

Early insights into the labor market effects of Generative Artificial Intelligence

Emma Wiles & John J. Horton

October 2024

As long as there has been an economy, it has been disrupted by changes in technology. Rapid advances in Artificial Intelligence, particularly in Large Language Models, have already begun to change both the kinds of work that employers of their workers and the skills available to those same workers.

AI has the potential to affect just about every part of the labor market, from how workers and employers search for and match with each other. AI is likely to automate many tasks currently performed by humans, but it can also lead to the creation of new tasks for humans to do — and it can complement workers by making the ones who use it more productive.

Beyond the automation debate, AI is also likely to have uneven effects on workers' skills and income. It is likely to lower costs, but it perhaps will also lower the quality of some goods and services.

This article explores the current state of research on AI and labor, examining evidence from both laboratory experiments and real-world studies, and it uses economic intuition where studies do not yet exist to make predictions for the labor market of the near future.

Automation vs augmentation

When new technologies affect the labor market, they tend to have two big direct effects.

First, they can have an automating effect, where workers are replaced by the technology. This effect is visible and very salient, for example it is obvious when robots take over for workers in auto factories.

The second effect, which we will call the augmentation effect, is when the technology creates new tasks and jobs.

AI can also have a third, indirect effect — boosting the productivity of workers who use it.¹

There is evidence for all of these effects playing out with generative AI. Some people who need copy-editing work done will no longer hire a copy editor; instead, they will use an LLM. This is a task that would have previously gone to a person and a clear example of automation.

But what about the new tasks and jobs that AI creates through augmentation? This effect is harder to see — but it is there. Every year, about 0.5 percent of the U.S. workforce finds themselves in occupations that didn't even exist the year before.² That's the equivalent of the entire population of Philadelphia working in a job that did not exist just a year earlier, and the effect repeats *every year*.

The direct impact of the technology on total demand for labor will depend on whether the automating effect (less for people to do) or the augmentation effect (more for people to do) will dominate.

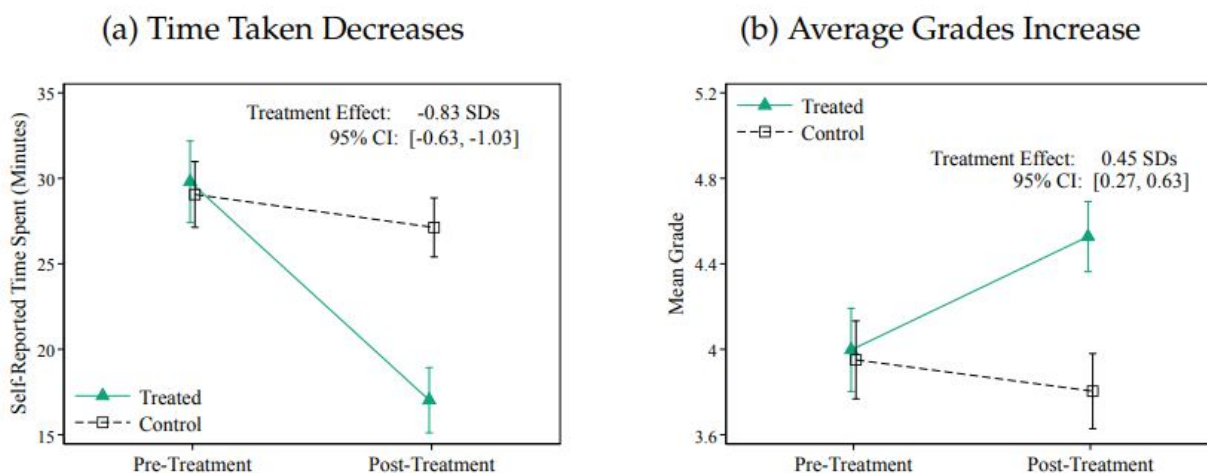
The economy's overall ability to employ people is a macroeconomic question that will depend on a variety of factors — monetary policy, whether job loss is concentrated in some particular places or skills, and much else. But even if the economy will tend towards a natural rate of employment in the long run, there will be large adaptation costs if AI directly automates away more than it creates new work. As we have seen from the China Shock literature, these adaptation costs can add up. It could mean some individuals suffer from lower lifetime earnings or leave the labor market early.

