How will AI affect the economy?

By Diego Comin

Professor of Economics

Dartmouth College

Both what is new and what we don't understand obsess us, and new things we don't understand obsess us doubly. The fact that the number of Google searches of “artificial intelligence” has multiplied by 10 in the last seven years is proof both of how new is AI and how little we understand it.

Artificial intelligence (AI) consists of the simulation of human intelligence in machines. This diffuse objective encompasses many techniques and technologies that allow computers and machines to carry out a broad range of tasks including decision making, learning, recognizing specific elements in images or videos, understanding human voice and reacting to particular stimuli.

Some applications of AI are very familiar. For example, virtual assistants like Siri or Alexa recognize natural language, process questions or requests that we ask them and respond to them. Streaming platforms make personalized recommendations based on our past choices. Fraud detection algorithms study deviations from spending patterns to detect fraud. Navigation apps use real-time traffic information to suggest the best routes. Ultimately, these AI applications have two common elements: they gather digital information and analyze it using sophisticated methods.

The potential to apply AI to many productive activities has led commentators to pose all sorts of questions: Will intelligent machines replace humans? Will AI revolutionize production processes bringing unparalleled growth in living standards? Will AI transform organizations, eliminating middle managers and replacing them with an army of technologists that train algorithms? Will AI make the income distribution even more unequal causing further social tensions?

As thought-provoking as these questions are, any specific answer at this point is mere speculation given the uncertainty that exists about the future evolution of AI, its applications and how we will use them. Furthermore, an added difficulty with AI is that, unlike other technologies whose use is obvious to the naked eye (for example, cars or wind turbines), most AI applications are not easy to observe when one runs in the park or sits in the office making it difficult to even measure the current diffusion of AI and its impact on the economy.

Despite these caveats, it is possible to make educated guesses about the answers to some of these questions because, like electricity, or computers, AI is what we call a "general purpose technology" (GPT). Caution. Do not confuse GPTs with the popular ChatGPT which stands for "Chat Generative Pre-trained Transformer." GPTs are technologies that have a wide range of applications across different sectors and have the potential to change production processes across the economy. The development and diffusion of GPT's follows some common patterns. Therefore, by looking at the past, we can understand what will occur in the future.

The first feature of GPT's is that they are not isolated technologies but a group of complementary technologies that provide much greater benefits when used together than when used individually. One consequence is that GPTs do not broadly diffuse until a significant number of applications have been developed. Hence, its long diffusion lags relative to non-GPT technologies. For example, the demand for electricity did not reach notable levels until a suite of home appliances that included the radio, the washing machine, the refrigerator, and the electric oven were invented in the first two decades of the 20th century. That was roughly four decades after Edison invented the first commercially viable lightbulb!

It is difficult and long to develop radically new technologies such as GPTs. For example, Thomas Edison tested more than 6,000 different materials to construct the filament for his light bulb in 1879. This was decades after the first models of bulbs were developed by other inventors, but their short lifespan did not make them commercially viable.

As a result of the complementarity of the technologies that make up a GPT and the difficulty of developing them, the impact of GPTs on the economy is very gradual, and only shows up in productivity statistics decades after the introduction of a GPT. The classical example here are computers. In 1987, 16 years after the commercialization of the first personal computer, Nobel prize winner Robert Solow formulated his famous paradox: "computers are everywhere except in productivity statistics." It was not until the mid-90s that U.S. productivity growth picked up and for a decade its rate was comparable to the golden 1960s.

Beyond their effect on productivity, much of the interest in GPTs concerns their potential distributional effects. Are there winners and losers from the diffusion of GPTs? and if so, can we anticipate who they will be? Economists have extensively studied the impact of specific GPTs on the relative demand of college- vs. non-college-educated workers. For example, firms took advantage of improvements in computing power and new software by incorporating into their production processes new tasks that required workers to use computers. Brand-new occupations such as programmer, chip designer, or IT consultant appeared, and the fact that in both new and old occupations, college-educated workers had an advantage in operating with computers fostered their relative demand causing an increase in their salary relative to non-college educated workers.

However, GPTs do not always increase the relative demand of skilled workers. Electricity facilitated the creation of larger, more efficient factories which enhanced the productivity of production workers which were relatively unskilled. GPTs related to transportation such as cars, trucks or planes allowed firms to reach new, more distant markets and increased the scale of their operations, bringing efficiency gains. Those impacted the productivity of all workers, skilled and unskilled, symmetrically. Therefore, the historical evidence does not support a consistent distributional bias of GPTs even though some GPTs can produce it.

What can we extrapolate from the historical regularities of GPTs to AI? What are the similarities and differences between AI and prior GPTs?

Let's start with the speed of diffusion and how long it will take for AI to show up in aggregate productivity data. By now it's safe to claim that the diffusion of AI resembles previous GPTs. The techniques used by AI to analyze data (e.g., machine learning, neural networks) have been around for various decades. As with electricity or computers, new AI applications are being developed increasing the attractiveness of incorporating AI into production processes.

Clearly, AI applications are still far from polished. For example, I prompted ChatGPT to write this article for me, and the result was so disappointing that here I am, typing.

