

## **INTRODUCTION**

### **What It's All About**

This is a book about AI disruption and the very real probability of millions of jobs and entire professions being replaced, but it's about more than that. It's about personal discovery.

Who are you, exactly, when the thing that you believe defined you is gone? No one wants to have that conversation, whether out loud or just in your head. The disruption to identities due to AI's replacement of what we do is accelerating the pace at which many are now coming face to face with that question.

As I was writing that it reminded me of the first line of one of my favorite Dionne Warwick songs, "What's It All About, Alfie?" I've asked myself some version of that question for most of my life.

I asked it again when copywriting, my profession for the past 15 or so years, was one of the first job fields to be taken over by AI. I suppose I'm somewhat lucky in that I've done so many different things that I never actually identified with one in particular. No, that's not true. I spent a year and a half living and working on a genuine paddle-wheeler steamboat on the Mississippi River. But that felt more like an adventure than a job, so I guess it doesn't count.

I do remember, though, standing on the dock in New Orleans and watching the boat leave without me the day I got off for the last time. There was a deep sense of loss, like a part of me was sailing away.

People who are being replaced by AI are feeling more than the loss of a paycheck. For all practical purposes, the role they played no longer needs them.

The question that was always waiting underneath is coming to the surface, whether we are ready or not.

Who are we without the title? What *is* it all about?

This book is an attempt to help you be ready for those questions. And the good news is that what you find underneath is almost always far more rewarding than the title itself.

## **CHAPTER ONE**

### **Something Big Is Happening And Most of Us Don't Know It Yet**

Think back to early 2020. If you were paying close attention, really close attention, you might have seen a few headlines about a virus spreading overseas. Most people weren't paying attention or maybe they thought it was another "Breaking News" headline that turned out to be anything but. We've certainly been conditioned that way.

Something about it did catch my attention. My brother was about to go on a cruise and I remember telling him that I didn't think he should be traveling out of the country. Of course, he ignored me and went anyway. Par for the course. But it wasn't long before it seemed like the world turned upside down almost overnight.

One day kids were in classrooms, restaurants had waiting lists and people happily packed together in theaters and sports arenas. The next day all of the toilet paper had disappeared from the grocery store shelves and there was a curfew.

The World Health Organization declared COVID-19 a Public Health Emergency of International Concern in January 2020. It was not officially over until May 5, 2023. Millions lost their lives and no one was left untouched.

So, why start a book about artificial intelligence like this? Because it is the last time in most of our lives that something arrived slowly enough for us to watch it coming and still be caught unprepared. Why? Because we simply have no frame of reference for this magnitude of disruption. Not anything recently with Covid and not anything *ever* with AI.

This is where we are again. What is coming moves slower than a virus, but it is gaining momentum and has the potential to change everything. The big difference is there won't be a "go back to normal" stage. This will stay.

## **But I Tried It and It Wasn't That Impressive**

I can almost hear some of you thinking: "I've seen AI. I've played around with it. It made some mistakes, got some things wrong, and honestly it wasn't all that mind-blowing." That is a completely fair reaction, and it's also a part of the problem. What most people have tried is a bit like judging the entire communications industry by using a rotary phone from 1965. The thing technically works. It just has almost nothing in common with what's in your pocket, or, more likely the palm of your hand, right now.

The free versions of these tools lag a year or more behind what paying users already have, and the firms adopting this in earnest are not using the free versions. The tools being used in those places right now are categorically different from the chatbot you may have experimented with and shrugged off. The gap between public perception and actual current capability is, at this moment, enormous. And that gap is part of why so few of us are prepared for what's coming.

## **The "Something Big is Happening" Man Himself**

Matt Shumer is an AI entrepreneur who has spent six years building an AI company and investing in the space. He knows this world from the inside, and he wrote about it with the particular urgency of someone who has watched these changes happen to himself first.

He described the moment he realized he was no longer needed for the technical work of his own job. Not because the work had got easier. Not because he had a capable tool at his elbow. Because the tool was doing the work completely, correctly and better, *without him*. He could describe what he wanted built and walk away for four hours, come back, and find it done. Not a rough draft. The finished job.

