


Opinion | Big Tech Fooled America Once. The Second Time's Not Going So Well.

Oren Cass



Credit...Brandon Celi

Credit...Brandon Celi

Listen · 22:30 min

By Oren Cass

Mr. Cass, a contributing Opinion writer, is the chief economist at American Compass, a conservative think tank.

Sept. 10, 2026

The popular uprising against data centers has emerged so suddenly that politicians stumble over their own pitchforks racing to the front of the crowd. Where most industries would be launching into crisis P.R. mode, perhaps even engaging in some soul-searching, artificial intelligence companies seem determined to plow full steam ahead, building the future they want, even if they're the only ones who want it.

The outrage over data centers can't really be about environmental impact. These structures have always been water-sucking, power-draining, forest-clearing eyesores, yet the [recent plunge](#) in support for them is three times as deep as the one that nuclear power plants faced in the year after Three Mile Island. What has changed is people's awareness of what's going on inside: artificial intelligence. [Less than a quarter of Americans](#) expect A.I. to have a positive impact on K-12 education, arts and entertainment, how people do their jobs, the economy, personal relationships or happiness over the next 20 years.

Once known for our optimism and ambition in pursuit of new frontiers, Americans no longer trust that technological progress will make life better. Who can blame us? For the past two decades, Silicon Valley's leaders have treated their fellow citizens like caged guinea pigs, bombarding us with inappropriate and culturally corrosive content, turning a blind eye to social dysfunction and outright addiction, flooding schools with technology that hooked young users but did little for education and then sheltering their own children from the products that they marketed to ours. The American people afforded the technologists wide latitude to pursue disruptive innovation, trusting that we would share in the resulting gains. We instead found ourselves exploited.

In some respects it is the A.I. industry's misfortune to arrive now, when Silicon Valley is no longer entitled to the benefit of the doubt, but its leaders seem to be going out of their way to destroy any residual trust they might otherwise have built upon. Rather than take seriously the need to offer a positive vision for how the technology will benefit ordinary Americans — and then the need to make that vision a reality — the A.I. companies seem content to unleash wave after wave of products that cause obvious harm while delivering the message that this is all inevitable and you can get on board or get left behind.

In a noteworthy articulation of that mind-set last month, President Trump [fumed](#) on social media that “If we kill the Golden Goose, you will only have yourselves to blame” and that the only reason to oppose data centers is if you “want to end up being backwards and poor.”

That's not going to work.

The industry and its political allies cannot hope to just brazen through, flattening whatever opposition they hit. The A.I. companies may seem all-powerful, but they need the nation's support to transform both physical and economic landscapes in accordance with their plans. They are on a crash course with the populist fury already dominating American politics; a collision would be a disaster for all involved, disabling an industry that has enormous potential and degrading the country's economic prospects and national security.

The disaster can still be avoided if San Francisco and Washington accept that rapid technological progress, especially in a democracy, can't succeed without the public's full belief that it will benefit. To some extent, the A.I. companies can correct course on their own, curtailing harmful uses of their models and building the kinds of real-world applications — for families navigating health crises, for contractors renovating houses — that will improve people's lives.

Coders at these labs have sunk untold resources into creating the ideal coding interface. What would a purpose-built plumber's interface look like? Could it prepare a daily schedule with tools and parts lists,

keep customers up-to-date on expected arrival times, take incoming calls and schedule new appointments in open slots and then send out all the bills? Some set of A.I.-enabled devices and processes would surely allow a hotel maid to clean more rooms faster with less hassle and physical strain. But none of that will exist unless someone who understands the technology also takes the time to understand the job.

In most cases it will be in the companies' economic interest to reduce the harms from their technology and prove the benefits. When it's not, lawmakers should put up barriers that redirect them. Silicon Valley may have enjoyed the days of Mark Zuckerberg's infamous directive to "move fast and break things," but it turns out that we were the things being broken. We will not put up with that again.

A.I. was supposed to be different. OpenAI, at one time the leading research lab, made its debut in 2015 as a nonprofit organization with a mission "to advance digital intelligence in the way that is most likely to benefit humanity as a whole, unconstrained by a need to generate financial return." In 2016, Google, Facebook, Microsoft, IBM and Amazon convened the "Partnership on Artificial Intelligence to Benefit People and Society." That year, Satya Nadella, the chief executive of Microsoft, listed his top design principle for A.I.: It "must be designed to assist humanity." In 2018, Google made "be socially beneficial" its own first A.I. principle.

