


OPINION
GUEST ESSAY

Silicon Valley Is Bracing for a Permanent Underclass

**By Jasmine Sun**

Ms. Sun writes about A.I. and Silicon Valley culture on Substack.

April 30, 2026

Most people I know in the A.I. industry think the median person is screwed, and they have no idea what to do about it. I live in San Francisco, among the young researchers earning million-dollar salaries and the start-up founders competing to build the next unicorn. While Silicon Valley has long warned about the risk of rogue A.I., it has recently woken up to a more mundane nightmare: one in which many ordinary people lose their economic leverage as their jobs are automated away.

Whether you talk with engineers, venture capitalists, founders or managers, or with doomers, accelerationists, lefties or libertarians, the so-called San Francisco consensus on the impact of A.I. for workers is bleak. Many are convinced that advanced A.I. will soon surpass human capabilities. This would produce tremendous growth and scientific achievement, but it would also displace millions of jobs as fewer humans are needed to make the economy run. The technology will depress economic mobility and exacerbate inequality, while ferrying power and wealth to the A.I. companies and the existing owners of capital.

This premonition is not a well-kept secret. It shows up in the Anthropic chief executive Dario Amodei's public pronouncements about a white-collar blood bath and in the disappearing-message Signal chats in which tech executives boast about the roles they plan to automate. You feel it in the fretting of recent college graduates who apply to hundreds of jobs without landing a single interview. You hear it in the gallows humor of the software engineers who joke about replacing themselves with Claude Code.

Some even believe that artificial general intelligence, or A.G.I., will create a permanent underclass. In the United States, the term "underclass" gained currency in the 1960s to describe the factory workers left behind by the postwar automation boom. Today, it has become repopularized as a viral term for a theory that posits that people have a limited window of time to build wealth before A.I. and robotics are advanced enough to fully replace human labor. At that point, we will get frozen in our current class positions: The rich will be able to deploy superintelligent machines to do their bidding, and everyone else will be rendered useless and

unemployable, left to live off welfare scraps.

Hyperbolic? Perhaps. But even those who view the idea of a permanent underclass as overblown tell me that the meme contains a kernel of truth. Yash Kadadi, a 23-year-old start-up founder and Stanford dropout, summarized the sentiment of his peers: “There’s only a matter of time before GPT-7 comes out and eats all software and you can no longer build a software company. Or the best version of Tesla Optimus comes out,” and can perform all physical labor as well. In that world, this year is a human’s “last chance to be a part of the innovation.”

Most economists and A.I. experts do not expect this scenario, but the persistence of the permanent underclass idea should concern all of us. First, because it signals how much collateral damage the A.I. companies will tolerate en route to A.G.I. And second, because the production of a social underclass is a policy choice. Instead of waiting for impact, we need to think seriously — now — about how we plan to support workers through A.I. disruption.

If left to its own devices, Silicon Valley may summon a permanent underclass through its own market logic. If you believe that human-substituting A.I. is inevitable, then every company should race to be the one to build it — and claim a market valuation the size of the economy and then some.

New A.I. models are assessed based on how well they do on a set of benchmarks — essentially standardized tests for the model. Increasingly, these evaluations emphasize real-world economic utility, which means that developers are aiming directly at replacing human capabilities.

The A.I. Productivity Index benchmark measures how frontier models perform across four jobs: investment banking associate, management consultant, Big Law associate and primary care physician. OpenAI established the GDPVal benchmark, which looks at 44 occupations, from real estate broker to news analyst. These measurements reflect A.I. progress but also direct it for researchers aiming for top marks.

“When we originally released GDPVal, which was just a few months ago, none of the models were yet on par with human experts,” said Tejal Patwardhan, who leads frontier evaluations at OpenAI. “Months later, we have over an 80 percent win rate compared to human professionals,” she said. As an example, she pointed to a research colleague who used to work as a banker, and who “keeps being shocked by how much of her old work the models can do.”

