IDEAS

The Man Who Saw AI Coming

Erik Brynjolfsson wants to talk with you about the future.

By Annie Lowrey

Illustration by The Atlantic. Source: Roy Rochlin / Getty.

JUNE 29, 2026, 7:31 AM ET

More than a decade ago, the economist Erik Brynjolfsson made a prediction: AI would change everything.

Humans began using tools millions of years ago. They cultivated grain and domesticated animals, and then developed written languages, iron tools, the printing press, gunpowder. Progress was slow and local until the mid-1700s, when modern society roared into being alongside machines and engines. Invention built on invention. Then, in recent decades, the pace of human progress slowed again. Productivity growth, the academic measure of how much better people are getting at generating outputs from inputs, collapsed.

“We are at a technological plateau,” in the midst of a “great stagnation,” the public intellectual Tyler Cowen argued in 2010. “Apart from the seemingly magical internet, life in broad material terms isn’t so different from what it was in 1953.” Shortly after, the economist Robert Gordon of Northwestern University released a best seller arguing that our era of unprecedented growth and life-changing innovation had come to an end.

Cowen and Gordon had the numbers on their side, but Brynjolfsson thought they were wrong. Along with the economist Andrew McAfee, he published two books arguing that the digital revolution was just getting started. AI was already besting human beings at cognitive tasks. Soon, it would make self-driving cars and hyper-smart computers look like nothing more than “warm-up acts,” the pair wrote in *The Second Machine Age*. Productivity growth would pick up, Brynjolfsson believed. Living standards would improve. Another new society would roar into being.

Last month, I visited Brynjolfsson at Stanford, where he teaches, with a few questions in mind: How had he known what AI would be capable of years before the release of ChatGPT and Claude Code? And what did he think it would do now? Would society trend toward debilitating mass unemployment or the joyful end of work?

I was not the only person seeking answers from the seer. Brynjolfsson is teaching overflow classes about AI and technological progress, as well as running a humming lab tracking AI’s effect on workers, firms, and human welfare. Academics are relying on his data. Silicon Valley luminaries such as Dario Amodei of Anthropic and Demis Hassabis of Google DeepMind are eating at his table. Senators and governors are calling him up.

Brynjolfsson—who is towering and gregarious, what I imagine the Brawny-paper-towel guy would be like if he’d become an academic rather than a lumberjack—is gratified that the rest of the world is finally “seeing what I was seeing,” he told me. AI is spurring advances in medicine, science, media, transportation, and energy. In time, it will do to white-collar work what machines did to blue-collar work: make people more productive, and push millions of them into new roles. He doubts that humans will fail to find new things to do. Yet he worries that algorithms could further cleave the 0.001 percent from everyone else, and that those elites could consolidate political and social power, along with precious capital. He added—joking, I pray—that a powerful AI could “do whatever it wanted, I mean, including get rid of us all.”

Could, he emphasized, not *would*. “We’re going to be scrambling with chaos, and I

don't know how we will handle it," Brynjolfsson said, before tossing on a leather jacket and loping off to give his students a pop quiz.

Brynjolfsson has always been optimistic about technological change. His father was a sheepherder on Iceland's northern coast who moved to Denmark, studied with Niels Bohr, immigrated to the United States, and worked on the Apollo missions. "It is a pretty amazing life to go from no electricity to putting people on the moon," Brynjolffson said. He became obsessed with AI when he was a kid reading Isaac Asimov instead of going to sleep, thinking, "This is going to be big."

As an academic, he came to understand why. Technological advances are different from those in many other industries. They are combinatorial, digital, and exponential. An improvement in machine learning boosts the effects of an improvement in chip speeds. (Thus, it is combinatorial.) An improvement in machine learning spreads quickly at low or even zero cost. (Because it is digital.) An improvement in machine learning sets up additional improvements in machine learning, instead of making marginal improvements harder to eke out. (In other words, it is exponential.)

