This past June, the most powerful AI ever released to the public was switched off for every person on earth at the same time.

It was three days old.

I'm going to tell you that whole story, because I've thought about it every day since. But first, the honest part. I hate writing. This letter took me a month, and I've almost deleted it more times than I can count. I wrote it for my family and my friends, and somewhere in the middle I realized I couldn't keep it that small. Part of it feels like God put this on my heart to share. The rest is simpler: there are real problems starting to form around AI, real opportunities too, and no one I talk to sees the full picture yet. Fair warning: this is not a feel-good letter. Parts of it should scare you. They scared me. This letter is long. Read all of it. If you just thought about saving it for later, you're exactly who I'm talking to. I hate writing and I wrote this anyway. Take the hint. And if you truly only have ten minutes, go read *Your job, specifically* and *The fork*, or go read the short version, then come back for the rest. You'll want to.

Most people are using AI to look things up, write emails, get through schoolwork. That's not a guess. OpenAI studied over a million ChatGPT conversations and found roughly three quarters were people looking up information, getting practical advice, or getting help writing something. Nothing wrong with that. I used it to help write this letter, and without it I probably would've put my laptop through a wall. But those are the smallest things it can do, and the distance between how the average person uses it and what it's actually capable of is getting enormous.

On June 9, a company called Anthropic released a model called Claude Fable 5. Their own announcement said its capabilities passed anything they had ever made available to the public, and the test scores backed that up. And this was not a toy. The guy who runs Anthropic, Dario Amodei, had already written a whole essay about what AI at this level could do. His prediction: take the medical progress we'd normally make over the next 50 to 100 years and cram it into 5 or 10. Eliminating most cancer. He puts the possible drop in cancer deaths at 95 percent or more. Preventing Alzheimer's. Preventing or outright curing most genetic disease, the kind that has wrecked families for generations. He wrote all of that back in 2024, and in the same essay he said AI at that level could arrive as early as 2026.

Three days after Fable launched, at 5:21 on a Friday evening, a letter arrived at Anthropic from the U.S. Commerce Department. It ordered the company to cut off every foreign national on earth. Not just people in other countries. Every foreign national anywhere, including Anthropic's own employees. The company got 90 minutes to comply. And since there is no way to check the nationality of hundreds of millions of users in real time, there was exactly one way to comply. Turn it off for everybody. By the end of the night, Fable was dark for every user on the planet. Not slowed down, not limited. Shut off completely.

One letter. Ninety minutes. Gone.

And be clear about what they turned off. Not a chatbot. The level of machine Amodei said could cram a century of medicine into ten years. Somebody you love is going to get one of those diagnoses. For three days, the thing with a shot at beating it existed. Nobody asked you.

Here's what actually pisses me off. The version they killed, the one you and I could use, was already the fenced-in one. Same brain, most powerful capabilities locked behind safety filters. The version without the fences is called Mythos, and it was never offered to the public at all. Anthropic says on its own website what Mythos can do: find and exploit software vulnerabilities better than any other model, and better than all but the most skilled human security experts alive. That one went to a couple hundred approved organizations. Federal agencies. JPMorgan. Apple, Google, Microsoft. You were never allowed near it. And that doesn't just piss me off. It could be the most dangerous threat to America there is: turning from a healthy capitalist country into a monopoly one.

The dates on the government's own letters spell it out. Both versions went dark in June, and they did not come back at the same time. On June 26, a letter from the same Commerce Secretary turned Mythos back on for the approved club. The public got Fable back on July 1. When the lights came back on, they came on for the club before they came on for you.

So yes, it's back. Nineteen days after everything went dark, Fable started rolling back out across the world. And I can hear the shrug from here: it was off for nineteen days and your life didn't change one bit. Right. Because you haven't built anything on it yet. In a few years this thing is answering your office phone, tutoring your kid, reading your mom's scans. You didn't notice this blackout. You will notice the next one.

But look at how it came back. It came back after Anthropic rebuilt its safety filters, and after the government, in the Commerce Secretary's own words, worked with the company to "analyze and approve" the model. The off switch runs through the government, and now so does the on switch. You don't get a vote. You might not even get a warning. It came back this time. Nothing says it does next time. And honestly, forget next time. While Fable was still dark, the White House asked OpenAI to restrict its newest models to a small circle of trusted partners too, before they even launched. OpenAI did it, grumbling publicly that this shouldn't become the norm. Two companies in one month. This is the pattern now.

And ask yourself why that switch exists at all. Nobody builds an emergency shutoff for something that doesn't matter. Governments don't hold emergency meetings over toys. So let me say it straight. AI is the single most important invention of our lifetime, and I think it ends up the most important one, period. That's not just a young guy with an AI company hyping his own product. The CEO of Google calls AI more profound than fire or electricity. Geoffrey Hinton, who won a Nobel Prize for the science underneath it, compares it to the industrial revolution. Take that comparison seriously for one second. The industrial revolution came for muscle. This one comes for the mind. Your mind. The thing your paycheck is written against. Put me next to those two and "biggest invention of our lifetime" makes me the calm one.

