

Technology alone is never enough for true productivity

Par M. Thirapounnho – GOAT LEARNING® Co-founding Partner



***Technology continues to change how we work,
and genAI is the latest, greatest disruptor. But
real productivity happens only when
companies holistically rethink how they do
business.***

As generative AI (gen AI)

grips the attention of business, corporate leadership can support their companies' growth most by prioritizing the needs of their workers while simultaneously standardizing the use of cutting-edge tech. In this episode of the *At the Edge* podcast, economist and Stanford professor Erik Brynjolfsson speaks with McKinsey's Lareina Yee about how gen AI differs from previous technological innovations, why it will likely augment more jobs than it replaces, and why keeping humans in the loop is essential.

- An edited transcript of the discussion follows. For more conversations on cutting-edge technology, follow the series on your preferred podcast platform.

Why gen AI is spreading faster than previous technologies

Lareina Yee: Erik, let's start with your research, which over the last three decades has looked at how digital technology and automation drive productivity growth. Right now, we're excited about generative AI. What's different about generative AI, and what might also be similar?

Erik Brynjolfsson: That's a great question. Let me start with what's similar, because going back over the decades we see some recurring patterns. When I first started working on this, economist Bob Solow asked me to look at the productivity paradox of some amazing technologies in the '80s, since computers weren't translating into real productivity gains. One of the things we learned was that awesome technology alone is not enough. What you really need is to update your business processes, reskill your workforce, and sometimes even change your business models and organization in a big way. This can lead to what we call a productivity J-curve, where initially, as you're doing all those changes, you don't see productivity gains, and may even experience a productivity loss. Over time, the second part of the J-curve kicks in, and you get these bigger benefits. I believe that we're likely to see a similar bit of a lag with generative AI, because it does ultimately require some rethinking of how businesses are run. But what's different with generative AI is that the lag is shorter; generative AI is happening a lot faster. And that's partly because generative AI is, frankly, one of the biggest, most effective technologies for changing the way work is done that's ever been invented. Also, it's just easier to implement than many earlier technologies, and a lot of it is something regular workers can get going on in a few hours or less.

The other difference is that we have a platform where we can distribute and implement generative AI—one that's been built out over the past couple of decades during those earlier technology revolutions like the internet, cloud, and infrastructure. This means that technologies like ChatGPT could go from zero to 100 million users in just 60 days, something we've never seen before.

The dawn of the second machine age

Lareina Yee: Erik, there are two phrases that get thrown around a lot. One is "Fourth Industrial Revolution," which is a really big concept. And then there's "knowledge workers." Can you help us understand how those two concepts come into play with generative AI?

Erik Brynjolfsson: Every time a new technology comes along, you need to rethink how the economy is run. If you simply pave the cow paths and put the

same technologies on top of the old way of working, you don't really get the business benefits. The earlier technological revolutions from the steam engine and electricity each triggered industrial revolutions. I actually like to compress them into two bigger, mega revolutions: the first machine age and the second machine age. I wrote a book called *The Second Machine Age* with Andrew McAfee, which focuses on automating and augmenting physical work, and then doing the same for cognitive work.

Whereas the first wave focused on us teaching the machines what to do, step-by-step, we're now in the second wave of that process. The new wave is based around machines learning how to solve problems on their own, which will probably play out over decades. And it will have a transformative effect at least as big as the First Industrial Revolution. And the reason—to your second point—is that it's affecting knowledge work. And most workers in the US economy are information workers or knowledge workers. They move around ideas—bits not atoms.

The physical component of the economy remains important. But now that we're addressing not just the physical but the cognitive, we are likely to have even bigger productivity gains than we saw in the past.

Automating and augmenting jobs, not replacing them

Lareina Yee: Let's talk a little bit about those productivity gains. You've done some research, and so have we. The number that people gravitate to is that these generative AI technologies have the potential to automate 60 to 70 percent of the activities for knowledge workers. That is a staggering number. Can you tell me what this means for jobs?

Erik Brynjolfsson: You talked about automating. I like to say both automating and augmenting, because it's rarely just replacing the entire job or an entire set of tasks. It's more often allowing people to do them better with more quality and effectiveness, augmenting their ability to do the job.

