LUISS

Degree Program in Economics and Business (Major in Finance)

Course of Macroeconomics

Artificial Intelligence and Labor Market Transformation: AI Exposure and Employment Dynamics Across U.S. Occupations

Prof. Giovanna Vallanti

SUPERVISOR

289081

CANDIDATE

Academic Year 2024/2025

Introduction

Artificial Intelligence (AI) has evolved from a narrow and specialized computer science field to a general-purpose technology. Rapid advances in machine learning (ML) and generative models have broadened the range of tasks that these systems can replicate, including, and most notably, tasks generally associated with high-skill cognitive work. As a result of this rapid evolution, the relationship between AI and the labor market has become increasingly relevant.

The economic literature frames AI as both an automation technology capable of replicating specific tasks previously performed by workers and a general-purpose technology that enables new products, processes, and complementary innovations. Recent research has shifted to task-level analysis, treating occupations as bundles of heterogeneous tasks and measuring AI exposure by technical feasibility. This work shows that whether exposure translates into displacement, augmentation, or negligible change depends on organizational choices, complementary skills, and economic incentives.

This thesis explores the relationship between AI exposure and employment growth, more specifically, whether more exposed occupations display systematically different employment growth patterns than less exposed ones. The central hypothesis is that AI exposure is associated with measurable differences in employment growth across occupations. To test this, the thesis uses a task-based AI exposure index along with occupation-level employment data from the U.S. Bureau of Labor Statistics across two periods: 2009-2014 and 2019-2024. These periods represent different stages of AI diffusion, allowing a comparative assessment of whether the exposure-employment relationship changes as AI capabilities increase. The empirical strategy includes cross-occupation regressions, distributional comparisons across exposure quartiles, and heterogeneity analysis across major occupational groups.

The thesis is structured as follows:

- Chapter 1 develops the conceptual and theoretical framework, defining artificial intelligence, reviewing its classification as a general-purpose and automation technology, and discussing task-based exposure and complementarity mechanisms.
- Chapter 2 examines AI adoption patterns, sectoral exposure differences, and diffusion dynamics across firms and industries.
- Chapter 3 presents the empirical analysis linking occupational AI exposure to employment growth, followed by an integrated discussion of findings, limitations, and policy implications.

Chapter 1 Foundations: Definitions, GPT, and Automation

1.1 Defining Artificial Intelligence and Its Technological Foundations

To understand clearly the economic impact that artificial intelligence might have on the job market, it is vital to know what AI means. The term "Artificial Intelligence" was first coined at the Dartmouth Conference in 1956. This was during the early years of computing, when research on symbolic, logic-based AI systems was conducted to imitate reasoning and problem-solving. Over time, this technology evolved from rule-based expert systems to more complex and dynamic applications such as chess-playing computers, internet search engines, and, most recently, the widespread adoption of data-driven machine learning algorithms (OECD, 2019a).

Because of the fast-paced and broad evolution of this technology, it is hard to find a definition that encompasses all; for this reason, there is no universally accepted definition of AI. Nevertheless, for this paper, I will use the OECD's working definition: AI refers to "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. It uses machine and/or human-based inputs to perceive real and/or virtual environments; abstract such perceptions into models (in an automated manner, e.g., with machine learning (ML) or manually); and use model inference to formulate options for information or action. AI systems are designed to operate with varying levels of autonomy"(OECD, 2019b). Simply put, these systems process data inputs like datasets, camera feeds, and sensors that enable them to perceive their surroundings (e.g., a self-driving car), and transform this data into complex models that make decisions based on inferred insights.

These models are generally developed through automated statistical inference, like ML or human coding. The unique ability of an AI system is to learn and improve from data over time, giving this technology multiple levels of autonomy. For example, some systems can run independently without human interaction, while others require supervision or rule constraints (Lane & Saint-Martin, 2021). This ability differentiates AI from previous technologies and is central to understanding its economic potential.

AI systems can be classified according to their learning framework (supervised, unsupervised, etc.), their task domain (image recognition, language models, etc.), and their function (classification, prediction, control, etc.). ML stands out for its effectiveness across a wide range of tasks, from natural language processing to medical image interpretation (Alnuaimi & Albaldawi, 2024).