One thing I refuse to let you walk away thinking: that this is too complicated for you, or that it's for people smarter or more technical than you. The biggest lie floating around about AI is that it's hard. It isn't. It's easy. The barrier to entry only looks high because nobody is teaching it. There's no class in school, your job isn't training you, nobody at the dinner table gets it, so everyone quietly assumes it must be complicated. It's not a technical skill. It's a thinking skill. Can you describe what you want? Can you hand over the context someone would need to do the job right? Can you look at the result and tell whether it's actually good? That's the entire job. If you could train a new employee, you can do this. The people who are great at this are not smarter than you. They just started earlier. And earlier is still available. Before I sign off, I'll tell you exactly where I'd start.

What you sell

Now think about what you actually sell for a living. Strip the job title off and it's your intelligence. The nurse deciding which patient can't wait until morning. The salesman reading a buyer. The teacher spotting the kid who's about to quit trying. The office manager who's the only person alive who knows how the place actually runs. We solve problems, we make decisions, we know things other people don't. That's the whole product.

What hasn't sunk in yet: on a growing list of those exact skills, the machine already beats the experts.

There's a test of graduate-level science questions built to be Google-proof. PhDs in the exact field score 65 to 70 percent on it. The top models score in the low 90s, and a PhD from a different field, given Google and half an hour per question, scores a 34. Math: in late 2024, mathematicians built original, unpublished problems specifically to stump AI, and Terence Tao, probably the greatest living mathematician, said they'd resist for "several years at least." They lasted eighteen months. Coding: the best model went from fixing under 2 percent of real software bugs in 2023 to roughly 95 percent today, and this February the industry started retiring the test because the machines maxed it out.

Bring all of that down to a smaller level. If the machine embarrasses the PhDs, the mathematicians, and the best programmers alive, what does it do with a stack of intake forms? A lesson plan? The quote you owe a customer by Friday? The routine half of your job was never a contest. And all of it gets better every few months, not every few years. Every few months the machine gets a raise. You don't.

Now, the machine can't do everything. Not yet. It's brilliant and lost at the same time. It doesn't know which problem is worth solving. It can't tell you what you actually want. It will hand you a confident, detailed answer that is flat wrong and never flinch. That gap, between raw horsepower and knowing how to aim it, is where you still live.

Don't relax about that gap. It's a job opening, and it's filling up fast. Your value doesn't disappear, it moves. It moves from being the smartest person in the room to being the one who can point the smartest thing on earth at the right problem and catch it when it's wrong. Hardly anyone has that skill yet. And if you're leaning on "people buy from people," careful. That's true, and by itself it will not save you. The people who start building it now are going to spend the next ten years so far ahead of everyone who waited

that it stops looking like a head start and starts looking like a different job market. And if you're 50 and quietly doing run-out-the-clock math, stop. This rewrite lands in the next few years, not at your retirement party. You get to choose which group you're in. You don't get to choose for very long.

I know all of that can sound abstract. So let me make it real, because I lived it.

The oh-no moment

I got into AI by using it to do my homework. That's the whole origin story. I thought it was the coolest thing I'd ever touched and I could not stop messing with it. At some point the homework stopped being the point and I just got curious about what else it could do. It kept surprising me. That curiosity has taught me more than school ever did, which is a sentence my teachers will be thrilled to read.

A while later I was interning at Colliers, a commercial real estate firm. My boss, Kent Willis, who's also a family friend, would have me "stack" properties near the Beltway in Houston, basically pulling together all the key details on a batch of buildings into one place. He'd ask me for five a day. Totally reasonable ask for an intern.

Using AI, I got to where I could do a hundred a day. Same work, twenty times the output.

It didn't feel like winning. It felt kind of ridiculous. I'm an unpaid intern, and I'm sitting there realizing that grown adults get paid full salaries to do the exact thing I'd just figured out how to do twenty times faster. That was the moment it really clicked. Not "AI is cool." More like, "oh no. The ground is moving under everyone, and almost nobody feels it yet."

So I called one of my best friends, Walker Deyo, and told him we had to move on this, and move fast. That's how Elevair started. Not from some polished business plan. From that exact feeling, that something enormous was happening and the window to do something about it was open right now.

That's a big part of why I'm writing this. That feeling wasn't a fluke, and it wasn't special to me. Your version of the five-a-day task is sitting somewhere in your week right now. Somebody is going to have the oh-no moment about your job. It should be you.

What an agent actually is

What I did with those properties wasn't some clever trick I invented. It has a name, and it happens to be the biggest shift in how work gets done right now. It's called an agent, and I'll be honest, the first time I watched one really work, it rattled me a little.

The AI you know is a chat box. You type a question, it types an answer, and then it just sits there waiting for you to go do something with it. Useful. Also the toddler version. An agent doesn't sit and wait. You hand it a goal and it leaves, breaks the job into steps, uses tools, searches the web, writes the code, checks its own work, and comes back when the whole thing is done. I still catch myself watching one run and thinking, I didn't have to do a single part of that.

That's what happened with my hundred properties. I stopped doing the work and started handing it off. Somewhere in there I stopped being the intern and became the manager of an intern that never gets tired, never complains, and works for about a nickel.

