From Artificial Light to Artificial Intelligence: Technological Change and Worker Perspectives

Kristen E. Broady*

October 2026

Abstract

Recent advances in generative artificial intelligence have generated considerable optimism about its potential to increase productivity, improve services, and change how work is performed, while also raising concerns about privacy, bias, job displacement, and the broader consequences of advanced AI. This paper asks: How should we understand the current AI transition in the context of earlier technological transformations? To address that question, the paper combines historical perspective and a synthesis of the automation and AI literature with a new analysis of nationally representative data from the 2025 Survey of Household Economics and Decisionmaking (SHED) and qualitative evidence from the 2026 Federal Reserve System Worker Perspectives Project. These sources distinguish technological capability and occupational exposure from actual workplace adoption and worker experience. Descriptive comparisons and multivariate analysis show substantial variation in generative AI use, with economic and educational characteristics being particularly important. Among users, perceptions are generally favorable regarding AI's usefulness for work, although concerns about employment remain. The findings suggest that generative AI is consequential, but its diffusion and effects remain uneven across workers and workplaces.

JEL Classification Codes: J11, J22, J24, O33

Keywords: Labor history, Artificial intelligence, Worker perspectives

1. Introduction

This introduction proceeds in two parts. The first considers contemporary expectations surrounding generative artificial intelligence, including its potential benefits, risks, and implications for workers and organizations. The second places these developments within a broader historical context, examining earlier technological transformations that altered production, work, skills, and everyday life.

1.1 Artificial Intelligence: Expectations and Historical Perspective

* Federal Reserve Bank of Chicago, 230 South LaSalle Street, Chicago, IL 60604-1413, USA. Email: kristen.broady@chi.frb.org. The views expressed are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

Generative artificial intelligence is now the focus of considerable research, conjecture, and forecasting. It holds the potential to raise productivity, stimulate innovation and alter the nature of work, but its adoption has so far been uneven across the global economy. Cerutti, et al. (2025) argue that differences in AI adoption reflect persistent structural disparities among advanced economies, emerging markets, and low-income countries. These differences suggest that the benefits and disruptions associated with AI may not be experienced uniformly across countries, industries, or workers.

At the same time its rapid development and expanding capabilities have raised important questions and concerns about the future of work, the skills workers will need, and the ways businesses and industries will operate. In urban planning, the use of AI raises concerns related to bias, transparency, accountability, privacy, and misinformation (Sanchez, et al., 2025). While AI is revolutionizing the educational landscape by increasing opportunities for personalized learning solutions, automated grading, and learning outcome assessment, there are concerns about privacy, learning loss, and equitable access (Alarcon, et al., 2025). In the healthcare sector, AI may improve the speed and quality of services provided to patients, buts its use also presents ethical and legal challenges related to patient privacy, biased results, and patient safety (Bataineh et al., 2024; Kooli & Muftah, 2022).

Beyond these sector-specific concerns, some researchers and organizations have raised much more far-reaching questions about the potential risks associated with advanced artificial intelligence. A letter released by the Future of Life Initiative, a nonprofit organization, addresses some of the potential harms of generative AI and asks, “Should we develop nonhuman minds that might eventually outnumber, outsmart, obsolete and replace us? Should we risk loss of control of our civilization?” (Clarke, 2023). A 2023 open letter from the Center for AI Safety sought to share concerns about severe risks associated with the rapid deployment of AI and stated that “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.” (Center for AI Safety, 2023).

The range of claims and predictions surrounding artificial intelligence is striking. AI is expected to improve productivity, expand the capabilities of workers and organizations, and change the delivery of services, while also raising concerns about bias, privacy, displacement, and even existential risk. These possibilities deserve serious consideration, but they should also be viewed in context. Technological change is not a new phenomenon. Historically, the transition from foraging for food to organized agriculture, industrialization, electrification, mechanization, and the development of mass transportation changed the structure of economic and social life in ways that were profound and enduring. My goal here is not to diminish the capabilities or potential impact of generative artificial intelligence. Rather, it is to place its development and implementation into historical context. Against that historical backdrop, the current AI transition is best understood not as an unprecedented break from the past, but as the latest stage in a much longer process of technological change.

Artificial light provides one particularly striking example of how technological change can alter both economic productivity and daily life. For most of human history, darkness imposed a hard constraint on what people could do after sunset. The earliest artificial light source was fire, followed by increasingly deliberate efforts to control and extend illumination. By around 500 BC, bamboo pipes were reportedly used in ancient China to transport natural gas for street lighting, while the Romans used oil lanterns to illuminate homes and public spaces (Brilliance LED, 2025). In later centuries, candles made from wax or tallow and oil lamps became common sources of household illumination, although they were relatively dim, smoky, and costly by modern standards (Brilliance LED, 2025).

The production of light itself was costly and labor intensive, relying on animal fat, oil, candles, and other combustible materials. In an episode of NPR's Planet Money Summer School on productivity, Stacey Vanek-Smith, Luigi Zingales and I discussed estimates by William Nordhaus showing just how dramatically the cost of light declined over time (Vanek Smith, Kestenbaum, and Goldstein, 2022; Nordhaus, 1996). In ancient Babylon, a full day's wages could purchase only about ten minutes of light equivalent to that produced by a modern bulb.

As explained by Ostrizhnaya (2022), the nineteenth century marked a major transition in the history of artificial illumination. The invention of the incandescent lightbulb came about through the work of several inventors. Experiments in 1840 by English scientist De la Rue involved running an electric current through a platinum wire in a glass cylinder. By the 1850s and 1870s, inventors including Heinrich Göbel, Alexander Lodygin, Joseph Swan, and others were extending bulb life and improving practicality. Thomas Edison's contribution came at the end of a long period of experimentation. Edison and his laboratory tested thousands of materials in search of a filament that could provide durable, commercially practical light. An early carbonized cotton-thread filament produced more than 13 continuous hours of illumination, and Edison filed a light-bulb patent in January 1880 (Hooper, 2022).

The economic consequences of these technological improvements were dramatic. By the nineteenth century, a day's labor could purchase several hours of kerosene light; by the late twentieth century, the same amount of labor could purchase roughly 20,000 hours of electric light (Vanek Smith, Kestenbaum, and Goldstein, 2022; Nordhaus, 1996).

The significance of that change extends well beyond the price of illumination. For much of human history, the setting of the sun sharply limited the activities that could be carried out after dark. Artificial light extended the usable day, making it possible to continue household production, craft work, food preparation, writing, study, religious observance, and other activities after sunset. Over time, improvements in lighting also increased the usefulness of homes, workshops, public spaces, and places of commerce beyond daylight hours. In that sense, artificial light did more than make illumination cheaper; it relaxed a basic constraint on how people could organize their time and activity. Much later, electrification magnified these effects and provided infrastructure for innovations far beyond lighting itself. Electric power transformed factories,

homes, transportation, communication, and eventually computing and digital technologies. As Zingales and I noted in our Planet Money discussion, technological progress often creates possibilities that could not have been anticipated before the underlying infrastructure existed.

The creation of artificial light is but one example of how technological innovation can fundamentally alter the conditions under which people live and work. The history of technological change also includes innovations that fundamentally altered basic constraints on human activity including: agriculture, which changed how societies obtained food and organized settlement; mechanization which changed how physical work was performed; and modern transportation, which increased the distance and speed people and goods could travel. The confluence of these transformations provides a broader historical context within which the development of artificial intelligence can be considered. The discussion that follows places AI within that longer history of technological change and considers what earlier technological revolutions can tell us about changes in work, skills, production, and the adoption of new technologies.

1.2 Industrial Revolutions – Technology in Context

Writer and futurologist, Alvin Toffler provided a useful framework within which to consider technological change and its effects on society. In his book, Previews and Premises, he discussed what he called a "restructuring of the entire technoeconomic base of society," as he described economies shifting from the First, Second, and Third Waves (Toffler, 1984). Toffler's waves do not correspond one-for-one with the numbered industrial revolutions commonly used in the contemporary literature. Rather, his framework describes broader socioeconomic transformations. His Second Wave encompasses the rise and maturation of industrial society, while his Third Wave captures the transition toward an information- and knowledge-based economy. Viewed in this way, Toffler provides a broader social framework within which the successive industrial revolutions can be considered.

