

# Print Edition

Monday 24 August 2026

EDITION NO. 32 · 1 PIECE · 10 PAGES

---

## IN THIS EDITION

1. The Infinite Canvas and the Agency Problem · Future History
- 

1

## The Infinite Canvas and the Agency Problem

BY DANIEL JEFFRIES · FUTURE HISTORY · 23 AUGUST 2026



Is anyone doing real work with AI agents?

AI agents can do almost anything. And that's exactly why most people have no idea what to do with them.

They're an infinite canvas, a wide open road. Where do you start? Anywhere. You just talk to the agent and ask and iterate and figure it out as you go.

But it doesn't feel to me like most people are doing this.

I had a friend ask me "How do I learn how to use agents?"

And I just thought, "why would anyone even ask this? You just talk to it and figure it out." You talk to it and ask and then keep iterating. That's it.

It's a bit like this meme:

## NORMAL PEOPLE

- I don't think I have the skills
- What if I fail?
- Will I look cringe?

## BUSINESS OWNERS



The truth is most people don't think "I'll just try and fail and then adjust until it works." They don't have a backlog of a thousand projects they always wanted to do. They don't have a million things they want to get done. They just want to know how to do one thing that is already worked out for them. And that's not how agents work today.

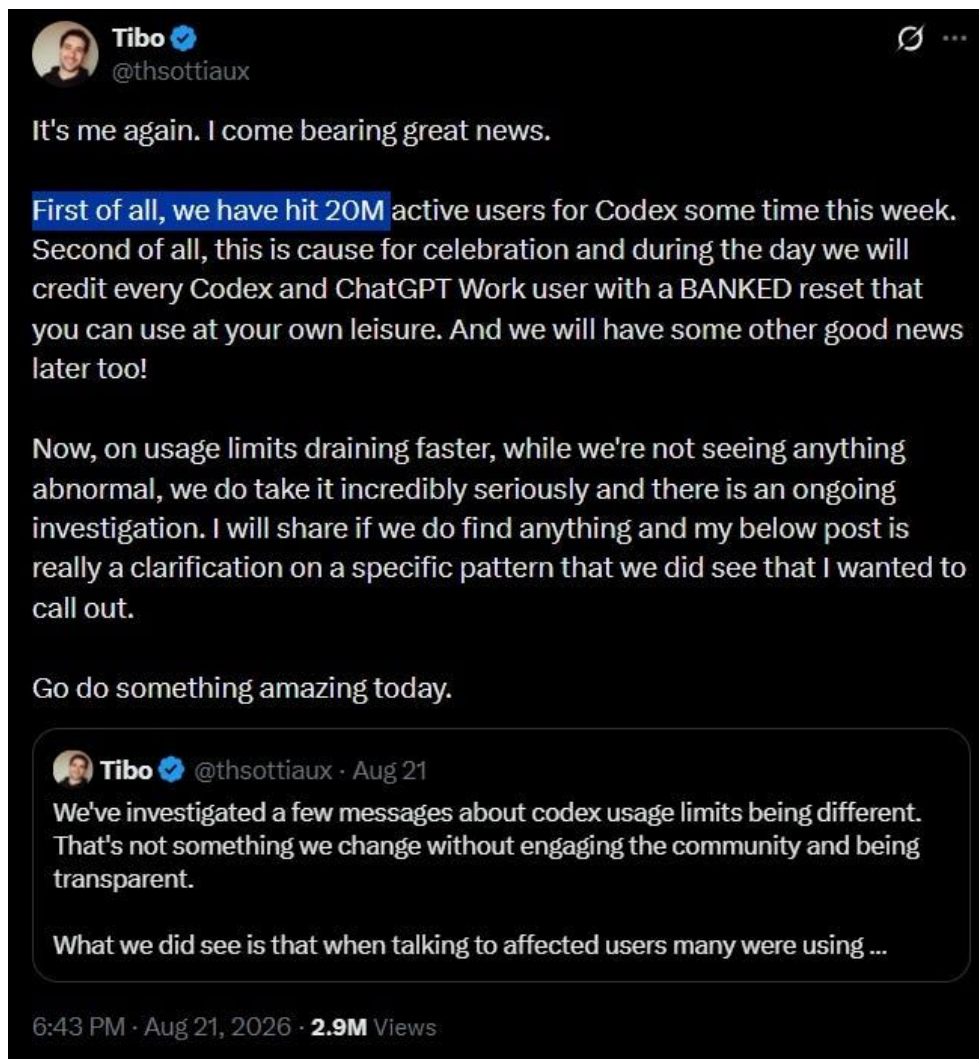
But if you're terminally online the way I am, people do seem to be doing something with agents, at least if you believe X and YouTube.

