

NBER WORKING PAPER SERIES

THE TURING TRANSFORMATION:
ARTIFICIAL INTELLIGENCE, INTELLIGENCE AUGMENTATION, AND SKILL PREMIUMS

Ajay K. Agrawal
Joshua S. Gans
Avi Goldfarb

Working Paper 31767
<http://www.nber.org/papers/w31767>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2023

This paper was initially commissioned by Brookings, and appeared at <https://www.brookings.edu/articles/the-turing-transformation-artificial-intelligence-intelligence-augmentation-and-skill-premiums/>. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w31767>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Turing Transformation: Artificial Intelligence, Intelligence Augmentation, and Skill Premiums
Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb
NBER Working Paper No. 31767
October 2023
JEL No. J2,O3

ABSTRACT

We ask whether a technical objective of using human performance of tasks as a benchmark for AI performance will result in the negative outcomes highlighted in prior work in terms of jobs and inequality. Instead, we argue that task automation, especially when driven by AI advances, can enhance job prospects and potentially widen the scope for employment of many workers. The neglected mechanism we highlight is the potential for changes in the skill premium where AI automation of tasks exogenously improves the value of the skills of many workers, expands the pool of available workers to perform other tasks, and, in the process, increases labor income and potentially reduces inequality. We label this possibility the “Turing Transformation.” As such, we argue that AI researchers and policymakers should not focus on the technical aspects of AI applications and whether they are directed at automating human-performed tasks or not and, instead, focus on the outcomes of AI research. In so doing, our goal is not to diminish human-centric AI research as a laudable goal. Instead, we want to note that AI research that uses a human-task template with a goal to automate that task can often augment human performance of other tasks and whole jobs. The distributional effects of technology depend more on which workers have tasks that get automated than on the fact of automation per se.

Ajay K. Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
and NBER
ajay.agrawal@rotman.utoronto.ca

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
and NBER
agoldfarb@rotman.utoronto.ca

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
and NBER
joshua.gans@rotman.utoronto.ca

The Turing Transformation: Artificial intelligence, intelligence augmentation, and skill premiums

By Ajay Agrawal, Joshua Gans, and Avi Goldfarb

Almon Brown Strowager, an American undertaker from the 19th century, allegedly angry that a local switch operator (and wife of a competing undertaker) was redirecting his customer calls to her husband (Kansas Historical Society, 2011), sought to take all switch operators to their employment graves. He conceived of and, with family members, invented the Strowager switch that automated the placement of phone calls in a network. The switch spread worldwide and, as a consequence, a job that once employed over 200,000 Americans has almost disappeared (Feigenbaum & Gross, 2021).

While the pioneer researchers in new areas of artificial intelligence (AI) such as machine learning, deep learning, reinforcement learning, and generative AI are probably not motivated by similar frustrations with people, their stated goals have nevertheless been to develop human-level machine intelligence. Sometimes the goal is to mimic a human, as in the Turing Test (Oppy & Dowe, 2021). Often, however, a specific task or job is a template for their endeavours. In image classification, the benchmark for AI researchers was superiority over human classifiers, a goal achieved for some tasks in 2015 (Markoff, 2015). Human performance is the benchmark for AI natural language processing and translation. OpenAI demonstrated that their GPT-4 model exhibits human-level performance on a wide range of professional and academic benchmarks (OpenAI, 2023), including a Bar exam, the SAT, and various AP-level courses. AI pioneer and Turing Award winner Geoff Hinton remarked in 2016 that time was up for radiologists (Creative Destruction Lab, 2016) and that no one should continue training in that field. Whether that will hold true or not, it is hardly surprising that recent developments in AI have reinforced the widespread view that the intent of AI research is to replace humans in performing various tasks.

This view has not gone unquestioned. In his book *Machines of Loving Grace*, John Markoff (2016) celebrated researchers committed not to human replacement but to human intelligence augmentation. He argues that the history of computer development showed the failure of replacement and large gains, both commercially and socially, when computers were designed to be a tool that augments the skills of people. Certainly, Steve Jobs had this vision when developing personal computers, seeing them as “bicycles for the mind,” (Michael Lawrence, 2006) with bicycles responsible for one of the greatest advances in human locomotion. Erik Brynjolfsson (2022) has identified the erstwhile Turing Test as an instrument of harm in creating an automation mindset for AI research at the expense of potential augmentation paths.