1.2 Why AI's Labor Market Impacts Are Distinct

Artificial intelligence represents a big step for automation potential, pushing through many of the limits that previous technologies like robotic machinery and software automation had. The ability to process unstructured data, improve with exposure to data, and generalize across task domains highlights the potential for this technology to be a transformative force within the economy. Unlike traditional automation, which was limited to routine, simple tasks like those generally found in manufacturing, AI extends its reach into non-routine, high-skill occupations that were previously considered to be immune to machine substitution (Aghion et al., 2017). Thus, the economic implications can be significant.

Many of AI’s capabilities are generally associated with white-collar work, for example: logical reasoning, pattern recognition, and problem-solving skills. Tasks within white-collar professions such as legal analysis, diagnostics, and data science could be automated in the future. More importantly, AI systems also have the potential to learn and improve over time, pushing the limit of what is automatable. This unique ability is especially present in ML systems, which are at the center of many AI applications (OECD, 2019a).

Despite rapid progress, AI still underperforms relative to humans in areas that require creativity, emotional intelligence, social reasoning, and the management of uncertainty. Tasks that we commonly associate with "humanity" are still a challenge for AI systems to replicate, so jobs in fields where empathy, moral judgment, and intuitive decision-making based on social cues, such as social workers, caregivers, and educators are still outside the reach of AI automation (Gries & Naudé, 2018). This limitation opens the door to AI complementarity rather than substitution, where AI augments human productivity rather than displacing labor outright. Workers in these fields may experience productivity gains by being able to focus on only those tasks that require human interaction, while AI supports the analytical and administrative parts of their jobs (Lane & Saint-Martin, 2021).

On one hand, AI can dramatically increase productivity by automating knowledge work and assisting decision-making. On the other, its diffusion could challenge traditionally protected occupational groups, namely, highly educated professionals, thereby intensifying public concerns about job security (OECD, 2019a).

These distinctive features, namely, the ability to self-improve, general-purpose nature, and applicability across a wide range of domains, are what make the impact on the labor market of this technology potentially more disruptive than previous waves of automation. Understanding the balance between complementarity and substitution is essential in assessing AI's macroeconomic effect.

There is consensus that AI is a general-purpose technology (GPT), a concept first introduced by Bresnahan and Trajtenberg in 1992 to describe technologies that can be applied across a broad range of sectors and occupations, are capable of continuous improvement, and augment complementary innovations. Historical examples of GPTs include the steam engine, electricity, and computing. AI meets the criteria of a GPT due to its core function of generating predictions, which serve as decision-making inputs in a wide range of fields such as education, medical diagnostics, and language translation (Agrawal et al., 2019).

ML systems can improve themselves over time and allow machines to interpret their environment, leading to waves of unprecedented innovation (Brynjolfsson et al., 2017). Similarly, Cockburn et al. (2018) describe machine learning as an "invention of a method of invention"; they argue that AI not only boosts productivity across sectors but also changes innovation itself. This is especially significant because scientific disciplines use classification and prediction, which are fundamental for new discoveries.

Understanding AI as GPT is fundamental in assessing the economic importance of this technology. As Brynjolfsson, Rock, and Syverson (2017) note, AI's broad applicability implies significant potential gains in output and overall welfare. However, this also means that its effects on the labor market, whether positive or negative, could be wide-ranging. As history shows us, automation has already displaced workers in many industries, but AI's broad applicability raises the possibility of job disruption across a much wider range. At the same time, AI's ability to boost innovation and productivity could potentially create new jobs and even new industries (Aghion et al., 2017).

Many economists also categorize AI as an automation technology, meaning a tool that can automate tasks that would have been otherwise performed by humans. This could reduce demand for certain types of workers, exerting downward pressure on their wages (Acemoglu & Restrepo, 2018; Aghion et al., 2017). AI, especially ML systems, could be more disruptive than previous automation technologies, such as industrial robots and machinery, because of its ability to solve more complex tasks, giving it a significantly broader applicability.

There is growing evidence that AI is enabling automation in domains previously considered safe, like high-skilled, white-collar jobs such as lab technicians, engineers, computer scientists, and actuaries, which are associated with non-routine, highly cognitive tasks. Traditional forms of automation have mainly targeted routine, repetitive tasks, often performed by low-skilled workers (Autor et al., 2003; Nedelkoska & Quintini, 2018). In contrast, AI's ability to process natural language, recognize patterns, and adapt to new data allows it to tackle non-routine, high-skilled tasks. AI may be a driving force behind the automation of occupations that were once considered securely human, such as those involving complex cognitive work (Aghion et al., 2017).