Almost no one understands this yet. Take advantage of that. The people pulling away aren't using AI, they're managing it, stacks of agents at once, one person running like a whole company. The intern who was proud of a hundred properties a day? Today I could build a system that does ten thousand in a day, one agent per property, another checking its work, thousands running at the same time. And before you file that under "kid on the internet exaggerating," the man who created Claude Code, the coding agent I use every single day, says there are days he manages tens of thousands of agents at once. He hasn't hand-written a line of code in eight months. One guy. That's not a prediction. That's his Tuesday.

The most credible skeptic in AI, Andrej Karpathy, a co-founder of OpenAI who this spring joined Anthropic to lead a research team (they are a way better company), says today's agents are glitchy, forget everything you tell them between sessions, and are probably a decade away from fully working. He's smarter than me, and my bet is the window is tighter than his. But look at what the same man posted the day after Christmas, two months later: "I've never felt this much behind as a programmer." The biggest skeptic in the field feels like he's getting lapped by the half-finished version. That should tell you everything about what happens when the finished one shows up.

So the gap means this. Your job isn't getting taken yet. It's getting rewritten. The value is sliding off doing the work and onto directing the thing that does the work, and catching it when it lies to you. That's a skill. It's learnable. I learned it as an unpaid intern who started out cheating at homework, and right now almost nobody has it. In that same Christmas post, Karpathy figured he could be ten times more powerful if he just properly strung together the tools that already exist, and said failing to claim that boost "feels decidedly like a skill issue." If falling behind is a skill issue for the guy who helped build OpenAI, what is it for the rest of us?

Pew runs the biggest survey in America on how people use AI. The 2026 edition doesn't contain a single question about agents. The number for how many regular people run them doesn't exist, because nobody thought to count yet. Empty means early. Start tonight at 22 or at 60 and you're early. And remember the club that got Mythos back first, the one you were never allowed in? This has no club. No list, no approval letter, nobody checking names at the door. It's just open. Walk in.

And it's already showing up at the front door of the job market. The most desirable first jobs in the country, the ones people spend four years of college grinding toward, don't look like they did even two years ago. Seven weeks ago the president of Goldman Sachs went on TV and described the firm as "a human assembly line," and said it's about to change the way factories did when the robots came. Goldman's CEO has already said AI can draft 95 percent of an IPO prospectus in minutes, work that used to take a six-person team two weeks. His words about the part that's left: "The last 5% now matters because the rest is now a commodity." The analyst classes college kids are applying into right now? Goldman, JPMorgan and Citi are shrinking them by as much as two-thirds.

It's already in the payroll data too. Stanford tracked millions of paychecks and found that since late 2022, employment for 22-to-25-year-olds in the jobs most exposed to AI has fallen 16 percent relative to everyone else, and the drop is speeding up, not slowing down. Same jobs, older workers: growing. The ladder isn't being pulled up. The bottom rung is snapping off.

And at the same time, the test to get in is quietly being rewritten. McKinsey now puts an AI in the room for final-round interviews and has candidates run it live. AI mentions in entry-level job postings nearly doubled in the past year. Shopify's CEO won't let a team hire a human until they prove an AI can't do the job. The jobs aren't gone. The test changed. And the people studying for the old test haven't noticed.

And the distance between the people who learn to run these things and the people still poking at a chat box like it's 2023 is not going to stay a polite little gap. It's opening into a canyon. What actually scares me is not the machine. It's how many people are watching all of this happen and just going through the motions. Scrolling past it. Nodding at dinner. Waiting for it to be somebody else's problem. The canyon doesn't care that you were busy.

Which drops us right in front of the question you've been dreading, the one everyone talks around at dinner. If the machine keeps getting better at the work, what happens to the people whose whole job was the work? I'll give you the real answer, both sides of it, and I won't flinch on either one. Next.

Your job, specifically

Yes. Jobs are going to be taken. I'm not going to dance around it, and I don't think it stays as slow as the experts keep promising. This spring, for the first time ever, AI became the number one reason American companies gave for cutting jobs, and it has held that spot every month since. The CEO of Salesforce cut his support team from nine thousand people to five thousand and his whole explanation was four words: "I need less heads." And you already saw the payroll data: it starts with people my age, and it's speeding up. Whatever today's numbers say, they're a photograph of a moving object.

We're not doomed, though. Everyone is just asking the wrong question. "Is my job safe" assumes jobs get taken by title, like a tornado that hits the accountants but spares the marketing department. That's not how this goes. Companies don't cut titles, they cut people, and inside every title the sorting has already started on one variable: who multiplies with the machine and who doesn't. The layoff list is secretly a fluency list. Two people with the same business card are about to have very different decades, and it won't be because one of them was smarter. It'll be because one learned to run the thing while the other kept calling it a fad.

That cuts the other way too, and I wish I could tattoo this on people: almost everyone would be better at their job if they made a real attempt to learn this. Better in a way bosses and customers notice. It costs less than your phone bill, it gets stronger every quarter, and most of your competition still hasn't touched it. There has never been a cheaper edge sitting in plainer sight.