They're running dozens of worktrees at the same time or snatching up 10K computers to run agents locally. NVIDIA just put out a workstation that costs 100K to run even bigger models. People will buy them.

YouTube is saturated with influencer videos about agents shouting about how "this is INSANE!".

OpenAI and Anthropic are the fastest-growing companies in history. They make the British East India Company look tiny. They've got billions of users paying them massive amounts of money. Companies are burning through their yearly AI budgets in months.

OpenClaw and Hermes have hundreds of thousands of stars and millions of downloads. Codex went from under a million users to 20 million in months.



Grok Bot is hot right now and promises to make inroads with regular folks. But only time will tell if it's easy enough to be the Instagram moment for agents.

But I ask you. If you really dig deeper and the question is:

“Is anyone actually doing any real work at all with these things?”

The answer is yeah, a few people. But mostly no. So I started to ask why? And as I dug deeper and deeper what I found is that too often it feels like working with agents is a sys-admin exercise. Every video on X and YouTube is basically:

Here's how to set them up in 100 easy steps!

Register this API key. And then this one. Set up this system. Check it. Install twenty plugins. Check. Fix. Run the doctor. Troubleshoot. Do it again.

It's a bit like running a marathon. You do it to say you ran the marathon, but what's actually at the end? What's the thing you were running towards or were you just running to run?

I've got dozens of agents running in parallel! Okay but what are they doing? The point seems to be getting them running in parallel instead of having an actual work task that requires agents running in parallel.

What's being done, other than coding and making more agents?

And the answer for most folks is: not much.

I know how I use agents. Research. A writing sparring partner. Editing. Website creation and deployment. Checking the background of an accountant. Researching travel. Language tutor. Coding. DevOps work.

But mostly I what I found from other people is that they're spending ninety percent of the time setting up the agents instead of doing anything real with them.

Early in 2026, Will Mandis wrote about this phenomenon in a scathing article called "Tool Shaped Objects" where he talks about a stunningly beautiful Japanese wood shaver from the 1700s, called a Kanna, that still gets made today. Its blades are forged from the finest steel and it takes dedicated and skilled craftsmen months to make one and set it up.

It's a fine machine and it shaves wood with unmatched precision. It's gorgeous, a machine work of art.

There's just one problem.

"It is beautiful. It is also, in the economic sense, worthless. A power planer does the same work in a fraction of the time. The Kanna exists so that the setup can exist."

So what does all this mean? Where is this all going if all we're doing is a painful exercise of setup for setup's sake?

Some folks see the fall of Babylon. Ed Zitron is no fan of AI and his writing practically salivates at the image of AI burning like Byzantium after its walls come crashing down. He digs into the scary numbers behind massive hyperscaler spending on AI and the spending is eye-watering:

"Microsoft has dedicated its Fairwater data centers (however much actually exists) entirely to OpenAI, much like Amazon has for Anthropic with however much of its massive Indiana-based Project Rainier has actually been turned on, and much like Google is in talks to backstop a \$15 billion data center project for Anthropic, along with data centers with Cipher Mining and TeraWulf and a \$35 billion private credit-funded Broadcom-backstopped deal where Google will sell Anthropic its TPU AI chips, put them in a Google-built data center, and rent them back to Anthropic."

Yeah, but that's a lot of demand obviously, right?

Maybe.

He goes on:

"I want to spell this out: when you remove Anthropic and OpenAI's compute spend, I am not confident that Google, Microsoft and Amazon have much of an AI business.

While many people believe — largely because the big three refuse to break out their actual AI revenues or disclose their customer concentration — that they have AI revenues coming from a diverse set

of different customers, the reality is that their largest cloud customers, let alone AI customers, are two companies that can literally not afford to pay them without a near-infinite flow of venture capital or debt.”

That is damning. Basically two companies driving all demand.

But Zitron seems to miss that those two companies mean massive demand from other companies, big and small, for those models. If nobody is using it then why are those companies growing faster than any other companies in history? Not just in revenue but in subscribers and API users?

So someone is using AI. But the question is, who? Like who is really doing anything with it, if at all?

And that’s where I’ve got some real nagging doubt.

I’ve been wondering whether anyone outside coding, the AI industry itself, and marketing is actually using agents to do real work? And I don’t mean asking an agent to rewrite an email. Or generating a picture of a cat dressed like Napoleon. I mean handing meaningful parts of your life and work to a machine and letting it help you do stuff autonomously.

