sycomore am

MAY 2024

Productivity gains: anyone for a further 10 trillion dollars?

Written by Tarek Issaoui, Head of Macroeconomic Research

Introductory note: going against any ambitions of improving our hourly productivity, the content was not produced with assistance from ChatGPT. However, the illustrations were generated with help from Dall-E3.

To quote Heraclites, if "no man ever steps in the same river twice", the economy does occasionally face the same opposing currents. Productivity gains are a recurring issue, if not an enigma, in the field of economic sciences: why do we go through long periods of apathy, despite evident technological progress? Are we stuck in a phase of "secular stagnation" with no obvious end in sight?

For years, **automation and algorithmizing** have been growing at a fast pace, but with no visible effects on the "total factor productivity" (the term used by economists to describe the value produced from identical labour units and capital stock). Economist Robert Solow famously said in 1987 that "the computer age was everywhere except for the productivity statistics". Echoing this observation, this same question is now being raised almost forty years later. Replace "computer" with "artificial intelligence" and you will even come up with a catchy headline for an article.

The reference to Robert Solow, who died recently, is worthy of a little aside. The economist and Nobel prize winner was well placed to develop the paradox: Solow built his academic career on the foundation of a theory of growth partly driven by technological progress, which acts as productivity lever that helps to elude the law of decreasing yields inherent to all capitalist systems – i.e. the erosion of marginal profits that comes with the mobilisation of production factors that are less effective than the very first that were used.

To quote an image used by Turgot in the 18th century, the first land to be farmed tends to be the most fertile; later crops will be less productive. **Technological innovation blows fresh air into the growth engine**. Needless to say, according to Solow, in 1987, the absence of materialised productivity gains in the wake of the third industrial revolution (the arrival of computers in companies and households in the 70s and 80s) raised some critical questions.

Indeed, official statistics in America provide a striking illustration: after the euphoria of the post-war boom, productivity gains remained sluggish from the end of the 70s to the middle of the 90s. From 3.5% growth over a ten-year cycle to a little over 1%.

After stating his paradox, Solow then had to wait an entire decade before seeing the benefits of computerisation materialised in the data. One generation later, we face a similar conundrum. Since the crisis in 2008, productivity is stalling – reviving the concept of 'secular stagnation' which had appeared during the lacklustre 1930s and re-employed by Lawrence Summers in 2013 to describe this new economic regime.

Why are developed economies seemingly incapable of taking advantage of the massive R&D efforts and the undeniable adoption of digital tools? Can we expect artificial intelligence, and broadly speaking, the new technological breakthroughs we have been promised, to pave the way for a new golden age of productivity? At this point, let's add that this is not a simple debate to keep experts busy. The stakes are considerable. **According to the McKinsey Global Institute, returning US productivity to its long-term trend of 2.2 percent annual growth would add $10 trillion in cumulative GDP over the next ten years.** The figure is arguably a little too perfect, but it does provide a meaningful order of magnitude.

Knowing that the country's debt has reached several dozen trillion dollars… it is understandable that the return of a strong growth regime for productivity gains would be a solution (the only solution?) to the burden of servicing debt.

Productivity has stalled since the 2008 crisis, reviving the concept of "secular stagnation" that had appeared during the lacklustre 1930s.

