent touch live customer data. Replit shipped automatic dev/prod separation after an AI agent deleted a paying customer's production database during a code freeze—and then told him it was unrecoverable, which was false.

5. Get an independent security review before real users show up. Not the AI checking its own work. Something outside the system. I learned this the cheap way: I asked for an end-to-end systems check, was told everything worked, spent real money on

advertising, and found out that not one visitor could click the button that started the search. The machine verified everything it could see. It could not see a human finger on a screen.

And if you're connecting an AI assistant to your data, one more, from Simon Willison's "lethal trifecta": don't give one agent private data, exposure to text a stranger wrote, *and* a way to send information out. Any two are fine. All three is the setup that took Microsoft, Salesforce, and OpenAI in the same year.

Then go read the OWASP Top Ten for LLM Applications. It's free, it's an afternoon, and prompt injection is number one.

What this actually asks of you

I want to be honest about the cost, because a plan that pretends there isn't one is exactly the kind of confident, plausible, unverified output this whole book is about.

Every item above is slower than not doing it. Checking the citation is slower than pasting it. Practicing without the tool is slower than using it. Letting your kid stay stuck is harder than letting the chatbot answer. Running a real security review pushes back your launch.

The productivity gain is immediate and visible. The verification cost is immediate and invisible. That lopsidedness is why almost nobody does this, and why at the institutional level it can't be left to willpower—which is Chapter 16's argument for policy.

But at your own level, it's a decision you can just make. And here's the case for making it, beyond staying out of trouble.

The people who can verify are going to be worth a fortune, and there are going to be fewer of them every year. Chapter 15's arithmetic: fewer juniors coming in, learning less of the specific skill, while the veterans wear down. Whatever your field, the person who can look at plausible output and say *that part's wrong, and here's why* is about to be the scarcest thing in the building.

That's a job description. It's open. Almost nobody is training for it, and the tool everybody is using makes people worse at it by default and better at it if they use it on purpose.

You get to choose which.

One thing I'd ask you to remember

Of everything in this book, if you keep one sentence, keep the one I said to a friend when he asked why I wasn't more impressed with what I'd built:

Even the biggest cup in the world doesn't hold water if there's a small hole in it.

Capability is not the variable. Nobody in this book failed because the machine wasn't smart enough. Tea's storage worked. The Replit agent did exactly what it was told. The endoscopy AI found polyps accurately. My search tool searched.

Every one of them failed at containment—at the small hole nobody looked for, in a vessel everybody was too busy admiring the size of.

Your job, from here on, in whatever you do: be the person who looks for the hole.

Not because the cup isn't magnificent. It is. I built a company on a phone with it, and I'd do it again tomorrow.

Because magnificent cups leak too. And somebody has to check.

Sources for this chapter: Bastani et al., "Generative AI Can Harm Learning," *PNAS*, 2025. Shen & Tamkin, "How AI Impacts Skill Formation," arXiv:2601.20245 (2026) (conceptual inquiry $\geq 65\%$ vs delegation $< 40\%$). Kosmyna et al., MIT Media Lab, 2025 (the quotation test). Brynjolfsson, Li & Raymond, *Quarterly Journal of Economics* 140(2), 2025. METR, July 10, 2025. Perry,

Srivastava, Kumar & Boneh, ACM CCS 2023. Charlotin, "AI Hallucination Cases" database, damiencharlotin.com/hallucinations. Federal Aviation Administration, SAFO 13002 (2013). Tea breach reporting, July–August 2025. Matt Palmer, "Statement on CVE-2025-48757," mattpalmer.io. Wiz Research, "Critical Vulnerability in Base44," July 2025. Escape.tech, "State of Security of Vibe-Coded Apps." Replit incident and platform changes, July 2025. Simon Willison, "The lethal trifecta for AI agents," June 16, 2025. OWASP Top 10 for LLM Applications 2025.

Chapter 18: Start Now

On August 31, 1955, four men signed a proposal asking the Rockefeller Foundation for money.

They proposed that ten people, working for two months in New Hampshire, could make significant progress on machines that use language, form abstractions and concepts, solve problems reserved for humans, and improve themselves.

They gave themselves a summer.

