

Andrei Rogobete: AI, Productivity and the Search for Meaningful Work

This paper is part of a series of essays that seek to explore the current and prospective impact of AI on business. A PDF version can be accessed [here](#).

The previous paper in this series looked at the impact of AI on work through the lens of Peter Drucker's concept of the 'Knowledge Worker'. In this paper we turn our attention to existing and emergent evidence on the impact of AI upon worker productivity. We contend that it is misguided to myopically focus on the perpetrated gains in productivity. Equal importance ought to be given to furthering our understanding of the impact of AI upon concepts of meaningful work, self-esteem and job satisfaction. The Judaeo-Christian framework discussed in the [first paper](#) in this series offers a moral basis that upholds the importance of human dignity and the intrinsic value of humanity as the sole bearers of the *imago Dei* (*image of God*).

The structure here is comprised of three parts. The first will look at both existing evidence and predictions for the impact of AI on productivity, highlighting the often-overlooked time delay between the arrival of new AI capabilities and their materialisation into beneficial productivity tools. The second section turns the attention to matters of meaningfulness, job satisfaction and employee wellbeing in relation to the use and integration of AI tools. The third and final section offers some concluding remarks on how we might begin to think about developing a morally robust symbiosis between AI and work.

'AI productivity gains may be smaller than you're expecting', reads the headline of a recently published report by ING Bank.¹ In May 2023, just 10 months earlier, the Brookings Institute published a research paper titled 'Machines of mind: The case for an AI-powered productivity boom'². What is the current state of AI when it comes to productivity? [Previously](#) we have seen how knowledge worker productivity, though important, presents us with challenges of measurability and accurate prediction. It is important to note here that when talking about AI we are referring primarily to generative AI rather than infrastructure AI which began spreading in the early 2010s and operates largely behind the scenes.

The Macro and Micro Economic Landscape

At almost two years from initial public release of ChatGPT we have an emerging story of two tales: there is a dichotomy of evidence between the Macro and Micro levels when it comes to AI-driven productivity gains. Let's briefly detail some of the existing evidence for each in part.

As of the first half of 2024 there is very little, if any, evidence of AI-driven productivity gains at the Macro level. This perhaps shouldn't come as much of a surprise since some economists, including Charlotte de Montpellier and Inga Fechner, argue that the biggest impact on productivity growth will be seen in 10-15 years' time. This assumes that AI does indeed lead to the much-needed complementary innovations that are expected to be dispersed across an array of different fields.³

The concept of 'complementary innovations' (i.e. innovations that follow and are enabled by the arrival of new technology) is an important one when it comes to gauging the potential impact of AI-driven productivity. General Purpose Technologies like electricity, the internet, personal computers and so on face what is known in productivity theory as a 'J curve' (note: 'GPTs' - not to be confused with ChatGPT which stands for Generative Pretrained Transformer).⁴ This holds that the arrival of new GPTs counterintuitively leads to an initial decrease in short-term productivity

measurements followed by a gradual increase in the medium-to-long term productivity – closely resembling a ‘J curve’. The J curve is largely due to difficulties in accurately measuring the initial GPT adoption investment in intangible capital, as economists like Erik Brynjolfsson *et al.*⁵ point out:

As firms adopt a new GPT, total factor productivity growth will initially be underestimated because capital and labour are used to accumulate unmeasured intangible capital stocks. Later, measured productivity growth overestimates true productivity growth because the capital service flows from those hidden intangible stocks generates measurable output. The error in measured total factor productivity growth therefore follows a J-curve shape, initially dipping while the investment rate in unmeasured capital is larger than the investment rate in other types of capital, then rising as growing intangible stocks begin to contribute to measured production.⁶

So it would be reasonable to assume a degree of delay between the period of initial investment, development and adoption of AI tools, and their derived productivity increases at the Macro level. Some long-term predictions remain optimistic: Goldman Sachs estimates a 7% (or almost \$7 trillion) increase in global GDP and a lift in productivity growth by 1.5 percentage points over a 10-year period – though it should be noted that this estimate is dependent upon AI’s future capabilities and adoption rates.⁷ Other predictions are more conservative: Daron Acemoglu, Professor of Economics at MIT, estimates that AI-driven GDP growth is unlikely to exceed circa 0.93% – 1.16% over the next 10 years, with a total factor productivity (TFP) of no more than 0.66% over the same period.⁸