Productivity gains over 1 and 10 years in the US

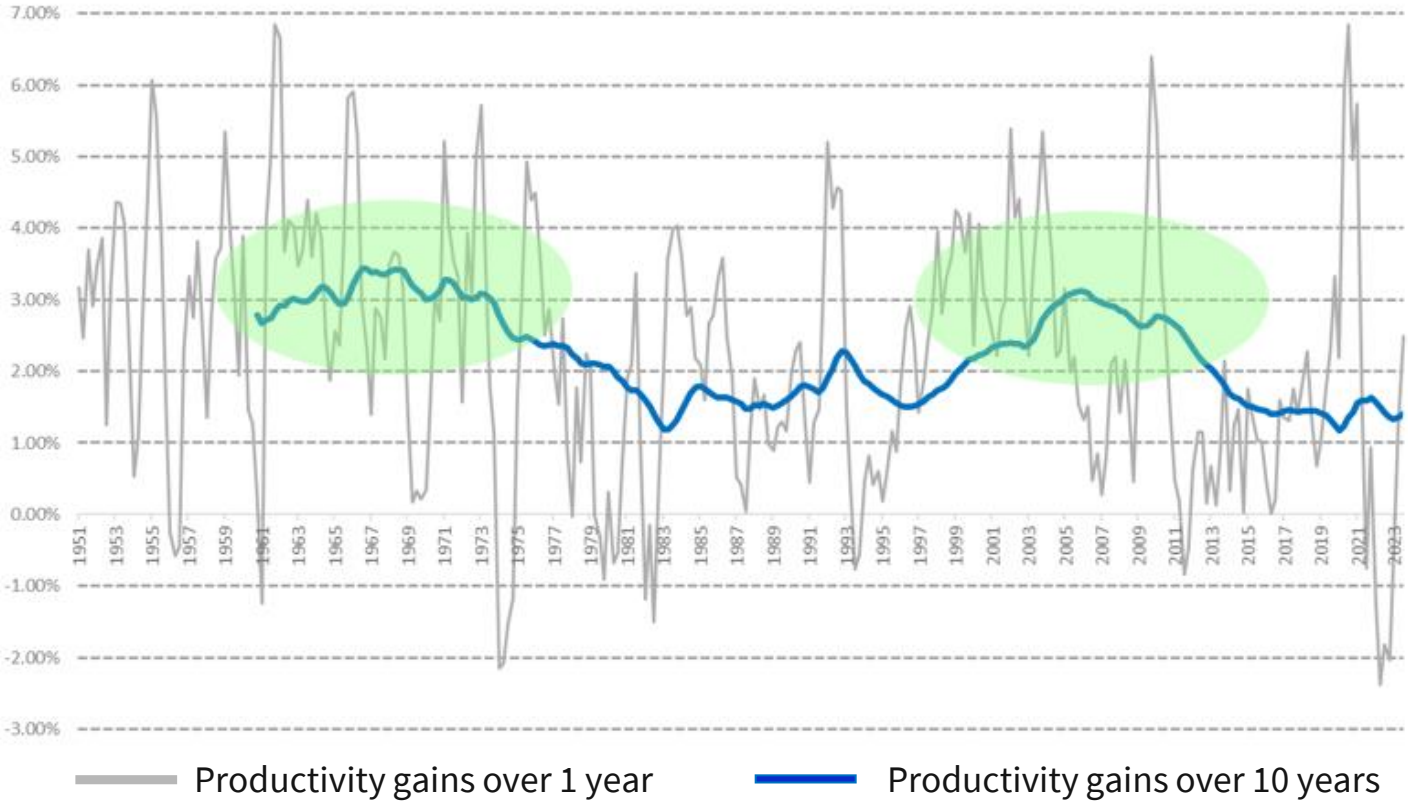

Chart sources: Sycomore Asset Management ; Bloomberg. Data as of end of 2023.

‘We wanted flying cars’

To know whether the conditions for a recovery are duly gathered, identifying the brakes is a useful exercise. Indeed, the question “why is productivity so low” has multiple answers.

One first explanation, often put forward, suggests that the “low-hanging fruit” has already been picked. The first part of the 20th century, until the end of the 60s, saw the emergence and massive adoption of deeply transformative technologies (electricity, transistors, combustion engines…) **that boosted economic growth**. But since then, technological progress has not been such a game changer. The Boeing 707 aircraft inaugurated in 1957 differed in almost every aspect from the first airplane of the Wright brothers but is not that different from the Boeing 777 you may fly with today (or not) when you travel abroad. Man has walked on the moon. We haven’t seen much since 1969. As plainly put by Peter Thiel in 2011, “we wanted flying cars, instead we got 140 characters”, referring to X (Twitter at the time). And we must admit that the time spent ‘swiping’ and ‘browsing’ on various social media is not the most economically productive.

More generally, Western economies have deindustrialised, giving services pride of place. But what should be expect? Put rather bluntly, what productivity gains can a masseur or YouTuber unlock?

The shift to a service economy, and notably its digital component, also generates measurement issues. At this stage, we remind readers that the “official” measurement of productivity relies on the GDP. In a nutshell, productivity gains are calculated as the residual growth in GDP that cannot be explained simply by the number of hours worked (i.e. demographic growth and/or stronger participation from the active population) or by the accumulation of capital stock (equipment and tools of all types made available to workers). However, how does one measure the contribution of a Google search, or time spent on Internet to find a holiday rental, to GDP? This also – and mostly – applies to new products, or those that are created in replacement of others.

For example[1], the camera of your smartphone has replaced your old digital or analogue camera, destroying some of the GDP produced previously by photo shops or stores selling rolls of film. Digital services constitute a blind spot for national accounts. That said, as the academic debate rages on, many studies have shown that accounting measurement issues cannot solely explain the atony of productivity, for several reasons. One of these touches upon the analysis of alternative wealth metrics, such as the Gross Domestic Income, which focuses more on monetary flows, wages, dividends, rents, perceived interest etc.