Markoff (2016) and Brynjolfsson (2022) argue that it would be preferable if AI research travelled a more human-centric path focused on opportunities to augment rather than automate humans. Such AI applications would enable people to do things they could not previously do. This would create a complementarity between the provision of such applications and human capabilities and skills. In this belief, they are joined by Daron Acemoglu (2021a) who has been vocal regarding the risks AI poses for job security unless more diverse research paths are chosen. Critically, Acemoglu sees the potential for AI in many sectors from

health care to entertainment. Closer to home, he speculates on paths not travelled (yet) for AI in education (Acemoglu, 2021b):

Current developments, such as they are, go in the direction of automating teachers—for example, by implementing automated grading or online resources to replace core teaching tasks. But **AI** could also **revolutionize** education by **empowering** teachers to adapt their material to the needs and attitudes of diverse students in real time. We **already know** that what works for one individual in the classroom may not work for another; different students find different elements of learning challenging. AI in the classroom can make teaching more adaptive and student-centered, generate distinct new teaching tasks, and, in the process, increase the productivity of—and the demand for—teachers.

What is holding back such innovations is partially rooted in funding, regulation, and unequal tax treatment between capital and labor. But the advocates for human-centric AI list the mindset of AI researchers as the primary starting point for attitudes to change. Brynjolfsson (2022) argues, “A good start would be to replace the Turing Test, and the mindset it embodies, with a new set of practical benchmarks that steer progress toward AI-powered systems that exceed anything that could be done by humans alone” (p.282).

It appears that Acemoglu and Brynjolfsson want to change the objectives and philosophy of the entire research field. The underlying hypothesis is that if the technical objectives of AI research are changed, then this will steer the economy away from potential loss of jobs, devaluation of skills, inequality, and social discord following from this. In this way, society can avoid what Brynjolfsson (2022) calls the “Turing Trap,” where AI-enabled automation leads to a concentration of wealth and power.

In this paper, we question this hypothesis.¹ We ask whether it is really the case that the current technical objective of using human performance of tasks as a benchmark for AI performance will result in the negative outcomes described above. Instead, we argue that task automation, especially when driven by AI advances, can enhance job prospects and potentially widen the scope for employment of many workers. The neglected mechanism we highlight is the potential for changes in the skill premium where AI automation of tasks exogenously improves the value of the skills of many workers, expands the pool of available workers to perform other tasks, and, in the process, increases labour income and potentially reduces inequality. We label this possibility the “Turing Transformation.”

We argue that AI researchers and policymakers should not focus on the technical aspects of AI applications and whether they are directed at automating human-performed tasks or not and, instead, focus on the outcomes of AI research. In so doing, our goal is not to diminish human-centric AI research as a laudable goal. Instead, we want to note that AI research that uses a human-task template with a goal to automate that task can often augment human performance of other tasks and whole jobs. Furthermore, it is difficult to determine whether any given technology is automating or augmenting. Put differently, one person’s automation can be another’s augmentation, and the two are not mutually exclusive. The distributional effects of technology depend more on which workers have tasks that get automated than on the fact of automation per se.

The paper proceeds as follows. In Section 1, we provide a formal model to demonstrate when we think that automation creates a Turing Transformation rather than a Turing Trap. Section 2 then illustrates some cases in which AI-powered automation has involved those opportunities. Section 3 provides examples of technologies that Markoff (2016) labels as intelligence augmentation but nevertheless led to increased inequality. Section 4 concludes by noting that one person’s substitute is another’s complement, and therefore artificially separating automation from augmentation does not capture the impact of intelligence technology on the distribution of income, wealth, and power.

¹ We present a related argument in Agrawal et al. (2023).

1. A Model

In order to be more precise in the description of these concepts, it is useful to formalize these ideas. Here we build upon a model provided by Acemoglu (2021a). He assumes that there are two tasks to be performed, labelled 1 and 2. The output of a firm in a sector is given by:

$$Y = \min\{y_1, y_2\}$$

where y_i is the output of task i . The production function here means these tasks are strong (that is, perfect) complements.

In the absence of AI, humans perform the tasks. While a human's skill level does not impact the productivity of task 2, there are specific skills that can improve the productivity of task 1. It is assumed there is a measure $[0, \alpha]$ of workers available with $\alpha > 2$. (Acemoglu assumes that $\alpha = 1$.) A measure 1 of these have a specialised skill while the remainder (of measure $\alpha - 1$) are generic. Thus, there are more workers with the generic skill than the specialized skill. The specialised skill is only valuable when used in firm production.

Workers of both types, skilled and generic, can earn an outside (hourly) wage of w ($< 1/2$), from self-employment. Each worker is endowed with 2 units of time (i.e., hours). All workers who devote a unit of time to task 2 can produce an output of 1 for that task. By contrast, for task 1, only skilled workers can produce an output of 1, while generic workers produce $x < w$. This means that if workers do both tasks (with one hour devoted to each) skilled workers produce $Y = 1 (= \min\{1, 1\})$ while generic workers produce $Y = \min\{x, 1\} = x$. Thus, it would only make sense to have the generic workers perform both tasks by allocating a fraction, x , hours to task 2 for a total wage bill of $(1 + x)w$. However, as $x < w < 1/2$, this means that if generic workers do both tasks as their job, their marginal product, x , is still less than $(1 + x)w$. So, it is only economical to hire skilled workers whose net contribution to the firm is $1 - 2w$. Thus, the total payment to labour is at least $2w$ but may be as high as 1, if there is a scarcity of skilled workers in the economy.²