In my opinion, there are three considerations we have to take into acount to assess how long we'll have to wait until we see AI reflected in the productivity statistics. AI applications may take longer to develop than applications for other GPTs for two reasons. First, machine learning and neural network techniques are very effective in identifying rich non-linear patterns present in the data. However, they need lots of data to work. For example, data that covers the financial transactions of several million borrowers is necessary to develop a neural networks algorithm that predicts credit default more accurately than traditional econometric models used by banks. This sort of data does not exist in most contexts, and it is difficult and costly to collect.

Second, the output of algorithms is opaque. For example, an algorithm may predict the probability that a potential client will default on a loan but does not shed light on how it reached this estimate and about the relevance of the different variables considered in the application. Furthermore, algorithms may be affected by confounding factors that have predictive power over the variable of interest. For example, the racial status of an applicant may play an important role in an algorithm even though it may not cause default per se because race is correlated with some causal drivers of loan default that the algorithm has not considered. Additionally, algorithms can be biased by the data used in the training stage A recent study[1] has shown that AI language models contain different political biases depending on the data used to train them, and that it is virtually impossible to clean training data ex-ante to prevent these ex-post biases. I anticipate that the opaqueness and potential biases of the algorithms that are the basis for all AI applications will create resistance to the development of applications slowing down the potential impact of AI in the economy.

There is however a countervailing force that will expedite the diffusion of AI applications relative to prior technologies. During the last 200 years, the speed

[1] "From Pretraining Data to Language Models to Downstream Tasks: Tracking the Trails of Political Biases Leading to Unfair NLP Models" *Shangbin Feng, Chan Young Park, Yuhan Liu and Yulia Tsvetkov.*

of diffusion of new technologies has continuously accelerated.[2] Technologies invented ten years later have on average diffused four years faster. This trend started with the Industrial Revolution and was not altered by the arrival of digital technologies. But, since those are more recent, they have diffused faster than any previous technology. AI applications are even more recent and will surely diffuse faster than any technology we have experienced in history. The example of ChatGPT, which had 1 million users two days after its launch, 100 million 9 months later and 200 million (expected) 13 months later, is consistent with this prediction.

A different question concerns the impact of AI on the long-run growth of the economy. Some commentators have conjectured that AI will change the way we innovate and that it will accelerate the long run growth of the economy. Their argument goes as follows. Artificial brains will replace human brains in developing ideas and as those will be more powerful and will not be subject to diminishing returns, the rate at which new ideas are created will increase bringing a new era of faster technological change and higher growth rates of productivity.

As appealing as this story sounds, I have doubts about its plausibility. Innovation is a complex process, and we are far from understanding it well. One thing we know about innovation is that it does not consist only in having new ideas. Ideas need to be transformed into prototypes that embody the ideas before they can be commercialized and used. But the most time-consuming and resource-intensive part of the innovation process is typically to tinker with the prototypes until they become viable machines, software, products, or processes. Tinkering is hard to automate and digitize and it is unlikely that AI will change that.

Additionally, good ideas do not just result from combining concepts in a reasoned fashion. Chefs do not create new recipes by orderly mixing ingredients, they rely on inspiration which is the "ability to understand something immediately, without the need for conscious reasoning." One thing is to learn what word comes next in a text or what concept mixes well with another based on what humans have done in the past. A much more demanding challenge is to develop the instinct that guides good researchers to produce

---

[2] "An exploration on technology diffusion" Diego Comin and Bart Hobijn, American Economic Review June 2010.

great innovations. I am therefore skeptical about the scope for AI to transform the innovation process and deliver the extraordinary riches dreamt by some.

The significant impact that information technologies have had on the skill premium and inequality have raised interest in the distributional effects of AI. At this early stage, there are already some relevant observations that hint at what we can expect in the future. Recent investigations have found that AI technologies are flattening organizations, eliminating middle-management layers, and increasing the demand for junior tech- and science-trained workers.[3] This incipient trend will impact the career paths of workers as middle management positions will cease to be a natural progression for young professionals.

The evidence also suggests that AI is changing the relative demand of college and post-graduate workers educated in STEM (i.e. Science, technology, Engineering and Math) relative to workers trained in social sciences and humanities. Additionally, companies that start with a larger share of their workforce trained in STEM are also more prone to adopt AI technologies having greater potential to grow and capture market share in their respective sectors from companies with lower reliance in STEM workers. As a result, we are starting to see a STEM premium in wages that raises the wages of skilled workers trained in STEM relative to those that are less capable to adapt to the new requirements in the workplace or whose job can be more easily replaced by AI. Along these lines, it is revealing that one of the key demands of the Hollywood writers to end their strike has been to limit the use of generative AI on script writing.

Despite the plausibility of these trends, it's important to bear in mind that they are based on projections made at the very early stages of AI diffusion and there is considerable margin of error. After all, in the most recent year for which I have seen a measure (2018), less than 0.1% of the US labor force was employed in AI intensive occupations. So, it may be still a bit too early to be obsessed with AI.

[3] Tania Babina, Anastassia Fedyk Alex X. He and James Hodson "Firm investments in artificial intelligence technologies and changes in workforce composition." 2023.