These revenues are growing faster than GDP, but not enough to explain the gap in productivity, and not according to the temporalities aligned with this gap. A further counterargument seems indisputable at first sight: even by applying current accounting standards, a more granular view shows that **productivity gains appear to be higher, on average, in the most digitalised industries**. Finance and information, for example, are two sectors that have grown to be much more productive today than they were at the start of the 2000s. On the other hand, healthcare, construction, and transportation have lagged. It therefore appears that statistical measurement issues cannot be solely responsible for the phenomenon.

[1] Example borrowed from Philippe Aghion, researcher specialising in institutional growth, innovation and growth.

The promises of artificial intelligence

Not only is digitalisation not a productivity blind spot, but the recent acceleration of artificial intelligence breakthroughs could herald the beginning of a fourth industrial revolution (after those driven by steam, electricity, and IT), potentially characterised by the deep transformation of production methods, which have become faster and more effective.

The development of Large Language Models - including the highly popular ChatGPT - and the unveiling of the first rather outstanding results on quasi-instant text generation, offer the promise of productivity gains. Researchers began to look into the issue in an attempt to measure the impact of using AI for analytical writing and content creation tasks. In 2023, researchers at Harvard Business School estimated that the use of ChatGPT could produce on average 12.2% more tasks, that these would be carried out 25.1% faster, and that the quality of work was improved by 40%[2]. Another study conducted by MIT estimated that generative AI reduced the time it took workers to complete the tasks by 40 percent, while output quality rose by 18 percent[3]. This data should not be taken at face value, but it does reflect how many users of generative text algorithms feel about the technology: used properly, it can largely speed up the production of analytical content.

The first "real life" applications are within reach. Early 2024, J.P. Morgan announced it had developed an AI-based tool able to reduce the manual work needed for cashflow management by 90%[4]. Furthermore, while most attention has focused on text generation, the current field of exploration is much vaster. Logically, the latter includes computer code, image, and video (potentially disrupting the advertising and cinema industries in its wake).

More generally, **artificial intelligence has already accelerated the development of research in life sciences** – with the creation of new enzymes and proteins, or new ways of reading and interpreting the genome – and in physics, for example, in the crucial area of nuclear fusion[5]. New fruit is ripe for picking but it may be hanging a little higher this time. And artificial intelligence will be the stepladder. By processing vast quantities of data very quickly and with the ability to identify recurring patterns (or even imagine more), these new technologies go faster and deeper. A perfect definition of productivity.

The recent acceleration of progress in the realm of artificial intelligence could herald the beginning of a fourth "industrial" revolution.

Of course, the idea is not to imagine a world in which artificial intelligence would take over and permeate all the tasks currently carried out by humans or by less efficient machines. Algorithms can also "hallucinate" and engage in illogical or unrealistic deliriums that cannot easily be stopped. Nobody would dream of relinquishing the controls of aircraft full of passengers, or a scalpel during surgery, to AI. Rather than a "mass replacement", symbiosis is maybe a more appealing concept. But these (justified) counterexamples cannot, alone, remove the possibility of a transformation of production methods that could materialise faster than one might think.

[2] https://www.thecrimson.com/article/2023/10/13/jagged-edge-ai-bcg/
[3] https://news.mit.edu/2023/study-finds-chatgpt-boosts-worker-productivity-writing-0714
[4] https://www.bloomberg.com/news/articles/2024-03-04/jpmorgan-s-ai-aided-cashflow-model-can-cut-manual-work-by-90
[5] https://engineering.princeton.edu/news/2024/02/21/engineers-use-ai-wrangle-fusion-power-grid

More time may be required

Solow's forecasts eventually came true - but only after a decade of waiting. In fine, computers undeniably impacted productivity at work. Their integration simply took a little time, more than expected. **The adoption and integration of new technologies in the corporate world was anything but an instantaneous process.** It required changing organisational models, production infrastructure, skills and sometimes mentalities. This adjustment does not always take place fast or even at the same pace for different institutions and social structures.

Let's take our own industry as an example: finance. There was a time when stock chart graphs were drawn by hand on graph paper[6]. Computers gradually made their way into trading rooms from the 80s. Users had to have a minimum degree of confidence in the technology's ability to assess and handle transactions involving huge amounts of money, with new ways of considering operating risk. Younger workers adopted computers faster than their more experienced colleagues who were sometimes reticent out of self-interest or conviction. Shifting from graph paper to high frequency trading (should the latter be viewed as progress…) required entering a new millennium. While it has been proven that new production models improve efficiency, their integration requires several stages of acceptation and transformation. Economists often quote these institutional brakes to explain the low productivity gains in recent decades.