Now zoom out with me, because the jobs story is one piece of something bigger. There was a long stretch of history when almost nobody could read. Reading was a professional skill. Monks had it, clerks had it, kings hired it, and everyone else took their word for what the book said. Nobody called that a literacy gap back then. It was just power. Mass literacy is what broke it. I think knowing how to use AI is this century's reading, and right now we are deep in the monks-and-kings phase. You watched this exact structure in June. Full strength for governments and giants, the safe copy and a shrug for everyone else. If nobody pushes back, that is the pecking order from here on.

Here's my theory, the one I'd defend in a room full of economists. The companies that get the best intelligence first get more than a head start, they get a compounding one. Better model, better decisions. Better decisions, more money. More money, earlier access to the next model. Run that loop for a decade and competition stops meaning anything. And the small firm loses, no matter how hard it works, because the other side gets to think with something it isn't allowed to buy.

This isn't hypothetical, either. That club from the start of this letter, the one whose lights came back on first? It has a name. Project Glasswing. Its stated job is defending the systems everybody depends on, power grids, hospitals, banks. And honestly, some of that is fair. Hospitals and power grids probably should jump the line. But notice what just became normal: somebody now ranks whose access to intelligence matters more, and the somebody is not you.

So here's where I stand, and you can hold me to it. The best intelligence in the world should be shared with all of humanity. With serious safety measures, absolutely, Fable proved how hard those are to get right. But "safe" and "reserved for the important" are not the same policy, and the second one gets dangerous quietly. China, who really holds the off switch, all of it is still coming in this letter.

Which brings the jobs question back to you, because the counterweight to all of this is almost embarrassingly simple: a country full of regular people who can run this technology themselves. Every plumber, teacher, paralegal, and shop owner who gets good at it is one more person the future can't quietly be decided for. That, and staying loud with our politicians while the rules of this era get written, because they are being written right now, the same way social media's rules got written while we were all busy posting. We get one setup phase, and we're in it. That's the real reason I keep saying learn it. The version of this that goes right needs a lot more of us in the room.

One more thing before you exhale. Everything I just said is about work that happens on a screen. Words, numbers, code. Most people hear that and quietly think, good thing I work with my hands. The people building this technology are spending billions betting they can build hands too.

About those hands

In 2021, Elon Musk announced Tesla would build a humanoid robot. He had no robot to show, so a guy in a white spandex suit came out and danced. People laughed, and honestly, the clip is still funny. Five years later he's telling shareholders this robot will be the biggest product of all time, bigger than the cell phone, and that there could eventually be tens of billions of them. And before you write that off as Elon being Elon, Morgan Stanley, about the most buttoned-up bank on Wall Street, projects a billion humanoid robots on Earth by 2050.

Slow down and actually picture a billion, because I swear, people are just asleep on this. A billion robots is three for every man, woman, and child in America. It's six entire American workforces, except this one doesn't sleep, doesn't call in sick, and gets smarter with a software update. And a billion is starting to look like the cautious guess. The bank that tracks this closest multiplied its own forecast by six in about a year, because robot costs were collapsing twice as fast as anyone modeled. The man building them says tens of billions. The price of the cheapest humanoid has fallen by about three quarters in two years. Every time the experts re-run the math, the number gets bigger. Forecasts that only get revised in one direction are telling you something.

And honestly, forget the billion. A number that big sounds far away, and far away is easy to ignore. The number I actually watch is five million. That's the point where I believe America starts feeling this for real, in wages, in hiring, in which businesses quietly stop needing a night shift. Five million robots is barely three percent of the American workforce, and three percent is plenty. It didn't take half the mortgages in America going bad to cause 2008. So when do we cross five million? Run the banks' own adoption curves and you get somewhere between the mid-2030s and the early 2040s, depending on whose math you trust. You already know which end I lean toward. My honest guess is the front edge, around 2035, and it moves closer every time the price falls faster than the forecast, which so far has been every time. Either way, that's inside your career, and right at the start of your kids'.

So back to Elon. How do you take a man like that seriously? There's a trick to reading him. This June, the New York Times went back through fifteen years of Musk's public promises, six hundred and two of them with actual deadlines, and scored the results. He delivered about one in five on time. Self-driving was "a year away" for close to ten straight years, and even the robot is late: five thousand promised for 2025, still in R&D this January. And yet. The people who laughed at reusable rockets watched a SpaceX booster land itself for the thirty-fifth time a few weeks ago. Electric cars were a punchline until every automaker on earth was chasing his. With Elon, the year on the label is negotiable, but the destination has a habit of arriving. My read, on the record: real humanoid robots, in our lifetimes, and probably earlier than feels comfortable.

Today, honestly? The robots doing actual paid work on actual shifts number in the low hundreds, worldwide. The famous robot fleet at BMW turned out to be one robot. Battery life dies before lunch, and a robot hand still can't come close to the seventeen thousand nerve endings in yours. But watch the direction, not the demos. China's government just ordered ten thousand humanoids into real jobs by the end of this year, and that's an order, not a forecast. You can buy a Chinese humanoid today for about the price of a used golf cart. The first American humanoid company just announced it's going public. The bodies are coming. The brains, as you've seen, are already here.