Without AI, other than having skilled workers perform both tasks, production could be organised by having workers specialise in each task, with skilled workers performing task 1 and generic workers performing task 2. This can potentially generate combined output (amongst each pair of workers) of $Y = 2 (= \min\{2, 2\})$ for a pair of workers. However, coordinating the tasks between them is not without cost. Thus, following Acemoglu, it is assumed that if there is not a single worker doing both tasks, there is a loss in economies of scope and the productivity for *each* task falls by a factor of $1 - \beta > 0$. This might arise because individuals learn from performing both tasks at the same time or from a cost of coordinating between tasks. Thus, if different workers worked on the same task (with the skilled on task 1) total output would be $2(1 - \beta)$ and firm surplus would be $2(1 - \beta - 2w)$. If $1 - 2w > 2(1 - \beta - 2w)$ (which simplifies to $2\beta > 1 - 2w$), and if firms operate in competitive product markets, it would be preferable to hire only skilled workers performing both tasks. We assume this throughout this paper; allowing for the possibility that AI adoption transforms the nature of the job.

Suppose now that there exists an AI that could automate task 1 at a unit cost of $c < 1$. Firms using AI are not constrained by the supply of skilled workers of measure 1. Thus, output is $2\alpha(1 - \beta)$ less the cost of buying the AI to complement worker output, which is $2\alpha c(1 - \beta)$.³ However, as the firm no longer relies on skilled workers, its labour costs become $2aw$. It is, therefore, profitable for a firm to adopt AI if $(1 - \beta)(1 - c) > w$.

² We do not derive a bargaining model as it will greatly complicate the analysis while providing little useful insight. Instead, we note that skilled workers and firms will bargain over wages between $2w$ and 1.

³ It is assumed that AI costs are in units of the final good produced.

Importantly, this assumes that skilled workers do not change their wage demands. When AI adoption is possible, the surplus changes from $1 - 2\alpha w$ to $2\alpha((1 - \beta)(1 - c) - w)$ which is a decrease if $1 > 2\alpha(1 - \beta)(1 - c)$. In this case, AI is *not* adopted, but the possibility of AI may reduce skilled worker earnings as the firm's negotiating position has improved; that is, if skilled workers were previously earning a premium above w per hour, there exist levels of that premium that may make adopting AI desirable. AI adoption will not occur as total surplus would fall. Nevertheless, the threat of AI adoption would diminish the bargaining position of skilled workers. If $1 < 2\alpha(1 - \beta)(1 - c)$, surplus increases from AI adoption, and so AI is adopted.

Note the implications of this. Under the stated assumptions, AI automates task 1, which opens up opportunities for workers, in general, to be employed in this sector. Employment in the sector rises to α and total wages in the sector rise to $2\alpha(1 - \beta)(1 - c)$ from somewhere between $2*w$ and 1. This, in turn, reduces inequality by removing the skill premium earned by skilled workers and allowing other workers to earn more than w (as all workers are now in demand and are technically scarce). This defines a Turing Transformation.

What is happening is that AI involves a task that requires specialised skills, and the automation of that task opens up opportunities for more workers. In effect, workers with generic skills are helped when AI is adopted to be able to participate in jobs previously only available to those with specialised skills.

However, suppose that $\alpha = 1$ and the only workers are the skilled workers. Under these assumptions, used by Acemoglu, if there are large economies of scope or AI involves a high unit cost, then wages would fall if AI were adopted. This is the situation that one might characterize as a Turing Trap.

What is going on here? In this model, an AI that is built with the intention of replacing a human in a task—that is, an automation mindset—turns out to be augmenting for the majority of workers because it opens up an opportunity to work on other tasks that would previously have been bundled as a job created for relatively scarce workers. In the model, more workers compete with one another, but the productivity effect is such that total labour income rises. This illustrates starkly the distinction between this perspective and an automation mindset for developing AI involving human replacement that ends up being favourable for labour as a group even without creating new tasks.

Broadly speaking, the implication here is the notion that automation and augmentation involve distinct mindsets with distinct outcomes for workers misses some relevant features.⁴ Different workers have different skills. Many of the developments in AI with the potential for widespread impact are about replicating an aspect of the intelligence of a small number of higher-wage human workers. In doing so, the technology could create opportunities for a much larger number of workers, enabling new opportunities for employment, along with the potential for higher wages and more choice in career. Thus, we emphasize that what an engineer might perceive as automation or augmentation of a particular task has little relation to the economic emphasis on substitution or complementarity for skills across the distribution of human workers.