Software development was his specialty, the thing he had spent years mastering, and an AI had taken it over. He wasn't bitter about it. He wrote about it with something closer to stunned matter-of-factness. Which, honestly, is somehow more unsettling than if he had been upset.

His essay, "Something Big Is Happening", published in February 2026, has been reported as viewed more than 80 million times.

Here's the part of Shumer's account that I think about the most. He tells a story about a lawyer friend he keeps urging to try using AI at his firm. The friend keeps finding reasons it won't work. It's not built for his specialty. It made an error when he tested it. It doesn't understand the nuance of what he does. And Shumer gets it. He's not being dismissive of his friend. But in the meantime, partners at major law firms have been reaching out to Shumer directly, because they've tried the current tools and can see where this is going.

When he spoke with Shumer, one managing partner, someone with decades of experience at a large firm, said he was using AI for hours every day. He described it as having a team of associates available instantly. And he said something that stayed with me: every couple of months, the tools get significantly more capable for his work, and he expects that, before long, they will be able to do most of what he does himself.

He's not panicking. But he is paying very close attention.

Not everyone was persuaded. Some read the essay as a pitch for paid AI tools, and one reviewer pointed out that Shumer has a history of overstating what a model can do. I think the specific things he describes are worth taking seriously anyway, and I'll show you why.

The lawyer and the managing partner are looking at the same technology. One of them is finding reasons it won't matter. The other is watching the change it is making in his business every sixty days. That gap is something we'll come back to, because it turns out to be one of the most human parts of the whole story.

## **So What Is Actually Happening?**

Let me try to explain artificial intelligence in a way that doesn't require a computer science degree, because I think the basics are genuinely important to understanding why this moment is different, and most explanations either talk down to people or assume too much.

For most of computing history, software worked by following explicit rules written by humans. A programmer would write: "if the customer enters X, do Y." If they entered something the programmer hadn't anticipated, the system would fail or produce nonsense. It was powerful within its rules and useless outside them. This is why early chatbots were so frustrating. They were basically elaborate sets of "if/then" instructions. The moment you went off-script, they fell apart.

Modern AI systems work differently. Instead of being given rules, they're given enormous amounts of text, images, audio and code and they learn patterns from that data. A large language model, which is the technical term for things like ChatGPT or Claude, has processed more written language than a person could read in a thousand lifetimes. From that ocean of language it learned how words relate to each other, how ideas connect, how arguments are structured and how questions get answered.

The result is a system that can have a coherent conversation about almost anything, write a legal brief, summarize a medical study, debug software, draft a marketing campaign or explain quantum physics in plain English. Not because it was programmed to do each of those things, but because it learned the underlying patterns of how humans think and communicate.

Now, here's where it gets especially meaningful for our conversation: these systems have been getting better, dramatically fast, and the rate of improvement has been accelerating. Not incrementally better. Categorically better. The difference between what existed two years ago and what exists today is not like going from a good phone to a slightly better phone, say an iPhone 15 to an iPhone 17. It's more like going from a bicycle to an airplane.

## **Why the Pace Matters So Much**

We're not good at thinking about exponential change. We're wired to think linearly. If something improved by ten percent last year, our instinct is to assume it'll improve by roughly ten percent next year. That serves us for most things in life.

That is not what's happening here. The researchers at an organization called METR have been measuring how long a task an AI can take on and still have a fair chance of getting it right. The tasks are real work, the kind that takes a skilled human hours. Their tasks mostly involve software and technical work, so treat the number as a floor rather than a ceiling. A few years ago the answer was measured in minutes. Then an hour. Then several hours. By the middle of 2026 they were publishing measurements above sixteen hours, which is the point at which they warned their own test could no longer measure it reliably. That figure has doubled roughly every seven months across the whole period, and closer to every four months since 2023.

By the time you're reading this, those numbers will be out of date, and I want you to take that as part of the evidence rather than a flaw. The pace of change is so fast that examples become obsolete before the ink dries.

## **It's Already Happening, Right Now, to Real People**

One of the things I want to resist is the framing that makes AI displacement sound like a future problem. The data on this is pretty clear. We're not making a forecast. This is happening right now.