At first I thought the problem was mostly technical. Agents are unreliable. They hallucinate. They need babysitting. Setting one up still feels like a sys-admin certification exam designed by a committee of sadists.

But that I realized it was something much deeper.

## The Agency Problem

Most people absolutely cannot, or will not, use agents in their current form.

But it’s worse than that.

The big labs keep warning that AI is coming for everyone’s jobs. America’s doomer movement talks about the machines as if Skynet is already filling out the paperwork. NIMBY coalitions fight data centers over water use while happily golfing and munching burgers and posting about it on their phone on social media, never noticing the irony. Gen Z and Gen Alpha workers quietly sabotage AI productivity mandates with feigned acceptance and performative use because helping the machine do more of their job feels a lot like volunteering to erase their own future.

This is not the Luddites smashing factory machines with hammers. It is subtler, quieter, more passive resistance. People sandbag the rollout. They use the approved tool for the approved little task, paste the result into a document, and change nothing fundamental about how they work. They comply with AI without ever really adopting it.

My wife is smart and perfectly capable of using GPT. She’ll ask it for travel recommendations or lists. She once used it to help write stories while she was learning a new language. She also used an online AI tool that translates TV shows to help her learn, and that felt a little like magic. Useful, real things. But it hasn’t become her interface to the world. It is still a tool she visits occasionally, not a collaborator woven deeply into her day. It’s like taking a two week vacation, not moving to a new country.

And honestly, why would it be? The dominant interface is an empty box staring back at you.

It's just blinking at you.

People ask me, "What can you do with AI?" My answer is: anything you can imagine.

That sounds liberating. In practice it is terrifyingly vague.

For high-agency people, "anything" is rocket fuel.

They have a million projects they have always wanted to try, companies they want to build, questions they want answered, trips they want planned, systems they want fixed, books they want written. Their problem is not finding a use for AI. For people like me, the problem is that we can't do all of it fast enough.

But most people do not walk around with a private backlog of impossible dreams. Their job is to do their job and go home. Their job is usually to be told what needs doing, master the process, complete it, and stop thinking about it when the day ends.

Give them an infinite canvas and they don't see freedom. They see another task: invent the task.



Most people aren't imagining anything at all. They don't need or want to do anything more. Give them the ability to do anything and they'd do absolutely nothing and be perfectly happy about it.


Most people dream of the beach and living by the beach. I live by one and almost never go. I'm bored out of my mind there unless it's for a quick swim or to watch my wonderful wife enjoy the waves. When she's happy, I'm happy. But I'm not happy because I'm at the beach. I'm happy because she is.

When I do go, I don't understand why everyone is so happy doing absolutely nothing, so I usually read or fall asleep. The beach is heaven to most of the world. I'm just thinking about all the other things I'd rather be doing.

The articles I want to write. The books I want to read. The business problems I could be solve. All the beautiful business ideas I'll never get to do because their just isn't enough time in the world. What's happening in the world and politics and whether I can do anything about it?

If you're building agents right now, you're not building them for the folks using Instagram or TikTok to look at funny videos. That interface/format/ease-of-use simply does not exist with where machine intelligence is today. It will take years to get there.

You're building them for a very small group of people. Think the red area on the diffusion of innovation curve:



We're not even on the early adopters part of the curve yet and so folks on the other part of this curve just won't be able to work with what we have now because it's too much setup and too many very rough edges. We're not at the iPhone 20 agents, we're at the PC in the 1980s version.

Some people really want to give agents a go and if they're anywhere outside of that red area it doesn't go very far, even if they're bright and clever.

Folks say to me, "So where do I start?"

The short answer is "I don't know."

And I don't know why they don't know.

Like I said earlier, my best friend is a smart guy and pretty tech-savvy. He said, "I want to learn what I can do with these things. How do I learn? What can I do with them?"

The answer is: you start talking to them, then try doing the things you've always wanted to do. They'll fail at some of it. You push back, iterate, figure out what's wrong, and eventually work it out together.

The problem is, there is nothing he really wants to do, nothing that's burning brightly in the back of his mind, nothing that he's wanted to do for forever but just couldn't do because he didn't have the time or the money or the skill or all of the above.

And if he could figure out what to do, he'd run right smack into the next problem.

## The Verification Problem

Sometimes agents don't know what they're doing, or they propose a stupid solution. You have to sniff that out. You have to smell it coming a mile away and say, "That's not right. What are you talking about? Let's do this instead or research it and find a better answer."