But as you say, it's affecting the majority of the tasks for knowledge workers. We pioneered this methodology, and the basic idea is that you can take a company and break it down into occupations, and then break those occupations down into individual tasks. Then, if you look at each of the individual tasks done by, for instance, software engineers, economists, or radiologists—who we learned perform 27 distinct tasks—you can evaluate each of them as to whether or not there's a gen AI solution that can help them. And by evaluating each of those individual tasks, you get a sense of how much of the economy is likely to be affected.

This is a moving target as technology improves. But one thing that was quite striking was the fact that there was not a single occupation we looked at that was entirely automated, although almost all the jobs had at least some affected tasks. Of course, it will take time for all that to happen, so this represents the potential, not the current reality. And one of the goals is to figure out those optimal places where the technology can be applied first, get it implemented, and convert the potential of technology into the reality of business value.

Lareina Yee: But there are also a lot of concerns in popular media that it could mean some jobs lost and some jobs gained. How should people in business today, at the start of this technological revolution, think about this?

Erik Brynjolfsson: You're absolutely right. **There will be jobs lost, and jobs gained.** But let's keep this in perspective. Technology's always been destroying jobs, and has always been creating jobs. It's actually more dangerous to try to freeze in place all of the jobs as they are today, or as they were ten years ago.

- The only safe option is to be continually reskilling and relearning and managing this dynamism that has made so many companies successful. That said, it's going to require a significant amount of retraining and flexibility in the workforce. Over time, it will, I think, tend to disproportionately eliminate some of the most routine, dangerous, boring kinds of jobs, and disproportionately leave more creative, human-centered work alone.
- That won't always be the case, but it's the general trend. And we're already seeing it complement a lot of human activity, allowing people to do their jobs more effectively than they could before. So in most cases, it's really not a replacement but more of an amplification.

Might AI empathize better than humans?

Lareina Yee: If I'm an average person at a company, what are the types of skills that I should develop to be a manager, a leader in this?

Erik Brynjolfsson: I want to be modest here because that's a tough question, and it's been evolving rapidly. When Andrew McAfee and I wrote *The Second Machine Age*, in 2014, we did make a list of skills, and, to be frank, it held up pretty well for five or ten years. But it's beginning to change. At the top of the list, we put creative work, interpersonal skills, and physical work. And those continue to be ones that machines have a hard time doing. But all of them are under fire right now.

- For example, there was a study where medical doctors were asked to answer some questions, and then an AI system was asked to answer the same questions. And most patients found the AI system's answers were actually somewhat more accurate, or the assessment was that the AI system's answers were somewhat more factually correct. So AI beat the doctors a bit on the IQ side. But more surprisingly, patients also found the AI system warmer and more personable. So that was a little scary to see the AI system beating doctors on both dimensions—and even more so on what we might have thought was the human advantage. We have to keep an eye on this and see how it's evolving.
- That said, I think there are some things where **humans continue to have a real advantage**, like large-scale planning and problem-solving, figuring out what needs to be done and prioritizing, and setting the goals that you want to achieve and pointing the system in the right place. Those are jobs for humans. **Another challenge for AI is dealing with exceptions.** We looked at call centers, and there are always some questions that come up a lot, and others that are very rare. And the AI system is great when it's seen the data and knows how to answer the question it's learned from the data. But when it's an exception, machine learning has trouble.

Gen AI can train workers and improve performance

Lareina Yee: You mentioned the call center study, which is something I've referenced a lot. I've oversimplified it by calling it "the triple win," because everyone was happy. Could you tell us about what your research found at this call center?

Erik Brynjolfsson: I love that name, “the triple win.” So let me just describe what we did first. We looked at the implementation of a large language model [LLM] that was to help with their call center operators. And the great thing for me as a researcher was that they phased it in. Some people had access to it, and some didn’t, so it was a natural experiment where we could really get causal estimates about how the technology performed. The other cool thing was that they used it to augment their workers, rather than try to replace them.