Thankfully, at the Micro level the picture is less murky. About half a dozen studies provide us with reliable data, three of which will be discussed here. The first is authored by E. Brynjolfsson, D. Li and L. Raymond and looked at the effects of using a generative AI conversational assistant (or AI chatbot) by 5,179 customer support agents.⁹ This likely represents the largest generative AI-workplace study of 2023 and its findings point to some positive outcomes for AI integration within this particular business scenario.

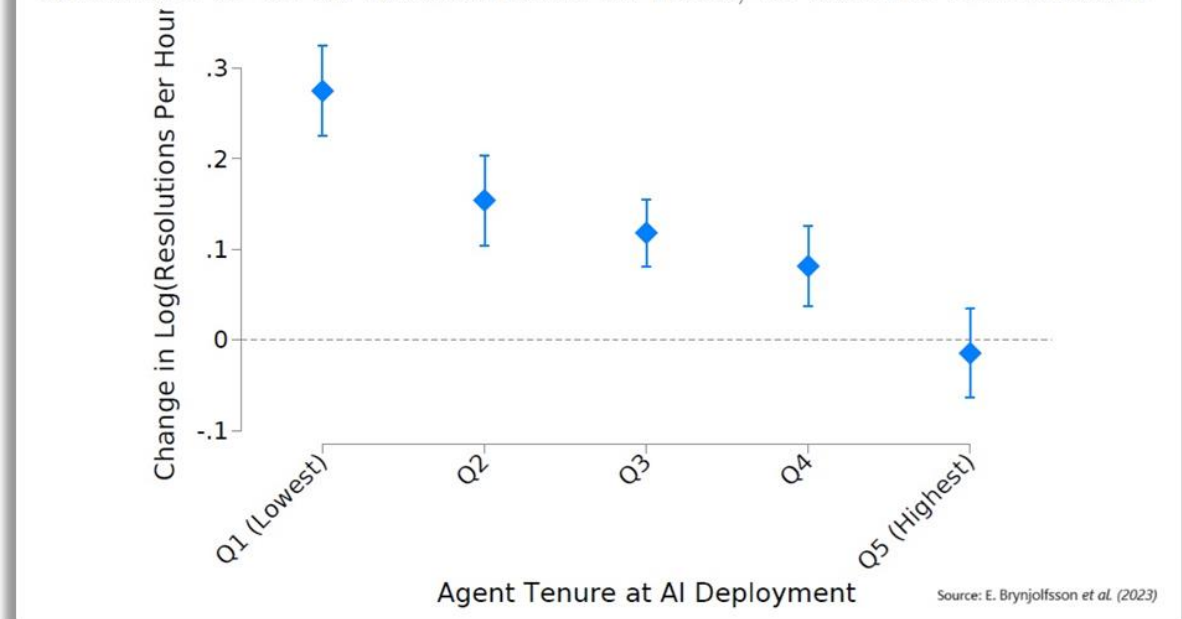
The productivity of each customer support agent was measured in resolutions per hour (RPH). Those that worked with the assistance of the AI chatbot completed on average 14% more RPH than those who didn’t.¹⁰ The research also found that ‘AI assistance improves customer sentiment, increases employee retention, and may lead to worker learning’.¹¹

What is even more interesting is the dispersion amongst high-skilled and low-skilled workers. Figure 1 illustrates the change produced in RPH (y-axis) following AI deployment to the lowest skilled workers (x-axis, Q1), through to the highest skilled workers (x-axis, Q5). The results point to a significant productivity gain of 35% for the lowest skilled workers (Q1), but negligible change in productivity for the highest skilled workers (Q5).¹²