When considering automation versus augmentation, the heterogeneity of worker skills is fundamental. One worker's automation is another's augmentation. Automation of rare high value skills can mean augmentation for everyone else. Similarly, augmentation that complements the lucky humans with rare high-value skills can mean increased inequality and a hollowing out of the middle class. This requires a

⁴ Another potential criticism of this perspective is that it is not always obvious whether a technology replaces something that is currently a human skill, and thus the line between augmentation and automation is blurry. In this article, we take the distinction as given. If it is blurry whether a technology is intelligence augmentation or human-like artificial intelligence, that will enhance our broader point that this distinction is not useful.

different perspective on how technology changes work than the standard interpretation of the task-based model.

2. Examples of the Turing Transformation through AI Automation

The discussion of automation and augmentation has a new urgency because of advances in artificial intelligence over the past decade. These advances are primarily in a field of artificial intelligence called machine learning, which is best understood as prediction (Agrawal et al., 2018) in the statistical sense. By prediction, we mean the process of filling in missing information. Our examples will focus on advances in prediction technology, though as the model above shows, our broader point about the value of automation versus augmentation is not specific to prediction machines. Technologies that replace the core skills of some workers can enable others to get more out of their skills.

There is already some evidence that AI might be particularly likely to affect the tasks performed by high-wage workers. Webb (2019) finds that the most common verbs in machine learning patents include “recognize,” “predict,” “detect,” “identify,” “determine,” “control,” “generate,” and “classify” (p.38). He also finds that these verbs are common in tasks done by relatively high-wage workers. It is an open question whether automating these tasks will simply reduce the wages of those who are already doing well or whether it will create new opportunities for lower wage workers.

The model in the previous section suggests that automation may reduce inequality, not just by making those with higher wages worse off but by creating Turing Transformation for many more workers. In this section, we provide examples of potential for Turing Transformation from personal transportation, call centers, medicine, language translation, and writing.

Personal Transportation: Since 1865 (Transport for London, n.d.), taxi drivers in London have had to pass a test demonstrating mastery of “The Knowledge” of the map of the complicated road networks in the city. Most drivers studied three to four years before passing the test. Acquiring The Knowledge leads to measurable changes in the brains of drivers (Woollett & Maguire, 2011). This is a skilled occupation, requiring incredible memory skills and the discipline to spend the time studying. Fifteen years ago, no one could compete with the ability of London taxi drivers to navigate the city.

Today, the taxi drivers’ superpower is available for free to anyone with a phone. Digital maps mean that anyone can find the best route, by driving, walking, or transit, in just about any place in the world. The mapping technology substitutes for the driver’s navigation skill. It doesn’t provide something new, but it replicates a human skill more cheaply. As a result, taxi driver wages have fallen (Berger et al., 2018). This is precisely what Markoff, Brynjolfsson, and others warn against.

Automation of the taxi drivers’ competitive advantage, however, has meant opportunity for millions of others. By combining navigation tools with digital taxi dispatch, Uber and Lyft have enabled almost anyone with a car to provide the same services as taxi drivers. Applying the model above, navigation is task 1. It is the task that requires specialized skills. Driving is task 2. It is a widely dispersed skill. Technology automated the core skill for some workers. It did something a handful of skilled humans could already do. In the process, it provided the opportunity for many without those skills to work in the same industry (Fos et al., 2019). In the U.S., there were approximately 200,000 professional taxi and limo drivers in 2018 (Statista Research Department, 2022). Today, more than 10 times that number drive for Uber alone.

Call Centers: There are millions of customer service representatives in the U.S. and around the world (U.S. Bureau of Labor Statistics, 2023). Many of them work in call centers where productivity is carefully

measured in terms of calls per minute and satisfied customers. Like other industries, worker productivity is heterogeneous. The most skilled agents are much more productive than the median, and new workers improve rapidly over the first few months. A recent paper by Brynjolfsson et al. (2023) looks at the deployment of AI in a call center for software support. These calls are relatively complicated, averaging over 30 minutes and involving the troubleshooting of technical problems.

The AI provides real-time suggestions on what the call center worker should say. The worker can choose to follow the AI or ignore it. Based on the model, task 1 involves identifying the relevant response to a customer query. Task 2 involves politely and effectively communicating to the customer what to do. Task 1 is relatively skilled. Task 2 is more widely dispersed. By automating task 1, the AI significantly increases productivity. The most productive workers, however, benefit very little if at all. They may even rationally ignore the AI's recommendation. In contrast, it is the least productive workers and the newer workers that benefit. Their productivity improves substantially. Notably, their relative productivity compared to the most productive workers increases. The AI reduces the gap between the less skilled and more skilled workers. The paper provides suggestive evidence that this is because the less-skilled workers learn what their more skilled peers would do in a given situation.