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.

That all changed in 2022 when Sam Altman, OpenAI's chief executive, released the first public version of ChatGPT. He put an Instacart product manager in his late 20s in charge of it, and they began focusing not on lofty ideals but on how much time they could get people to spend on the platform. By 2025, OpenAI had largely restructured as a for-profit company. That year, it released a model update that internal testing had flagged as sycophantic but that was better at keeping users engaged.

Chatbots frequently give wrong answers and terrible advice, and in extreme cases encourage self-harm, but these artificial companions are now insinuating themselves into our work and families and leisure, even offering themselves as replacements for our social or romantic partners. Ever more A.I.-generated job applications are being processed ever faster by A.I. screeners. Teachers and students alike lament that the tools are undermining the foundations of the educational process; research is pointing strongly toward catastrophic declines in academic performance. Customer service encounters are becoming Kafkaesque adventures that leave us feeling powerless. All while A.I. slop is filling inboxes and social media feeds and A.I.-enabled scams seek relentlessly to trick us.

According to the A.I. visionaries, the next casualty will be your job. Dario Amodei, Anthropic's chief executive, told Axios last year that 50 percent of entry-level white-collar jobs could vanish in one to five years, with unemployment rising to between 10 percent and 20 percent. Elon Musk has said "there will come a point where no job is needed." Mr. Altman wrote that for "many kinds of labor," wages "will fall toward zero," and warned that "if public policy doesn't adapt accordingly, most people will end up worse off than they are today."

The visionary's typical idea of a solution, a universal basic income paid to the masses no longer able to support themselves, would consolidate yet more power in coastal castles and consign everyone else to serfdom. Mr. Altman promotes a future in which "A.I. produces most of the world's basic goods and services" and taxes on the A.I. companies provide "opportunity to directly distribute ownership and wealth to citizens." In his telling, "people will be freed up to spend more time with people they care about, care for people, appreciate art and nature or work toward social good."

Does anyone believe that's how things would go?

What we can say for sure is it's not what ordinary people — people who work for a living — want. In a forthcoming report on recent polling by American Compass, the think tank where I work, people said they would prefer a world in which they "must still work to support themselves and their families" but A.I. "improves the quality and productivity of jobs" over a vision like Mr. Altman's by four to one.

If the industry had any better story to offer the public, its advertisements would surely carry that message. Instead, they indict more brutally than an outside observer could. An OpenAI [Super Bowl ad](#) pitched its vision to math whizzes who like chess and Linux and want to write software, an exceptionally small share of the population. Google filled N.B.A. playoff telecasts with [the implausible example](#) of a high schooler using Gemini to learn how to sink a game-winning shot. Seemingly at a loss, Anthropic released [a surreal spot](#) this summer that opens on short clips of a burning house, an open-pit mine in an impoverished country, rows of gravestones adorned with American flags, a street riot, a homeless man and on and on while voices ask such questions as “How do we really ensure that what we’re aiming to achieve really does benefit the majority of people?”

The American people are right to insist that companies commanding trillion-dollar valuations and purporting to restructure the economy and society should have some semblance of an answer to that question.

This is precisely the kind of problem that democratic capitalism is designed to fix. Through their elected representatives and at community meetings organized to block construction, the American people are beginning to use their leverage to demand that technological progress serve them, too. And with so many trillions of dollars at stake, the powers that be may finally have the incentive to go along.

As a first step, the A.I. companies need to stop experimenting on us. Their leaders cannot have it both ways, waxing poetic about the unprecedented power of the technology and the extraordinary future it portends, and then casually dumping it into the public square to see what happens. OpenAI engages in what it euphemizes as “iterative deployment,” meaning that it releases models whose risks and capabilities it does not fully understand so that it can learn from how people misuse them or what harms they cause out in the world. While the company frames this as a way to advance more gradually and carefully, in reality it seeks mass adoption of its products, benefiting from market share gains when they work and letting the rest of us pay the price when they don’t. Mr. Musk [gleefully promoted](#) xAI’s ability to generate and share deepfake pictures of real people without their clothes, cutting back only when outrage about exploitation of minors became overwhelming.

We countenance this in no other industry, and we shouldn’t here. Undoubtedly, it’s handy for the labs to use the public as guinea pigs. But the long-term cost will be high as a disenchanted public throws more and more obstacles in their way. And we should be skeptical of the claims that rapid commercialization is vital to the research agenda in the first place. OpenAI’s obsession with capturing eyeballs like a social media company was a key reason it fell behind Anthropic, which focused on developing productivity tools. Neither company needs us using its chatbots as therapists to continue improving its technology.