Corporate executives accelerate layoffs and slow hiring because they don’t want to be the firm lagging behind. After laying off nearly half of his company’s employees in March, the Block chief executive Jack Dorsey told Wired that coding agents such as Anthropic’s Opus 4.6 and OpenAI’s Codex 5.3 “presented an option to dramatically change how any company is structured, and certainly ours.” Investors responded with a 25 percent stock price surge in after-hours trading.

Sometimes, layoffs happen even before executives know how or whether A.I. will replace those roles. When chief executives are “saying they’re cutting jobs because of A.I., other people feel like they have to too,” explained Zoë Hitzig, an economist who previously worked at OpenAI. “That dynamic could make the changes happen sooner than efficiency would dictate.”

Tech workers, for their part, are scrambling for lucrative A.I. jobs in hopes of securing financial freedom — even when they harbor ethical hangups. “People feel like there are not that many

opportunities to make money in the future,” said Steven Adler, a former employee on OpenAI’s safety team who now writes a newsletter on A.I. policy. “Even if someone thinks it is personally distasteful to make money from building technology that companies say may literally kill everyone, many people are just cogs in the machine.”

Sign up for the Opinion Today newsletter Get expert analysis of the news and a guide to the big ideas shaping the world every weekday morning. [Get it sent to your inbox.](#)

This apparent dissonance can be justified if you believe that the arc of technological progress is fixed. For instance, the founders of Mechanize, a once buzzy start-up with a mission to “enable the full automation of the economy,” argued in a blog post that “the only real choice is whether to hasten this technological revolution ourselves, or to wait for others to initiate it in our absence.”

Many A.I. employees are ultimately motivated by visions of a beautiful future: a promised land where goods are cheap, diseases are cured, and abundant machine labor liberates humans to enjoy lives of infinite leisure. But increasingly, they also worry about triggering a jobs apocalypse along the way. “There are some people who care about jobs and inequality because they really care about people. There are others who think this is going to lead to instability, insurrection and revolution, and that’s bad for business,” said a researcher who has worked at two frontier A.I. labs, and who spoke on the condition of anonymity because of fear of professional retaliation. (In general, tech industry sources expressed more extreme concern about the labor market impacts of A.I. in private conversation — but suddenly became optimists once I turned on the mic.)

The three leading A.I. labs — OpenAI, Anthropic and Google DeepMind — have set up new teams to measure and communicate about the economic impacts of the technology. All three are planning to take a more active policy stance in the coming year. But when I spoke with the technical researchers, economists and policy experts charged with this task, I was not reassured. What I found was a well of worry, good ideas and limited commitments from corporate actors whose core business model relies on the very disruption they are warning about.

Since its early years, OpenAI believed that A.G.I. would transform the global economy and generate untold wealth for its creators. The leadership held that government action would be critical for helping people navigate the disruption that A.I. caused. In a 2021 blog post, the company’s chief executive, Sam Altman, predicted that within decades, “unstoppable” A.I. systems would be able to do almost any job a human could, and thus would shift power from labor to capital. His proposed solution was to aggressively tax assets: land and A.I.-company shares. “If public policy doesn’t adapt accordingly, most people will end up worse off than they are today,” Mr. Altman wrote.

But when the veteran lobbyist Chris Lehane joined OpenAI in April 2024, he spun a sunnier economic story. He and his team appeared to deprioritize research projects that could produce unflattering results, including studies on the environmental impacts of A.I., on the gender gap and the urban-rural divide in ChatGPT usage, on how ChatGPT guides users’ career decisions

and on long-run economic forecasting, according to multiple sources. Instead, Mr. Lehane focused the company's economic messaging on A.I.'s concrete benefits, such as the new jobs and the growth in the gross domestic product that OpenAI's data center investments would create.

"Whenever someone wrote a paper which talked about some negative aspect of A.I., he would say, 'We're not going to release something about a problem until we have a solution for it,'" said an employee who worked with Mr. Lehane, and who spoke on the condition of anonymity to discuss internal deliberations. Mr. Lehane characterized his approach differently: He wanted the economists on OpenAI's global affairs team to "inform smart public-policy making," not conduct "niche" academic research. "We want to do applied physics, not theoretical physics," he said when we spoke in March.