The study found ‘... suggestive evidence that the AI model disseminates the best practices of more able workers and helps newer workers move down the experience curve’.¹³ In other words, the AI chatbot proved to be an effective tool at learning from the best resolutions for certain problems and distributing this knowledge at greater pace and with higher accuracy to the most novice and low-skilled employees. It is important to note that the AI chatbot in this particular study was designed to augment and assist each particular issue and resolution. The final decision of whether to adopt or reject the AI’s suggestions remained entirely at the discretion of the customer support agent.¹⁴

FIGURE 1: HETEROGENEITY OF AI IMPACT, BY SKILL AND TENURE

A. IMPACT OF AI ON RESOLUTIONS PER HOUR, BY SKILL AT DEPLOYMENT



A second notable study by S. Peng *et al.* looked at GitHub’s ‘Copilot’, an AI assistant utilised in computer programming.¹⁵ A group 95 programmers recruited via Upwork, a freelance jobs platform, were tasked with implementing an HTTP server in JavaScript as quickly as possible (though the technical details are not essential for the lay reader). Of the 95 programmers, 45 were in the treated group and 50 in the control group. Performance was measured by (A) task success and (B) task completion time. The results revealed no difference of statistical significance in (A) task success – in other words, both groups completed the challenge with a high rate of success. However, the results did show a 55.8% decrease in (B) completion time for the treated group compared to the control group. This translates to 71.17 minutes versus 160.89 minutes – a net reduction in completion time of 89.72 minutes for the treated group of programmers that utilised GitHub Copilot.¹⁶ It is important to note however, that the study did not evaluate the quality of the code produced by the two groups, and discrepancies here may be significant for the real-world impact of relying on AI tools in programming.¹⁷ So programmers that utilised GitHub’s Copilot finished the challenge an average of 1h 30min quicker than those who did not.

The third study worth mentioning is entitled ‘Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence’, authored by S. Noy and W. Zhang, both from MIT.¹⁸ As the title suggests, the research took an empirical look at the effects of using ChatGPT for a variety of mid-level business-related writing tasks.¹⁹ The study recruited 444 professionals with a higher degree of experience from fields such as data science, human resources, consultancy and marketing. They were all tasked with completing 20–30-minute assignments such as writing a more important email, a short report, press releases, an analysis of various bits of data and so on – all encounters designed to resemble a real-world work environment.²⁰

Between task 1 and task 2, 50% of the participants (i.e. the treatment group), were given the possibility of using ChatGPT for their second task (neither group used AI for the first task). The results in productivity were measured in earnings per minute, with each piece of final documentation being independently evaluated for content quality, writing and originality, and assigned a score. The results reveal a substantial increase in productivity by reducing the average task completion time from 27 minutes to 17 minutes. What is perhaps more interesting is that

the blind evaluations in quality produced reveal an improvement of 4.54 with ChatGPT versus 3.79 without (on a scale of 1-7).²¹

AI and Perspectives on Meaningful Work

The evidence presented thus far broadly points to the adoption of generative AI tools having a positive impact on productivity. However, myopically focusing on productivity gains at the expense of other factors that are relevant to work such as meaning, self-esteem and job satisfaction, risks giving us a distorted and incomplete understanding of the multifaceted implications of adopting and integrating generative AI within the workplace. Indeed, a closer look at some of the relevant studies reveal a more complex picture. Let's start with the concept of meaning and self-esteem.

In philosophy the relationship between work and meaning is well-established, with notable studies by Diddams and Whittington,²² J.B. Ciulla,²³ C. Michaelson²⁴ and others. Within the social sciences we also find a convoluted landscape that encompasses meaningful work, drawing upon contributions from organizational studies, psychology, economics, political theory, and sociology.^{25 26} What exactly does it mean for something to take on the adjective 'meaningful'? The etymology of the word 'meaning' expresses the importance or value of something.²⁷ To become 'meaningful' is to give significance, intentionality and a purpose that pervades the action or the subject in question.