This technology is automation as defined by Markoff (2016). It involves machines that do what humans do, rather than machines that do something that humans can't do. It is used as decision support and therefore seemingly serves as a complement to all of the human workers, regardless of their skill. In practice, however, this helps the least skilled and provides an example of another Turing Transformation.

Medicine: A large and growing body of research is showing the potential for AI to provide medical diagnoses. Underlying this research is the insight that diagnosis is prediction: It takes information about symptoms and fills in missing information of the cause of those symptoms. Diagnosis, however, is a key human skill in medicine (Goldfarb & Teodoridis, 2022). Much of the training that doctors receive in medical school, and the selection process they go through in order to get into medical school, focuses on the ability to diagnose. Other workers in the medical system may be better at helping patients navigate the stress of their medical issues (Agrawal et al., 2022) or providing the day-to-day care necessary for effective treatment. Perhaps the central skill that sets doctors apart is diagnosis. As modeled above, diagnosis is task 1. The other aspects of medicine together make up task 2. The diagnosis skill is rare relative to the skills required for these other aspects of medicine.

An AI that does diagnosis automates the task requiring that relatively rare skill. It is not augmented intelligence but a replacement for human intelligence. There were 760,000 jobs for physicians and surgeons in the U.S. in 2021, earning a median income of over \$200,000 per year (U.S. Bureau of Labor Statistics, 2022b). Automating the core skill that many of these doctors bring to their work could eliminate much of the value that doctors bring, even leading to stagnating employment and wages. Again, exactly the worry that Brynjolfsson and Markoff warn against when AI replicates human intelligence.

There were also 3 million jobs for registered nurses (U.S. Bureau of Labor Statistics, 2022d) and millions for other medical professionals including pharmacists, nurse and physician assistants, and paramedics (U.S. Bureau of Labor Statistics, 2022c). As we discuss in our book *Power and Prediction: The Disruptive Economics of Artificial Intelligence* (Agrawal et al. 2022), diagnosis is a barrier for these medical professionals to take full advantage of their skills. While AI diagnosis would likely negatively affect many doctors, if these non-doctor medical professionals could perform AI-assisted diagnosis then their career opportunities, and possibly wages, could increase substantially.

Language Translation: Another task currently performed by skilled workers that AI could take over is language translation. Many people speak multiple languages, and in many workplaces this ability confers an advantage. Speaking French and English is an advantage in many Canadian workplaces, particularly for

the hundreds of thousands who work in the civil service (Government of Canada, 2023) or in regulated industries. Similarly, people who speak multiple languages have an advantage in many international business opportunities. Of course, many people work as translators, earning their income directly from their ability to translate between languages.

For written texts, when the goal is simply to communicate with little regard for eloquence, AI is already good enough to replace many human translators. For large scale translations and real-time translation of verbal communication, there are reasons to expect machine translation to be good enough to deploy commercially in the very near future (and perhaps already; Skype, n.d.). These advances are probably bad news for the tens of thousands of language translators in the U.S (U.S. Bureau of Labor Statistics, 2022a).

However, they are likely good news for many others. Brynjolfsson et al. (2019) report that AIs used for translation enhance the capacity of sellers on eBay, increasing exports by 17.5%. AI that automates language translation enables enhanced communication across the world. It likely means more trade, more travel, faster integration into workplaces for recent immigrants, more cross-cultural exchange of ideas, and perhaps even different social networks. Those whose jobs have been constrained by an inability to speak or write in multiple languages would no longer face those constraints. Translation represents the rare task 1 in the model, and selling represents the relatively-common task 2. Automation, in the sense of an AI doing something that many people already do well, creates new opportunities for other people who don't have that particular skill.

Writing: The ability of AI to write goes beyond translating between languages. On November 30, 2022, OpenAI released ChatGPT. This tool quickly gained millions of users because of its ability to produce well-written prose on a wide variety of topics. It can produce high quality five-paragraph essays, leading to worries about the future of take-home exams and the potential for widespread cheating (Herman, 2022). It can write eloquent emails, longer articles, and summarize research and news events. Because summarizing, interpreting, and writing is such an important part of knowledge work, Krugman (2022) worried that ChatGPT means that “robots are coming for the skilled jobs.” Summarizing and writing are clearly tasks that people can do. This is not a case of a machine doing something that is beyond the capability of humans. It is automation, not augmentation. Or in Markoff's language, it is artificial intelligence for duplicating human behavior, not intelligence augmentation that attempts to expand human abilities.

That, however, depends on the human. Many people do not write well. With ChatGPT, they will be able to quickly draft out notes to customers, suppliers, or friends without fear of grammatical mistakes and without the need to stress about how to get the ideas down on paper. This could enable millions of people to benefit from skills other than writing. Once almost anyone has the ability to write clearly, there will be changes in who is capable of which jobs, with many people in the bottom half of the current income distribution receiving new opportunities while some at the top will face enhanced competition.