Toffler characterized the First Wave, the Agricultural Revolution, as the period when nomadic and hunter-gatherer societies increasingly settled in permanent locations and agricultural production became central to economic life (Toffler, 1984). This transition altered where people lived, how labor was organized, the size and structure of populations, and ultimately the range of economic activities communities could support. Weisdorf (2005) reviews a broad economic literature attempting to explain why societies made this transition, including theories emphasizing environmental change, population pressure, technological knowledge, and changes in the relative productivity of farming and foraging. One important implication of this literature is that agriculture changed the allocation of labor. Sedentary agricultural communities could support larger populations and, as agricultural production expanded, could also support people who did not themselves produce food. Weisdorf discusses the emergence of non-food-producing specialists as one consequence of the transition: food producers could support individuals engaged in producing other goods and activities, permitting a greater degree of occupational

specialization. Agriculture therefore changed not only how food was produced, but also what other kinds of work a society could sustain.

That transition was neither instantaneous nor uniformly beneficial. Weisdorf also describes models in which farming could require more labor time than foraging and in which population growth following the adoption of agriculture could eventually reduce standards of living. The importance of the Neolithic transition, therefore, does not rest on the proposition that farming immediately made everyone better off. Rather, it fundamentally changed the relationship among land, labor, population, production, and specialization and created conditions for increasingly complex economic organization.

Agricultural production itself continued to change substantially over the centuries that followed. Economic historians disagree over exactly when these later transformations should be dated and whether they should be treated as one or several agricultural revolutions. Thompson (1968), for example, argues that it is more useful to distinguish different kinds of technical and economic change rather than regard agricultural development as a single unitary revolution. He describes a progression from subsistence production toward market-oriented farming, followed by changes in crop rotation, livestock improvement, enclosure, the intensity of cultivation, and eventually more capital-intensive and mechanized production.

From the farmer's perspective, Thompson characterizes part of this transformation as a managerial revolution. New systems of husbandry required farmers to manage land, labor, livestock, and other resources more systematically in order to obtain higher yields and greater income per acre. Later changes also increased dependence on purchased inputs and outside markets. Seeds, fertilizers, feeding stuffs, drainage systems, and purpose-built farm structures became increasingly important, while farmers devoted greater amounts of capital to production. Thompson argues that this process made farming increasingly resemble manufacturing: rather than operating largely as a closed system, farms increasingly purchased inputs that were processed in order to produce goods for sale.

Mechanization brought another major change in the relationship between labor and production. In Thompson's later stage, human labor was increasingly replaced by mowers, reapers, binders, threshers, and other machinery, while tractors replaced much of the work previously performed by horses. At the same time, increased fertilizer use and other technological changes allowed greater specialization in agricultural production.

Kerridge (1969), by contrast, cautions against imposing too neat a chronology on these developments. He disputes the conventional argument that the British Agricultural Revolution occurred primarily between approximately 1750 and 1850 and argues that many important changes in agricultural practice began considerably earlier, including innovations in crop rotation, livestock production, and the cultivation of new crops. For Kerridge, much of what later historians grouped into an eighteenth-century revolution had roots stretching back into the sixteenth and seventeenth centuries.

The disagreement between Thompson and Kerridge is useful because it illustrates a broader point about technological revolutions: their boundaries are rarely as precise as the labels later assigned to them. New technologies and production practices can emerge gradually, coexist with older methods for long periods, and diffuse unevenly across regions and producers. Thompson himself notes that different stages of agricultural change were conceptually distinct but did not occur in perfectly separate chronological periods.

Regardless of the precise periodization, agricultural innovation fundamentally changed what could be produced and how production was organized. Greater agricultural productivity supported larger populations, production became increasingly oriented toward markets rather than household subsistence, and technological change altered both the amount and type of labor required to produce food. Over time, specialization, commercialization, capital investment, and mechanization transformed agriculture from an activity that occupied much of the population into one capable of producing substantially more food with a smaller share of workers.

The experience of agricultural transformation is relevant to contemporary discussions of artificial intelligence for reasons beyond the eventual substitution of machinery for labor. New agricultural technologies changed who performed particular tasks, the skills required to perform them, and the amount of output that could be produced with available labor. They also spread unevenly. Some farmers adopted new production methods, machinery, and purchased inputs earlier than others, while older methods continued alongside newer ones for long periods. Current adoption of generative AI may follow a similarly uneven path. Even as AI tools become more widely available, workers do not necessarily encounter, use, or benefit from them at the same rate. Differences in access to these technologies, employer encouragement, and workers' ability or willingness to incorporate them into their jobs may therefore be as important as the technology's technical capabilities.

Toffler's Second Wave, the Industrial Revolution, brought another major transformation in the organization of production. New machinery and factory production increasingly replaced the agricultural way of life across parts of the world (Toffler, 1984). Production became more mechanized and concentrated in factories, changing both the scale at which goods could be produced and the kinds of work people performed. Rather than simply producing more of the same goods, industrialization altered where production occurred, how workers were organized, which skills were demanded, and the relationship between workers and machines.

The Industrial Revolution provides another useful comparison. Mechanization did not affect all workers or occupations in the same way, nor did technological change simply eliminate labor. New production methods changed the tasks workers performed, increased the importance of some skills, reduced demand for others, and created new occupations and forms of work. As was the case with mechanization, AI is not being integrated into all industry sectors equally. The impact of AI depends on how applicable it is to particular job tasks, whether workers actually use it, whether employers support or encourage its use, and whether the technology is experienced

primarily as a complement to the worker's skills or as a potential substitute for them (Broady, et al., 2026).

In *The Third Wave*, Toffler described the movement toward what he viewed as a technological and information-based economy. He predicted that education would "become interspersed and interwoven with work" and believed that accountability, adaptability to change, sensitivity to coworkers, and independent thinking would become increasingly important in the workplace (Toffler, 1980). He also anticipated diversified renewable energy sources, advances in electronics and genetics, and production processes that would increasingly move beyond traditional assembly-line systems. His broader argument was that technological change would not simply introduce new machines, but would continue to alter the organization of production and the skills expected of workers.

Toffler's emphasis on education and adaptability is particularly relevant to the current development of generative AI. Unlike earlier forms of mechanization that substituted machines for physical tasks, many generative AI applications require workers to interact directly with the technology: asking questions, evaluating responses, revising output, and determining when and how the technology should be used (Tomlinson et al., 2025). The ability to benefit from these tools may therefore depend in part on workers' existing skills, education, experience, and opportunities to learn new ways of working. At the same time, workers with different levels of education may encounter AI in very different types of jobs and may have different expectations about whether the technology will improve their work or threaten their employment.

More recently, Włodarczyk (2022) characterized Industry 4.0 as a multifaceted transformation of manufacturing that extends beyond the adoption of new technologies to include changes in production processes, organizational structures, and business and operating models. Central to the concept is the integration of digital and physical systems, with technologies such as artificial intelligence, robotics, cloud computing, cyber-physical systems, and the Internet of Things enabling greater automation and connectivity across machines, systems, workers, and production processes. The term *Industry 4.0* was first used publicly in 2011 at the Hannover Messe trade fair, where a German working group introduced *Das Zukunftsprojekt Industrie 4.0*, a national initiative focused on the digitization of manufacturing. Włodarczyk places this development within a broader evolutionary process in which interconnected technologies increasingly combine the digital, physical, and biological worlds and transform both production and business organization.

Tallat et al. (2023) describe the movement from machine and information-system interoperability under Industry 4.0 toward a closer relationship between humans and robots under Industry 5.0. This human-machine relationship is one of the most frequently cited characteristics of Industry 5.0. Nahavandi (2019) distinguishes the Fourth and Fifth Industrial Revolutions by describing Industry 4.0 as primarily concerned with automation and Industry 5.0 as a synergy between humans and autonomous machines.

The concept of a Fifth Industrial Revolution, or Industry 5.0, is still developing, and the literature does not offer a single universally accepted definition. In their review of the Industry 5.0 literature, Barata and Kayser (2023) note disagreement over whether Industry 5.0 represents a distinct industrial paradigm or an advancement of Industry 4.0. They describe it broadly as a humanized approach to technological transformation that attempts to balance the needs of workers and society with technological advancements, economic competitiveness, and the sustainable use of resources. Their review identifies three concepts that have increasingly come to characterize Industry 5.0: human-centricity, sustainability, and resilience.