Human beings began making thinking machines in the 1950s. Progress continued, if slowly and imperceptibly, from that point forward. Brynjolfsson was never sure when progress would become fast and perceptible, he told me in his Stanford office, which is dominated by a humanoid robot, now defunct. But he was "pretty confident" that it would. With exponential advances, "barely anything changes, and then at some point the function crosses this point of salience. Technically, the percentage growth rate is the same the entire time. But it's off a small base, then once the base is big enough … "

Brynjolfsson got his Ph.D. at MIT and spent 30 years at the institution, studying how humans put the technologies they invent to use and the effects on output, employment, and efficiency. The process is often messy and slow, he found. Just out of graduate school, he coined the term *productivity paradox* to describe why the IT revolution took decades to show up in national statistics. (Computers raised productivity in the 1990s, a big blip on a middling trend.) A 50-pound desktop arrived in Dave from accounting's cubicle in 1982. Dave stopped doing calculations by hand and quit hunt-and-peck typing in 1987. He helped implement a digital ordering system in 1989. His firm's profitability started to pick up only in 1994.

Sometimes, life-changing technological advances do not affect marquee economic statistics at all. Smartphones, social networks, GPS, tap to pay, ride-sharing apps and

gig-work platforms, costless global videoconferencing—they have yet to charge American productivity or boost GDP growth. They likely never will, Brynjolfsson thinks. Digital services tend to be low-cost or zero-cost. Payroll reports and GDP data never capture these services' true value or give a sense of how much they have changed modern life.

Brynjolfsson has kept track of AI throughout his career. The first course he ever taught, at Harvard's Extension School in 1985, was about "expert systems"—hand-coded, if-then decision-making proto-AIs. He watched as algorithms revolutionized search, as self-driving cars took the wheel, as software became capable of image recognition. In 2012, he decided to "evangelize" his belief that technology had hit an inflection point, in response to Cowen and Gordon and many other big-name academics.

He and McAfee wrote *Race Against the Machine*, arguing that the digital revolution was accelerating and that technological unemployment could soon afflict white-collar, middle-skill workers. In 2014, they beefed up the argument in *The Second Machine Age*. Thinking machines would generate enormous bounty, they argued. Human minds would have to ensure that that bounty got distributed in socially beneficial ways.

When Brynjolfsson arrived at Stanford, in 2020, his ideas still seemed outlandish. Productivity remained close to scratch. The start-up economy was in a recession. Meta, Google, and Apple hadn't put out an exciting new innovation in eons. Their core products were enshittifying, when they weren't destroying attention spans and crushing democracy. "Everybody was talking about going to Austin or Miami," Brynjolfsson said. "I said: *Guys, be honest with me. Did I miss it?*"

Yet that same year, Brynjolfsson got access to GPT-3, a progenitor of ChatGPT. He fed it one of his papers, asked it for comments, and was "blown away" by the response. He prompted the chatbot to do it again, but this time channeling Taylor Swift. "What you gonna do when the knowledge factory has no floor?" it wrote. "How you gonna run when the knowledge factory is in the cloud?" He "thankfully" did not try to sing the lyrics, but he kept staring at them, thinking, "Holy shit." Soon after, OpenAI and Anthropic released products to businesses and consumers. Financiers started to push hundreds of billions of dollars into algorithms and data centers.

When I asked Brynjolfsson's colleagues and collaborators to describe his strengths, I expected them to talk about his skill at coding or data collection. Instead, they invariably brought up his people skills. A former student marveled that they had "never—literally never—seen Erik be mean to anyone." Daniel Rock of the University of Pennsylvania remembered watching him bound past a bunch of eminences at a conference to beeline to a table of students talking about AI. "Erik is exceptionally warm," Nela Richardson, the chief economist at ADP, the payroll processor, told me—unprompted, as soon as we started a call. She described meeting "an economist who gives you eye contact" as a "special" thing.