1.3 Occupational Exposure to AI and the Transformation of Work

To better understand the impact of AI on the labor market, it is vital to comprehend which occupations are most likely to be affected by AI's capabilities, and we measure this by assessing to which degree skills needed for certain occupations overlap with what AI technology can replicate. However, high exposure does not necessarily translate to substitution; it merely means that many of an occupation's tasks can be carried out by AI models. Whether those tasks are automated, transformed, or augmented depends on other factors such as cost, regulatory acceptance, and organizational choices (Lane & Saint-Martin, 2021).

Recent studies provide different methods to measure AI exposure:

- Webb (2019) links AI patents to occupational tasks in the U.S. O*NET database using verb-noun pair matching. This allows for a quantifiable comparison of AI research with the linguistic structure of job descriptions.

- Felten et al. (2021) use a skill-based approach by linking ten core AI applications (including image recognition, language modelling, translation, and speech recognition) to the 52 O*NET occupational abilities using crowd-sourced expert judgments, constructing an AI Occupational Exposure (AIOE) index that measures how closely an occupation's ability profile aligns with the most prevalent AI capabilities.

- Brynjolfsson et al. (2018) construct a suitability framework to assess which tasks are more likely to be performed by AI. This includes criteria such as routineness, frequency, and the degree to which a task can be described by formal rules.

Using each of these three methods, a pattern emerges: most occupations considered to be highly exposed to AI are generally white-collar jobs that were previously considered to be shielded by the displacement effect of automation. Felten et al. find that almost all jobs with high AI exposure require higher education. At the top of the table, we find engineers, actuaries, mathematicians, epidemiologists, etc. While Webb and Brynjolfsson et al. find more variety in exposure levels across the skill distribution, they similarly find the same relationship, high-skilled occupations are more exposed.

These results are in contrast with earlier research on risks of automation. Studies by Nedelkoska and Quintini (2018) and Frey and Osborne (2017) found that low-skilled, routine occupations such as cleaners, cashiers, and industrial machine operators were the most vulnerable to automation. Intuitively, these results highlight the difference between earlier technology and the more recent AI. Earlier robotic automation technology generally involved physical manipulation and could complete complex manual tasks, while AI technologies are complex software-based models that rely on iterative learning and perception (Felten et al., 2021). Thus, newer studies concentrate narrowly on tasks solvable by statistical AI systems, which may not capture tasks embedded within older automation technologies (Lane & Saint-Martin, 2021).

1.3.1 AI Exposure ≠ Job Elimination: The Role of Complementarity

While exposure is great at determining the likelihood that AI will be used in certain occupations, it does not tell us whether AI will replace human labor or augment it. Almost all occupations are comprised of both skills that can be replicated by AI models and skills that are still resistant to replication. For example, a financial analyst can use AI for building quantitative models, but they still rely on their judgment for decision-making, such as which stocks to buy or sell. This task heterogeneity within occupations suggests that AI will transform jobs rather than eliminate them.

This is reinforced by the concept of task-level complementarity: AI systems used for certain tasks within jobs, the long, repetitive, and mentally tiresome tasks are done by AI, leaving workers to focus on the higher value, more human-centric tasks. Fossen and Sorgner (2019) argue that occupations will not vanish but rather change structurally as AI improves specific functions within them. Similarly, AI enhances productivity by performing information and time-intensive tasks, allowing workers to redistribute time to those activities that require emotional intelligence, discretion and creativity (Agrawal et al., 2019).

For example: Alibaba's customer service chatbot handled over 95% of inquiries during a major 2017 sales event, enabling human employees to manage more complex and emotionally sensitive requests (Zeng, 2018). In healthcare, AI is aiding radiologists and clinicians in diagnostics, but the delivery of care, especially involving empathy or human judgment, remains irreplaceably in the hands of humans (Davenport & Kalakota, 2019).

Business surveys support this trend. According to Agrawal et al. (2019), most AI startups do not aim to replace labor entirely. Instead, they design systems to complement and expand human capabilities. Similarly, in firm-level surveys: companies report productivity improvement and service personalization as primary drivers of AI adoption, not cutting costs associated with labor (Accenture, 2018; Bessen et al., 2018).