These developments make the worker's experience of technological change especially important. Much of the research on automation and artificial intelligence has focused on what technologies are capable of doing, which occupations are exposed to them, and how employment may change as a result (Autor, 2015; Autor & Salomons, 2018; Broady, 2017). Those questions remain important. But the rapid diffusion of generative AI also raises a more immediate set of questions about what is happening inside workplaces now. Are workers actually using these technologies? Are employers encouraging them to do so? Do workers believe AI helps them complete their jobs more efficiently or perform tasks they could not previously perform? And do they see AI as improving their career opportunities or threatening their employment?

This paper focuses on those differences. The remainder of the paper proceeds as follows. Section 2 reviews the evolving literature on automation, artificial intelligence, and occupational exposure, with particular attention to the distinction between technological capability, task-level exposure, and realized labor-market outcomes. Section 3 describes the 2025 Survey of Household Economics and Decisionmaking (SHED), the Federal Reserve System's Worker Perspectives Project, and the empirical strategy used to examine workplace generative AI use. Section 4 presents the quantitative results on AI adoption, frequency of use, worker perceptions, and differences across demographic and economic characteristics, while drawing on the Worker Perspectives findings to provide additional context for how workers experience AI in employment and job-search settings. Section 5 discusses these findings in light of the historical and contemporary literature on technological change, including what the evidence suggests about the distinction between AI's technical capabilities and its actual diffusion across workers and workplaces. Section 6 concludes.

2. Automation, Artificial Intelligence, and the Changing Nature of Occupational Exposure

Concerns about the effects of automation on employment predate the recent expansion of generative artificial intelligence. Much of the modern literature has focused on the extent to which new technologies can substitute for workers, the occupations most exposed to technological change, and the degree to which automation alters tasks rather than eliminating entire jobs. Over time, however, both the technologies themselves and the methods used to assess their labor-market implications have changed substantially. Earlier studies largely evaluated whether occupations were technically susceptible to computerization or automation. More recent

work has increasingly focused on the specific tasks workers perform, the extent to which artificial intelligence can assist with or complete those tasks, and the ways technological exposure may lead to augmentation, restructuring, or substitution rather than straightforward occupational disappearance.

Frey and Osborne (2017) provided one of the most influential modern assessments of occupational automation risk. Using a machine-learning model, they estimated the probability that 702 U.S. occupations could be computerized and concluded that approximately 47 percent of U.S. employment was in occupations facing a high probability of automation. Their estimates suggested a clear socioeconomic pattern: lower-wage occupations requiring less education were generally more susceptible to automation, while higher-wage occupations requiring greater levels of education were comparatively less exposed. Importantly, however, their framework was based not simply on whether machines could perform routine work. Frey and Osborne identified three major engineering bottlenecks that limited the range of tasks susceptible to automation: perception and manipulation, creative intelligence, and social intelligence. These constraints helped explain why many occupations involving complex physical environments, creativity, interpersonal interaction, persuasion, or judgment appeared relatively insulated from automation at the time.

The technological environment has changed considerably since that analysis. Advances in robotics, machine learning, multimodal systems, foundation models, and generative AI have narrowed several of the technical limitations Frey and Osborne identified. With respect to perception and manipulation, robots historically struggled to navigate unstructured environments, handle irregularly shaped objects, and recover from mistakes. More recent systems combine visual, tactile, acoustic, and proprioceptive information to improve environmental perception and physical interaction. Advances in sensor fusion, reinforcement learning, three-dimensional machine vision, and foundation models have allowed robots to operate more effectively in dynamic or cluttered environments and to adapt their behavior when conditions change (Patel and Reynolds, 2025; Naderi et al., 2024; Novo et al., 2024; Singh et al., 2024; Lahemer and Rad, 2024; Sorall, 2025). Other developments have improved robotic failure recovery, locomotion, and reconfiguration, further reducing the distinction between fixed, task-specific automation and more adaptable machines (Berdugo, 2025; Knight, 2025; Tu et al., 2025). These developments do not imply that human perception and dexterity have been fully replicated, but they substantially alter the technical boundary around work once considered difficult to automate.

Changes in generative AI have similarly challenged earlier assumptions about creative intelligence. Frey and Osborne (2017) treated creativity as difficult to automate in part because creative processes were difficult to specify through conventional computer programming. Generative AI systems operate differently. Large language and foundation models can generate text, images, music, computer code, and other forms of content in response to natural-language prompts. Research examining AI-generated humor, for example, has found that people

sometimes rate jokes produced by ChatGPT as comparable to or funnier than those produced by humans (Gorenz and Schwarz, 2024). Other research has examined AI-generated music and found that listeners distinguish between human and AI-generated compositions in complex ways rather than simply rejecting machine-produced creative output (Fišer, Martin-Pascual, and Andreu-Sánchez, 2025). These developments suggest that creative tasks cannot be treated as categorically beyond the reach of automation simply because they require the production of novel content.

At the same time, the ability to generate creative output does not eliminate the role of human knowledge, judgment, or expertise. Generative AI systems require users to formulate prompts, evaluate responses, identify errors, revise output, and determine whether the resulting material is appropriate for a particular purpose. The usefulness of these technologies therefore depends partly on workers' existing knowledge and their ability to interact effectively with the systems. Moreover, problems including hallucinated information, copyright and licensing concerns, and resistance from industries whose workers seek to protect creative ownership continue to constrain the use of generative AI for some tasks (Gorenz and Schwarz, 2024). Thus, advances in creative capability may change the division of labor between workers and machines without necessarily eliminating the need for human participation.

A similar evolution has occurred with social intelligence. Frey and Osborne (2017) argued that tasks involving persuasion, negotiation, emotional interpretation, and other forms of interpersonal interaction would be difficult to automate because machines were unable to recognize and respond to human emotions in real time. More recent evidence suggests that this boundary has also become less clear. Large language models can identify broad emotional patterns in written material and produce responses that reflect elements of emotional reasoning. Finet, Kristoforidis, and Laznika (2025), for example, found that AI tools could provide relatively consistent assessments of the general emotional tone of written material, although their performance was weaker for more nuanced emotional states. Fatahi et al. (2025) similarly found systematic differences between human and ChatGPT responses in medicine and finance: human responses more frequently incorporated hopeful or reassuring language, while AI responses tended to be more neutral, cautious, or probabilistic. These findings suggest that AI systems can perform some tasks associated with emotional interpretation without necessarily reproducing human emotional experience itself.

The narrowing of these engineering bottlenecks has important implications for how occupational exposure to technology is measured. Frey and Osborne's framework asked whether an occupation was technically susceptible to computerization given the technological capabilities available or foreseeable at the time. As those capabilities expand, occupations once viewed as comparatively resistant may contain tasks that are now readily assisted by artificial intelligence. This does not mean that the original estimates were incorrect; rather, it illustrates the importance of distinguishing a technology-specific estimate of feasibility from a permanent classification of occupational risk. The technological frontier itself moves.

Tomlinson et al. (2025) provide a more recent approach that reflects this shift. Rather than estimating automation probability primarily from occupational characteristics, they examined approximately 200,000 anonymized interactions between workers and Microsoft Copilot during 2024 and linked those interactions to work activities defined in O*NET. Their resulting AI applicability measure reflects both whether AI successfully assists with particular work activities and how broadly those activities occur within an occupation. Occupations with higher applicability scores include work involving writing, editing, programming, sales, customer service, and clerical activities. Occupations with lower scores tend to involve physical interaction with people, operation or monitoring of machinery, and other manual activities.

This distinction represents an important evolution in the literature. Frey and Osborne's automation probabilities and Tomlinson et al.'s AI applicability scores measure related but different concepts. The former estimates the technical susceptibility of occupations to computerization; the latter draws on observed worker-AI interactions to identify activities for which generative AI appears applicable. Neither measure should be interpreted as a direct forecast that a particular occupation will disappear. Rather, both provide information about the types of work most likely to be affected by technological change. As noted in earlier work, these measures are particularly useful for understanding which occupations and workers may experience changes in job tasks, even when employment itself does not decline.