And that's the part that should hold your attention, because it plugs straight into everything I've told you about access. In June we watched the government switch off the smartest software on earth with a letter. Annoying, scary, survivable. It came back. Now picture that same power applied to a workforce. Machines that think and move, and somebody deciding who gets them. Whoever holds the keys at that point doesn't just out-think the competition, they out-build, out-repair, out-work everyone who wasn't on the list. Because a gated model is one thing, but a gated workforce is a different kind of country.

Which is why I need you to understand the fork we're standing at. All of this goes one of two ways, and the difference between the two futures is not the technology. It's who gets it. That part of the story is next, and it's the reason I started writing this letter in the first place.

The fork

Back to June, because I owe you the rest of that story, and I'm not going to be calm about it. The legal authority behind that letter had never been used that way in its history, not once. There was no published standard the company failed, no hearing, no appeal, no way to contest any of it. When nearly eighty security executives signed an open letter demanding a transparent process for decisions like this, they got silence. And for nineteen days, the best research tool ever built sat dark for the exact people using it to chase the cures from the start of this letter. Those are the receipts. Now the theory they prove.

But first the government's side, because I promised you honesty and their side isn't nothing. Researchers at Amazon had found a way to trick the model past its security guardrails. The model's own maker had spent months telling everyone how dangerous these capabilities could be, and as one security researcher put it, if you describe your product like a weapon in every press release, eventually a government takes you at your word. The White House says it asked for a fix first and got refused. The company disputes that flat out.

Fine, reasonable people can argue about the jailbreak. What nobody can defend is how it was done. No standard, no process, ninety minutes, worldwide. And the detail that actually keeps me up: two days before the letter, Dario Amodei published an essay asking the government to build exactly this kind of safety authority, done properly, with testing and written standards, including a guarantee that citizens keep access to AI when a government acts badly. He asked for a referee. He got a switch.

You want to know how asleep people are? Nobody even polled it. The most important technology humans have ever built got switched off for the entire world and switched back on nineteen days later, and not one polling firm thought Americans would care enough to ask. Three quarters of us tell pollsters we worry about AI concentrating too much power. Then the concentration demonstrated itself in broad daylight, and the country scrolled past.

Now the danger I named at the start, the monopoly thing, because it's time to prove it. Everything that works in America works because of competition. Anyone can beat anyone. The plumber who answers his phone takes the job from the one who doesn't. The startup nobody heard of eats the giant that got lazy. That's the whole reason things here get better and cheaper every year, and we're so used to it we think it's automatic. It's not automatic. It only works if everybody gets to compete with roughly the same tools.

Now add June. The best intelligence on earth, the thing that already beats the experts, sits behind a government letter, an approved list, and a price. Remember my loop from the jobs section: better AI, better decisions, more money, first in line for the next model. If only certain companies are allowed the best intelligence, winning stops being about who's better. You win because you're on the list, and you lose because you're not, no matter how hard you work. Who wins, what things cost, what you get paid, more and more of that gets decided by who has access instead of who does the best job. That's not capitalism anymore. It just still has the name on the door.

And there's no working around this one, because intelligence isn't one tool. It's the thing that makes every other tool better. Repossess a man's truck and he buys another truck. Repossess his access to the best intelligence, and June proved they can, ninety minutes, no appeal, and he isn't behind in the race. He's out of it. That's what a monopoly country is. Not cartoon rich guys in a smoky room. An economy where the list matters more than the work, one list at a time.

And this is already being built, in the open. Ten days before the letter, an executive order quietly created a system where AI companies volunteer their newest models to the government for review, up to thirty days before the public sees them, with the capability bar set by a classified process. And that OpenAI grumble I mentioned earlier deserves its full quote. Their newest models went, in their own words, to "a small group of trusted partners whose participation has been shared with the government," and their protest, published the same day, said the arrangement "keeps the best tools from users, developers, enterprises, cyber defenders, and global partners who need them." The gatekeeper wrote that about its own gate.

And it isn't just governments. The strongest tiers at Google and xAI run two and three hundred dollars a month. Meta, which spent a decade preaching open AI for everybody, just shipped the first closed model in its history, and its CEO now talks about being "careful about what we choose to open source." Government gates, member gates, price gates. Stack them and you get an org chart for intelligence itself.

But the thing to see about that chart is not where you sit on it. It's that everyone on it has a row above them. The companies answer to the letter. The letter answers to whoever holds office. Anthropic sat on top of every tier that existed, the safety reputation, the government contracts, all of it, and still went dark on ninety minutes' notice. When access decides everything, nobody owns their spot. Everybody rents. So this isn't the 99 percent

versus the 1 percent, and I'm not interested in that fight. Nobody, at any size, wants to bet their success on who gets access to the best intelligence. Not a plumber, not a hospital, not a bank, not a trillion-dollar lab. The moment access beats effort, every business in America is one relationship away from irrelevant, and the giants most of all.

