AI Agents Are Here. Nobody's Ready.

Matteo Wong, Charlie Warzel

Close your eyes, and imagine a frictionless day.

You wake up to a notification that a grocery-store employee will deliver all of the ingredients for the three recipes you saved last night to your “Weeknight Dinner” folder. While you brew your coffee, a program browses online retailers to compare toilet-paper prices, and another solicits competitive bids for Sheetrock work you need done in the garage. During your commute, a bot moves your upcoming dentist appointment to avoid a scheduling conflict. On your lunch break, you get an alert that computers acting on your behalf noticed some unclaimed funds and class-action payouts in a state database and claimed them. Also, the machine has negotiated a cheaper internet plan. By the time you get home, the groceries are on your doorstep, and you’ve gotten back $150.

This is the future that OpenAI, Meta, xAI, Google, and the rest are selling with a spate of always-on AI assistants that work for you. “Imagine every person on earth with a second mind beyond their own,” Meta’s chief AI officer, Alexandr Wang, wrote of the company’s new AI helper, Muse. A few weeks after Muse’s launch, OpenAI launched its competing product, Dots, which Sam Altman claimed would let you “get more of your time and attention back.”

Read: The singularity is not what it seems

These AI agents promise to bulldoze the inefficiencies of daily life. Think about all of the opportunities you’re currently leaving on the table: You don’t change cell-service providers for a flash sale because of the paperwork involved. You haven’t fiddled with your 401(k) investments in years. You have kids to pick up, a date to dress for, sleep to catch up on. But your AI agents are different—they’re relentless, tireless, and shameless. They work in the background, without supervision, eagerly helping with all of these tasks and a bunch you didn’t even ask for. They use every tactic and loophole at their disposal to get you the best price and the highest-quality service because they’re able to pay attention to everything at once.

The AI companies see a massive opportunity for hyper-personalized agents to become a new intermediary layer for the internet and the physical world. This future is already unfolding: People are giving Muse access to their phone number, email, credit cards, and bank accounts to make dinner reservations, contact customer service, and book vacations. Mark Zuckerberg says that he is using Muse to coach his martial-arts form and help him bake with his daughter (he isn’t very familiar with baking, being “normally into smoking meats”).

Silicon Valley’s vision of boundless efficiency is shortsighted, or at least incomplete. Near the end of his blog post about Muse, Wang writes that “we’ve never seen humanity with every single person’s agency fully switched on.” He paints a utopian vision of laptop-class optimization. But such frictionless commerce, if it ever does come to pass, could bring about something else entirely: chaos.

The era of the agentic web started less than 24 months ago. Anthropic’s Claude Code, introduced in early 2025, allowed agents to write software and use other programs on your laptop. OpenClaw, launched later that same year, could navigate the web on your behalf. In the spring of 2026, Instinct, an invite-only app that let users text requests to a bot over iMessage or WhatsApp, began capturing the imagination of Silicon Valley insiders. It could book museum tickets, purchase gifts, and find in-network doctors.

But it wasn’t until the launch of Meta’s Muse in early September that personalized agents hit the mainstream. Muse has been downloaded more than 5 million times in the past month and is especially popular among early adopters—entrepreneurs, investors, software developers. These users are now

reporting back on their experiences and offering previews of our potential agentic future, which sounds both strange and jagged.

When we reached out to Rajan Lukose, a start-up founder in the Bay Area, to discuss his use of AI assistants, he had his Muse set up a calendar invite. In the late 1990s and 2000s, Lukose worked for corporate research-and-development teams in Silicon Valley and did early research on web agents. Now, he told us, “they’re actually, finally becoming sort of real.” Lukose has been using Meta’s program for “annoying things that I would never do”: tracking travel sites over time to book the cheapest flights, submitting receipts for DoorDash credits, sorting out problems with a magazine subscription. “That’s something that’s a real pain in the ass that just gets done,” he said.

Lukose said that the primary upside of using an agent is clawing back time and money. A few years ago, for instance, Lukose installed an automated tool that went through his Gmail, looked at various transactions, and filed claims under price-purchase guarantees. “That was like $600 that I just would not have gotten myself,” he said. The agentic era promises to deliver those kinds of savings to everyone.

Decades ago, the internet was supposed to create perfect markets, with users able to compare the cost and quality of products and services from anywhere in the world. But it failed: People continued to make lots of decisions based on habit. That, in turn, is good for business. “A lot of things online are premised on limited human attention,” Andrey Fradkin, an economist at Boston University, told us. Digital advertisements, recommendation algorithms, brand reputation, and celebrity endorsements steered consumers toward making suboptimal purchases.