The automation and augmentation effects of LLMs will take time to emerge. Firms don't immediately replace their entire workforce with robots.

The productivity effect, however, can be studied as soon as the technology is available to researchers.

Productivity effects and evidence from the lab

There is evidence from a whole slew of lab experiments that having access to LLMs makes workers faster at task completion. This effect was found in one experiment in which workers were asked to complete a variety of writing tasks. The workers given access to ChatGPT completed the writing tasks faster and with no loss to quality.³ In another experiment with GitHub Copilot, coders with access to AI performed coding tasks 50 percent faster, also with no loss to quality.⁴



Source: "[The impact of ai on developer productivity: Evidence from github copilot.](#)" by Peng, Kalliamvakou, Cihon, and Demirer. (2023).

This early evidence shows some of the promise of how AI can boost worker productivity, but the experiments have important limitations. They are generally in lab settings where the tasks that workers are hired to do are crafted by researchers. In many cases, the LLM is able to do the task at hand with little or no input from the worker. To understand if and how productivity gains will come to workers in real jobs, for which tasks are more complex and specific, we need to look at evidence from field experiments.

Evidence from the field: more negative

Field experiments and natural experiments are run “in the wild” rather than in a lab. They take place in a real market, in this case a real labor market, online platform, or firm.

In natural experiments, the randomness in the study occurs by chance (or for some purpose other than research). In field experiments, which are harder to run but easier to analyze, the randomness of the experiment is specifically for the purpose of the research at hand.

On the question of how AI affects productivity, it is not obvious that the productivity effects we see in lab settings will hold up in the real labor market. The impact on speed or quality could be a construct of the tasks that the researchers are writing for these workers to do. These tasks may not be representative of the complex, context-specific work that people perform in real-world settings.

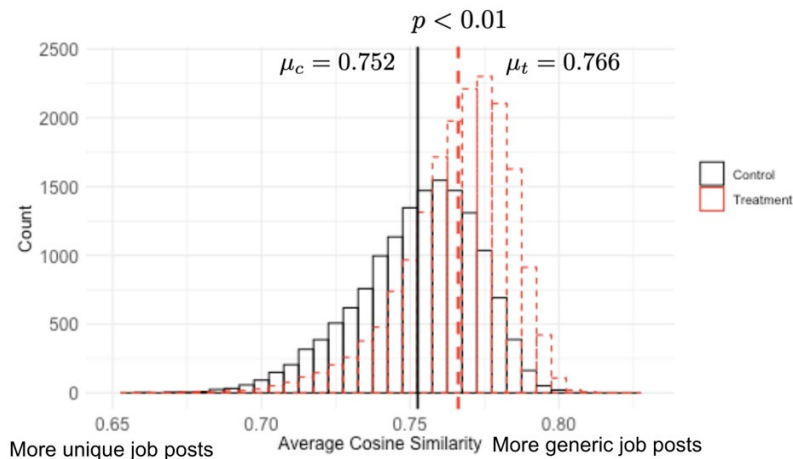
While research is early, in some cases, the positive impact of AI on quality of work in lab settings do appear to hold up in the real world. One study shows that when surgeons were paired with robots in health practices in England, the patients of AI-enhanced surgeons saw lower post-operative length of stay and decreased likelihood of death.⁵ In another study, customer service agents who were given access to LLM suggestions performed better at their jobs by resolving more customer complaints.⁶

Despite these positive examples, in most cases the relationship between AI and output quality in real-world settings is more negative than those from the lab, so far at least.

Yet the average effect of AI on productivity does not tell you if, in the long run, AI will be utilized in ways that increase productivity. Important to understanding why this could be is the economics concept of “free disposal”, which implies that consumers can’t be hurt by getting more things, even if they didn’t want some of those things. For example, those who love period films but hate horror movies are not harmed when Netflix adds a horror movie to its catalog. They can simply not watch it. Similarly, if AI is worse at some task than the human using it, the human could simply not use the AI for that task.

Yet free disposal does not always apply in reality. When Apple “gifted” the new U2 album to every iPhone users’ library, it sparked a massive backlash despite the free disposal. Likewise, research suggests that in some cases, when people are given suggestions from AI that are worse than what they would have done without it, they still take the bad suggestions.