I'll take it one step further, and this part is mine. Every unequal stretch of history had a floor under it, and the floor was need. The people at the top needed the people at the bottom. Fields, factories, armies, customers. That need is why every gilded age eventually bent back toward the middle. Strikes worked because labor was needed. Votes carried weight because taxpayers and soldiers were needed. Now AI does the thinking, robots do the moving, and for the first time ever the need itself starts to disappear. Take away being needed, and whatever the top still sends down isn't owed to anybody anymore. It's charity, and June showed you what happens to things that only flow by permission. They can be shut off in ninety minutes.

I'm not predicting doomsday. I'm pointing at the floor, and telling you this technology is pulling it out from under us faster than laws have ever moved. The people building it can see the same thing, and they're already buying insurance: Anthropic just put two hundred million dollars into studying what AI does to workers, a sitting senator filed a bill to pay every American a dividend from AI companies, and Sam Altman has started talking less about basic income and more about everyone owning a piece. When the most optimistic people on earth start buying insurance, believe them.

Here's the urgent part. Both futures in the next section are built from parts that already exist today. The lists exist. The order is signed. The gates are half-built. Nothing about which future we get is decided yet, and everything about it is being decided now, by whoever shows up. The whole reason I wrote you this letter is that I'd rather it be decided by a country full of people who can run this thing than by a room you and I will never be in.

What's truly coming

Everything I've shown you so far already exists. This part is about what the people building it say comes next, because that's what decides how much time you have. Their own words, on the record, with dates.

Sam Altman, who runs OpenAI, wrote last summer that "we are past the event horizon; the takeoff has started." His company has put its next deadlines in public: an AI that works as a real research intern by this September, two months from now, and a fully automated AI researcher by March 2028. An AI that is itself the researcher, not a tool that helps one. They admit they might miss those dates, but look at what they think is close enough to put a date on.

Dario Amodei, whose predictions opened this letter, wrote in January that powerful AI, his term for a machine smarter than a Nobel Prize winner across most fields with millions of copies running at once, could be as little as one to two years away. He's careful to say it could take longer, and just as careful to add there's "a very strong chance it comes in the next few." He calls what that adds up to a country of geniuses in a datacenter. Mark Zuckerberg says Meta has "begun to see glimpses of our AI systems improving themselves," improvement he calls slow for now but undeniable, and that superintelligence is "now in sight." And the cautious one in the group, Demis Hassabis at Google DeepMind, who has a Nobel Prize for his AI work, says five to ten years. Five to ten years. That's the careful answer now.

You don't have to trust a single CEO, though. An independent research group called METR measures one thing: how long a stretch of human work an AI agent can finish on its own. In early 2025 that number was doubling about every seven months. When they re-measured this January with newer models, the doubling had sped up to about every four months. The best models now run around five hours of human work at a stretch, at a coin-flip success rate, and yes, that caveat is real. But their own extrapolation puts agents finishing month-long projects before this decade is out, if the trend holds.

And notice what OpenAI's deadlines are really about: AI doing AI research. Google says on its own blog that three quarters of its new code is now written by AI and approved by engineers, up from half last fall. Amodei says the feedback loop "may be only 1-2 years away from a point where the current generation of AI autonomously builds the next." Once that loop closes, the machines stop waiting on human researchers. That's the mechanism behind every date you just read. A research team that copies itself.

The money is saying the same thing. The four biggest tech companies spent around four hundred billion dollars on AI infrastructure last year and have told investors this year lands somewhere between six and seven hundred billion. Seven datacenter sites are under

construction for OpenAI alone, planned to draw about as much power as New York City at its peak. They are restarting the nuclear plant at Three Mile Island just to feed Microsoft's share of it.

Could they all be wrong together? Sure. They hedge it themselves, and the loudest forecaster who bet on all of this landing by 2027 has already pushed his own date into the 2030s. So don't take the bold numbers. Take the slow ones. Karpathy's decade before agents really work. Hassabis's five to ten years. The robot math from the last section. The cautious estimates, from the most careful people in the field, all land in the same window: the early-to-mid 2030s. That's why the two futures below are set in 2035. The year comes from the careful numbers, not the wild ones.

So here's 2035, both ways it can go. Not a think tank's forecast. Mine. And I'll stand behind every word of it.

The bad one first. It's 2035 and nothing dramatic ever happened. No crash, no uprising, no headline. The gates just never opened. The best intelligence on earth runs three tiers deep behind approval lists, and everything important runs through it: medicine, money, law, energy. Your business competes against companies that think with something you're not allowed to buy, and it loses politely, every quarter, forever. Your kid's brilliant idea dies in a notebook, because the intelligence to build it rents by permission and permission has a waitlist. The cancer drugs exist. They arrive in order, and your family's place in line depends on tiers nobody voted on. Effort stopped deciding anything years ago. Standing decides, and everybody just knows it, the way everybody once knew not to question the king. And the people at the top didn't win either. They spend their lives defending their row, watching the company above them and the letter above that. A whole country, top to bottom, organized around not losing access. Nobody chose this. It compounded while everyone slept.