1.3.2 Bottlenecks and Limits to Automation

As we have seen, some jobs are more resistant than others to AI's automation despite its rapid evolution. These more resistant jobs have in common the requirement for manual dexterity and human interaction, for example, massage therapists, educators, and social workers. In all major studies, we observe that these types of jobs have consistently lower exposure. One of the explanations can be Moravec's Paradox, which identifies the mismatch between complexity and automation, simulating high-level reasoning is computationally easier than simulating basic perceptual/motor tasks. A finding by Gries and Naudé (2018) highlights that AI systems find sensorimotor complex tasks challenging. Cooking, gardening, or cleaning are hard to replicate, even if they are considered "simple" by humans.

Another explanation is the Polanyi Paradox. Essentially, it captures the difficulty to "teach" computers "tacit knowledge", which describes everything we know how to do but may not be able to describe. "We can know more than we can tell" (Polanyi, 2009). Almost all occupations rely in part on this type of knowledge and skills such as intuition, interpreting social cues and adapting to dynamic work environments.

AI systems are just question-answering tools, not question-asking tools. This means that those who make decisions about what to work on, entrepreneurs, innovators, and problem-solvers will always be needed even in an AI-rich world (Brynjolfsson & McAfee, 2017).

Conclusively, AI exposure helps us understand which occupations are comprised of multiple technical tasks that can be replicated by AI systems, though it does not translate to AI substitution but more to AI complementarity because of those tasks that are still irreplicable and essential to all kinds of occupations. Ultimately, the question of what the labor market will look like in the future will depend on how we deploy technologies, how we reorganize work, and how institutions manage this transition.

Chapter 2 AI as a Labor Market Force: Exposure, Adoption, and Diffusion

Historically, the introduction of General Purpose Technologies (GPTs) such as electricity, the internet, and computing has brought about significant economic transformations, not only by raising productivity but also by reshaping the structure of labor markets (Agrawal et al., 2019; Lipsey et al., 2005; Varian, 2018). As AI evolves and increasingly reveals more core characteristics of a GPT, such as broad applicability, self-improvement, and complementarity with other technologies, it too will have a significant impact on the labor market (Filippucci et al., 2024).

2.1 Sectoral Differences in AI Exposure

Looking at the exposure levels across sectors, there is a clear trend: sectors reliant on cognitive tasks, especially routine, data-driven, and analytically structured tasks, consistently exhibit higher exposure, while sectors where tasks are centred around physical inputs, adaptive motor coordination, and interpersonal skills lag behind in exposure. This heterogeneity between sectors is explained by the capabilities of AI and its weaknesses. Knowledge-intensive service (KIS) industries, including Finance, Information and Communication Technologies (ICT), Publishing, and Professional Services, are the most affected by AI because they depend on highly skilled professionals and advanced technologies. KIS contains a bigger share of tasks technically exposed to AI systems, estimated to range from 50% to 80%, depending on exposure from baseline AI capability (e.g., rule-based systems and predictive models such as spam filters), either baseline AI capability or advanced AI capability (e.g., large language models, generative AI) (Eloundou et al., 2024; Filippucci et al., 2025). High exposure in KIS is particularly related to core activities in KIS industries, which are focused on: data interpretation, structured information synthesis, pattern identification, and logic-based decision-making.