This spring, as fears of A.I.-induced job losses were becoming impossible to ignore, OpenAI started to share solutions. In April, the company released a white paper outlining an "Industrial Policy for the Intelligence Age" that declares the necessity of ambitious New Deal-style policies to combat the concentration of wealth and power in firms like OpenAI. In Mr. Lehane's telling, industrialization "really threw off that relationship between capital and labor" and facilitated the rise of "fascism and communism."

Many of the ideas listed in OpenAI's white paper are radically progressive: a 32-hour workweek, higher taxes on corporations and capital gains and a "public wealth fund" that provides all citizens an equity stake in A.I. companies. Others more clearly cohere with company interests, such as accelerating energy grid expansion and establishing a national "right to A.I." that would give foundation models to schools and libraries.

Still, the document is vague on implementation mechanics and whether OpenAI will advocate the policies listed. In an emailed statement, an OpenAI spokesperson declined to provide examples of specific legislation the company supports, but said that it has talked to members of Congress and the Trump administration about their intent to contribute to a public wealth fund, among other ideas.

OpenAI has not always lived up to its idealistic promises. In 2025, the company removed a profit cap that had previously limited investors' and employees' returns to 100 times their initial investment. The pro-A.I. super PAC Leading the Future, funded in part by OpenAI's president, Greg Brockman, has spent over \$2 million on ads against the New York congressional candidate Alex Bores, who introduced safety regulation for large A.I. developers and released a plan to fund direct payments to Americans by taxing A.I.

I spoke with Mr. Adler, the former OpenAI employee, who shared feedback with Mr. Altman on his early proposals for a public wealth fund and a land value tax. "I hope OpenAI is willing to fight for these prosocial ideas with policymakers," he said in reference to the new white paper. "The A.I. industry is engaged in cutthroat competition over truly world-changing technology. Unless we change their incentives, we shouldn't be surprised when companies cut corners, even if they've said the right things."

And then there is Anthropic, which fashions itself the industry Cassandra. Mr. Amodei spent much of the past year on a nonstop media circuit predicting that 50 percent of entry-level white-collar jobs may disappear by 2030.

But his longer-term concerns are about deeper matters than job losses. In a roughly 20,000-word essay about A.I. risks posted to his personal blog in January, he warned that A.I. may create “an unemployed or very-low-wage ‘underclass’” for people with “lower intellectual ability.” That group would grow to encompass more of the population as A.I.’s capabilities allow it to outpace more humans. In that world, what’s at risk is not only wages but democracy itself. “The balance of power of democracy is premised on the average person having leverage through creating economic value. If that’s not present, I think things become kind of scary,” Mr. Amodei said last year to Axios.

At the same time as A.I. erodes ordinary workers’ leverage, it may concentrate power and wealth in large companies and the U.S. government — two entities whose interests are increasingly linked. A.I.-related investments such as software and data centers accounted for 39 percent of U.S. economic growth in the first three-quarters of 2025, per an analysis by the St. Louis Fed. That gives the federal government a vested interest in sustaining the A.I. boom. Mr. Amodei acknowledges that this concentration can lead to “the reluctance of tech companies to criticize the U.S. government, and the government’s support for extreme anti-regulatory policies on A.I.”

In March, the company started the Anthropic Institute to house its teams working on economics, societal impact and frontier safety. The institute is led by Jack Clark, the affable British journalist turned A.I. billionaire and Anthropic co-founder, who seems to be replacing Mr. Amodei on the media tour of late. When we spoke, I asked Mr. Clark if he, too, expects A.I. to create a permanent underclass.

“This is basically a societal choice,” he replied. Like Mr. Altman and Mr. Amodei, Mr. Clark sees the default path for A.I. as dire: one where we “let technology rip, and don’t think about the social effects until later.” But he also feels optimistic that sufficiently conscientious A.I. builders and policymakers can steer the ship away from the storm.