This task-oriented interpretation is consistent with a broader literature in labor economics emphasizing that technological change does not operate solely through the elimination of occupations. Autor, Katz, and Krueger (1998) connected the spread of computer technology to skill-biased technical change and widening educational wage differentials. Autor (2015) subsequently emphasized the importance of distinguishing jobs from tasks: technologies may substitute for particular activities performed within an occupation while complementing workers in others. The employment consequences of automation therefore depend not merely on whether technology can perform existing tasks but also on whether new forms of work emerge and whether workers are redeployed toward activities in which human labor retains comparative advantage.

Acemoglu and Restrepo (2019) develop this distinction through the concepts of displacement and reinstatement. Automation can reduce labor demand by allowing capital or technology to perform tasks previously undertaken by workers, while the creation of new tasks can reinstate labor into the production process. Net employment outcomes depend on the relative strength of these forces. Autor and Salomons (2018) similarly show that productivity growth may reduce employment within particular industries even when broader economic adjustments and the emergence of new tasks or industries offset some of those losses at the aggregate level. Autor (2022), however, raises the possibility that digital technologies may not be producing new labor-intensive tasks at the same rate as earlier technological transformations, potentially increasing the range of workers exposed to displacement.

Overall, this literature cautions against interpreting technological capability as equivalent to realized labor-market displacement. An occupation may contain tasks that can be automated without the occupation disappearing. AI may allow workers to perform some activities more quickly, shift time toward other tasks, change the skills employers demand, or alter the way work is organized. In other cases, technological adoption may reduce labor demand. The eventual employment consequences depend not only on what a technology is technically capable of doing, but also on whether firms adopt it, how they integrate it into production, whether complementary tasks emerge, and how workers adapt.

Recent evidence reinforces that distinction. In *Rethinking Automation Risk: AI Applicability and Occupational Outcomes, 2019–2024*, Broady, Dunson, and Barr (2026) compare employment and wage changes across occupations with differing levels of Frey and Osborne automation risk and Tomlinson et al. AI applicability. Their results complicate a simple displacement interpretation. Occupations with high AI applicability experienced overall employment and wage growth between 2019 and 2024, while occupations classified as having high or moderate automation risk experienced weaker employment performance than low-risk occupations. Wages nevertheless increased across all automation-risk categories. The authors therefore argue that exposure measures are better interpreted as indicators of potential occupational restructuring than as direct forecasts of employment decline.

That distinction is especially important for generative AI because the technology has diffused quickly while remaining unevenly integrated into work. Technical capability can advance more rapidly than workplace adoption, and occupational exposure can exist long before measurable employment effects appear. The relevant questions are therefore no longer limited to whether AI can perform particular tasks. They also include whether workers are actually using the technology, how frequently they use it, whether employers encourage its use, whether workers perceive it as improving their productivity or expanding the tasks they can perform, and whether those experiences differ systematically across groups of workers.

This progression from estimates of technical susceptibility, to task-based models of technological change, to observed interactions between workers and generative AI provides the foundation for the empirical analysis that follows. The existing literature provides increasingly sophisticated measures of what automation and AI may be capable of doing to occupations and tasks. The 2025 Survey of Household Economics and Decisionmaking adds a different perspective: it allows us to observe whether workers report using generative AI in their jobs and how they assess its effects on their work and careers. The Federal Reserve System's Worker Perspectives Project adds still another dimension by documenting how workers in lower- and moderate-income households describe encountering AI in employment and job-search processes. Read alongside one another, these sources allow the analysis to move from technological capability and occupational exposure toward actual adoption and worker experience.

3. Data and Methodology

Since 2013, the Federal Reserve Board has conducted the Survey of Household Economics and Decisionmaking (SHED) to assess the economic well-being of U.S. adults and identify emerging risks to household finances (Board of Governors of the Federal Reserve System, 2026). The survey covers a broad range of topics related to financial well-being, including income and expenses, credit access and use, savings, retirement, housing, economic hardships, education, and student loans. The 2025 SHED was fielded from October 17 through October 28, 2025, and the final analytic sample included 12,934 respondents ages 18 and older living in the United States. The Employment and Job Quality module added particular attention to workers' use of generative artificial intelligence (Board of Governors of the Federal Reserve System, 2026).

The 2025 report treats generative AI as an increasingly important feature of workers' job experiences. One-quarter of workers reported using generative AI for their jobs in the prior month (Board of Governors of the Federal Reserve System, 2026, pp. 17). The survey also asked workers whether AI saves time, improves the quality of their work, enables them to perform new tasks, whether their employers encourage its use, and whether they expect AI to improve their careers or replace their jobs. Among workers who had used generative AI, 81 percent agreed that it saved them time, while smaller majorities said it improved the quality of their work or enabled them to perform new tasks (Board of Governors of the Federal Reserve System, 2026, pp. 18). AI users were also considerably more likely than workers overall to view the technology positively and to expect it to improve their careers.

The report also highlights substantial differences in workplace AI use across workers. Use increased sharply with education: 43 percent of workers with a graduate degree reported using generative AI in the prior month, compared with 10 percent of workers with a high school degree or less. Workers with more education were also more likely to report that their employers encouraged AI use. Differences appeared by age, disability status, and job autonomy as well. Among employees who often or always had control over how they completed their work, 30 percent reported using generative AI, compared with 20 percent of those with less autonomy (Board of Governors of the Federal Reserve System, 2026, pp. 18).

The Federal Reserve System's Worker Perspectives Project provides a complementary qualitative source for understanding workers' experiences during roughly the same period as the 2025 SHED. The project was designed to speak directly with workers and job seekers from low- to moderate-income households and to provide insight into the "why," "what," and "how" underlying economic and labor market trends that may be difficult to observe in aggregate quantitative data (Miller, et al., 2026). The 2025 research consisted of 11 focus groups conducted between October and December 2025 with 63 participants from across the United States, including Spanish-language sessions. All participants lived in households with annual income of $85,000 or less, a threshold selected to approximate the income level of roughly half of U.S. households.

The Worker Perspectives Project used a semi-structured focus-group design and an inductive qualitative approach to identify recurring experiences, attitudes, and themes across participants. A common facilitation guide provided consistency across sessions while allowing moderators to ask follow-up questions based on participant responses and group discussion. The research team analyzed the qualitative evidence to identify themes and findings, and a subsequent validation session with a subset of participants was used to confirm that the themes and descriptive language accurately reflected participants' experiences. The present study draws on the findings reported in the published Worker Perspectives report rather than conducting a separate reanalysis of the underlying focus-group transcripts.

The focus groups examined how workers and job seekers were experiencing changes in household costs, employment opportunities, job quality, education and skill-building, and confidence in the labor market. Participants described substantial financial pressure, limited economic mobility, unstable employment, and persistent job searching. Of particular relevance to this study, participants also discussed the growing role of artificial intelligence in the job-search process. Many participants engaged in active job searches believed that AI-based screening systems were excluding them from consideration for jobs, and some reported feeling disadvantaged when interacting with AI during hiring and interview processes. These concerns contributed to broader uncertainty about future employment opportunities and confidence in the labor market.

The empirical analysis proceeds in two stages. First, weighted descriptive statistics are used to compare the prevalence and frequency of workplace generative AI use and workers' perceptions of its benefits and risks. All estimates use the SHED cross-sectional survey weight. Second, logistic regression is used to examine whether differences in workplace generative AI use persist after accounting for observable worker characteristics. Let AI_i equal one if worker i reported using generative AI at work during the prior month and zero otherwise.

The baseline model is specified as:

$$\Pr(AI_i = 1 \mid X_i) = \Lambda(\beta_0 + \beta_1 Female_i + \beta_2 Race_i + \beta_3 Age_i + \beta_4 Education_i + \beta_5 Income_i + \beta_6 Region_i). \quad (1)$$

where $\Lambda(\cdot)$ denotes the logistic cumulative distribution function. Race and ethnicity, educational attainment, household income, and Census region are entered as categorical variables, while age is measured continuously. The reference categories are male, White, high school diploma or GED, household income below $25,000, and the Northeast. Because coefficients from nonlinear probability models are not directly interpretable as percentage-point changes in probability, results from the baseline specification are reported as average marginal effects (AMEs). For age, the AME represents the average change in the predicted probability associated with a one-year increase. For categorical variables, the reported effect represents the average discrete change in predicted probability relative to the reference category.