His conviviality powers his work. At Stanford, he set up his lab in the computer-science building instead of the economics building, making him a bridge between the number crunchers arguing that AI was a "fad" and the code writers arguing that "the world was going to end," Houda Nait El Barj, a former student who now works at OpenAI, told me. He hosts and convenes, avidly and constantly—billionaires, Nobel laureates, garage-dwelling start-up types, obscure researchers, students too young to rent a car. And he runs his classes like business-school classes, with industry figures giving frequent lectures. "Although some academic economists are a bit snooty about it, it's a real strength," the economist Diane Coyle of the University of Cambridge told me; it gives him and his students insight into what AI can actually do, how businesses are actually using it, and how it is changing the tasks that workers perform.

His lab is dedicated to studying those changes analytically: figuring out how quickly firms are adopting AI, and whether the technology is acting as a complement to human workers (making them better at their jobs, boosting productivity) or as a substitute for them (rendering them obsolete, suppressing employment) or all of the above, as seems to be the case right now. A recent paper Brynjolfsson co-authored with Danielle Li and Lindsey Raymond found that generative AI is boosting customer-service workers' output by as much as 30 percent. Another study found it is increasing the employment of some workers and decreasing the employment of others.

Measuring the impact of technological change is always "very difficult," Bharat Chandar, one of the researchers in Brynjolfsson's lab, told me. "We don't have an experiment where we can compare one world versus the other." That leaves academics to pore over surveys that are too slow and too broad to be useful, or to devise new ones. In 2021, Brynjolfsson and a few other academics collaborated with the Census Bureau to collect data from thousands of industrial firms. AI had experienced a productivity paradox, they found, but that paradox had resolved at warp speed. Firms

using AI robots were less efficient and profitable at first, and more efficient and profitable soon after. Stanford also teamed up with ADP to produce measures of how AI is affecting employment, wages, hiring, and occupational tasks. The resulting "canaries dashboard" indicates that workers in their early 20s are, indeed, canaries in the coal mine. AI has put them into a hiring recession.

To many AI executives, and to many ordinary people, Brynjolfsson's predictions seem modest. More than half of adults are worried that AI will put someone in their household out of work. While in the Bay Area, I heard smart people argue that productivity growth would quintuple to 10 percent; we would descend into neo-feudalism; we would cure cancer; we would become even more ensorcelled by our smartphones and even less happy; the GDP report would become meaningless. The forecast I found most moving, a tentative one, came from a Stanford undergraduate. They worried that knowledge jobs might not be available when they graduate. I wanted to give them a bear hug.

But plenty of academics believe that Brynjolfsson is still overestimating AI's impact. As his own work shows, people take time to learn how to use new tools. Companies take time to rejigger their workflows and update their systems. Right now, firms are finding that AI "really speeds up or increases" the output of some workers, Coyle, of Cambridge, said—"but that doesn't translate to *organizational* gains. It may be that one department is way ahead, but if the other departments are not catching up, the whole process won't improve."

Other people emphasize the rules and regulations slowing down diffusion and stymieing change. The FDA has to approve each iteration of the AI software used to analyze medical images, for instance. Radiologists still have to read each scan and collaborate with their colleagues to diagnose and treat patients. Plus, AI isn't cheap; companies have to spend millions to get benefits, diverting capital from other resources.

Gordon, of Northwestern, is among the most prominent economists arguing that AI won't improve the underlying economy much at all, whatever his longtime buddy at Stanford says. In 2020, the two made a friendly public bet. If the annual labor-productivity rate in the coming decade averaged out to less than 1.8 percent, Gordon would win. If it came in higher, Brynjolfsson would. (The wager is worth $400, and the loser will give the money to charity.)