In a forthcoming study, we give evidence from an experiment on an online labor market, where half of new employers interested in posting a job were given first drafts, written by AI, of the job post.⁷ We find that these employers are much more likely to post a job, but those jobs are much less likely to find someone to hire. These employers spend 40 percent less time writing the job posts than those without the AI-written draft, and the jobs they post have text that is much more generic than the texts of other job posts. It looks like in this case, access to the AI draft “crowded out” effort that the employer would have otherwise put into writing a job post that accurately reflected what type of worker they were looking to hire.



Source: "[More, but Worse: The Impact of AI Writing Assistance on the Supply and Quality of Job Posts](#)" by Wiles and Horton. (2024).

Another example finds that when recruiters have access to applicant recommendations by very high-quality AI, they take the AI's suggestions *even when the suggestions are not correct*.⁸ But when they have access to suggestions from a lower-quality AI, they perform better. Both of these studies point to the perils of mindlessly trusting AI despite free disposal.

Lowering costs

Whether AI automates or augments humans, we would expect the cost of the goods or services that they produce to go down. So far, the literature comprehensively suggests that AI drastically reduces the time it takes workers to do a task. This large decrease in time spent per task should dramatically decrease the costs of production for firms, even if it comes with some loss to quality.

One way to think about this is to compare photography pre and post the smartphone era. In the 20th century, photography was an expensive hobby requiring specialized equipment, which was inaccessible to the masses. Photos were perhaps more carefully composed, and the average person might only buy a few rolls of film a year to dole out only when the payoff of a photograph was high. Now most of the people on earth own a smartphone with a camera, and can take close to unlimited photographs. Photography is accessible to nearly everyone, and there are exponentially more photographs in existence. My (Emma's) camera roll has a few excellent shots I never would have had without my smartphone, but mostly it just contains screenshots, memes, selfies, photos of sandwiches, photos I took by accident, etc. Despite the huge increase in the cameraphone technology, it is still likely the quality of the average photo in 2024 is much lower than it was in 2000.

Generative AI could have a similar impact in some domains, decreasing the cost of creating certain goods and services, making them cheaper and more accessible — but with a possible decline in their average quality. AI-generated articles, logos, tweets, or even legal documents, for example, might be lower in quality than human-created versions, but their relatively lower cost should make them accessible to a much wider range of people.

Importantly, lower costs can also benefit workers and not just consumers. Research by James Bessen shows that if lower costs drive prices down, and if consumers are "elastic" (responsive) enough, this could cause an increase in consumer demand large enough to grow employment — even though the labor required to produce a given unit of output declines.⁹

Comparative advantage

For now, people seem to be better than LLMs at a good number of things, but that is not always guaranteed to be the case. LLMs are getting trained on larger and larger datasets and the scope of tasks they can do better than humans is increasing every day. It is possible that one day, AI will be better than humans at literally every type of task. In this case, will everyone be out of a job?


It's unlikely, because of what economists call "comparative advantage". While AI might be better at everything than people, there will be some things that people are relatively better at. To give a concrete example of this, imagine that AI is better at gene sequencing and accounting than people. Humans and AI both have a finite amount of energy (AI is limited by computing power). So if AI is much better at gene sequencing than people, but only slightly better at accounting, it is likely that we will hand over gene sequencing to be done by AI and we will do more accounting.

Winners and losers

As AI gets cheaper and more widespread, there will inevitably be people who benefit from it and people who are hurt by it. There is a fair bit of evidence from the economics of AI literature that when you give people access to AI tools, the lowest-skilled people are the ones who benefit the most. For example, in Wiles and Horton (2023), we find that giving job applicants access to AI while they write their resume causes the worst writers to improve the most from the software.¹⁰

The economists Kyogo Kanazawa, Daiji Kawaguchi, Hitoshi Shigeoka, and Yasutora Watanabe find that giving taxi drivers AI suggestions for route improvements only resulted in time saved for the worst drivers, and at best had no impact for the most skilled drivers.¹¹