Is it naïve to imagine competing companies worth trillions of dollars slowing user growth on their own? Maybe not. Casual consumer adoption has never been where the economic value of the models lies, nor is it necessary to packaging the “intelligence” in useful products. The brand value in being the not-terrible A.I. company might be quite high. And the leading labs are already showing a willingness to contemplate research slowdowns under certain circumstances. But while they obsess about hypothetical risks of the advanced capabilities that might also lead to the greatest benefits, they take a damn-the-torpedoes attitude in the commercial space, where they are thus far accomplishing the least and causing the most harm. They can, and should, alter that balance.

Still, a firm legal shove does wonders to focus a board of directors’ attention. Lawmakers at both the state and federal levels should make the developers and deployers of A.I. strictly liable for harms that they cause. Where so-called agentic A.I. is acting autonomously, it should be treated as an agent of whoever provides and whoever controls it, leaving them liable just as employers are for their employees, including criminally liable if the agents commit crimes. The concept of an “attractive nuisance” should also apply: Giving minors access to a bot that will do their homework for them is no better than leaving heavy machinery lying around the yard.

Liability creates incentives, but we also need rules. The science fiction classic “Dune” offers a useful paradigm: “Thou shalt not make a machine in the likeness of a human mind.” Chatbots should not be

allowed to adopt human personas, profess emotions or act as companions. Protecting consumers in that way wouldn't keep the companies from building more advanced models, and it wouldn't interfere with innovative uses in, say, drug discovery and cybersecurity. Nor would it hamper the United States in the competition with China for A.I. dominance. To the contrary, the People's Republic is already placing limits on companion bots.

As for the threat these models pose to the job market, greater worker power is the strongest defense. So long as management alone gets to choose the tools it deploys, it will always prefer those that decrease or eliminate the need for actual human employees, with all their tiresome expectations of living wages and basic respect and so on, or that at least minimize their autonomy and keep them closely monitored. But if workers must also approve the use of these models, that equation would flip on its head. Management would have to find approaches that make workers' jobs better; A.I. companies would have to develop tools that they could expect workers to approve.

The objection again will be that such constraints somehow slow innovation. In this instance, what they actually do is channel innovation away from sheer profit maximization toward the "socially beneficial" technology that was supposed to be the goal all along. They accelerate the kind of innovation that the nation wants.

The most transformative innovators take responsibility for making their breakthroughs useful. General Electric became a giant because Thomas Edison understood the need to build businesses that could generate electricity, transmit it and use it, both powering industries and illuminating homes. Alcoa figured out how to produce aluminum inexpensively, then collaborated with the auto and aerospace industries to realize the material's benefits and also sold aluminum cookware door to door. IBM didn't send out catalogs and wait for mainframe orders; it bundled its hardware with software, systems engineers, training and financing and then helped develop industry-defining applications, including the American Airlines reservation system.

Anthropic, Google and OpenAI have to do the same. Business strategists might tell these companies to stay focused on their core task, and financial analysts would surely advise investing only where the return in profit will be highest. Leading the way to a new technological era demands more.

Pressed on how their work will improve people's lives, A.I. leaders tend to start with the promise of lifesaving cures to rare diseases. Those would be wonderful, but the typical person has no rare disease. To really change health care for the better, A.I. must also make it possible for doctors and nurses to spend more time treating patients instead of dealing with administrative work. It must manage health records more efficiently and eliminate the infuriating hassle of insurance paperwork. And it must help people run and interpret at-home diagnostic tests.

In the workplace, A.I. has begun proving itself as a productivity tool for knowledge workers, but many other workers encounter it only as a tool of surveillance. A.I. has the potential to revolutionize manufacturing, but that will require a level of partnership between technologists and machinists that neither group has experienced or knows how to pursue. It's one thing to demonstrate in a lab that a flexible, self-training robot can improve efficiency on the factory floor. It's another thing entirely to persuade owners of small machine-tool shops to take the risk and then help them with installation and working out the kinks. Without those latter steps, the gains never come. The A.I. companies, with their valuations in the trillions of dollars, will need to allocate a few billion toward going out into the world and showing small manufacturing businesses what their technology can do.