To examine whether gender differences in workplace AI use vary across racial and ethnic groups, a second specification allows the association between gender and AI use to differ by race and ethnicity.

The interaction model is:

$$\Pr(AI_i = 1 \mid X_i) = \Lambda[\beta_0 + \beta_1 Female_i + \Sigma_r \gamma_r Race_{ir} + \Sigma_r \delta_r (Female_i \times Race_{ir}) + Z_i'\theta]. \quad (2)$$

where Z_i contains age, educational attainment, household income, and Census region. The δ_i terms represent the gender-by-race/ethnicity interactions, and White men form the omitted gender-race category. Because interaction coefficients in a nonlinear model are not themselves direct probability differences, the analysis also reports adjusted predicted probabilities for men and women within each racial and ethnic group.

A joint Wald test is used to determine whether the gender-by-race/ethnicity interaction terms are collectively distinguishable from zero. The null hypothesis is:

$$H_0: \delta_1 = \delta_2 = \delta_3 = \delta_4 = 0. \quad (3)$$

The Wald statistic is calculated as:

$$W = \hat{\delta}'[\mathrm{Var}(\hat{\delta})]^{-1}\hat{\delta}. \quad (4)$$

and is evaluated against a chi-square distribution with four degrees of freedom. Rejection of the null would indicate that the relationship between gender and workplace AI use differs systematically across racial and ethnic groups. The interaction coefficients, adjusted probabilities, and joint Wald test are reported together so that the individual and collective evidence from the interaction specification can be interpreted separately.

4. Worker Perspectives and Views on AI

The analysis begins with the weighted descriptive statistics described in Section 3, which provide an initial picture of workplace generative AI use and workers' experiences with the technology. Table 1 provides an initial overview of both the prevalence of workplace generative AI use and how workers who use the technology experience it. The table distinguishes between whether workers use generative AI at all, how frequently users rely on it, and what users believe the technology does for their work and careers. Data reported in the three panels suggest that generative AI remains far from universal in the workplace, but among workers who do use it, engagement can be fairly frequent and perceptions of its immediate usefulness are generally positive. At the same time, workers are less uniformly positive when asked about broader career implications, suggesting an important distinction between finding AI useful for completing work and assessing what the technology may mean for longer-term employment.

Panel A shows that 24.7 percent of workers reported using generative AI at work during the prior month, compared with 75.3 percent who did not. Thus, despite the rapid expansion and visibility of generative AI, workplace use remained concentrated among a minority of workers in 2025.

This distinction is important because much of the public discussion surrounding AI focuses on the technology's capabilities or potential effects, whereas these data measure whether workers are actually using it in their jobs. The results therefore suggest a substantial gap between the broad availability of generative AI technologies and their incorporation into workers' day-to-day employment.

Panel B shows considerable variation in the intensity of use among workers who had adopted generative AI. Approximately 30 percent used it daily or multiple times per day, and another 24 percent used it several days per week. At the other end of the distribution, 29.1 percent reported using generative AI less than one day per week. Weekly users accounted for the remaining 17.3 percent. The distribution therefore does not suggest a single typical pattern of workplace AI use. For some workers, generative AI appears to have become a routine part of work, while for others it remains an occasional tool. This variation in frequency may be particularly important when considering the extent to which AI is changing how work is actually performed rather than simply whether workers have experimented with the technology.

Panel C provides additional insight into what workers believe they gain from using generative AI. Time savings are by far the most widely reported benefit: 81 percent of users agreed that AI saves them time, while only 3.6 percent disagreed. Other perceived benefits are positive but less universal. A majority reported that generative AI enables them to perform new tasks (55.5 percent) or improves the quality of their work (51.8 percent). Employer encouragement is similarly divided, with 50.8 percent agreeing that their employer encourages AI use and 36.5 percent neither agreeing nor disagreeing. These results suggest that the strongest perceived benefit of generative AI may currently be efficiency rather than a wholesale expansion or improvement of what workers are able to produce.

The responses concerning careers and job replacement are more mixed. Slightly less than half of users agreed that having generative AI available would improve their career, while 41.8 percent neither agreed nor disagreed. At the same time, 22.1 percent expressed concern that AI could replace their job, compared with 47.2 percent who disagreed and 30.7 percent who were uncertain. This pattern is notable because workers appear considerably more certain about AI's immediate usefulness than about its longer-term implications for their careers and employment. The results therefore point to a distinction between workers' assessments of what generative AI does for them now and their expectations about what it may mean for them in the future.

The Worker Perspectives focus groups add another dimension to these findings. Whereas the SHED captures workers’ experiences using generative AI in their own jobs, focus-group participants also described experiencing AI as a system used by employers. Among participants actively searching for work, many believed AI-based screening was excluding them from consideration, and some felt disadvantaged by AI-mediated interview processes (Miller et al., 2026). These accounts suggest that workers may encounter AI both as a tool they use and as a technology used to evaluate or sort them in the labor market.

The qualitative evidence also helps contextualize the relatively mixed SHED responses on careers and displacement. SHED respondents who currently use generative AI were more likely to describe immediate benefits than to express concern about job replacement. By contrast, Worker Perspectives participants, many of whom were navigating job searches or less stable employment circumstances, described greater uncertainty about AI's effects on future employment opportunities. The difference is not necessarily contradictory. Rather, it suggests that perceptions of AI may depend heavily on whether workers encounter the technology as a useful tool within an existing job or as part of a more uncertain process of competing for employment.

Table 1. Generative AI Use and Perceptions among Workers

2025 Survey of Household Economics and Decisionmaking (SHED)

Panel A. Use of generative AI at work in the prior month

Response	Unweighted N	Weighted percent
Yes	1,850	24.7%
No	5,638	75.3%
Total answering	7,488	100.0%

Panel B. Frequency of generative AI use among users

Frequency	Weighted percent
Multiple times per day	15.0%
Daily	14.6%
Several days per week	24.0%
Weekly	17.3%
Less than one day per week	29.1%
Total	100.0%

Panel C. Perceptions of generative AI among users

Statement	Agree	Neither agree nor disagree	Disagree
It saves time	81.0%	15.4%	3.6%
It improves quality	51.8%	35.6%	12.5%
It enables me to do new tasks	55.5%	31.6%	13.0%
My employer encourages me to use it	50.8%	36.5%	12.8%
Having it available will improve my career	48.1%	41.8%	10.1%
I am worried about AI replacing my job	22.1%	30.7%	47.2%

Notes: Panel A includes the 7,488 workers who answered the generative-AI-at-work question. Panels B and C describe workers who reported using generative AI at work in the prior month. Percentages are weighted using the SHED cross-sectional survey weight.

Table 2 presents unadjusted differences in workplace AI use by race and ethnicity and gender. AI use was highest among Asian workers, 38.6 percent of whom reported using generative AI at work in the prior month. Use was 24.6 percent among White workers, 22.2 percent among Hispanic workers, and 20.7 percent among Black workers. Overall differences by gender were small, with 25.1 percent of men and 24.4 percent of women reporting workplace AI use. Within the Black worker population, however, the unadjusted difference was somewhat larger: 22.3 percent of Black men reported using generative AI compared with 19.3 percent of Black women.

Table 2. Unadjusted Generative AI Use by Race/Ethnicity and Gender

Weighted shares of workers reporting generative AI use at work in the prior month

Panel A. Race/ethnicity

Race/ethnicity	Weighted share using AI
White	24.6%
Black	20.7%
Hispanic	22.2%
Asian	38.6%
Other	23.5%

Panel B. Gender

Gender	Unweighted N	Weighted share using AI
Men	3,993	25.1%
Women	3,495	24.4%

Panel C. Black workers by gender

Gender	Unweighted N	AI users, unweighted N	Weighted share using AI
Men	398	89	22.3%
Women	454	85	19.3%

Notes: Percentages are weighted using the SHED cross-sectional survey weight. Values are unadjusted descriptive statistics.