Figure 1 AI Exposure Is Highest in Knowledge Intensive Services

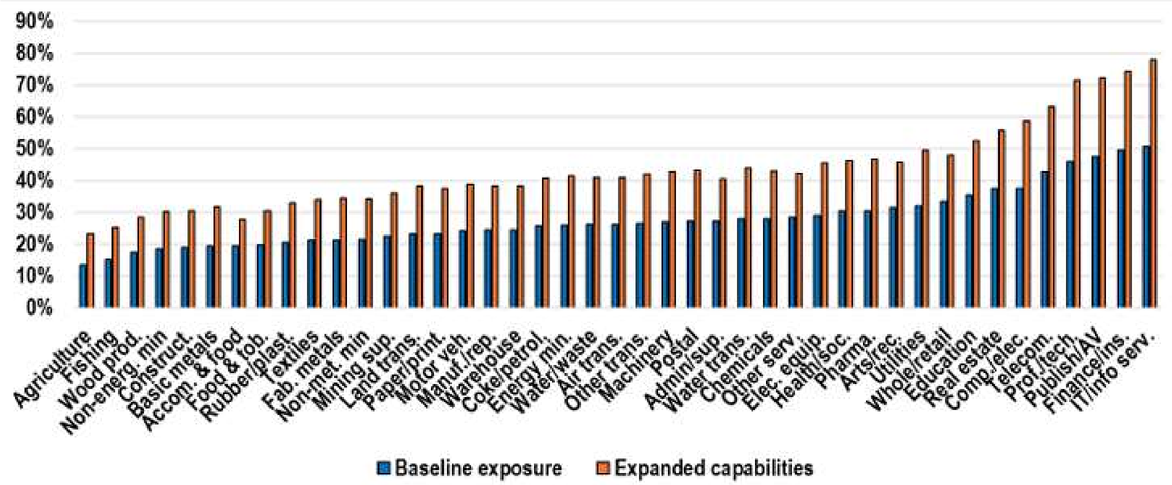

Note. The share of tasks in each sector that are exposed by AI, under current and expanded capabilities (average across G7 economies). Aggregated from tasks and occupations to sectors. Baseline exposure relies on the main task-level measure of Eloundou et al. (2024), while Expanded capabilities rely on the task-level numbers in Eloundou et al. (2024) that are obtained when AI is combined with additional software.

Source: "Macroeconomic Productivity Gains from Artificial Intelligence in G7 Economies" (p. 14), by F. Filippucci, P. Gal, K. Laengle, and M. Schief, 2025, OECD Artificial Intelligence Papers. (https://doi.org/10.1787/a5319ab5-en)

Also, financial analysts, legal professionals, and software developers are occupations that require information processing activities that AI systems can replicate. It is not surprising that Large Language Models (LLMs) can now draft reports, summarize case law, and produce computer code. The deep potential for integration is exactly why these roles are vulnerable to transformation or even displacement.

In contrast, like agriculture, mining, and construction experience much lower AI exposure, somewhere in the range of 10% and 30% (Filippucci et al., 2025). Not surprisingly, these jobs are mainly physical and rely on sensory inputs in real-time. For example, harvesting fruits and vegetables or navigating a construction site requires complex mobility dependent on a variable environment and the ability to rapidly adapt to changes in conditions and situations.

This asymmetry in exposure leads to the conclusion that the impact of AI on the labor market will be highly heterogeneous. While highly exposed sectors will experience major structural changes from the number of tasks redesigned to the reconfiguration of jobs and skills, lower exposed sectors will probably experience a limited disruption. Connecting this to productivity, we can also expect gains in economic benefits to be concentrated in highly exposed sectors. These profit and income gains will probably also accentuate inter-sectoral income and employment discrepancies.

To some extent, uneven exposure raises questions about how the government and other policymakers should respond to the undesired need to reallocate workers to more productive roles in the economy, as well as support their education and income support. Workers in the high-exposed sectors will need significant reskilling to help them transition to roles that augment that technology, while workers in low-exposed sectors may see greater job security, but lower wage growth on account of the appearance of productivity-enhancing technology. As such, understanding inter-sectoral exposure is particularly relevant for assisting the design of targeted policy and efficient labor market policy interventions.

2.2 Defining and Measuring AI Adoption: From Tools to Transformation

To assess how AI affects employment and productivity, it is vital to understand what AI adoption is. In this thesis, occasional and isolated use of AI tools by employees does not constitute AI adoption, we will refer to AI adoption as the structural and strategic integration of this technology into the core processes of a firm (Filippucci et al., 2025). Excluding occasional and isolated use is important because only widespread use of this technology is likely to trigger organizational restructuring and task reallocation, which in turn are measurable impacts on the labor market.

Therefore, using AI tools like chatbots or recommendation engines occasionally for responding to emails, as an example, would not qualify as full adoption. Differently, automating a firm's credit-risk assessment framework using AI systems, or introducing an AI-driven screening and filtering system to optimize the hiring processes, is AI adoption, because it changes workflows, reduces or shifts labor inputs, and redefines organizational capabilities. These subtle differences are critical to distinguishing real transformation from superficial technological experimentation and ultimately assess macroeconomic implications.

Firm-level data on AI adoption remain fragmented globally. Business surveys conducted by national statistical offices including Statistics Canada, Eurostat, ISTAT, INSEE, and ONS show that adoption can vary widely depending on the country and the size of the firms. In the European Union, around one in five businesses with at least ten employees reported using AI in 2025, but among large firms (>250 employees) the share exceeded 50% (Eurostat, 2025). National surveys reflect these differences: adoption stood at 16.4% and 53.1% for larger firms in Italy (ISTAT, 2025), 10% and 33% in France and 42% for firms in ICT (INSEE, 2025), 23% in the United Kingdom (ONS, 2025), while Canada reported 12.2% in 2025 with a peak of 35.6% for information and culture industries (Statistics Canada, 2025). Germany stands out with higher reported penetration, at around 40.9% of firms in 2025 with peaks of 84.3% for companies in advertising and market research and 73.7% for the ICT sector (ifo Institute, 2025). Across all sources, adoption is consistently more prevalent in data-intensive industries such as ICT, finance, and telecommunications where the potential returns are more immediate and tangible.