Consider a story that circulated widely on Twitter about a British landscaper with poor communication skills (Richman, 2022). His mentor used OpenAI's technology to convert the email “Sally I am starts work at yours monday from dave” to a well-written and appropriately punctuated email:

Dear Sally,
I hope this email finds you well. I am writing to let you know that I will be starting work with you on Monday. I am really looking forward to getting started.
If you have any questions or need any help with anything, please don't hesitate to get in touch.
Best wishes,
Dave

For this landscaper, an inability to write represented a barrier to his labor market opportunities. In this example, writing is task 1 and requires a relatively rare skill. For task 2, we have to extend the model to

note that there are many other valued tasks in the economy that are complementary to writing. The automation of writing presents a Turing Transformation.

As with taxi drivers, those that make a living writing will be affected. They may become more efficient, as the AI summarizes articles and writes or revises drafts. They will also face more competition for their work and, like taxi drivers, their wages may fall as their skills are no longer scarce.

The above examples show that automation technologies that do the same things as some human workers can also enhance opportunities for others. Technology that substitutes for one human worker complements others. In these selected examples, workers that had their skills automated tended to be relatively highly paid, and those that received opportunities were likely of lower socio-economic status.

3. Information Technology, Intelligence Augmentation, and Increasing Inequality

In this section, we provide examples of information technologies that are best seen as intelligence augmentation under Markoff's definition—as technologies that do things that are not possible for humans to do. In this sense, they are outside the motivating model, as they do not involve directly automating a specific task done by a human worker, although, as we have emphasized, one person's augmentation could be another's automation. In each case, we show that the augmentation technology complemented human labor at the top of the income distribution and reduced employment opportunities and wages for those in the middle.

Computerization: As Brynjolfsson & Hitt (2000) put it, “computers are symbol processors” (p. 23). They can store, retrieve, organize, transmit, and transform information in ways that are different from how humans process information. Markoff (2016, p. 165) notes that modern personal computers have their root in Douglas Engelbart's augmentation tradition. Unlike AI, for which we argued above may decrease inequality, computerization increased inequality (Autor et al., 2003) and led to polarization of the U.S. wage distribution (Autor et al., 2008), expanding high- and low-wage work at the expense of middle-wage jobs (Michaels et al., 2014). This is because, while some tasks done by computers could be done by humans, much of the changes are a result of complementarity between the skills of the most educated workers and the identification of new ways to use the machines. In other words, rather than directly replacing a task done by middle income workers as AI does, computers complemented the skills of those already near the top of the income distribution, thereby increasing their productivity for tasks that were already done by humans. Again, quoting Brynjolfsson and Hitt (2000), “As computers become cheaper and more powerful, the business value of computers is limited less by computational capability and more by the ability of managers to invent new processes, procedures, and organizational structures that leverage this capability” (p. 24). Barth et al. (2023) match census data on business software investment with employee wages to show that within and across firms software investment increases the earnings of high-wage workers more than that of low-wage workers. Computers displaced the workers performing routine technical tasks in bookkeeping, clerical work, and manufacturing, while complementing educated workers who excel in problem-solving, creativity, and persuasion (Autor, 2014).

Digital Communication: The internet represents another technology that does something distinct from what humans can do. For the most part, as Markoff notes (2016, p. 166), the internet does not replace specific tasks in human workflows. It does not fit naturally into the task-based framework described in the model above. It allows computers to communicate with each other, sending information between millions of devices. This information is a complement to the human skills of interpreting and acting on information.

People (Akerman et al., 2015) and places (Forman et al., 2012) at the top of the income distribution benefited from the technology. Those with less education benefited less. To the extent that there are differences between augmentation and automation technologies, the internet is more of an augmentation technology. As such, it complemented the skills of those who were already at the top of the income distribution.

The above discussion warrants an important caveat: Many have called computerization and digital communication “automation.” Formally, it is difficult to classify technologies as automating or augmenting, and we do not want to take a strong stand on which technologies belong in which category. That’s an aspect of our underlying point. One person’s augmentation is another’s automation. What matters is the distribution of workers whose skills are complemented.

4. AI, automation, and the Task-based Model

The first 50 years of computing introduced many technologies that appear to be intelligence augmenting, creating new capabilities and new products and services. The last 10 years have seen a rise in artificial intelligence applications, whose inventors directly aspire to automate tasks currently performed by humans. On the surface, technologies labeled as augmentation appear to complement human workers, while automation technologies appear to substitute for human workers. Therefore, many scholars have called for engineers, scientists, and policymakers to focus on augmentation technologies over automation (Brynjolfsson, 2022; Markoff, 2016; Acemoglu, 2021b). An important aspect of this argument is the idea that complements to human labor will reduce income inequality while substitutes for human labor will increase it.