The descriptive differences in Table 2 do not account for differences in workers' other characteristics. Hence, Table 3 moves beyond the unadjusted comparisons by presenting results from the baseline logistic regression specified in Equation (1). The model examines the association between worker characteristics and workplace generative AI use simultaneously, controlling for age, gender, race and ethnicity, educational attainment, household income, and Census region. Consistent with the empirical strategy described in Section 3, the results are reported as average marginal effects, which indicate percentage-point differences in the predicted probability of workplace generative AI use. The estimates reveal substantial differences in the strength of these associations. Education and household income stand out most clearly, while gender and region show little independent association with AI use after the other included characteristics are taken into account. Age is also strongly associated with use, and some racial and ethnic differences remain after adjustment.

Age is negatively associated with workplace generative AI use. Each additional year of age is associated with an estimated 0.36-percentage-point decline in the probability of using generative AI, holding the other included characteristics constant. Because age is entered continuously, the effect accumulates across larger age differences. For example, the model implies a substantially

larger difference between workers separated by ten or twenty years than the single-year estimate might initially suggest. The result is consistent with workplace AI adoption being more prevalent among younger workers even after differences in education, income, gender, race and ethnicity, and region are considered.

Gender, by contrast, has little association with AI use in the baseline model. Women are estimated to be about one percentage point less likely than men to report workplace generative AI use, but the confidence interval includes zero and the estimate is not statistically significant. This result is particularly notable in light of the nearly identical unadjusted use rates for men and women shown in Table 2. In other words, both the descriptive and multivariate results provide little evidence of a broad gender difference in workplace AI use when gender is considered across the workforce as a whole.

The results for race and ethnicity are more varied. After adjustment, the estimated differences between Black and White workers are small and statistically insignificant. The estimated difference for Hispanic workers is positive and approaches, but does not reach, the conventional 5 percent level of statistical significance. By contrast, Asian workers are approximately 6 percentage points more likely than otherwise similar White workers to report using generative AI, and workers in the residual “other race/ethnicity” category are estimated to be approximately 6.7 percentage points more likely. These results are also informative when compared with the unadjusted statistics in Table 2. In particular, the lower unadjusted rate of AI use among Black workers does not persist once differences in the other included characteristics are taken into account. That suggests that the raw Black–White difference reflects, at least in part, differences in characteristics associated with AI use rather than an independent racial difference captured by this specification.

Educational attainment exhibits one of the clearest gradients in the model. Relative to workers with a high school diploma or GED, workers with some college or an associate degree are about 4.1 percentage points more likely to report workplace AI use. The difference increases sharply at higher levels of educational attainment: workers with a bachelor's degree are approximately 14.5 percentage points more likely to use generative AI, while those with a master's degree or higher are about 20.1 percentage points more likely. In contrast, the estimated difference between workers without a high school diploma or GED and the reference group is negative but statistically insignificant. The pattern is therefore not simply one in which each additional educational category produces a uniformly larger probability of AI use. Rather, the largest differences emerge among workers with bachelor's and graduate-level education.

Household income displays a similarly pronounced gradient. There is essentially no adjusted difference between workers in households earning less than $25,000 and those earning $25,000 to $49,999. The difference becomes considerably larger for household incomes between $50,000 and $99,999, with an estimated increase of 6.7 percentage points in the probability of workplace AI use. Workers in households earning $100,000 or more have by far the largest income-related difference, with an estimated probability nearly 20 percentage points higher than workers in households earning less than $25,000. The sharp increase at higher household income levels suggests that workplace exposure to and adoption of generative AI is disproportionately concentrated among workers with greater economic resources.

The strong associations between workplace AI use, education, and household income are also consistent with themes that emerged from the Worker Perspectives focus groups. Participants

from households earning $85,000 or less frequently described education and skill-building as important to employment mobility, but many also reported that the cost of education, work and family obligations, and uncertainty about the returns to additional credentials limited their ability to pursue further training (Miller et al., 2026). These qualitative findings do not establish the mechanisms underlying the SHED estimates, but they illustrate some of the constraints that may affect workers' ability to acquire the skills, credentials, or workplace opportunities associated with new technologies.

Finally, there is little evidence that workplace AI use differs systematically across Census regions after the other characteristics are considered. Relative to workers in the Northeast, estimated differences for workers in the Midwest, South, and West are small and statistically insignificant. Geography, at least at the broad regional level captured here, appears considerably less important than education, income, or age in explaining differences in workplace AI use.

Overall, the data in Table 3 suggests that the uneven adoption of generative AI is associated more strongly with workers' education, household income, and age than with broad differences by gender or geography. Race and ethnicity present a more nuanced picture: some adjusted differences remain, while others that are visible descriptively largely disappear. This distinction is important because it shows why the multivariate analysis adds something beyond simply comparing group percentages. The characteristics of workers who use generative AI overlap, and some apparent group differences look quite different once those characteristics are considered simultaneously.

Table 3. Weighted Logistic Regression Predicting Generative AI Use

Average marginal effects (AMEs). Dependent variable equals 1 if the respondent used generative AI at work in the prior month and 0 otherwise.

Characteristic	Average marginal effect	95% confidence interval	p-value
Age, each additional year	-0.36 pp	-0.43 to -0.29	<0.001
Gender			
Male	Reference		
Female	-1.05 pp	-2.90 to 0.79	0.262
Race/ethnicity			
White	Reference		
Black	+1.05 pp	-2.06 to 4.15	0.509
Hispanic	+2.68 pp	-0.02 to 5.38	0.052
Asian	+6.04 pp	2.46 to 9.62	0.001
Other race/ethnicity	+6.70 pp	0.53 to 12.88	0.033
Education			
High school diploma/GED	Reference		
No high school diploma/GED	-3.01 pp	-6.87 to 0.86	0.127
Some college/associate degree	+4.13 pp	1.63 to 6.63	0.001
Bachelor's degree	+14.48 pp	11.76 to 17.21	<0.001
Master's degree or higher	+20.08 pp	16.79 to 23.37	<0.001
Household income			
Less than $25,000	Reference		
$25,000-$49,999	+0.55 pp	-2.53 to 3.62	0.726
$50,000-$99,999	+6.71 pp	3.82 to 9.59	<0.001
$100,000 or more	+19.83 pp	16.83 to 22.83	<0.001
Region			
Northeast	Reference		
Midwest	+0.77 pp	-2.11 to 3.66	0.600
South	+2.13 pp	-0.48 to 4.74	0.110
West	+1.60 pp	-1.24 to 4.44	0.270

Notes: Weighted logistic regression estimated on the 7,488 workers who answered the generative-AI-at-work question using the SHED cross-sectional weight. Age is entered as a continuous exact-age variable. Gender, race/ethnicity, education, household income, and Census region are categorical indicators. AMEs are percentage-point differences in the probability of using generative AI, holding the other included characteristics constant. Reference categories are shown in the table.

The small overall gender difference may conceal variation across racial and ethnic groups. Table 4 therefore allows the association between gender and workplace AI use to differ by race and ethnicity while continuing to account for age, education, household income, and region. Adjusted probabilities were nearly identical for White men and women and for Hispanic men and women. The largest estimated gender differences appeared among Black and Asian workers. Among Black workers, the adjusted probability of workplace AI use was 27.5 percent for men and 22.0 percent for women, a difference of 5.5 percentage points ($p = .058$). The interaction results should nevertheless be interpreted cautiously. The female-by-Black interaction coefficient was statistically significant ($p = .037$), indicating that the gender difference among Black workers differed from the corresponding gender difference among White workers in this specification. However, the four gender-by-race interaction terms were not jointly statistically significant (Wald $\chi^2(4) = 5.48$, $p = .241$). The results therefore do not provide evidence of systematic gender differences across racial and ethnic groups as a whole, although the Black male–female difference warrants further examination.

Table 4. Gender Differences in Generative AI Use by Race/Ethnicity

Weighted logit with gender x race/ethnicity interactions and controls for exact age, education, household income, and Census region.