The World Economic Forum, not exactly a fringe alarmist organization, is the group that runs the annual gathering in Davos where world leaders and CEOs talk about the state of things on a global basis. Their 2025 Future of Jobs report found that 41 percent of employers expect to reduce their workforce in areas where AI can do the work. The same report also projects more jobs created than lost overall. Goldman Sachs has estimated that generative AI could expose the equivalent of 300 million full-time jobs worldwide to automation. The Inter-American Development Bank puts average exposure to these tools at 28 percent of jobs within a year and 44 percent within ten.

And the jobs going first are not the ones most of us expected. The old assumption was always that automation would come for physical labor, factory workers, assembly line workers, people doing repetitive, manual tasks. That's been the story of every previous wave of automation, going back to the Industrial Revolution.

What's actually happening this time is close to the reverse. It's the people with degrees and training who are being hit first. Accountants, lawyers, writers, software engineers, analysts. Work that happens on a screen. Work we spent decades telling our kids was the safe harbor, where they would be fine if they just got an education and developed cognitive skills.

It turns out that what makes AI systems particularly powerful is that they're good at language and reasoning and pattern recognition, your basic job description for how most office workers spend their days. And they're getting better at it faster than almost anyone anticipated. We have always told ourselves that people would do the things machines couldn't. We are getting close to running out of things machines can't do.

To give you a sense of the specific roles most affected right now: computer programmers, accountants and auditors, legal and administrative assistants, customer service representatives, content writers and market research analysts are among those facing the highest near-term disruption.

There's also a gender story here that gets far less attention than it deserves. A 2023 analysis from the Kenan Institute at the University of North Carolina found that about 79 percent of employed women in the United States work in roles where a quarter to half of the tasks are exposed to generative AI, against 58 percent of men. The International Labour Organization's 2025 study found that in high-income countries, 9.6 percent of women's jobs sit in its highest risk category, against 3.5 percent of men's.

We should note that this is about which jobs are most exposed, not a count of people already replaced. Think of it more like the writing on the wall that is warning that this wave of displacement is likely to hit women harder and faster.

## **The Part That's Hardest to Explain**

I've been talking about AI as if it's a fixed technology that companies are now deploying. But there's a development that changes the picture in a significant way, and I want to try to explain it clearly even though it sounds a little like science fiction.

Some AI systems are now being used to help improve AI systems.

That sounds simple and actually makes sense, don't you think? The implications are far-reaching and anything but simple.

A Chinese company called MiniMax reported that one of its models ran more than a hundred rounds of a loop on itself, analyzing its own failures, changing its own code and re-testing, and lifting its scores on internal evaluations by about 30 percent. MiniMax says the same model now handles between 30 and 50 percent of one research team's daily work, with humans stepping in only for the crucial decisions, and that it is helping to build the next version of the tool. That account comes from MiniMax's own announcement, so treat it as a company describing its own work, but the direction is what matters here.

This is what researchers call the recursive self-improvement problem, and it's been a theoretical concern in AI safety circles for years. Part of the concern is a system getting capable enough to be dangerous on its own. It's that once an AI system can contribute to improving AI systems, the pace of development can accelerate beyond what humans can meaningfully oversee or direct. Each improvement makes the next improvement easier, and, once started, needs less human effort to keep going.

OpenAI's own announcement for its coding model described it as "our first model that was instrumental in creating itself", meaning it debugged its own training, managed its own deployment and diagnosed its own test results and evaluations. That's not a prediction about some future milestone. That's a company telling you, in their own words, that the model they just released helped to build itself.

Then there is this. Stuart Russell is a professor at UC Berkeley and the co-author of "Artificial Intelligence: A Modern Approach", the standard textbook in the field, used in over 1,500 universities across 135 countries. Dr. Russell has been writing and speaking about the dangers of pursuing Artificial General Intelligence, AGI, which is a hypothetical type of AI that can learn, reason, and perform any intellectual task that a human can. In a 2023 lecture at Berkeley called "How Not to Destroy the World with AI", Russell said "You should not deploy systems whose internal principles of operation you don't understand, that may or may not have their own internal goals that they are

pursuing and that you claim show sparks of AGI. If we believe we have sparks of AGI, that's a technology that could completely change the face of the earth and civilization. How can we not take that seriously?"