- **So what did we find?** First off, we quickly found very large productivity gains, double-digit gains, sometimes as much as 30 or 35 percent. Within four or five months, the people using the AI system were significantly outperforming their colleagues who’d been on the job for a year or more.
- Second, we found that customers were happier and customer satisfaction scores were higher. We also examined customer sentiment by looking at millions and millions of transcripts, using natural language processing to search for happy and angry words. And there were a lot more happy words from customers when the LLM was used, so they seemed to be better off.
- **Last, but not least**, I was a little worried this would turn into some sort of an electronic sweatshop to squeeze the employees. But we found the opposite to be true. The employees were actually happier, and there was significantly less employee turnover as well.

So it is a triple win, like you said. The shareholders gained because the system was creating more **productivity and business value**. The customers gained more satisfaction and better sentiment. And the employees were happier, with less turnover.

Lareina Yee: It’s pretty remarkable. This was all with humans and machines working together in a way to create a better result. I love that the individual call center agents were happier and more willing to stay with their organization. Do you know why that is?

Erik Brynjolfsson: We found that the least-skilled workers, the same people who had previously been doing worse than anyone in the company, actually benefited the most. That’s because they were basically being coached and learning from the system. And the system was identifying the best answers from across the organization, which often came from the best workers, making them accessible to the least-skilled and least-experienced workers and really bringing them up a lot faster.

- So they achieved the biggest gains, up to 35 percent, whereas a lot of the most-skilled workers had basically zero gains. They were already doing everything close to optimal. And so the system is doing something we hadn’t seen before in the history of technology, which is finding a way to capture the tacit organizational and individual knowledge of the best workers and making it accessible to other workers.
- And the last thing I’ll mention is that, every once in a while, when the system was down for a couple hours, they still had to answer customer calls. We were worried they’d been de-skilled and they’d kind of used it as a crutch. But the opposite turned out to be true. In fact, the workers who had been using the system continued to give better answers because they had internalized them and continued to pass them on to the customers.

Support, not supplant, humans

Lareina Yee: One piece of research we're doing at McKinsey involves the concept of experience capital, or, rather, the knowledge gained with firsthand experience. You learn a lot during your formal education, but what you learn on the job is equally important. The idea of being able to accelerate that learning is the essence of reskilling, and CEOs and senior executives are eager to accelerate the spread of best practices so frontline workers reach the expert level faster. We think there's **about a \$4.4 trillion value in increased productivity**. So if I'm a company trying to capture maybe a couple hundred million of that growth, how do I get started?

Erik Brynjolfsson: Technology alone is not enough. You need to convert it into changes in the way you run your business, and that starts with identifying what the opportunities are.

Lareina Yee: You're pushing the boundaries of academic thinking and even challenging the way in which we understood things. Another piece is this concept of the "Turing trap." Could you tell us about that?

Erik Brynjolfsson: Alan Turing, the renowned computer scientist and mathematician, came up with this iconic test of intelligence, which came to be known as the Turing test. The basic idea was to suppose you had a computer in one room, a human in the other, and you interact with each of them without knowing which is which. And if the user can't tell the difference between the computer and the human, Turing said it has passed the Turing test, and you can consider it intelligent.

- I remember reading about it as a kid and thought, "Oh, that's a great test. It makes a lot of sense to me." But I've since concluded that it's actually not a good test at all. The sad truth is humans are not that hard to fool, and even very simple AI systems can trick them. So it's not really a test of intelligence. But the more important thing for me as an economist is that it's a bad goal.
- If engineers, AI researchers, or business executives are trying to replace humans with machines, that creates two really serious problems. The first one is that it's actually not ambitious enough. Machines can do so many things better than humans. The point is not to try to mimic exactly what a person does. It's to lean into the strengths of technology, whether it's reading millions of books or synthesizing lots of texts, since these are things that humans really can't do.
- Second, and more profoundly, it could lead to a really **dystopian society**. Because if you substitute a machine for human labor, it *tends to drive down the value of wages*, make human labor less valuable, and concentrate wealth and power among capital or technology owners.