In Mr. Clark’s future utopia, society can choose to “expand the share of human labor” in relational roles like teaching and nursing, even while A.I. displaces jobs in other sectors. For example, someone who might have become a customer service agent could train as a teacher’s assistant instead — a job that he expects to be more fulfilling for many workers, and in a setting where human presence matters more.

Unlike cash-only safety nets such as a universal basic income, Mr. Clark’s approach preserves work as a source of both individual leverage and personal purpose, even if it favors different occupations. “What A.I. should allow us to do is pay these jobs way more and massively multiply the number of them,” Mr. Clark said, adding, “Of course, we and the other companies have to deliver on the money side.”

The money will come from selling enterprise A.I. agents, a product category in which Anthropic is the current market leader. Agents are large language models that can undertake sequences of actions in pursuit of a goal — like a remote co-worker who lives inside your computer. Because agents like Claude Code can work on projects independently for hours without human prompting, they are at the forefront of concerns around job displacement. Anthropic’s enterprise agents are so popular that the company’s annualized revenue has surged to \$30 billion at its current pace, up from \$9 billion at the end of 2025.

However, Anthropic's coffers probably won't be emptied in the service of public work-force programs unless politicians compel the company to do so. Anthropic has not yet released a set of economic policies that the company supports, either in broad strokes, as with OpenAI's white paper, or by endorsing specific legislation, as Google did when it picked a list of 15 A.I. work-force assessment and education bills. When I asked Mr. Clark if the Anthropic Institute planned to lobby for the redistributive measures he alludes to, he demurred, describing policy advocacy as "the end of a very, very long chain of work." (Anthropic has, however, contributed \$20 million to a political group backing Mr. Bores.)

The mood inside Anthropic is uneasy. The company has become one of the most desirable employers in town, pairing a rocket-ship business model with high-minded ethical principles. Yet in conversations with employees, I also hear a palpable sense of existential vertigo about the magnitude of the societal changes they are bringing forward. Many engineers run several Claude Code agents simultaneously, giving them tasks to complete overnight so that someone — human or machine — is always on the clock. They muse about the postwork future while pulling 80-hour workweeks. Even their own berths may not be safe, implies their boss: "It may be feasible to pay human employees even long after they are no longer providing economic value in the traditional sense. Anthropic is currently considering a range of possible pathways for our own employees," Mr. Amodei wrote.

Compared with those at OpenAI, Anthropic's research teams seem less afraid to highlight the bad alongside the good: what its researchers call the "light and shade" of A.I. In January, Anthropic published a paper revealing that a small but increasing fraction of Claude users are delegating their most personal and consequential decisions to A.I. — a choice they often later regret. "You made me do stupid things," one such user told Claude. In another experiment, Anthropic researchers found that junior engineers who relied on A.I. coding agents not only didn't complete tasks much faster; they also understood their work less when quizzed about it afterward. The labor market implications are grim. At the same time that early-career workers are competing with A.I. for jobs, they may be stunting their own skill development by overusing A.I. tools.

Some employees are looking beyond their day jobs to alleviate the harms of A.I., from job loss to bioterrorism. Anthropic staff members have pre-committed billions of dollars in individual donations to nonprofits they choose, including many organizations dedicated to preventing catastrophic A.I. outcomes.

So while Anthropic employees insist that positive A.I. futures are possible — or else they wouldn't be building it — they often seem uncertain about whether that world is likely, or whether they personally are bringing it about.

On the evening of Feb. 25, several dozen A.I. employees and civil society advocates gathered in a converted warehouse in San Francisco's sleepy Dogpatch neighborhood to hear the Democratic pollster and strategist David Shor. The event was titled How to Prepare Our Politics for A.G.I., and doubled as a fund-raiser for a new "six-to-nine-month sprint" to rally Democratic politicians around the campaign issue of A.I. job displacement.