Data highlights that adoption is not only not distributed equally within the economy, but it also depends on firms' size, digital infrastructure and access to skilled labor. Even within AI-using companies, most rely on off-the-shelf products that have little organizational impact. Only a minority of companies move from experimenting to reengineering internal systems in response to AI capabilities, as OECD denotes as "transformational" adoption (Filippucci et al., 2025). Moreover, this kind of adoption depends on complementary investments in cloud infrastructure, cybersecurity and education. These investments constitute a significant entry barrier for small companies with low margins.

AI's productive revolution depends greatly on how intensively companies incorporate this technology. Superficial integration will result in marginal benefits and a meager reshuffling of the labor market. So, future research needs to clearly define where the line between superficial adoption and structural integration is drawn for it to capture AI's impact on the economy in any substantive manner.

2.3 Patterns of AI Adoption and the Labor Market Diffusion Curve

Adoption rate curves of GPTs are generally S-shaped, meaning slow early adoption is usually concentrated around a handful of innovators and technology early adopters, followed by mass adoption driven by positive results such as productivity increments, improvements in functionality, ease of use, and lower costs. Eventually, diffusion rates reach a plateau as technology reaches maturity and penetration (Geroski, 2000; Hall, 2006; Rogers, 1983). Historical examples of this diffusion are electricity, personal computers, mobile phones, and the internet, all of which have revolutionized the economy and the labor market. Electricity and computers reached 23% and 40% adoption rates, respectively, ten years after their commercial breakthrough in the American market, while mobile phone use reached over 60% penetration in the same period (Comin & Mestieri, 2018; Tankwa et al., 2025).

Early evidence suggests that AI may follow a similarly shaped, but significantly steeper, adoption curve. Firm-level estimates pointed to growth rates of around 50% year-on-year from a low baseline of 3–5% adoption in the early 2020s (Filippucci et al., 2025). More recent official surveys reinforce this pattern: in the European Union, adoption among enterprises with at least ten employees rose from 8% in 2023 to 13.5% in 2024 and 19.95% in 2025, with 50.03% among large firms (Eurostat, 2025). National data show similar momentum: Italy increased from 5% to 8.2% and 16.4% in 2025, France from 6% to 10% between 2023 and 2024 (INSEE, 2025; ISTAT, 2025), Canada nearly doubled from 6.1% in mid-2024 to 12.2% in mid-2025 (Statistics Canada, 2025), and the United Kingdom reported 9% adoption in 2023, which increased to 23% in 2025 (ONS, 2025). Germany, too, stands out with around 40.9% of firms using AI in 2025 up from 27% in 2024 (ifo Institute, 2025). Taken together, these figures indicate that AI adoption may have already entered the steep phase of diffusion, particularly among large and data-intensive enterprises, echoing the trajectory of past GPTs.

Moreover, the cost of AI systems is declining rapidly. Recent forecasts suggest that AI-construction and deployment expenses have declined by as much as 80% in the past couple of years alone, as open-source models are enhanced, hardware efficiency upgrades improve, and computing infrastructure is accessed over the cloud (André et al., 2025). This reduction in costs resembles previous computing and storage cost reductions that were central to the widespread adoption of earlier technologies.

Despite these accelerating trends, high adoption might not translate into immediate and homogenous impacts on employment. As highlighted before, the integration of AI into core processes requires complementary investments into data quality, digital infrastructure, cybersecurity, and AI-specific expertise; otherwise, firms may remain at the "experimentation" stage of adoption using AI only peripherally (Brynjolfsson et al., 2020). On one hand, big firms with established digital infrastructures are generally quicker to adopt impactful AI applications, while Small and Medium Enterprises (SMEs) often lag behind because they lack the scale and ability to implement new technologies. This phenomenon can potentially widen already large productivity gaps across firms. Sector is also an important variable: finance, telecom, and ICT are pioneers in adopting AI because of their data-rich environments and relatively repetitive operations, while sectors such as construction, agriculture, and mining face more physical and coordination-related constraints.