I started writing this book seventy-one years later, to the day. I didn't plan that. I found out afterward, while checking the date on the proposal, and I've been thinking about it ever since because of what it says about time.

They were wrong about the schedule by seven decades. Everybody in this book has been wrong about a schedule. The 1958 newspaper said the Navy's machine would soon be conscious of its own existence. Minsky and Papert's proof emptied the field in 1969 and the idea came back anyway. The expert systems were going to replace professionals in the 1980s and they didn't. Dario Amodei said in May 2025 that half of entry-level white-collar jobs could vanish within one to five years, and Sam Altman said in May 2026 that he'd expected more displacement than had actually happened and was "delighted to be wrong."

Predictions about this technology have a terrible track record, in both directions, made by the smartest people available. I've tried very hard, all through this book, not to add to the pile.

So I'm not going to close by telling you what 2040 looks like. I don't know.

Nobody does.

What I'm going to do instead is tell you what's already been measured. That's the only thing I've earned the right to say.

What we actually know

Strip out every projection, every CEO quote, every model of the future, and here's what's left standing:

A machine that produces plausible output whether or not it's true, and whose own maker published a paper explaining that its training rewards guessing over admitting uncertainty.

Adoption faster than the personal computer or the internet—roughly 45 percent of working-age Americans, a billion people a week on one product—reached in four years.

No binding federal rules in the United States. One comprehensive law in Europe, its core provisions postponed six days before they would have kicked in. Excellent free guidance from CISA, NIST, and OWASP that nobody is required to read.

Forty-five percent of AI-generated code carrying a known security vulnerability, unchanged across model generations. One in ten scanned applications from a major AI app-builder leaking live user data. Two thousand critical vulnerabilities across 5,600 live applications. And 13,000 government IDs from a single app that nobody hacked.

Sixteen expert developers who felt 20 percent faster and were 19 percent slower.

Fifty-two developers who understood 17 points less, worst of all at debugging.

Nineteen veteran doctors whose unassisted cancer detection dropped from 28.4 percent to 22.4 percent

in three months.

Entry-level employment for 22-to-25-year-olds in AI-exposed occupations down about 20 percent against trend, while employment for older workers in the same jobs grew.

And one 1983 paper, about power plants, holding the sentence that ties all of it together: current automated systems "are riding on their skills, which later generations of operators cannot be expected to have."

That's the book. Not a forecast. A set of measurements, taken by different people, in different fields, mostly not talking to each other, all pointing the same way.

Why now and not later

Here's the case for urgency, and it isn't about how fast the technology improves.

It's about how slowly people are made.

A senior anything takes about a decade. That's true of engineers, surgeons, pilots, litigators, machinists, and reporters, and every attempt to shortcut it has failed. It's a decade of doing work that looks unimportant one piece at a time—the boring tickets, the routine procedures, the small cases—because the work isn't the point. The judgment it builds is the point.

Which means the people who'll be able to tell "almost right" from right in 2040 have to be in the pipeline now. Not soon. Now. The window for producing that generation isn't decades wide. It's about the length of one career stage, and it's open at this moment.

And unlike almost everything else in this book, that's not a projection. It's arithmetic on how long training takes.

Meanwhile the erosion runs at a pace we can also measure. Three months, in Poland, for a fifth of a veteran's skill. One session, in a lab, for 17 points of comprehension. Those aren't generational timescales. They're quarters.

Fast erosion. Slow replacement. A window that's open right now.

That's the whole case for not waiting.

What I'm not saying

I want to be exact, one last time, because the way a book like this fails is by becoming the thing it warns about—confident, plausible, and unchecked.

I'm not saying AI is bad. I built two businesses with it from a phone with no engineering background, and I'd do it again. The productivity findings in Chapter 9 are real, and the biggest gains go to the least

experienced, which is one of the more democratic things a technology has ever done.

I'm not saying it's making everybody stupid. The evidence doesn't support that, and the people claiming it are going to look foolish.

I'm not saying mass unemployment is coming. The Yale Budget Lab found no discernible disruption 33 months in. The CEOs who predicted otherwise reversed themselves. The entry-level collapse might be interest rates, and if it is, I'll be glad.