Panel A. Adjusted probabilities

Race/ethnicity	Men	Women	Female-male difference	p-value
White	23.6%	23.4%	-0.2 pp	0.862
Black	27.5%	22.0%	-5.5 pp	0.058
Hispanic	25.8%	26.8%	+1.0 pp	0.681
Asian	31.6%	26.8%	-4.9 pp	0.145
Other race/ethnicity	31.2%	29.3%	-1.9 pp	0.757

Panel B. Gender x race/ethnicity interaction terms

Interaction term	Logit coefficient	Robust SE	z	p-value
Female x Asian	-0.251	0.199	-1.26	0.207
Female x Black	-0.392	0.188	-2.09	0.037
Female x Hispanic	-0.022	0.156	-0.14	0.890
Female x Other race	-0.053	0.339	-0.16	0.876

Notes: Panel A reports adjusted probabilities implied by the interaction model. Panel B reports the corresponding gender x race/ethnicity logit interaction coefficients. White workers are the reference race/ethnicity category. The joint Wald test of the four gender x race/ethnicity interaction terms is chi-square(4) = 5.48, p = .241.

5. Discussion

The results presented here suggest that generative AI is already having a meaningful impact in the workplace. However, its current effects remain more uneven and limited than some of the broader rhetoric surrounding the technology might imply. The technological capabilities of generative AI have advanced rapidly, and recent research shows that machines can now perform or assist with tasks once viewed as resistant to automation, including activities involving writing, programming, creativity, and elements of social and emotional interpretation. At the same time, the worker-level evidence presented in this paper shows that the existence of those capabilities does not mean that they have been adopted uniformly or that they have transformed work to the same degree across the labor force.

That distinction is important in historical perspective. Earlier technological transformations altered fundamental constraints on economic and social activity. Artificial light extended the usable day; agricultural innovation changed how societies produced food and supported occupational specialization; mechanization transformed physical production; and electrification provided infrastructure that enabled subsequent waves of technological development. These changes eventually affected large portions of economic and social life, but their diffusion was neither immediate nor uniform. Generative AI may ultimately prove comparably consequential, but the available evidence does not yet suggest that its effects are of the same magnitude as technologies that changed basic physical or organizational constraints on human activity. Rather, the current period appears more consistent with an early and uneven stage of technological adoption in which capabilities have advanced faster than their incorporation into everyday work.

The SHED results illustrate this distinction clearly. In 2025, only about one-quarter of workers reported using generative AI at work in the prior month, and frequency of use varied substantially even among users. For some workers, generative AI has become part of the regular workday; while for others, it remains an occasional tool. This pattern suggests that the relevant question is not simply whether generative AI can perform particular tasks, but whether workers have access to it, whether employers permit or encourage its use, whether its applications are relevant to the tasks workers perform, and whether workers have the knowledge and opportunity required to incorporate it effectively.

The results also suggest that the current benefits of generative AI may be more incremental than revolutionary for many workers. Time savings are the most widely reported benefit, while smaller majorities report improvements in work quality or the ability to perform new tasks. That pattern is consistent with AI functioning, at least for many current users, as a productivity-enhancing tool rather than as a wholesale replacement for existing work. The distinction matters because improvements in efficiency can still have substantial economic effects without immediately producing large-scale occupational displacement. Workers may complete familiar tasks more quickly, shift time toward other activities, or alter how particular tasks are performed without experiencing an immediate change in occupation or employment status.

At the same time, the results point to an uneven distribution of opportunities to use and benefit from AI. Educational attainment and household income are among the strongest correlates of workplace AI use in the multivariate analysis. Workers with bachelor's and graduate degrees are considerably more likely to report using generative AI than workers with only a high school diploma or GED, and the probability of use rises sharply among workers in higher-income households. These relationships remain even after accounting for age, gender, race and ethnicity, and region.

There are several possible interpretations of these patterns, and the data do not permit a definitive test among them. Workers with more education may be concentrated in occupations containing a greater share of tasks to which generative AI is currently applicable. Higher-income workers may also be more likely to work in organizations that provide access to AI tools, encourage experimentation, or have already incorporated AI into workflows. Education may additionally affect workers' ability to evaluate AI-generated output, identify errors, formulate useful prompts, or determine when the technology is appropriate for a particular task. These possibilities are consistent with the broader literature on task-based technological change, but they should not be interpreted as evidence that education or income themselves cause AI adoption.

The Worker Perspectives findings reinforce the importance of this distinction. Participants from lower- and moderate-income households frequently described education and skill development as important to economic mobility, while also reporting financial, family, and work-related constraints that limited opportunities to pursue additional training. These accounts provide useful context for the SHED results. If access to new technologies increasingly depends on the kinds of

jobs workers hold, the training employers provide, or the resources workers can devote to acquiring new skills, then uneven AI adoption may reflect broader inequalities in opportunity rather than simply differences in personal willingness to use the technology.

The findings also have implications for educators and workforce-development institutions. If workers increasingly need to assess AI-generated information, determine when it is reliable, and understand how to integrate AI into existing work processes, then education and training may need to place greater emphasis on judgment, evaluation, and application rather than simply on technical familiarity with a particular platform. This is consistent with the broader historical pattern discussed earlier in the paper: technological change often alters the skills workers need without eliminating the need for workers themselves.

For policymakers and government officials, the results suggest caution against treating AI exposure as equivalent to labor-market displacement. Measures of occupational susceptibility and AI applicability are useful for identifying where change may occur, but they do not by themselves reveal how workers or employers will respond. The current evidence suggests a more complicated process in which adoption varies widely, worker experiences differ, and immediate productivity benefits coexist with uncertainty about longer-term employment consequences. Public policy aimed at responding to AI may therefore need to distinguish among several separate issues: access to technology, worker training, employer adoption, job-search and hiring practices, and actual employment displacement.

The relatively modest concern about job replacement among current AI users should also be interpreted carefully. Only a minority of users reported worrying that AI would replace their jobs, but that does not necessarily imply that displacement concerns are unimportant. Current users may be more likely to occupy jobs in which AI is complementary to their work, may feel confident in their ability to use the technology, or may simply perceive the immediate benefits more clearly than longer-term risks. The Worker Perspectives participants, particularly those engaged in job searches, offer a different perspective in which AI is associated more directly with uncertainty and exclusion. The contrast underscores that workers' views of AI may depend heavily on their position in the labor market and on whether they encounter the technology as a tool, a complement, or a barrier.

The evidence suggests that generative AI should be understood neither as a technology whose importance is overstated nor as one whose economic effects are already fully realized. Its capabilities are substantial, and the range of tasks to which it can be applied continues to expand. But the historical record suggests that transformative technologies become economically consequential through diffusion, complementary investments, changing work practices, and broad incorporation into everyday activity. By that standard, generative AI appears to be in an important but still incomplete stage of diffusion.

Several limitations should be considered when interpreting these findings. The SHED data are cross-sectional, and the multivariate estimates identify associations rather than causal

relationships. The models also do not directly account for occupation, industry, employer AI policies, or the specific task composition of workers' jobs, all of which may contribute to observed differences in workplace AI use. In addition, the Worker Perspectives findings are qualitative and are drawn intentionally from workers and job seekers in households with annual incomes of $85,000 or less. They provide important context for understanding how some lower- and moderate-income workers experience AI and the labor market, but they are not nationally representative and should not be interpreted as estimates of the prevalence of those experiences across the broader workforce.

6. Conclusion

Generative artificial intelligence is an important technological development, but its historical significance should be judged by more than its technical capabilities. Earlier technologies became transformative through widespread diffusion, complementary investments, changes in work and production, and incorporation into everyday economic life. Artificial light ultimately altered a fundamental constraint on human activity and, through electrification, became part of the infrastructure of the modern economy. Generative AI has not yet produced a transformation of comparable scale. That is not to say that its long-run impact will be limited. However, the evidence presented here suggests that the current AI transition is still underway. Its eventual significance will depend not only on what the technology is capable of doing, but also on how broadly it is diffused, how firms and workers incorporate it into production, and how institutions respond to the opportunities and disruptions that follow.

References

Acemoglu, Daron, and Pascual Restrepo. "Automation and new tasks: How technology displaces and reinstates labor." Journal of economic perspectives 33, no. 2 (2019): 3-30.

Alarcón, Claudia Myrna Méndez, Olusiji Adebola Lasekan, Margot Teresa, and Godoy Pena. "AI concerns among educators: A thematic analysis." *International Research Journal of Medical Sciences* 6, no. 1 (2025): 1031-1045.

Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." Journal of Economic Perspectives 29 (3): 3–30.

Autor, David. The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty. No. w30074. National Bureau of Economic Research, 2022.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. "Computing inequality: have computers changed the labor market?" The Quarterly Journal of Economics 113, no. 4 (1998): 1169-1213.