Now the good one, because it's just as real and pieces of it already exist. Same year, 2035, and intelligence got treated like electricity. Not free, not unregulated. Just available, the way power and roads and the internet were available, to anyone with something to build. Nobody wins electricity. Everybody builds on it. Your kid has the best tutor on earth in her pocket, and so does every kid in every zip code, and the difference in what they become goes back to being effort, not access. The town doctor walks into every appointment with a specialist in her phone. Your fifty-person company fights like a five-hundred-person company, and I can tell you that part works, because building those companies is what I

do all day. The cancer drugs still arrive, except they arrive the way vaccines did, everywhere, in the same decade, and the line to get them is the line at the pharmacy. The giants do better here too. Bigger markets, faster science, customers who can afford them, and no empire hostage to a Friday letter. None of this is fantasy. Switzerland already runs a fully open national AI model. Singapore just made months of frontier AI access a government benefit for workers. China put AI class in every school in the country. The pieces exist. They just aren't assembled here yet.

Same technology in both endings. The only difference is the architecture, and the architecture question is one question: how many hands hold the thing? The real lesson of June isn't that a government did something scary, it's that it could, and it could because the best intelligence on earth lives inside a couple of companies, which means it can be switched off with a couple of letters. You don't fix that by asking gatekeepers to be nicer. You fix it by making sure copies exist that nobody can recall. That's what open source means: models whose weights, the actual brains, are published for anyone to download, run, and build on. Once the weights are public, there is no off switch. Not for a company, not for a government, not for anyone.

And before somebody tells you that's naive, America already ran this experiment. In the nineties, the government classified strong encryption as a munition and tried to stop it at the border. People printed the source code in a book and shipped the book, because books are protected speech. A federal court ruled code is speech. By 2000 the controls had collapsed, encryption went everywhere, and history's verdict isn't even controversial: the open side won, American industry won, and every single one of us got safer. You bank on your phone today because openness beat control the last time they fought.

And this time there's a scoreboard. While our best models sit behind letters and lists, China is flooding the world with free ones. Don't mistake that for generosity. It's strategy, and it might be the sharpest play they've run against us yet: give the world your models free, let every country and company build on them, and you own the defaults for a generation. It's the same playbook you just read about, turned back on us while we charge admission. On the biggest open marketplace for AI models, three out of every five tokens already run through Chinese models. Malaysia's national AI runs on one. Beijing pitches free models to the developing world as the cure for American gatekeeping, and the pitch is landing. There's no beating free-and-everywhere with locked-and-approved. The only counter is better, and just as open. Even the White House's own AI plan admits open models carry "geostrategic value" and says the government "should create a supportive

environment" for them. Those are the government's written words. June was the government's actual behavior. One of them has to win, and the default settings of the entire free world ride on which.

So yes, safety measures, real ones, written down, with standards and hearings and appeals, the way this country regulates every other powerful thing it builds. I told you earlier where I stand, that the best intelligence in the world should be shared with all of humanity. Now you've seen why.

Which leaves one big question standing, because none of this is free. These labs burn money like rocket fuel, the papers keep calling it a bubble, and if the bubble people are right, everything in this letter softens. They're not right. Why they're not right is the most misunderstood part of this whole story, and it's next.

The bubble

Let me make the skeptics' case for them first, because the losses are real and I'm not going to pretend otherwise. Leaked financial statements verified by the Financial Times show OpenAI brought in about thirteen billion dollars last year and spent thirty-four billion doing it, a twenty-one billion dollar operating loss, and their own investor documents don't project profitability until the end of the decade. The market has noticed. The Nasdaq corrected this spring, and in late June Oracle, one of the companies building the datacenters, had its worst week since the dot-com era. Even Sam Altman, asked point blank whether investors are overexcited about AI, answered "my opinion is yes" and predicted someone is going to lose a phenomenal amount of money. So when your buddy says it's a bubble, don't argue with him. Parts of it probably are.

Argue with what he does next, which is shrug and go back to his phone. Because that same Altman answer had a second half. Asked whether AI is the most important thing to happen in a very long time: "my opinion is also yes." People treat "bubble" and "real" like opposites. History says they usually arrive together.

In the late nineties, telecom companies buried more than five hundred billion dollars of mostly borrowed money in fiber optic cable. Twenty-three of them went bankrupt, and in 2002 only about three percent of that fiber was even lit. The investors got flattened, the fiber stayed in the ground, and the dirt-cheap bandwidth it left behind is a big reason streaming and the modern internet exist. Railroads, same story a century earlier: a

quarter of the country's rail mileage went into receivership after the Panic of 1893, and the country kept the tracks and the cheap freight. Amazon fell ninety-four percent in the dot-com crash, the crash that supposedly proved it was all fake, and it was quietly the most important company of its era the entire way down. Wrong money, right thing. That's the most common ending in the history of new technology.

And the demand side of this one is not smoke. OpenAI's revenue went from thirteen billion to a twenty-five billion dollar annual pace by March. Anthropic told investors in May, in a setting where lying is securities fraud, that it had crossed a forty-seven billion dollar run rate. Half the population started using this stuff within three years of it existing, faster than the PC spread, faster than the internet. And the price of yesterday's best AI keeps collapsing, roughly forty-fold per year by some measures. The losses aren't the smell of a scam. They're the cost of building the factory while the line out front keeps doubling.