Another key factor affecting the adoption-employment relationship is how AI is implemented. Data suggest that in most cases, AI is augmenting or replacing some tasks within jobs; this leads to job functions being altered rather than jobs being displaced entirely. Workers may spend less time on repetitive and administrative tasks, leaving more time to focus on more complex and creative tasks requiring judgment, communication, or oversight.

In conclusion, while AI has a higher rate of diffusion compared to previous GPTs, the economic and labor market consequences depend not only on diffusion, decreasing costs, and task-level integration but also on the firms' absorptive capacities, workers' and institutions' adaptability.

2.4 Implications for Labor Demand and Occupational Change

Given the differences in AI adoption and exposure between firms and sectors, the effects on the labor market will be multi-dimensional. AI will transform the workforce both through the automation of occupations and skills and the creation of new jobs and skill requirements.

AI exposure will increasingly determine the needs of the job market. Cognitive, analytic, and administrative jobs are most exposed; jobs that involve systematic processing of data, routine and repetitive decision-making, rule-based workflows, and data input are most likely to be partially or fully automated by AI. These tasks represent exactly the technical capabilities that predictive and generative AI models excel at. In sharp contrast, jobs that are more reliant on skills such as emotional understanding, dexterity, context-sensitive thinking, like personal care, social services, or craftsmanship, are less exposed (Felten et al., 2021).

Another important aspect is AI adoption; this might cause polarization within the job market, increasing the gap between high-skilled and low-skilled workers. Jobs that are more technical and require professional skills stand to gain from AI, increased productivity and the possibility of growing demand. Careers where AI helps in diagnosing, modelling, summarizing, and writing without fully taking over human discretion, like software developers, doctors, project managers, and legal professions might experience gains. For example, existing proof indicates that lawyers/judges are extremely exposed to AI while also being extremely complementable, thus unlikely to be replaced but rather complemented to any significant extent (Cazzaniga et al., 2024).

On the other hand, jobs that are centered around repetitive tasks, especially in the middle-skilled levels like entry-level office work, accounting, and customer services, are more likely to experience reshaping and even elimination. This is in line with past trends that were observed in past waves of automation, where repetitive, codable tasks were the first to be replaced by software or machines (Lassébie & Quintini, 2022). Different from before, though, is that automation is taking place “up the skill ladder” affecting high-skilled jobs that were previously considered immune.

Really important is how employers choose to integrate this technology. If deployment of AI is done complementarily, for example, with decision aid tools, individualized learning plans, or extended diagnostic tools, then employment can remain stable or increase in some industries, since individuals are made more efficient and shifted upwards. Here, AI serves the role of a force multiplier and not a labor displacer.

Instead, if businesses pursue cost-minimizing strategies, using AI models to replace workers, then the displacement effect can be greater and more immediate. As an example, customer service is increasingly using AI chatbots and voice interfaces for simple interactions, thus decreasing the need for human operators and at the same time not creating new corresponding jobs.

Adequate policy responses, such as investments in education, active labor market programs, and social protection, are vital, as without these, many displaced workers might experience prolonged periods of unemployment. This is particularly concerning within highly exposed industries with limited channels of reskilling or mobility. The distributional effect of change promoted by AI can exacerbate inequality not just between regions and occupations but also between labor and capital, as those businesses that can employ AI at full efficiency acquire disproportional productivity and profit leads (Aghion et al., 2017).

Summarizing, the effect of AI adoption on the labor market depends greatly on firm strategies, sectoral characteristics, and on how public institutions respond. Whether AI will be a complementary tool augmenting productivity or a displacer will depend on the institutional frameworks, firm-level deployment strategies, and policy interventions that shape how this technology is integrated into the workplace.

2.5 Demand Reallocation and AI-Induced Shifts in Sectoral Employment

While much of the focus on AI and how it will impact the labor market was put on task automation and substitution, equally important are the indirect effects AI has on the demand side of the economy through changes in consumption patterns and sectoral reallocation of labor. More in depth, these effects are caused mainly by increases in productivity, which in turn lower the relative price of goods and at the same time increase real wages, especially in the more exposed sectors. Consequently, both households and firms adjust their expenditure accordingly in ways that can reshape the structure of employment across the economy (Aghion et al., 2017).