Work is therefore not just a means of economic survival but also a fundamental source of self-identity, worth, and purpose. Work carries repercussions that move beyond the mere intellectual or physical act itself. C. Cordasco from the University of Manchester highlights two broad categories from which work derives meaning and self-esteem: intrinsic and extrinsic. Intrinsic factors involve pride in one's unique personal or collective skills, a genuine interest and enjoyment in the work itself (be that physical or cerebral) and contributions to an organisation or indeed a wider field. Extrinsic factors include the ability to provide for oneself and one's family, the recreational freedom that work provides, the affiliation with certain groups and social networks, and so on.²⁸

The [first paper](#) within this series we considered a Judaeo-Christian approach to AI and work. We highlighted how this implicitly raises wider questions of purpose, meaning and a sense of calling that pervades the mere temporal dimension of work. The Judaeo-Christian perspective therefore seeks to re-evaluate of the gift and place of human agency and responsibility within creation. The foundational texts can be found in Genesis 1:28 and 2:15 where humanity is called to 'Be fruitful and increase in number; fill the earth and subdue it. Rule [...] over every living creature that moves on the ground. [...] The Lord God took the man and put him in the Garden of Eden to work it and take care of it.'²⁹ The command here is here is one of teleological reflection through human capabilities of that which is divine: humanity is given freedom and authority to order, create, steward, and against the backdrop of original sin, also to destroy.

Judaeo-Christian teaching therefore places the concept of work as a key part of what it means to be made in the *Imago Dei* (the image of God), and to actively partake in the eschatological realisation of creation. Work is thus an integral element of Christ's redemptive transformation of the individual and indeed the world. Meaning therefore, finds its ultimate source in the creator God, and this of course encompasses meaning within the realm of work. It is a distinctly human pursuit – no other species on earth searches for meaningful work. Indeed, no other species even reaches a point of asking the question: 'Why am I doing what I am doing?'. As David Atkinson rightly points out in his commentary on Genesis: 'To be in his image is to be aware of ourselves as his creatures'.³⁰

This ability for profound self-reflection is a core characteristic of what it means to be image bearers of the divine. It informs and shapes the meaning of work: if humanity has been gifted with intellectual abilities such as creativity, problem-solving skills, discernment, a capacity to learn new skills and to avoid past mistakes, and has been entrusted with these abilities to care for and steward over creation, then anything that risks compromising these qualities warrants careful attention and scrutiny. The Judaeo-Christian perspective on meaningful work is in some sense dualistic: on one hand God is the ultimate source of purpose and meaning, and on the other, human capabilities play a role in fulfilling and partaking in the larger narrative of God's redemption of creation.

If we turn back to AI, what is the likely impact going to be on meaningful work and job satisfaction? The evidence, while still in its infancy, is patchy. Emergent studies point to both positive and negative outcomes. S. Noy and W. Zhang found that augmentation with ChatGPT in the variety of common office tasks, '...increases job satisfaction and self-efficacy and heightens both concern and excitement about automation technologies'.³¹ The study points out that the recorded increases in job satisfaction are likely due to a heightened sense of achievement when completing a more difficult or tedious task with the assistance of ChatGPT, and in a shorter amount of time than would have otherwise been possible.³²

However, another study by P.M. Tang *et al.* cautions against an overdependence on AI systems as a leading factor in social disconnection and worker loneliness:

This coupling of employees and machines fundamentally alters the work-related interactions to which employees are accustomed, as employees find themselves increasingly interacting with, and relying on, AI systems instead of human coworkers. This increased coupling of employees and AI portends a shift towards more of an "asocial system" wherein people may feel socially disconnected at work.³³

Similarly, C. Cordasco points out that while AI development poses a significant threat to traditional sources of self-esteem derived from work, halting AI is neither feasible nor the best solution. Instead, society should explore new ways of cultivating self-esteem that align with the evolving technological landscape.³⁴

A report by Boston Consulting Group's (BCG) Henderson Institute investigated how people can 'create and destroy' value with Generative AI and found that, '...it isn't obvious when the new technology is (or is not) a good fit, and the persuasive abilities of the tool make it hard to spot a mismatch. This can have serious consequences: When it is used in the wrong way, for the wrong tasks, generative AI can cause significant value destruction'.³⁵ The study had access to over 750 BCG consultants as subjects and found that in areas such as creative product innovation, AI tools boosted productivity by 40%, but in areas like business problem solving, generative AI actually led to a 23% reduction in productivity. The report also highlighted an important trade-off when it comes to collective creativity. Whilst individual performance may be boosted by 40%, collective diversity of ideas may fall by 41%. This is largely because AI chatbots tend to produce the same or similar responses to the same specific prompts – resulting in positive outcomes at the individual level but repetitive and less diverse outcomes at the collective level.³⁶ The potential impact of AI tools on human creativity also seems to be an issue of concern: out of a group of 60 BCG consultants, 70% expect a negative impact on creativity, 26% do not anticipate a negative creative impact, and 4% are unsure.³⁷