But set all of that aside, because you might still doubt it, and it turns out not to matter. Run the skeptics' ending yourself. The crash comes, valuations get cut in half, some funds get wiped out. Now walk back through this letter and find one thing a falling stock price undoes. The model that beats the PhDs still exists at zero dollars a share. The agents still run. The switch still sits exactly where it sits. Jeff Bezos, whose own company was the ninety-four percent crash, calls this an industrial bubble: when the dust settles, the investors eat the loss and society keeps the infrastructure. A crash changes who gets rich off this, not what exists, and not one warning in this letter. The bubble argument is about money, and mostly other people's money at that. Your question is what you do about all of it, and I promised you an exact answer.

Your move

First, quick, the world you're walking into, because it changed while I was writing this. Since February, ChatGPT shows ads in its free version. Roughly half of the new articles on the internet are now written by AI. Merriam-Webster's word of the year was "slop." When Google answers your question with AI at the top of the page, clicks to the actual sources get cut roughly in half. The internet is filling with machine-written everything, the machine that reads it for you is becoming the front door, and the front door sells ad space now. So nobody is short on information anymore. What's scarce is a person who can tell good from bad, and real from confident but wrong. Judgment.

That's not philosophy, it's in the hiring data. The number one skill employers say they want is analytical thinking. AI literacy is the fastest-growing skill in American job postings. And Harvard Business Review just named the opposite problem, workslop: AI output that looks like work and says nothing, and takes about two hours to untangle every time it lands on someone's desk. In an office drowning in machine output, the person who can direct the machine and then judge what it hands back is the person everyone suddenly needs. The typing was never the job; the thinking was, and it still is.

So, the promise: exactly where I'd start. Two moves, one evening.

Find your five-a-day. Somewhere in your week is the task you do over and over. The quote, the schedule, the report, the follow-up emails, your version of my five properties. Tonight, open the AI and make it do that task with you. Not a poem. Not a trivia question. Your actual work. Hand it the same context you'd hand a new hire, what the task is, what good looks like, one finished example, and when it gets something wrong, and it will, tell it what's wrong and make it go again. The first version will be eighty percent. The third one will scare you.

And pay the twenty bucks. ChatGPT Plus, Claude Pro, Google AI Pro, they all run about twenty a month. The strongest models regular people are allowed to touch live in the paid tiers. The free version is the loaner car, and it has ads on the doors now. Most people who tell you AI is dumb have only ever driven the loaner. Twenty dollars is what it costs to stop guessing about the thing this entire letter is about.

Do that one evening this week, and somewhere in the next month you'll hit your version of my hundred properties, the moment you stare at the screen and feel the ground move. Mine came as an unpaid intern. Yours might come on a Tuesday night at the kitchen table. That moment is the entire assignment. The tools, the tricks, the prompts, you'll chase all of that on your own afterward, because after it you won't be able to stop. And every single person who gets there is one more of us in the room while the rules of this era get written.

And when the moment hits, don't stop at the chat box, because the chat box is the smallest version of this. The real power is agents. Stop asking the machine questions and start handing it whole jobs. Brief one the same way you briefed your five-a-day, then let it run: it leaves, does the work, checks it, and comes back done. The tool I'd learn is Claude

Code, the same one I run my company on. A month of that and you're one person directing a staff that never sleeps, while your competition is still typing questions into a box.

When it happens, I want to hear about it. Which brings me to the last thing.

The door

One last idea and I'll let you go. It's the one I'd bet on. Remember Karpathy's biggest complaint about agents, that they start every session from scratch and remember nothing you tell them? Everybody is racing to make AI smarter. Watch the other race, the one to fix the forgetting. My theory: the next winner isn't a smarter AI, it's the AI that knows you. Your customers, your prices, your history, which jobs you lost and why, what your best month looked like and what actually caused it. The day the forgetting gets fixed, this stops being a tool you pick up and put down and becomes an employee whose tenth year knows everything its first year learned. The first business in a town whose AI truly knows the whole operation is going to be brutally hard to catch.

Building exactly that is what my company does. Elevair wires a business's entire customer engine together and gives it a brain that knows the operation. So yes, I'm talking my own book here. I'm also all in on that book, my money, my twenties, all of it, which is more than most people talking their book can say.

Now the door. Text me: 713 826 9548. A question, an idea, the thing you're building, the thing you're stuck on, send it. I will help anyone who actually wants to learn this, because everything in this letter comes down to how many regular people get good at it in the next few years. You won't be bothering me. This is what I do all day, and the people who move early are exactly the people I want to know.

So don't just read this and nod along. Build something. Fix your business with it, start the one you've been sitting on, learn the skill that changes your next ten years. The window is wide open right now, and I want the people I know walking into this decade awake. Honestly, I want some of you building it with me.

That's the letter. I hate writing, and I wrote you all of this anyway. Take the hint.