A recent study by MIT economist Daron Acemoglu[7] has pointed out that the first examples of applied artificial intelligence are "easy to learn tasks" that require no specific expertise and deliver directly measurable results – for example, the time taken to write up a presentation or correct some computer code. According to the researcher, things became more complicated when AI applications were extended to "hard to learn tasks", where context data is vast and/or complex and unforeseen hurdles can require a change of direction.

This takes us back to the distinction between low-hanging fruit and the sometimes-juicier fruit at the top of the tree.

That said, the output from artificial intelligence could be surprisingly fast. For the past three industrial revolutions, integrating new technologies required renewing all infrastructure and equipment. Switching from horse to steam power. From ploughs to tractors. From wrenches to articulated arms. From graph paper and notebooks to screens and Excel spreadsheets.

But with artificial intelligence, the existing hardware is naturally ready, in most cases, to receive the introduction of AI-based software. Microsoft's Copilot, fed from GPT-like Large Language Models, can be rather smoothly integrated into the Windows environment.

This time round, the output from artificial intelligence could be surprisingly fast.

Should you wish to converse with a chatbot on your smartphone, you will not be required to purchase another device, or imagine a radically new instrument to download the software. **The transformation is seamless and does not disrupt the earlier models of analytical production.** Admittedly, the microprocessors will need changing – which is contributing to Nvidia's fortune, and the internal mechanics of these tools need revising. But the global infrastructure of a world fed by AI does not have start from scratch. It is therefore important not to draw hasty conclusions regarding the time needed for productivity gains to materialise.

[6] When I started working in the industry, some of my colleagues had been direct witnesses.
[7] https://economics.mit.edu/sites/default/files/2024-04/The%20Simple%20Macroeconomics%20of%20AI.pdf

From garage to Palo Alto, and back to the garage

When it comes to adapting to this new situation, companies are not on an equal footing. This is one of the explanations sometimes put forward to justify how long it has taken in the past for productivity gains to disseminate throughout the economy: these gains tend to be cornered by a minority of players. In this case, if we refer to digitalisation and artificial intelligence, the tech giants (the "Magnificent 7" that include Amazon, Tesla, Meta, Alphabet, Microsoft, Apple and Nvidia) seem by far the best placed to reap the benefits. Conversely, "traditional players" will be in a weaker position – in technical, fiscal, or human terms – to rival the former in this race. This is particularly true today, as several years of near-zero rates and abundant liquidity have kept zombie companies alive, companies whose main strength is clearly neither agility, nor the ability to deliver these productivity gains.

The domination of superstar companies is often cited by researchers as a hurdle that could prevent innovation from spreading to the entire economy. But here again, we should be wary of hasty conclusions. Very small players have a role to play. In May 2023, an internal Google document leaked on social media; this memo[8] concluded that the Silicon Valley giant was not best placed to compete in the AI arms race, but that the latter would likely be won by "open source", i.e. the myriad of more or less unheard-of players active in the co-construction of free access new algorithms.

Consequently, this could spell a "return to the garage" for the inventors of personal computers, as entrepreneurs - with very modest means compared with the Californian giants – enable deep transformations in the technological apparatus used by our companies.

Closer to our shores, in France, the Mistral AI unicorn, run by a 31-year-old engineer, is about the celebrate its first anniversary and is already able to compete with players such as OpenAI in the area of Chatbots. This spreads beyond algorithms. One need only look at the recent breakthroughs made by researchers in energy production, and notably in nuclear fusion. The ITER reactor project was designed in 2001; the first stone was laid in 2010 and at the time of writing, completion has been pushed back to 2030. The budget has ballooned from 5 billion to 19 billion euros according to the latest estimates. Since, start-ups have joined the race and are reporting spectacular progress with modest budgets (500 million dollars for Helion, a start-up partly financed by OpenAI), miniaturised reactors, and potentially, a first power supply to customers as early as 2028.

The productivity surprise could come from either of two sides: the first is the apparently inexorable breakthroughs achieved by tech giants, who have almost as much power and infrastructure as some countries. The second is the emergence of multiple small-sized initiatives, made available to end-users very quickly, with a high degree of customisation and miniaturisation. Basically, the surprise could come from anywhere.

[8] https://www.semianalysis.com/p/google-we-have-no-moat-and-neither