Conclusions

It is important to note that when attempting to draw conclusions about the impact of AI upon the world of work, we are (whether we like it or not), operating along several core variables, or axes.

The first would be the level of automation (high) versus augmentation (low). The second represents the level of skill of the employee or group of employees in question. Here it is becoming increasingly apparent that there seems to be a positive reduction in productivity inequality, with at least in these nascent stage, low-skilled workers standing to benefit the most from AI tools. There is also a challenge of AI discernment, what some authors have called a 'jagged technological frontier', whereby the most successful employees and managers will learn to distinguish which tasks are best suited for AI assistance and which aren't. ³⁸ The third is perhaps less a variable than a recognition that the business world represents a plethora of highly distinct work contexts and scenarios where AI implementation may or may not play an important role.

All of these factors are essential when attempting to understand the impact that generative AI has upon work. Broadbrush conclusions about the impact of AI are at best generic, and at worst, inaccurate. Therefore, at least in these early stages, we have to operate on a case-by-case basis and seek to identify and understand areas where AI is a net contributor, and not a hindrance, to both productivity and matters surrounding meaningful work.

Central to the Judaeo-Christian framework is the importance of humanity as the sole image bearer of the divine, tasked with responsibilities of stewardship over nature. In fulfilling the stewardship command, humanity also has the duty of recognising and protecting distinct human attributes such as meaning, purpose, self-esteem and creativity. Emergent technologies therefore ought to be developed and harnessed in harmony with the qualities conferred by humanity's uniqueness, not against them.

Bibliography

- ¹ Charlotte de Montpellier, Inga Fechner, 'AI productivity gains may be smaller than you're expecting', ING Bank, April 2024, <https://think.ing.com/articles/macro-level-productivity-gains-ai-coming-artificial-intelligence-the-effect-smaller/>.
- ² Martin Neil Baily, Erik Brynjolfsson, Anton Korinek, 'Machines of the Mind: The Case for an AI-powered Productivity Boom', Brookings Institute, May 2023, <https://www.brookings.edu/articles/machines-of-mind-the-case-for-an-ai-powered-productivity-boom>.
- ³ Charlotte de Montpellier, Inga Fechner, 'AI productivity gains may be smaller than you're expecting', ING Bank, April 2024, <https://think.ing.com/articles/macro-level-productivity-gains-ai-coming-artificial-intelligence-the-effect-smaller/>.
- ⁴ Erik Brynjolfsson, Daniel Rock, Chad Syverson, 'The Productivity J-Curve: How Intangibles Complement General Purpose Technologies', *American Economic Journal: Macroeconomics*, Vol. 13(1): 333-72, (January 2021), DOI: 10.1257/mac.20180386.
- ⁵ Ibid. p. 1
- ⁶ Ibid. p. 3
- ⁷ Goldman Sachs, 'Generative AI could raise global GDP by 7% ', April 2023, <https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7-percent.html>
- ⁸ Daron Acemoglu, 'The Simple Macroeconomics of AI', paper prepared for *Economic Policy*, Massachusetts Institute of Technology, (April 2024), p.4.
- ⁹ Erik Brynjolfsson, Danielle Li, Lindsey R. Raymond, 'Generative AI at Work', *National Bureau of Economic Research - Working Paper 31161*, <https://www.nber.org/papers/w31161>.
- ¹⁰ Ibid. p. 10
- ¹¹ Ibid.
- ¹² Ibid. p. 15
- ¹³ Ibid.
- ¹⁴ Ibid. p. 9
- ¹⁵ Sida Peng, Eirini Kalliamvakou, Peter Cihon, Mert Demirer, 'The Impact of AI on Developer Productivity: Evidence from GitHub Copilot', *arXiv Accessibility Forum*, (February 2023), [arXiv:2302.06590 \[cs.SE\]](https://arxiv.org/abs/2302.06590).
- ¹⁶ Ibid. p. 5
- ¹⁷ Ibid. p. 8
- ¹⁸ Shakked Noy, Whitney Zhang, 'Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence', *Science*, Vol. 381(6654): 187-192, (July 2023), DOI: 10.1126/science.adh2586.
- ¹⁹ Ibid. p. 1
- ²⁰ Ibid. p. 2
- ²¹ Ibid. p. 4
- ²² Margaret Diddams, J.Lee Whittington, Daniel T. Rodgers, Joanne Ciulla, 'Book review essay: Revisiting the meaning of meaningful work', *Academy of Management Review*, Vol. 28(3):508-512, (June 2003), DOI: 10.2307/30040737.
- ²³ J. B. Ciulla, *The working life: The Promise and Betrayal of Modern Work*, (London: Times Books, 2000), pp. 266.
- ²⁴ Christopher Michaelson, 'Meaningful motivation for work motivation theory', *Academy of Management Review*, Vol. 30(2): 235-238, (April 2005), <https://doi.org/10.5465/amr.2005.16387881>.
- ²⁵ Ruth Yeoman (ed.), Catherine Bailey (ed.), Adrian Madden (ed.), Marc Thompson (ed.), *The Oxford Handbook of Meaningful Work*, (Oxford: Oxford University Press, 2019), pp. 544.
- ²⁶ Catherine Bailey, Marjolein Lips-Wiersma, Adrian Madden, Ruth Yeoman, Marc Thompson, Neal Chalofsky, 'The Five Paradoxes of Meaningful Work: Introduction to the special Issue 'Meaningful Work:

Prospects for the 21st Century”, *Journal of Management Studies*, Vol. 56(3): 481-499, (May 2019), <https://doi.org/10.1111/joms.12422>.

²⁷ Cambridge Dictionary, ‘Meaning’, (July 2024), <https://dictionary.cambridge.org/dictionary/english/meaning>.

²⁸ Carlo Ludovico Cordasco, ‘Should We Halt AI to Protect Meaningful Work?’, *ResearchGate*, (December 2023), DOI: 10.13140/RG.2.2.22893.77288, p.7 -18.

²⁹ The Holy Bible, (NIV Translation).

³⁰ David Atkinson, *The Bible Speaks Today Series: The Message of Genesis 1—11: The Dawn of Creation*, (Westmont, IL: InterVarsity Press, 1990), p. 37.

³¹ Shakked Noy, Whitney Zhang, ‘Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence’, *Science*, Vol. 381(6654): 187-192, (July 2023), DOI: 10.1126/science.adh2586. p.1.

³² *Ibid.* p. 9

³³ Pok Man Tang, Joel Koopman, Ke Michael Mai, David De Cremer, Jack H. Zhang, Philipp Reynders, Chin Tung Stewart, and I-Heng Chen, ‘No Person Is an Island: Unpacking the Work and After-Work Consequences of Interacting with Artificial Intelligence’, *Journal of Applied Psychology*, Vol. 108(11): 1766–1789, (2023), <https://doi.org/10.1037/apl0001103>.

³⁴ Carlo Ludovico Cordasco, ‘Should We Halt AI to Protect Meaningful Work?’, *ResearchGate*, (December 2023), DOI: 10.13140/RG.2.2.22893.77288, p. 35-36.

³⁵ François Cadelon, Lisa Kraye, Saran Rajendran, and David Zuluaga Martínez, ‘How People Can Create—and Destroy—Value with Generative AI’, *Boston Consulting Group Henderson Institute*, (September 2023), pp. 21.

³⁶ *Ibid.* p. 15

³⁷ *Ibid.* p.16

³⁸ Fabrizio Dell’Acqua, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Kraye, François Cadelon and Karim R. Lakhani, ‘Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality’, *Harvard Business School*, Working Paper 24-013, (June 2024), p. 2.