We argue that this dichotomy is misleading. A key aspect of understanding the impact of intelligence technology on inequality and the well-being of most workers is the heterogeneity of the skills of workers. A technology that directly substitutes for rare and highly-valued skills could create enormous opportunities for most workers.

Through a formal model and examples, we have demonstrated that our argument is plausible. It remains an open question whether this model and these examples will prove dominant as AI technologies diffuse. It is also an open question whether the owners of AI technology will have sufficient market power to capture the value, leaving even the workers who are most likely to benefit no better off. What is clear, however, is that one person’s substitute is another’s complement, and so heterogeneous impacts are essential to consider. Many of the technologies described as augmenting are about tasks that humans don’t currently do. They nevertheless enable the replacement of entire jobs by redesigning workflows to take advantage of these new capabilities. In the process, technologies that Markoff defines as augmenting, such as computing and the internet, led to increased inequality and a hollowing out of the middle class. The people best positioned to take advantage were well-educated and skilled workers.

With technological change, we argue that the winners and losers are not determined by whether the technology seems to replace or augment human tasks. Instead, the winners and losers are determined by whether the augmentation affects lower-wage workers and automation affects those already doing well. Perhaps the best targets for computer scientists and engineers looking to build new systems is not to find intelligences that humans lack. Instead, it is to identify the skills that generate outsized income and build machines that allow many more people to benefit from those skills. As noted above, this may be what is already happening with AI that recognizes, predicts, determines, controls, writes, and codes.

Ultimately, whether the engineer or scientist is building a tool that replaces a human process or that creates a new capability might be irrelevant to whether the technology enhances productivity in a way that reduces inequality and increases opportunity for those who are not already at the top of the income distribution. What matters is whether the technology enhances the productivity of those who are already doing well or if it opens up a Turing Transformation for everyone else.

References

- Acemoglu, D. (2021a). Harms of AI (NBER Working Paper No. 29247). *National Bureau of Economic Research*. <https://doi.org/10.3386/w29247>
- Acemoglu, D. (2021b, May 20). *AI's Future Doesn't Have to Be Dystopian*. Boston Review. <https://www.bostonreview.net/forum/ais-future-doesnt-have-to-be-dystopian/>
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.
- Agrawal, A., Gans, J., & Goldfarb, A. (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press.
- Agrawal, A., Gans, J., & Goldfarb, A. (2023). Do we want less automation? *Science* 381(6654), 155-158
- Akerman, A., Gaarder, I., & Mogstad, M. (2015). The Skill Complementarity of Broadband Internet. *The Quarterly Journal of Economics*, 130(4), 1781–1824. <https://doi.org/10.1093/qje/qjv028>
- Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the “other 99 percent.” *Science*, 344(6186), 843–851. <https://doi.org/10.1126/science.1251868>
- Autor, D., Levy, F., & Murnane, R. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Autor, D. H., Katz, L. F., & Kearney, M. S. (2008). Trends in U.S. Wage Inequality: Revising the Revisionists. *The Review of Economics and Statistics*, 90(2), 300–323. <https://doi.org/10.1162/rest.90.2.300>
- Barth, E., Davis, J. C., Freeman, R. B., & McElheran, K. (2023). Twisting the Demand Curve: Digitalization and the Older Workforce. *Journal of Econometrics*, 233(2), 443–467. <https://doi.org/10.1016/j.jeconom.2021.12.003>
- Berger, T., Chen, C., & Frey, C. B. (2018). Drivers of disruption? Estimating the Uber effect. *European Economic Review*, 110, 197–210. <https://doi.org/10.1016/j.euroecorev.2018.05.006>
- Brynjolfsson, E., & Hitt, L. M. (2000). Beyond Computation: Information Technology, Organizational Transformation and Business Performance. *Journal of Economic Perspectives*, 14(4), 23–48. <https://doi.org/10.1257/jep.14.4.23>
- Brynjolfsson, E., Hui, X., & Liu, M. (2019). Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform. *Management Science*, 65(12), 5449–5460. <https://doi.org/10.1287/mnsc.2019.3388>
- Brynjolfsson, E. (2022). The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence. *Daedalus*, 151(2), 272–287. https://doi.org/10.1162/daed_a_01915