And over time, that can lead to a trap that I've called the **Turing trap**, where the people who have lost their economic bargaining power also lose their political bargaining power. So the vast majority of people are left behind, or perhaps left to the charity of the people with all the wealth and power, which isn't a very stable outcome either.

So instead, I propose a different approach, which is **making machines that complement and extend human abilities, allowing them to do things they never could have done before**. If you have machines that augment humans instead of replacing them, you not only create a lot more value, but you have a higher ceiling, and you're also more likely to share prosperity. I think that's the path going forward. We can avoid the Turing trap by complementing instead of substituting human labor.

‘Technology alone is not enough.
You need to convert it into changes
in the way you run your business,
and that starts with identifying
what the opportunities are.’

Lareina Yee: There’s a lot to unpack there, Erik. How should business executives think about the Turing trap and their own strategies and implementation?

Erik Brynjolfsson: I think the easy path and, frankly, the lazy one is to just look at what’s already being done and say, “How can we get a machine to do the same thing as this person?” The greater value lies in asking yourself, “How can I do something new that hasn’t been done before?” or “How can I augment this work to create new kinds of quality?” If you’re looking to automate the workers, they’re not going to help you. If you’re looking to augment them and help them do their job better—with higher quality in a faster, more productive way—they’re going to help you with those tough questions of how to be more creative. Ultimately, that’s the much bigger win, and that’s what most of the successful companies are doing.

Thinking beyond GDP to determine digital value

And it also requires different markers. **One thing you and I have talked about is our shared love of data and the fact that some of our traditional metrics don’t work in a digital economy.** One of my earliest memories of studying economics is learning the concept of GDP, but you have a very different take on it. Can you tell us how we need to evolve our thinking?

Erik Brynjolfsson: **First, let me praise GDP.** It was one of the great inventions of the 20th century when Simon Kuznets and his team devised it in the 1930s. That said, great as it was for the 20th century, I don’t think it’s the best way to measure things in the 21st century.

We’ll want to keep GDP around for certain things. But we’re developing a new metric called GDP-B. The “B” stands for benefits, which measures the value consumers derive from these goods and services, as opposed to what they cost. So while GDP measures all the things that are bought and sold in the economy, that means that if something has zero cost, with a few exceptions, it has zero weight in calculating GDP.

But this podcast is going to be free, right? The average American spends about nine hours a day looking at screens of various sizes. That means they’re mostly consuming digital information. And a lot of that digital information is just invisible in the GDP statistics.

‘If you’re looking to automate the workers, they’re not going to help you. If you’re looking to augment them and help them do their job better—with higher quality in a faster, more productive way—they’re going to help you with those tough questions of how to be more creative.’

The idea of GDP-B is looking at how much value you get from something even if it’s free. So maybe you get \$10, \$20, or \$50 a month of value from Wikipedia, and get some value from Google search, Google Maps, and all the other digital goods and services. If you add all that up nationwide, or globally, it amounts to trillions of dollars in value. And as that digital part of the economy grows, it’s going to be more important for us to measure where the value is coming from, not just what we’re paying.

Lareina Yee: So if I’m at a company and am about to launch a set of big bets on generative AI solutions and would like to incorporate the concept of GDP-B into my work, how would I start to think about that?

Erik Brynjolfsson: We typically do a massive number of online-choice experiments that involve offering people compensation to stop using a particular good or service. So we might ask some people, **“If I paid you \$10, would you stop using Facebook for a month?”** Some people will say yes, and others no. You then vary the amounts you offer by as little as \$1 or as much as \$100. And if you do this with enough people, you start getting a range of answers, and you start generating a demand curve where you’ll find that more people will say yes to the larger amounts and fewer people will say yes to smaller amounts.

That produces a **downward-sloping demand curve**, so you now have a sense of what the demand is for Facebook. By doing this at scale with hundreds of thousands or millions of people, you start getting a sense of what parts of the company’s goods and services are actually being valued by the customers.



27-33 Rue du Colonel Pierre Avia
75015 Paris – FRANCE
www.goat-learning.com

Maintaining optimism, thoughtfully

Lareina Yee: Erik, you're fundamentally an optimist about technology and the future. And I think you said once, **"Technology is not destiny. We shape our destiny."** I love that quote. As you look around the corner and envision a positive future with generative AI, what do you see?