AI agents are now here to finally fulfill the promise of perfect commerce. “The agent knows your preferences; they can go to every single platform, find the best products, the best deal for you, and then help you make that purchase,” Fradkin said. The agent won’t be fooled by sponsored posts or give up on shopping when they get distracted or hungry. “The current attention-based economy—if you gain the eyeball, you have a better chance to sell certain goods—that may change,” Ken Huang, who studies agentic AI at the University of San Francisco and runs an AI consultancy, told us. The bots are impervious to human-targeted nudges. Their attention is infinite.

In the short term, however, this automated quest for efficiency is creating some unsettling dynamics. Over the past month, X, Reddit, and other forums have been overrun with examples of agents haggling customer-service representatives and contractors. Take the case of the Muse user Luke Whiting, who sent his agent to check on a prescription at Walgreens. Whiting, who works for an open-source-database company in San Francisco, told us that the pharmacist at the other end of the call refused to speak to an AI system. His Muse suggested that the best way forward would be for Whiting to abandon using the agent altogether: “Honest take: other stores might say the same thing,” it said. Having the Muse call other pharmacies “could just burn time. The fastest path is probably you calling yourself.”

On a Subreddit for call-center employees, a recent thread is titled, in all caps, “I AM NOT AI.” “I have been asked 3 times this week ‘are you human?’” one person replied. Other users who identify themselves as call-center workers discuss their frustrations. One says that when customers question whether they’re talking to AI, they like to respond, “I’m a real boy,” like Pinocchio. Another, fed up with rude callers, vented, “Customers deserve ai.”

Companies, meanwhile, are starting to tap AI systems to squeeze more out of their customers. According to Fradkin, many hospitals don’t charge as much as they could, simply because medical-billing codes are countless and labyrinthine. But using AI to document and code patient visits reliably allows hospitals to charge for more, and more expensive, services. That likely contributed to patient costs increasing by $650 million across Blue Cross Blue Shield from 2023 to 2025. And that’s only the beginning; health payments are projected to rise another 9 percent next year, largely because of AI billing. McDonald’s is reportedly using AI to analyze millions of consumer transactions and adjust prices at each location, for each menu item, based on local demand. (McDonald’s has said that its pricing portal is “a tool, not a mandate” to help franchisees.) It’s not hard to imagine vendors using AI to develop personalized, on-demand inventory management and dynamic pricing to maximize profits for streaming subscriptions, sports tickets, drug

prescriptions, consumer electronics, insurance policies, and loans.

This optimization assumes that the agents work as intended. So far, there have been plenty of stumbles: In a paper published earlier this year, a group of AI researchers found that agents, given assistant-style tasks in a virtual environment, leaked Social Security numbers and were easily hijacked by other people. Agents powered by Chinese models refused to complete tasks that involved politically sensitive topics. Two weeks ago, a tech executive's AI agent randomly posted his financial information in a workplace chatroom for employees to see.

Things could get weirder still: Tech companies market their agents as perfectly rational, when, in actuality, they can be extremely erratic and unpredictable. Sophia Kazinnik, who studies AI agents at Stanford, found some models to be overly generous and others to be exceedingly greedy. Across the board, they were far less patient than humans. When two personal AI agents are haggling over prices in the wild, the outcomes might not be maximally profitable. And there's no assurance that all future agents will be created equal: More powerful models may have an upper hand in negotiations.

Even if the agents are able to effectively shop and negotiate for us, important rules governing the privacy rights, liability, and economic allegiance associated with these tools will have to be codified. Does your Muse have a fiduciary responsibility to you, or a legal responsibility to Meta? If Muse pays for a $2,000 espresso machine instead of a $200 one, can you be refunded? If Muse leaks your account and routing numbers, is Chase required to refund any fraudulent payments?

At some point, everyone's agents could reach parity. "Just like today, every business has a website and social-media accounts and email addresses and all these things, every business is going to have agents too," Zuckerberg said in an interview this week.

In this future, communication and commerce function almost entirely via computer proxies. Your agent calls a fancy restaurant for a dinner reservation, and a digital maître d' picks up. Your agent and your insurer's agent and your hospital's agent all negotiate the price for your recent appointment. Prices fall, buyers and sellers are better matched, and businesses are incentivized to deliver higher-quality, customized services and products. This is "what economists dream about," Fradkin said.

But friction can serve a purpose for businesses. Banks make great sums when customers leave money in low-interest checking or savings accounts, and studies show that sizable portions of retailer subscription revenue are driven purely by "inertia," another way of saying that people forget to cancel. (Friction can help consumers too: Inattention and incomplete information at times lead landlords, pharmacies, and other businesses to undercharge.) Our entire economy and basically every company are structured around the notion that humans have mental and physical limitations. Without them, existing market structures could collapse.