Figure 1: AI effects by skill

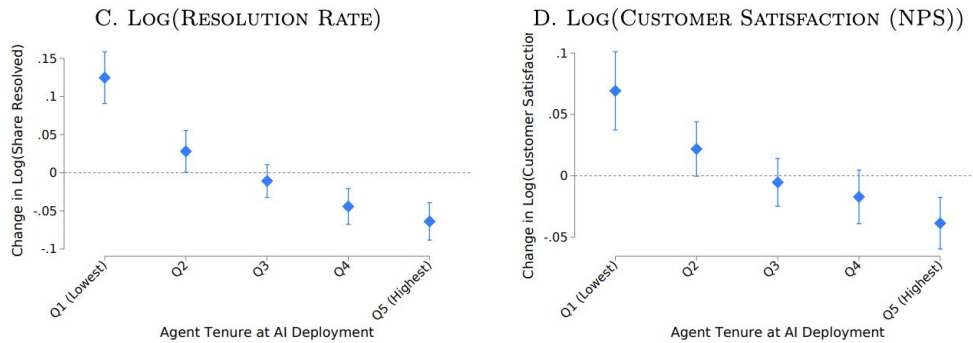


Note: This figure plots the estimated value of AI Navi effects by skill with 90% confidence intervals in the shaded areas. The Weibull hazard regression, where the interactions of AI usage dummy and skill index and square of skill index are added to equation (1), is estimated. See column (1) of Table A1 for the corresponding estimates. The outcome is cruising time. The negative estimate indicates that AI usage reduces the search time. Higher skill index indicates more skilled drivers. Figure A3 displays the distribution of the skill index. Drivers whose skill index ranges from -1.5 to 1.5 cover 92.4% of all cruises. The "full" sample with all drivers is used.

Source: "[AI, Skill, and Productivity: The Case of Taxi Drivers](#)" by Kanazawa, Kawaguchi,

Shigeoka, and Watanabe. (2022).

In some cases, AI not only makes the lowest-skilled people's output better, but it can also make the highest-skilled people's output worse. In the example earlier of the customer service agents with access to AI, the worst customer service agents benefitted from having access to the technology, but the best customer service agents' performance was lower when they had access to the AI.¹² This result holds for both performance (measured as what fraction of customer disputes they resolved) and their customer satisfaction ratings after the fact.



Source: "[Generative AI at Work](#)" by Brynjolfsson, Li, and Raymond. (2023).

Similarly, in a study by Jonathan H. Choi and Daniel Schwarcz, law students were randomly given access to GPT4.¹³ Students at the bottom of the class performed much better on law school exams if they had access to the AI, but the best students actually did worse if they had access to the AI.

These studies suggest that widespread use of AI could cause a compression of the skill distribution, which in turn could lead to a narrowing of the wage distribution within occupations.

One might infer from this that overall income inequality in the economy would decrease as a result. But if the impacts to productivity are larger for higher wage jobs — say, because software engineers and lawyers have the largest productivity impacts — then the demand for high-wage jobs will increase relative to lower-wage jobs and could dominate the within-occupation wage compression. Luckily, what types of jobs workers hold are not fixed.

Upskilling

In many ways, what happens to workers will depend on how they respond to changes in the demands for skills. The half-lives of skills are short already, and will only get shorter as AI improves.

As the scope of what AI can do increases and the quality of its output improves, this will call for different skills from workers. Firms can hire and fire new workers every time their demand for skills change, but this is expensive, organizationally complicated, and bad for morale. And firms seem to believe in the value of retraining current employees to some extent — a BCG study from 2023 suggests that such investments in upskilling or reskilling represent approximately 1.5 percent of organizations' total budgets.

One way for workers to upskill is to learn how to work with and manipulate the new technology. One thing that makes LLMs a unique technology is that they can both potentially do parts of your job better than you, but this same technology can be used to teach you how to learn new skills. LLMs are the most patient tutors out there.¹⁴ This can be used to make previously lower skilled people higher skilled and command higher wages.

Ideally, AI tools will be accessible to workers across socioeconomic backgrounds and income distribution. But it's hard to believe this is the case. In a survey from Pew Research Center, 76 percent of people in high-income

brackets reported knowing “a lot” or “a little” about ChatGPT, but this number was only 44 percent for low-income people.