That's the verification problem in action. It's the biggest problem in AI right now, not alignment, no matter what the doomers tell you. If you don't know how to do the thing, or you're not the kind of person who can figure it out by asking questions, thinking critically, and doing your own research, then you have no idea whether the AI's answer is any good. If you have no taste, it doesn't matter if AI barfs out a design. You don't know enough about negative space, kerning, or how fonts affect people emotionally to tell whether it works.

The biggest irony in AI today is that the people getting the most out of AI right now are the ones *who can already do many of the things it does*, or who can recognize when it gets them right.

How do you know it did the right thing?

I just know and I don't know how I know because I can already do it.

It's not common sense, and I'm not sure it can be taught. It might be wisdom, taste, experience, or some sixth sense or the 10,000 hour rule. It's a bit like noticing the uncanny valley: you just know something is wrong, even if you can't quite pin it down and so you just don't accept the result and try again. You're going to pull that loose string in the carpet until it comes apart or disappears.

The other thing is that a lot of AI work today feels like managing people. Machine people, but managing people nonetheless and not hands on work.

Management is juggling and context switching constantly.

This person is stuck and doesn't know what to do, so you point them forward. This person is done, but half the work is wrong, so you explain what needs fixing and send them back. Someone else doesn't even know where to start, so you give them a direction.

If you're not a manager and have never led people, that can feel totally alien.

Put all that together—no ideas, no agency, no burning desire to make things happen, doing things only because you want them over with—and you've got a real adoption problem for the current generation of AI agents.

And that's okay.

That's how diffusion works.

The products for the early part of that curve are not the products for the later parts of the curve.

I still know how to take apart a computer and put it back together because I'm old enough to have lived in a world where that was still necessary. But today's people who grew up after the 80s and 90s PC revolution? They don't need to know. Their Mac is sealed. The memory is soldered on. When it breaks down it's a shock, because it just works 99.9999% of the time. When it breaks down they take it to the shop and it gets fixed somewhere off stage.

That's a late stage product.

Agents are not there and they can't be there yet because machine intelligence needs to get better and the interfaces themselves need to get melted down and reforged into something totally new. That will take time.

With today's agents, most people log in, look around for two minutes, and feel like a writer staring at a blank page with a terminal case of writer's block. They don't know where to start, so they give up.

And the bigger problem is: there's really nothing they urgently want to do.

AI isn't a "do this one tiny thing" tool. It's a "do whatever you can imagine" tool, and most people have no idea what to imagine.

That's different from software of the past. You log in to Instagram and it does one thing well: it gives you a dopamine hit. You understand it in minutes. Its utility is baked into the design. Its purpose is dead simple: look at things you like, find more of them, skip the rest. It's a slot machine. It's not complicated. It's not even checkers. It's tic-tac-toe.

### **Agents Amplify Agency. They Don't Create It.**

Most people aren't trained to think and create and question and wonder. They're not thinking about all the things they could do. They're thinking, "What do you want me to do?" or "What do I have to do?"

The disconnect is that there are a million little things I want to do with agents and that I'm doing as many as I can.

Here's what I did with agents in just a single day last week (and I'm probably forgetting something).

- They helped me create my website for my new product and published it.
- They acted as sparring partners for the website write-up and marketing copy and refined it and iterated on it.
- They wrote 95% of the code.
- They deployed the systems and wrote the operator guides.
- They made me a graphic this morning to illustrate an idea I had.
- They researched some trip I am planning with my wife based on criteria I gave it after another place we wanted to go fell through because the Airbnbs just sucked.
- They helped me research an accountant who proposed something to me that seemed shady and it turned out it was because the dude had a previous firm and court cases against him.
- They helped me brainstorm on this article and corrected my typos and did research for quotes and articles to link.

On and on and on.

I can't even remember all the things I did with agents yesterday, much less a week ago. And what I did yesterday is more than what most people will do with this generation of agents in an entire year.

I think the thing is that most people don't really need to get things done at all. They get things done because they have to get things done because they have no other choice but given infinite freedom they would do nothing. I can't imagine doing nothing.

The people getting crazy absurd leverage from them today are not necessarily the smartest people in the room. But they do have a brilliantly burning desire to do things and a massive backlog of things they've always wanted to do. And they're the people with taste, unfinished plans, and an almost pathological need to make something happen.

And until AI goes from a 1980s era computer that you have to set up and wire yourself to the iPhone where all you have to do is take it out of the sleek little box and touch it and it does all the work for you and abstracts away all the pain and complexity, then we aren't crossing the chasm to the next phase of AI adoption and the Singularity is not near.