Under a disco ball, with La Croix cans in hand, tech workers perched on benches and beanbags while Mr. Shor presented a slew of public opinion polls spotlighting Americans' economic fears

about A.I. One slide showed that 79 percent of voters are worried about the “government not having a plan to protect workers,” and 72 percent are concerned that A.I. “drives down wages for people like you.”

While the American public ordinarily hesitates to support left-wing policies like a jobs guarantee or single-payer health care, A.I. seems to expand the political Overton window. “Right now, the argument is, ‘You’re all about to lose your jobs, and the choice is either you get nothing and starve, or we do something fair,’” Mr. Shor said. “People don’t want to be members of the permanent underclass.”

Not all policies are created equal, however. A universal basic income is unpopular, but a federal jobs guarantee has legs, Mr. Shor found. American voters don’t care about beating China, but they are excited about A.I. curing diseases. And, crucially, populism sells. In one of the top-performing political ads that Mr. Shor’s data firm tested, the nameless narrator declares: “We make the corporations and billionaires who profit from A.I. pay their fair share.” The ad concludes: “They work for the bots. We work for you.”

The presentation ended with a pitch to the audience: “\$700 billion a year is being spent” on A.I. transformation, said Mr. Shor. “For less than what the industry spends in one hour,” donors can equip Democratic politicians with winnable campaign messaging — slogans, ad concepts, banner policies — for the job disruption that he believes is coming.

During the Q&A portion, an audience member asked about the risk of “crying wolf” on A.I. job disruption. What if it happens more slowly than predicted? Mr. Shor scoffed. “People’s bar is way too high on this. The reality is, if one concentrated industry with 1,000 people loses their jobs, it’s going to be the biggest story of the century.”



Ben Denzer

If A.I. companies and American voters are waiting for policymakers to act, many still seem paralyzed by the data (or lack thereof). Nobody can predict how far A.I. capabilities will progress or how fast A.I. will spread. Economists also disagree on whether wage inequality will rise or fall, whether consumer demand is elastic or capped and whether economic growth will be linear or exponential. As a result, many are hesitant about making aggressive forecasts, even in scenarios in which advanced A.I. rapidly surpasses human ability.

Yet there are a few predictions that most analysts agree on. We are already seeing some labor market impacts now, with employment declining for young workers in highly A.I.-exposed occupations like software engineering and customer service. More knowledge-work roles will be automated by A.I. over the next five years — first as more organizations learn to adopt A.I. tools, closing the gap between theoretical capability and observed usage, and second as the models themselves improve.

If current trends continue, A.I. models and agents will be capable of performing a wider range of knowledge-work tasks at higher levels of complexity. At that point, A.I. shifts from automating single tasks to taking over entire roles. Hiring may slow in accounting, marketing, design, administrative work and other white-collar professions.

The work force will shift toward less automatable jobs where humans retain a comparative advantage — such as entrepreneurship, care work, the skilled trades and entertainment like sports and the performing arts. We will also see new jobs we haven't imagined yet, in numbers we cannot predict. Many displaced workers will struggle to retrain, as they have in past

automation waves. Education, health care and tax systems will require an overhaul if white-collar employment is no longer a reliable path to middle-class stability.

At a societal level, the result of mass automation is a decline in worker bargaining power and the labor share of income. This conclusion is supported by the majority of economic research. Leaner A.I.-native firms with a small number of human employees could outcompete those with more workers, much the way the success of technology-intensive superstar firms propelled the decline of U.S. labor share around the turn of the 21st century. A.I. model developers and A.I. infrastructure companies will most likely explode in value, earning a cut of every transaction.

Some analysts, like the economist Anton Korinek, of the University of Virginia and the Anthropic Institute, suggest that *no* human job is invulnerable in the long run, once A.I. can outperform humans at everything. Others, such as the M.I.T. economist David Autor, argue that new industries will emerge to meet infinitely unfolding consumer demand, just as our ancestors could not have fathomed the modern roles of flight attendants and software salespeople. Ultimately, the severity of disruption depends on how fast and how far automation goes.