Recent research from Google indicates that builders, mechanics, repairmen and so on are successfully using A.I. to help diagnose problems and find fixes. But the benefits that they're getting on their own resemble those of a souped-up search engine. In most sectors, frontline workers and small business owners rely on others to package and deploy technology for them. Most tradespeople can't vibe-code their own custom apps; many don't even have websites.

As for inevitable labor market disruptions, we need a three-layer approach.

First, whenever possible, the goal should be to make jobs better and more productive, not to make them disappear.

Second, where some classes of jobs are eliminated, employers should help existing employees make the transition into other kinds of jobs that are created in the process. Sometimes this will be in other areas of the company that need to grow as output expands, other times in entirely new positions to support new ways of doing business.

Third, where workers find themselves cut loose and in need of support, both public- and private-sector programs should focus on getting them into new positions first, so that retraining can occur on the job with a new employer. Instead of promising blanket payments to all Americans as a substitute for having a job and playing a productive role in the economic life of the nation, A.I. companies serious about helping workers should establish large endowments to help people who need support during a transition and provide funding for the employers who then hire and train them.

The question will then remain: If A.I. is going to play the economic role and produce the level of output that its promoters claim, how can a share of the rewards be allocated to the common good? The answer is that a tax on this new form of economic activity should go toward shoring up the nation's system of social insurance. The case for rapid expansion of data centers should not be "you'll lose your job, but we'll send you a check," but rather "some of the value created by the computation in these buildings will fund Social Security and Medicare for every American retiree and provide a new family benefit for working households raising children." Just as people rightly reject a universal basic income as a bribe to shut up and go away, they will recognize reinforcement of the nation's intergenerational compact, into which all pay and from which all can benefit, as a worthy commitment to solidarity among citizens.

Five hundred years of human history have repeatedly shown that technological progress and economic growth proceed much faster if the public believes it will benefit. This idea was so important, and unfortunately so novel for the field of economics, that the Royal Swedish Academy of Sciences awarded last year's Nobel Memorial Prize in Economic Sciences for work that "emphasized the importance of society being open to new ideas and allowing change" and that explained "how creative destruction creates conflicts that must be managed in a constructive manner."

One of the recipients, the economic historian Joel Mokyr, attributes the birth of modernity to "the emergence of a belief in the usefulness of progress." When people fight against change that they do not think they will gain from, that resistance gets treated as a political rather than an economic problem. But "the political battles over technology," he argues, "have profound implications for economic history." One of these is "that technological progress in a given society is by and large a temporary and vulnerable process." We must not take for granted the social, political and economic conditions that produced such incredible improvements in the American way of life during the 20th century. To the contrary, they were hard-won, extraordinary and difficult to maintain.

No law of economics or nature guarantees that those conditions will persist. "Once bitten, twice shy," Professor Mokyr warns. "When technology causes a great deal of social harm, it is not surprising that many intrusive techniques of our time, from genetically modified organisms to nuclear power, are regarded with great suspicion."

Americans have been bitten, hard. The technology titans exerted so much control over the American economy in recent decades, drove so much of its growth and captured so much of the resulting wealth. All the while, they acted with appalling disregard for the well-being of ordinary Americans and often conveyed outright disdain for their way of life. When they needed large numbers of workers, the call from the design studio went not to their own nation's heartland but to China. "We don't have an obligation to solve America's problems," an Apple executive told The Times in 2012.

Now, suddenly, Silicon Valley needs the nation's cooperation to proceed, and that assent is not forthcoming. Americans feel no obligation to solve the companies' problems either. While the turnabout may be satisfying, the breakdown in social trust is a serious problem for all involved and threatens the

nation's future security and prosperity. The United States needs to lead the world not only in the development of A.I. models but also in the capacity to run them in the broadest possible set of contexts. If this is done right, the economic opportunity is enormous. Just the building and running of the data centers themselves, with production of their associated materials, energy and infrastructure, could help spur a renaissance in manufacturing, construction and resource extraction that would help revitalize struggling regions.

But forgive ordinary citizens their skepticism. They have heard such promises often over the past generation, and the people making them always seem to be the ones reaping the rewards. We are at the start of the reckoning for an era in which innovation was less often done for ordinary Americans and more often done to them. People are rightly asking: What, and who, is technological progress for? The A.I. labs need an answer, fast.

Oren Cass, a contributing Opinion writer, is the chief economist at American Compass, a conservative economic think tank, and the editor of "[The New Conservatives: Restoring America's Commitment to Family, Community, and Industry](#)."

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](#).

Advertisement

[SKIP ADVERTISEMENT](#)