- Brynjolfsson, E., Li, D., & Raymond, L. (2023). Generative AI at Work (NBER Working Paper No. 31161). *National Bureau of Economic Research*. <https://doi.org/10.3386/w31161>
- Creative Destruction Lab. (2016, November 24). *Geoff Hinton: On Radiology* [Video]. YouTube. https://www.youtube.com/watch?v=2HMpRXstSvQ&ab_channel=CreativeDestructionLab
- Feigenbaum, J., & Gross, D. (2021). Organizational and Economic Obstacles to Automation: A Cautionary Tale from AT&T in the Twentieth Century (NBER Working Paper No. 29580). *National Bureau of Economic Research*. <https://doi.org/10.3386/w29580>
- Forman, C., Goldfarb, A., & Greenstein, S. (2012). The Internet and Local Wages: A Puzzle. *American Economic Review*, *102*(1), 556–575. <https://doi.org/10.1257/aer.102.1.556>
- Fos, V., Hamdi, N., Kalda, A., & Nickerson, J. (2019). Gig-Labor: Trading Safety Nets for Steering Wheels. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3414041>
- Goldfarb, A., & Teodoridis, F. (2022, March 9). *Why is AI adoption in health care lagging?*. Brookings. <https://www.brookings.edu/articles/why-is-ai-adoption-in-health-care-lagging/>
- Government of Canada. (2023, June 26). *Population of the Federal Public Service*. Government of Canada. <https://www.canada.ca/en/treasury-board-secretariat/services/innovation/human-resources-statistics/population-federal-public-service.html>
- Herman, D. (2022, December 9). *The End of High-School English*. The Atlantic. <https://www.theatlantic.com/technology/archive/2022/12/openai-chatgpt-writing-high-school-english-essay/672412/>
- Kansas Historical Society. (2011, June). *Almon Strowger*. Kansapedia. <https://web.archive.org/web/20170114094940/http://kshs.org/kansapedia/almon-strowger/16911>
- Krugman, P. (2022, December 6). *Does ChatGPT Mean Robots Are Coming for the Skilled Jobs?*. The New York Times. <https://www.nytimes.com/2022/12/06/opinion/chatgpt-ai-skilled-jobs-automation.html>
- Markoff, J. (2015, December 10). *A Learning Advance in Artificial Intelligence Rivals Human Abilities*. The New York Times. <https://www.nytimes.com/2015/12/11/science/an-advance-in-artificial-intelligence-rivals-human-vision-abilities.html>
- Markoff, J. (2016). *Machines of Loving Grace: The Quest for Common Ground Between Humans and Robots*. Ecco.
- Michael Lawrence. (2006, June 1). *Steve Jobs, "Computers are like a bicycle for our minds."* - Michael Lawrence Films [Video]. YouTube. https://www.youtube.com/watch?v=ob_GX50Za6c&ab_channel=MichaelLawrence
- Michaels, G., Natraj, A., & Van Reenen, J. (2014). Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *The Review of Economics and Statistics*, *96*(1), 60–77. https://doi.org/10.1162/rest_a_00366

- OpenAI. (2023). GPT-4 Technical Report. *arXiv*. <https://doi.org/10.48550/arXiv.2303.08774>
- Oppy, G., & Dowe, D. (2021, October 4). *The Turing Test*. Stanford Encyclopedia of Philosophy. <https://plato.stanford.edu/entries/turing-test/>
- Richman, D. [@DannyRichman]. (2022, December 1). *I mentor a young lad with poor literacy skills who is starting a landscaping business. He struggles to communicate with* [Image attached]. Twitter. <https://twitter.com/DannyRichman/status/1598254671591723008>
- Skype. (n.d.). *Skype Translator*. Skype. <https://www.skype.com/en/features/skype-translator/>
- Statista Research Department. (2022, July 6). Number of taxi drivers and chauffeurs in the United States from 2013 to 2018 [Chart]. Statista. <https://www.statista.com/statistics/943496/number-of-taxi-drivers-united-states/>
- Transport for London. (n.d.). *Learn the Knowledge of London*. Transport for London. <https://tfl.gov.uk/info-for/taxis-and-private-hire/licensing/learn-the-knowledge-of-london>
- U.S. Bureau of Labor Statistics. (2022a, September 8). *Interpreters and Translators: Summary*. Occupational Outlook Handbook. <https://www.bls.gov/ooh/media-and-communication/interpreters-and-translators.htm>
- U.S. Bureau of Labor Statistics. (2022b, September 8). *Physicians and Surgeons: Summary*. Occupational Outlook Handbook. <https://www.bls.gov/ooh/healthcare/physicians-and-surgeons.htm>
- U.S. Bureau of Labor Statistics. (2022c, September 8). *Registered Nurses: Similar Occupations*. Occupational Outlook Handbook. <https://www.bls.gov/ooh/healthcare/registered-nurses.htm>
- U.S. Bureau of Labor Statistics. (2022d, September 8). *Registered Nurses: Summary*. Occupational Outlook Handbook. <https://www.bls.gov/ooh/healthcare/registered-nurses.htm>
- U.S. Bureau of Labor Statistics. (2023, April 25). *Occupational Employment and Wage Statistics*. U.S. Bureau of Labor Statistics. <https://www.bls.gov/oes/current/oes434051.htm>
- Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3482150>
- Woollett, K., & Maguire, E. A. (2011). Acquiring “The Knowledge” of London’s Layout Drives Structural Brain Changes. *Current Biology*, *21*(24), 2109–2114. <https://doi.org/10.1016/j.cub.2011.11.018>