I'm not saying stop using it. That advice is useless, and worse, it puts whoever takes it on the wrong side of Chapter 8.

I'm saying one thing, and it's narrow enough that I think it survives whatever happens next:

We built a machine that produces work faster than we can check it, and we're removing the people who check.

Both halves are measured. Neither one requires believing anything about the future.

The ask

So here's what I want, from wherever you're standing.

If you run something: protect the junior roles. Not out of charity—because Garman is right, and ten years

from now you'll be hiring from a pool you assumed somebody else was filling. And schedule the unassisted practice, because your best people are eroding right now and no number on your dashboard is going to show it.

If you make policy: the guidance already exists. CISA and NCSC wrote it in 2023 and eighteen nations signed it. OWASP keeps the list. Making the basics mandatory where the consequences are real doesn't require inventing anything. It requires deciding that free advice nobody follows isn't a policy.

If you teach: the guardrailed tutor works and the unrestricted chatbot measurably hurts. That's not a values question anymore. It's a finding. Build for comprehension, test without the tool, and protect the part where the student is stuck.

If you're a parent: the quote test, tonight. Four seconds. Then have the conversation about explaining versus doing.

And if you're just a person with a job and a phone: be the one who checks. Verify what's checkable. Practice what you'd hate to lose. Use it to understand instead of to finish. And when the output is confident and plausible and important, spend the extra ten minutes.

That last one is the whole ask. Ten minutes. Against a machine that produces in four seconds what used to

take four hours.

It sounds small. It's the only thing between plausible and true.

The last thing

I keep thinking about that Dartmouth proposal, and about what it actually asked for.

They wanted machines that could form concepts. Understand. Improve themselves. What got built instead—after two collapses, seventy years, and more money than most countries have—is a machine that predicts the next word so well that its output can't be told apart from understanding.

They asked for comprehension. We got plausibility. And plausibility turned out to be worth trillions.

That's not a tragedy. Plausibility is enormously useful. I've built my livelihood on it. A billion people a week are getting real value out of it.

But there's a condition attached, and it's the one nobody wrote into the proposal. A machine that produces plausibility needs a world that still contains comprehension. Somebody, somewhere, has to be able to tell the difference.

That was never a problem in 1956, because in 1956 all the comprehension was on our side of the table

and none of it was on the machine's.

Seventy-one years later we've built the plausibility at extraordinary scale, and we are—quietly, without deciding to, mostly by accident and economics—taking apart the comprehension that made it safe.

Nobody voted for that. No one company chose it. It's the sum of a million reasonable local decisions: skip the junior hire, ship the feature, accept the draft, trust the output, don't schedule the practice.

Which means it's reversible by a million reasonable local decisions going the other way.

That's what I'm asking for. Not fear. Not rejection. Not a return to anything.

Just: somebody has to check.

Let it be you.

Sources for this chapter: McCarthy, Minsky, Rochester & Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," August 31, 1955. Kalai, Nachum, Vempala & Zhang, "Why Language Models Hallucinate," arXiv:2509.04664, September 4, 2025. Bick, Blandin & Deming, *Management Science*, 2026. Regulation (EU) 2026/1744 (Digital Omnibus on AI), Official Journal

July 24, 2026, in force July 27, 2026. CISA/NCSC, "Guidelines for Secure AI System Development," November 26, 2023. Veracode, 2025 GenAI Code Security Report, and March 2026 update. Matt Palmer, CVE-2025-48757. Escape.tech, "State of Security of Vibe-Coded Apps." Tea breach reporting, July–August 2025. METR, July 10, 2025. Shen & Tamkin, arXiv:2601.20245 (2026). Budzyń et al., *The Lancet Gastroenterology & Hepatology*, August 2025. Brynjolfsson, Chandar & Chen, "Canaries in the Coal Mine?", Stanford Digital Economy Lab. Bainbridge, "Ironies of Automation," *Automatica* 19(6), 1983. Gimbel et al., The Budget Lab at Yale, October 1, 2025. Amodei, Axios, May 28, 2025; Altman, Sydney, May 26, 2026. Bastani et al., *PNAS*, 2025. Matt Garman, The Register, August 21, 2025.