Autor, David, and Anna Salomons. Is automation labor-displacing? Productivity growth, employment, and the labor share. No. w24871. National Bureau of Economic Research, 2018.

Barata, João, and Ina Kayser. "Industry 5.0–past, present, and near future." *Procedia Computer Science* 219 (2023): 778-788.

Bataineh, Abdallah Q., Alaa S. Mushtaha, Ibrahim A. Abu-AlSondos, Saeed Hameed Aldulaimi, and Marwan Abdeldayem. "Ethical & legal concerns of artificial intelligence in the healthcare sector." In *2024 ASU International Conference in Emerging Technologies for Sustainability and Intelligent Systems (ICETSIS)*, pp. 491-495. IEEE, 2024.

Berdugo, Sophie. "Scientists taught an AI-powered 'robot dog' how to play badminton against humans — and it's actually really good." Live Science. (September 3, 2025).

Board of Governors of the Federal Reserve System. 2026. *Economic Well-Being of U.S. Households in 2025*. Washington, DC: Board of Governors of the Federal Reserve System. May 2026.

Brilliance LED. (2025). "The Evolution of Artificial Light Sources—A Walk through Time." *Brilliance LED*. Accessed October 9, 2026. Artificial Light Sources | History and Evolution of Artificial Light

Broady, Kristen. "Race and jobs at high risk to automation." Joint Center for Political and Economic Studies, December 18 (2017).

Broady, Kristen E., Darlene Booth-Bell, Anthony Barr, and Aisha Meeks. "Automation, artificial intelligence, and job displacement in the US, 2019–22." *Labor History* 67, no. 2 (2026): 257-273.

Broady, Kristen, Caleb Dunson, and Anthony Barr. "Rethinking Automation Risk: AI Applicability and Occupational Outcomes, 2019-24." (2026).

Center for AI Safety. "Statement on AI risk." *https://www. safe. ai/statement-on-ai-risk* (2023).

Cerutti, Eugenio, Eugenio M. Cerutti, Antonio I. Garcia Pascual, Yosuke Kido, Longji Li, Giovanni Melina, Marina Mendes Tavares, and Philippe Wingender. *The global impact of AI: Mind the gap*. International Monetary Fund, 2025.

Clarke, Laurie. "Call for AI pause highlights potential dangers." *Science* 380, no. 6641 (2023): 120-121.

Fatahi, Somayeh, Mehdi Khalaj, Shahabeddin Abhari, Julita Vassileva, and Chanchal Roy. "Beyond Words: Exploring Emotional Contrasts Between Human and ChatGPT Responses in Medicine and Finance." In International Conference on Human-Computer Interaction, pp. 136-151. Cham: Springer Nature Switzerland, 2025.

Finet, Alain, Kevin Kristoforidis, and Julie Laznicka. "The Limits of AI in Understanding Emotions: Challenges in Bridging Human Experience and Machine Perception." F1000Research 14 (2025): 582.

Fišer, Nikolaj, Miguel Ángel Martín-Pascual, and Celia Andreu-Sánchez. "Emotional impact of AI-generated vs. human-composed music in audiovisual media: A biometric and self-report study." PloS one 20, no. 6 (2025): e0326498.

Frey, Carl Benedikt, and Michael A. Osborne. "The future of employment: How susceptible are jobs to computerisation?" Technological forecasting and social change 114 (2017): 254-280.

Gorenz, Drew, and Norbert Schwarz. "How funny is ChatGPT? A comparison of human-and AI-produced jokes." Plos one 19, no. 7 (2024): e0305364.

Hooper, Charles. "Important Facts About Thomas Edison & the Invention of the Light Bulb." *Sciencing*. Last modified March 24, 2022. https://www.sciencing.com/important-thomas-edison-invention-light-bulb-6305/

Kerridge, Eric. "The agricultural revolution reconsidered." *Agricultural History* 43, no. 4 (1969): 463-476.

Knight, Will. "This Robot Only Needs a Single AI Model to Master Humanlike Movements." Wired. (September 3, 2025).

Kooli, Chokri, and Hend Al Muftah. "Artificial intelligence in healthcare: a comprehensive review of its ethical concerns." *Technological Sustainability* 1, no. 2 (2022): 121-131.

Lahemer, Elfituri SF, and Ahmad Rad. "An Audio-Based SLAM for Indoor Environments: A Robotic Mixed Reality Presentation." Sensors 24, no. 9 (2024): 2796.

Miller, Sarah, Merissa Piazza, Elizabeth Bogue Simpson, and Kristen Broady. "When every dollar counts: Inside the economic struggles of workers who earn low to moderate incomes." *Notes from the Field* 20260624 (2026).

Naderi, Hossein, Alireza Shojaei, and Lifu Huang. "Foundation Models for Autonomous Robots in Unstructured Environments." arXiv preprint arXiv:2407.14296 (2024).

Nahavandi, Saeid. "Industry 5.0—A human-centric solution." *Sustainability* 11, no. 16 (2019): 4371.

Novo, Alvaro, Francisco Lobon, Hector Garcia de Marina, Samuel Romero, and Francisco Barranco. "Neuromorphic perception and navigation for mobile robots: a review." ACM Computing Surveys 56, no. 10 (2024): 1-37.

Ostrizhnaya, Polina. "Let There Be Light: The History of Artificial Illumination." *Scientific Russia*, February 19, 2022.

Patel, Marcus, and Ethan Reynolds. "Robotic Perception and Manipulation in Unstructured Environments." International Journal on Mechanical Engineering and Robotics 13, no. 1 (2025): 16-22.

Sanchez, Thomas W., Marc Brenman, and Xinyue Ye. "The ethical concerns of artificial intelligence in urban planning." *Journal of the American Planning Association* 91, no. 2 (2025): 294-307.

Singh, Chahat Deep, Botao He, Cornelia Fermüller, Christopher Metzler, and Yiannis Aloimonos. "Minimal perception: enabling autonomy in resource-constrained robots." Frontiers in Robotics and AI 11 (2024): 1431826.

Sorall, Pavel. "2D or 3D? Finding the Perfect Machine Vision Fit For Your Application." Photone. (January 21, 2025)

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. "Working with AI: Measuring the Occupational Implications of Generative AI." *arXiv preprint arXiv:2507.07935* (2025).

Tu, Yuxiao, Guanqi Liang, Di Wu, Xinzhuo Li, and Tin Lun Lam. "Locomotion and self-reconfiguration autonomy for spherical freeform modular robots." The International Journal of Robotics Research (2025): 02783649251360360.

Weisdorf, J.L. (2005), From Foraging To Farming: Explaining The Neolithic Revolution. Journal of Economic Surveys, 19: 561-586. https://doi.org/10.1111/j.0950-0804.2005.00259.x

William D. Nordhaus. The Economics of New Goods (University of Chicago Press, 1996), chap. 1, https://www.nber.org/books-and-chapters/economics-new-goods/do-real-output-and-real-wage-measures-capture-reality-history-lighting-suggests-not.

Tallat, Raiha, Ammar Hawbani, Xingfu Wang, Ahmed Al-Dubai, Liang Zhao, Zhi Liu, Geyong Min, Albert Y. Zomaya, and Saeed Hamood Alsamhi. "Navigating industry 5.0: A survey of key enabling technologies, trends, challenges, and opportunities." *IEEE Communications Surveys & Tutorials* 26, no. 2 (2023): 1080-1126.

Thompson, Francis Michael L. "The second agricultural revolution, 1815-1880." *The Economic History Review* 21, no. 1 (1968): 62-77.

Toffler, Alvin. *Previews and premises*. Black Rose books, 1984.

Toffler, Alvin, and Toffler Alvin. *The third wave*. Vol. 484. New York: Morrow, 1980.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. "Working with AI: Measuring the occupational implications of generative AI." *arXiv preprint arXiv:2507.07935* (2025): 1-41.

Vanek Smith, Stacey, David Kestenbaum, and Jacob Goldstein. 2022. “Planet Money Summer School 8: Productivity & Getting Lit.” *Planet Money*, NPR, August 31, 2022.

Włodarczyk, Estera. *Industry 4.0: The Role of IIoT in Digital Transformation of the Manufacturing Sector*. White paper. Comarch, 2022.