Those who don't adapt

A reasonable fear about AI is that for those who don't upskill, their wages will decrease. But this is not necessarily the case. After Steve Jobs announced the Death of Flash, and that Apple would not be supporting Flash, the first order effect might have been a huge decrease in the number of jobs that require Flash and a hit to the wages of the many workers previously working with Flash. But that's not what happened. Jobs which required Flash did in fact go down. But many workers, the younger and maybe more nimble ones, transitioned away from Flash, such that the workers for whom Flash was really all they could do ended up seeing no changes to their wages. Yes, there were fewer Flash jobs. But as workers upskilled or moved away from the dying technology, it left fewer Flash workers and, therefore, no hit to their wages.

Conclusion

We are in the early days of these technologies. In this essay, we have provided evidence from the first round of research on the impacts of generative AI. These studies suggest that widespread use of AI could cause a large increase in productivity and a compression of the skill or even wage distribution.

An important caveat is that these studies are all measuring impacts in “partial equilibrium”. These results are evidence of what happens to an individual who starts using AI in a market in which other people are not using it. We do not yet know, however, what will happen as AI adoption becomes more widespread across the market. It remains uncertain how the demand for different types of work will shift — how AI will alter the relative demand for high-skilled versus low-skilled labor, for instance, or the demand for white-collar jobs compared to more technical roles. The general equilibrium effects of AI on the labor market are still unknown, and these shifts in demand will have significant implications for workers, firms, and the macroeconomy.

Emma Wiles is an Assistant Professor at Boston University's Questrom School of Business in the Information Systems group. John J. Horton is an Associate Professor of Information Technologies at the MIT Sloan School of Management.

Explore the Economic Innovation Group's American Worker Project [here](#).

¹ Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3-30.

² Lin, Jeffrey. “Technological Adaptation, Cities, and New Work”. *The Review of Economics and Statistics* 93, no. 2 (2011): 554–74. <http://www.jstor.org/stable/23015954>.

³ Noy, Shakked and Zhang, Whitney, Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence (March 1, 2023). Available at SSRN: <https://ssrn.com/abstract=4375283> or <http://dx.doi.org/10.2139/ssrn.4375283>.

⁴ Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirel. "The impact of ai on developer productivity: Evidence from github copilot". arXiv preprint arXiv:2302.06590 (2023).

⁵ Tafti, Elena Ashtari. Technology, skills, and performance: the case of robots in surgery. (2022).

⁶ Brynjolfsson, Erik and Li, Danielle and Raymond, Lindsay. Measuring the Productivity Impact of Generative AI. (2023) NBER Working Paper No. 31161.

⁷ Wiles, Emma and Horton, John. More, but Worse: The Impact of AI Writing Assistance on the Supply and Quality of Job Posts. (2024) Unpublished manuscript.

⁸ Dell'Acqua, Fabrizio and McFowland, Edward and Mollick, Ethan R. and Lifshitz-Assaf, Hila and Kellogg, Katherine and Rajendran, Saran and Krayer, Lisa and Candelon, François and Lakhani, Karim R., Navigating the Jagged Technological

Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality (September 15, 2023). Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013.

⁹ Bessen, James. AI and jobs: The role of demand. No. w24235. National Bureau of Economic Research, 2018.

¹⁰ Wiles, Emma and Munyikwa, Zanele and Horton, John. Algorithmic Writing Assistance Increases Hires. (January 2023) NBER Working Paper No. 30886.

¹¹ Kanazawa, Kyogo and Kawaguchi, Daiji and Shigeoka, Hitoshi and Watanabe, Yasutora. AI, Skill, and Productivity: The Case of Taxi Drivers (October 2022) NBER Working Paper No. 30612.

¹² Brynjolfsson, Erik. And Li, Danielle. And Raymond, Lindsay. Measuring the Productivity Impact of Generative AI. (2023) NBER Working Paper No. 31161.

¹³ Choi, Jonathan H. and Schwarcz, Daniel, AI Assistance in Legal Analysis: An Empirical Study (August 2023). 73 Journal of Legal Education (forthcoming, 2024).

¹⁴ Mollick, Ethan R. and Mollick, Lilach, New Modes of Learning Enabled by AI Chatbots: Three Methods and Assignments (December 13, 2022).