We're not going to go much farther describing the dangers of the unregulated and unchecked growth of AI. You can find information and predictions pretty much anywhere, if you want to look for it, and much of it is better informed than anything I could add.

But I do think it's important for you to know that the people building these systems, including the chief executives of the major AI companies, signed a public statement in 2023 saying that the risk of extinction from AI should be a global priority alongside pandemics and nuclear war. And still they continue building.

The explanation I keep coming back to is simple: they can't stop. The competition, the investors and the sheer momentum of it mean that any company slowing down for safety hands the lead to a competitor. It's the logic of an arms race.

Nobody set out to create this. It's something closer to a force of nature that we have set in motion and are now trying to steer.

If, like me, you've been a fan of the Jurassic Park movies, you may feel there's a similarity to all of this. Jeff Goldblum's character, Ian Malcolm said, "Your scientists were so preoccupied with whether or not they could, they didn't stop to think if they should". His question came too late. The dinosaurs were already in the park.

## **What You're Probably Feeling Right Now**

Anthropic, the company behind Claude, did something at the end of 2025 that I found genuinely useful, and published the results in March 2026. They asked 80,508 of their users across 159 countries and 70 languages how they actually feel about AI, beneath all the noise of the public debate.

Two things came out of it. People named real hopes and real fears in the same breath, and jobs and the economy sat right at the top of the worry list, alongside AI getting things wrong.

What that says for you, personally, is that if you picked this book up feeling a mix of fascination and dread, you're not confused. You are reading the situation accurately. The people paying close attention are not divided into a hopeful camp and a frightened one. It's the uncomfortable experience of holding opposing emotions at the same time.

That's who this book is for. Not the person who has decided AI is going to save the world and not the person who is convinced it's the end of everything. The person who suspects something extraordinary is happening and is looking for a way to think about it that goes deeper than the headlines.

Most of the what is coming out about AI right now does a good job of sounding the alarm. The articles and the podcasts will make the ground shake beneath you and keep you awake at night. They tell you the water is rising. They don't hang around and help you work out what to do when it's lapping at your door.

That's what I'm trying to do here. Go beyond the headlines.

## **The Question Under the Question**

The economic question (what will I do for work? how will I earn a living? what skills will still matter?) is urgent and very real. But I've come to believe that there is a deeper question underneath that AI displacement is bringing to light for many people.

When you meet someone, one of the first things you will hear is "What do you do?" The answer will almost always be some variation of "I am a (insert job title)". This could be anything from corporate CEO to working mom.

Before we go further, I used the word "job" but we don't actually have a good word for what people are losing to AI, and that's part of why the whole conversation feels so difficult. "Job" is too small, too transactional, too tidy. "Work" is too abstract, as though what we do for eight or ten hours a day doesn't shape the way we move through the world. And "profession" only fits the people who see themselves that way, which leaves out the millions who show up every day, not because it fulfills them, but because it keeps the lights on and the kids fed. None of these words capture the artist who never sold a painting, the gardener whose tomatoes are the talk of the neighborhood or the retiree who volunteers teaching adults to read.

The truth is that AI isn't just automating tasks. It's destabilizing the roles we've built our lives around. Some of those roles were chosen with passion; others were chosen from necessity; many were chosen simply because we stayed somewhere long enough for it to become comfortable. Whether we loved the work or merely endured it, it still gave our days structure and added depth to our relationships. When AI becomes able to do what we've done for years, what we have foolishly built our sense of self-worth on, we lose more than income. We lose orientation. And because we don't have a single word that captures all of that, we end up talking around it instead of naming it.

Research on AI-driven job displacement has found a cluster of psychological responses that look remarkably like grief. Shock and disorientation first. Then a slow erosion of professional identity. Anxiety about the future. Social withdrawal as friendships that were built at work slip away. And often a profound sense of betrayal, the feeling that you held up your end of the deal, that you worked hard, you improved your skills and you reliably showed up for years and the deal was broken anyway.

What comes through clearly in almost every account is how unprepared people were. Not for the money or the loss of the job itself, but for how much of their self-worth was tied up in the answer to "what do you do?"

That is what this book is most interested in. The economic side of it matters, and I am not going to pretend otherwise. But the inner disorientation is the deeper loss, and it is the one almost nobody is talking about.