So far, Brynjolfsson is winning by a few tenths of a percentage point. But he's not

really winning, Gordon told me, because AI isn't really boosting productivity. To explain why, he divided American workers into three groups: industrial workers (stevedores, roofers); people who can work from home, whether they do or not (accountants, writers); and people who have to go to work in person (baristas, nurses). "By far, the biggest adoption of AI has been in the working-from-home category," he said, yet all of the "productivity improvement in the last eight years is outside that category." Baristas and roofers have gotten more efficient, and accountants have gotten *less* efficient—thanks a lot, ChatGPT.

If AI does make white-collar workers more efficient, the country's productivity rate will still not go up, he said. A Google engineer produces something like $215 of labor value an hour, a waiter something like $27 an hour. If high-productivity firms shed workers and low-productivity firms gain them, he said, it would "completely offset the benefit." The United States is also facing gale-force productivity headwinds, such as the aging of the population and the relentlessly rising cost of health care, education, and housing. What's AI going to do about that?

Gordon is impressed by ChatGPT, which he uses to look up information about the Great Famine and the panzer and the like. (It's not going to tell *him* anything about economics.) Still, he fears AI's potential to undercut wages, increase unemployment, and immiserate young people, particularly the ones who pay a fortune for a useless college degree. And he delights, at least a little, in the tech optimists learning how unpopular their optimism is.

When we spoke earlier this month, college students were booing commencement speakers who talked up AI. His "beloved colleague"—Joel Mokyr, who just won a Nobel Prize for his work on "sustained growth through technological progress"—was due to give a convocation address at Northwestern soon after. Gordon wasn't invited to the address, he said, but he had commissioned a research assistant to text him and let him know if Mokyr got heckled. In turn, he would let *me* know, he said, beaming as if he had won the lottery and not some piddling $400 bet.

Mokyr groaned when I called him to ask about the speech, mock annoyed at his colleague's caper and unfazed by the prospect of being booed. Asked for his thoughts on our AI future, he cited the Talmud: "Ever since our Second Temple was destroyed, the art of prophecy was given to the fools."

At the time of the destruction of the Second Temple, the productivity of human workers took perhaps a millennium to double. At the height of the Industrial Revolution, it took a decade. As of five years ago, it would have taken a century. If Brynjolfsson is right about AI, it might take only a generation.

He believes in AI because he trusts in history (it would be "weird" if this kind of technology didn't lead to exponential growth) and because he is witnessing the revolution himself. "A lot of economists look at data sets and numbers and try to figure out what's going on," he said. Too few visit companies, talk with managers, and speak with actual human beings.

He gets why students are booing commencement speakers. (Mokyr did not end up being one of them, by the way. He delivered a podium-pounding pro-immigration stem-winder and got a standing ovation.) He sympathizes with people who feel like they're on a digital roller coaster and can't get off. Not long ago, a student came to his office hours. She was graduating. She didn't have a job lined up. Nor did many of her friends. "She was worried that her generation was doomed," Brynjolfsson told me. "That's heavy." Yet human beings find it easier to worry about what may be lost than to imagine what could be created, he said. We don't want AI to do everything humans can do. We want AI to change what humans can do. Beyond the corpus of current human achievement, he said, exists "this much bigger space, bigger than you know, as big as you can imagine, of things that have never been done before."

Brynjolfsson wants people to see that, while also understanding the social risks AI poses. Right now, "human skills are inherently decentralized," he said. "No matter how smart Bill Gates is, he can't make all the decisions in Microsoft." But if AI is performing thinking tasks and executing on them, this limitation disappears. AI could create significant unemployment if humans do not adapt fast enough; it could create malign new forms of inequality too. Society needs to act now, or "a lot of people could be made worse off."

The politics that can help get us there are beyond his remit, he said; policy isn't his area of expertise. So he is doing what he always does. He is building the body of human knowledge. And he is acting as a social network. He is convening. He is talking. He is getting people together to figure it out.

ABOUT THE AUTHOR

Annie Lowrey

Follow

Annie Lowrey is a staff writer at *The Atlantic*.