"There's this strange protection provided by the inefficiencies in our economy," Rory Van Loo, a lawyer and a professor at the Wharton School, told us recently. In 2019, well before the rise of large language models, Van Loo envisioned a phenomenon he dubbed "digital market perfection." Imagine everyone has a digital agent that is always looking for a lower price or a newer, better service. In this world, Van Loo argues, agents could, all at once, direct millions of consumers toward a new product or service—a process he calls "hyperswitching." That would in theory be good news for consumers, but it would be sheer hell for companies. The nightmare scenario that Van Loo and others have conjured is a style of agentic bank run, where similarly attentive agents all pull money from a bank at once to get their clients better interest rates elsewhere.

Maybe this sounds good or exciting to you—a kind of consumer's revenge—but in practice, stable commercial markets would begin to look and function like the stock market. "If, all of a sudden, people could purchase whatever they wanted at the lowest price with the perfect ability to see every possible option, there'd be a lot of volatility in the world. It would seem highly efficient from an economic standpoint, but it could have devastating effects," Van Loo said. Picture 1 million personal agents all switching their consumers to a new brand of paper towels on Amazon because a markdown made them

three cents cheaper. The same Sweetgreen salad is $18 one week and $38 the next. Your electric bill is $150 one month and $450 the month after, and you're uncertain with each payment whether your local utility is about to go bankrupt.

"Stocks are bought and sold fast; markets change and plummet—such that we have a stop button for it, because it can go off the rails," Van Loo noted. "But we don't have those kinds of protections in most other markets." There's potential for collateral damage everywhere. If millions of people switching to higher-yield accounts causes a bank run, that's a disaster for any person or institutional investor with money locked up at that bank. If half of Verizon's customers switch to AT&T in one week to take advantage of a flash sale, the business could collapse overnight, affecting the other half of its customers. Sometimes, this rapid switching could signal healthy competition. Other times, it could cause organizations to implode.

A world where AI proxies do the planning and talking could also mean a future where humans are constantly surprised by things they once were able to anticipate: a birthday present for your spouse ordered in a preposterous size that you must try to explain; a plumber showing up out of the blue.

Or perhaps none of this will happen. It's possible that an assistant-dominated economy turns out to be far more mundane, defined by tacked-on fees to keep the agents in check. Instead of agents allowing you to negotiate and bring down health-care, insurance, and electrical bills, the companies responsible for all of these services charge your agents every time they request to pull your data. Banks, to bolster their revenues in the event of a bank run, institute new and higher fees on checking accounts. Amazon, unable to profit from dynamic pricing and ads in the face of perfectly rational agents, raises prices for consumers and takes a bigger cut from manufacturers. "Generally, if one fee goes away, you get another," Fradkin said. In other words, everything could be faster, easier, and more efficient for the consumer—but potentially more expensive.

Van Loo said that the bigger AI companies will eventually find ways to strike deals with platforms via these protocols, negotiating access, commissions, and partnerships in ways that benefit the companies, not necessarily the consumers. This could result in a sort of preferential treatment like the deals already in place between Meta and Walmart, Sephora, and other retailers, which will effectively cause Muse to prioritize their products. Taking another approach, Amazon has blocked Muse and encourages visitors to browse with its own agent, "Alexa for Shopping," meaning the company is basically negotiating on both sides of the deal. Meanwhile, this week, JPMorgan announced that it would not allow Instinct, Muse, or other agents into customer accounts. Ultimately, this might all lead to "the same old heavily obstructed, inefficient world we have now," Van Loo said. "But it will look far more efficient, because it will save some people time and feel much more convenient." Very powerful companies will design new structures for markets and the internet, determining what can be sold, by whom, and for how much—as they always have.

For all the time and money the agentic future could save us, the technology entails another sort of cost: an additional layer of abstraction and disengagement. We've largely outsourced our decision making to a small number of platforms: Amazon guides our shopping; Google Maps steers us toward certain restaurants and primary-care providers; Instagram dictates our news consumption. We are told that these recommendations and ads are, based on troves of personal data, already optimal. The companies pushing agents assure us that the decision making will now be even more rational and the deals even better. But checking these claims for ourselves will be harder, because the logic behind those recommendations will be buried in large-language-model chains of reasoning. Eventually, the skeptical instinct to check agents' work could atrophy altogether. Over time, humans will play a smaller and smaller role in their own life.

In the age of the algorithm, it's easy to throw out cliché terms such as *late-stage capitalism* in an attempt to describe the extractive, inhumane dynamics of the economy. But what is an agentic future if not the furthest extension of our current market logic? A future of capitalism so perfect, so efficient, that it doesn't need us at all.