Erik Brynjolfsson: Let me connect the first two things you said. I'm an optimist, but I also don't think that technology shapes our destiny. And that makes me what I call a **"mindful optimist."**

There are some people who are just unconditional optimists, who run into friends and always say, "Hey, don't worry. It's all going to work out in the end. Just sit back." Then there are some people who are equally pessimistic. Both of those groups make the same mistake, which is taking away any agency and thinking, "OK, the technology's going to do something to us, and we just have to see what that is."

My approach is different. I believe that we decide what happens with the technology. I'm a mindful optimist because I think **we have the opportunity to do some great things with technology.** And if we do, we're not just going to create trillions of dollars of value. **We're going to have the best decade ever in human history.**

- But I also want to acknowledge that **it's not automatic.** And if we do the wrong things, this could be perhaps the worst decade in human history. There are some awesome and even frightening technologies currently emerging that can spread misinformation, economic disruption, and pandemics, so we have to be careful how we manage them and make sure we do the right thing.
- So the mindful part is as important as the optimist part. And I think by creating a **vision of a positive future**, one that uses technology to create shared prosperity, that uses it as a complement rather than as a substitute, and that involves workers and helps people become more creative, we can build that future together.
- But it's something that's going to take a fair amount of hard, but fun, work by managers, policy makers, researchers, and technologists. And by working together, we have more powerful tools than we've ever had before. So we're going to have more power to change the world in, I hope, a good way in the coming years.

AI as imperfect critic and collaborator

Lareina Yee: Your agency point, the ability to decide for ourselves, is so critical. All the stakeholders do have to come together, since it requires some sort of collective action. That said, **generative AI is one of those technologies that feels incredibly personal.** So maybe on a lighter note, with all you know about generative AI, how do you use it for your own personal productivity?

Erik Brynjolfsson: I try to use a cluster of different tools: Gemini, Claude, and ChatGPT. They all have different strengths and weaknesses, and they're all constantly evolving. I even use Pi for a little coaching, but I try to use a lot of different ones.

I do think the technology is a great critic, and a fun thing I discovered is that it can read my papers and then critique them. It's fun because it doesn't always get it right, but it often has some good insights. I think this highlights the fact that you want to keep a human in the loop. **Combine the technology with a person to make the final call.** So I encourage people to make it a collaborator, to make it a partner. Don't try to outsource your work to it. But if you work together with the system, I think you can often improve your work.

Lareina Yee: I've heard many people use the **analogy of AI as your hardest-working intern**. But I actually much prefer your frame of it perhaps being your most caring critic helping you with your work.

Erik Brynjolfsson: It can be, exactly.

Lareina Yee: A couple more fun questions, given that you're an economist. If you had an hour with any economist in the history of the world, who would you want to spend your hour with?

Erik Brynjolfsson: I mentioned Bob Solow earlier. He was a kind of mentor of mine at MIT [Massachusetts Institute of Technology], a Nobel Prize winner who's not around anymore. He just had this way of cutting right to the core of things and simplifying them in a totally friendly and helpful way. A brilliant guy, but also one of the nicest guys you'd ever meet. And it would be awesome to have an hour to show him what technology is doing today. He came up with what was called the Solow paradox, or the productivity paradox, which showed that early computers weren't creating value. And now we want to understand how AI can create more value. I'm sure he would have had some interesting insights, and I'll try to channel him as I do my research going forward.

Lareina Yee: That would be an amazing conversation to hear. Whenever we get together, oftentimes you're carrying books and I'm carrying books. What's on your book list these days?

Erik Brynjolfsson: I'm just rereading *The Geek Way*, by Andrew McAfee. He spent a lot of time visiting companies and understanding what makes them successful. Geek, by the way, is a compliment at MIT, where we were both working and he still is. So teaching people the geek way is teaching them how to be cool and effective for a successful company.

Lareina Yee: Erik, thank you so much for your time. It's been an absolute joy.

Erik Brynjolfsson: Always good to talk to you, Lareina.