But the debate over the most extreme scenarios conceals a more immediate threat: Even in the most limited case, A.I. will break the career ladder for millions of current and future workers, a prospect often waved away with euphemisms like “transitional friction.” The Oxford economist Carl Benedikt Frey puts it plainly: “Most economists will acknowledge that technological progress can cause some adjustment problems in the short run. What is rarely noted is that the short run can be a lifetime.”

Powerful A.I. may look alien, but the political dilemmas it raises are not. Some economic policy experts predict that A.I. will look like an accelerated and expanded version of deindustrialization. But rather than companies outsourcing jobs to overseas workers, they will be outsourcing them to A.I. agents. “The China shock unfolded over several years, whereas this could happen over two years,” said Bharat Ramamurti, a former deputy director of the National Economic Council in the Biden White House. “These companies have spent so much money developing models that there’s going to be immense pressure on them to generate revenue through quick adoption.”

“I’ve interviewed so many college students who are super fearful about what the future means, and their narrative is exactly the same as those blue-collar guys in the heartland,” said Molly Kinder, a senior fellow at the Brookings Institution who studies work and automation. In Ms. Kinder’s view, A.I. companies’ narratives about abundance repeat the same flawed promises of globalization. “Our economy grew extraordinarily and prices went down, but there were clear losers.”

In this sense, A.I.’s broad capabilities foster a rare class solidarity between white-collar and blue-collar workers. When 20-something software engineers in San Francisco talk about escaping the permanent underclass, I hear them projecting concerns about their own precarity: What happens if the invisible hand of the market decides that my skills are no longer valuable? Who will catch me if I fall? For once, a rarefied class of employees — those used to being the automaters, not the automated — is reckoning with their potential obsolescence.

It is not as if the United States has never before seen problems of wealth inequality, a declining

labor share of the economy or technological shocks to jobs. But this time we might finally do something about it, now that some of the most privileged are vulnerable.

“I think you’re going to see a battle of ideas in the next presidential election,” said Ms. Kinder. A.I. has risen in importance to voters faster than any other issue in the past year, per Mr. Shor’s polling data. And Democrats ought to be especially alert: Their younger and more college-educated voters are more exposed to A.I. than Republicans are. Senator Mark Kelly and Representative Ro Khanna have announced sweeping A.I. agendas. The technology is an opportunity for gutsy politicians — especially populist candidates vying in a crowded 2028 presidential primary — to push ideas that are usually too radical for moderate voters to swallow.

Society’s ability to cushion A.I.’s disruption may determine whether we get to reap its gains at all. Without a safety net and a transition plan, blunt protectionism is workers’ rational response to automation. If you hear that A.I. will entrench a permanent underclass, you’ll do anything to stop it. Across the United States, there are new proposals for bans on data center construction, on self-driving cars and on chatbots for broad consumer uses like therapy and law. In the extreme, populist rage can metamorphose into violence. In April, an attacker tried to firebomb Mr. Altman’s home, and another is accused of targeting an Indianapolis city councilman who approved a local data center project.

And what if we don’t act? What if we “let technology rip”? What if millions of people do lose their jobs to A.I., and nobody puts up the money or policy solutions to help them? In March, the Palantir chief executive, Alex Karp, spoke on a panel with the Teamsters president, Sean O’Brien. “The biggest challenge to A.I. in this country is political unrest,” Mr. Karp said. “If I were sitting here in private with my peers, I’d be telling them the country could blow up politically and none of us are going to make any money when the country blows up.”

Jasmine Sun writes about A.I. and Silicon Valley culture on Substack and as a contributing writer at The Atlantic. She has previously worked in startups and in A.I. policy.

The Times is committed to publishing a diversity of letters to the editor. We’d like to hear what you think about this or any of our articles. Here are some tips. And here’s our email: letters@nytimes.com.

Follow the New York Times Opinion section on Facebook, Instagram, TikTok, Bluesky, WhatsApp and Threads.

A version of this article appears in print on , Section SR, Page 6 of the New York edition with the headline: Silicon Valley Is Bracing For a Permanent Underclass