At the center of this phenomenon is a modern reinterpretation of Baumol's cost disease, the tendency for wages in jobs that have experienced little or no increase in labor productivity to rise in response to rising wages in other jobs that did experience high productivity growth. In turn, these sectors of the economy become more expensive over time, because the input costs increase while productivity does not (Baumol, 1967). Typically, this affects sectors that require human presence, like services more than manufactured goods, and in particular health, education, arts, and culture. When extended to the context of AI, the concern is that high productivity growth industries such as ICT, finance, and business services can shift consumption and demand for workers toward non-automatable sectors, low productivity growth sectors, like education, healthcare, and hospitality services. These sectors are generally characterized by labor-intensive tasks that require human interaction and are irreplaceable by current AI technology. Accordingly, where aggregate productivity is larger in the short period, the growth effect is dampened in the longer period if labor is redirected toward structurally low-growth industries (Bessen, 2018).

This reallocation may not necessarily translate into a net negative total employment growth; it could remain unchanged or even grow. However, what is important is the nature of new employment. Employment in lower productivity growth industries may have lower rates of pay, career development opportunities, and work conditions than employees in higher productivity growth sectors. So, even if the general level of employment is stagnant or even increasing, employment quality and income distribution may fall, and thus increase labor market segmentation and inequality.

Moreover, this demand reallocation has different effects in the short and long run. For instance, in the short run, lower prices in combination with higher wages in higher productivity sectors can increase discretionary spending, resulting in increases in the demand for goods and services outside these sectors. This increase in demand may improve employment and wages in lower productivity sectors. However, in the long run, without any improvements in these sectors' productivity growth, it could lead to a structural trap where an ever-growing share of the workforce migrates and concentrates towards low productivity and stagnating sectors.

Past precedents reveal more about this pattern. Transitioning from industrial-centred to service-based economies in the past four decades resulted in a steady brake on developed economies aggregate productivity growth (Stéphane Sorbe et al., 2018). This trend could be worsened by AI, creating a paradox: while AI can significantly boost productivity within a few industries, it can end up driving the growth of low-productivity jobs and therefore nullify its long-run macroeconomic benefit.

Another important aspect to consider is the elasticity of demand in AI-exposed sectors in the long run. Sectors like finance or law, considered high productivity growth, can realize productivity gains that in turn reduce costs and thus increase demand for those services and consequently employment, creating a sort of virtuous cycle that, after a certain point, reaches saturation and additional employment stalls. If demand in these sectors is inelastic, then even substantial cost reductions will not lead to equivalent employment growth. Instead, gains in disposable income might be redirected towards those sectors with lower productivity and consequently increase employment in these low-productivity sectors.

The speed of AI innovation introduces another layer of complexity. If institutions cannot match the speed of AI development like the educational system, workers may be increasingly pushed towards lower productivity sectors, giving rise to even higher polarization between industries. In this case, even if total employment remains resilient, real wage stagnation and the erosion of middle-income jobs are an increasing possibility.

In conclusion, the long-term effect that AI will have on the labor market also depends on whether productivity gains in highly exposed sectors are matched by employment growth in highly productive sectors. Without such a match, an extensive part of the workforce risks being stuck in low-growth, low-productivity occupations. Proactive policy frameworks are vital for addressing this problem, investments in education, training, tax, and welfare policies are needed to transition workers efficiently across sectors.

Chapter 3 Empirical Analysis: AI Exposure and Employment Dynamics

3.1 Research Question and Empirical Strategy

This chapter focuses on the relationship between AI exposure and employment growth in the US across occupations using two comparative periods. The objective is to test whether more exposed occupations display different employment dynamics relative to the less exposed ones.

The empirical analysis uses AI exposure scores along with employment data from the U.S. Bureau of Labor Statistics, from which employment growth is calculated using the logarithmic difference in total occupational employment between the end and start of the two time frames:

2009-2014, representing an earlier phase with limited AI diffusion in production environments

2019-2024, representing a later phase characterized by broader AI capability and adoption

This two-period structure allows for a comparative assessment of whether the exposure-growth relationship strengthens, weakens, or remains stable across time. If AI exposure plays a meaningful role in shaping employment dynamics, one would expect the exposure coefficient, dispersion patterns, or group differentials to change substantially between the earlier and later periods.

Occupational AI exposure is measured using an externally constructed task-based exposure index mapped to SOC (Standard Occupational Classification) occupations. This measure captures the degree to which tasks performed within an occupation are technically exposed to AI capabilities. Importantly, it does not measure AI adoption, but rather potential exposure based on task content. The index draws on the methodology developed by Felten et al. (2021), which links AI application capabilities to occupational ability requirements.

The empirical strategy proceeds in several layers:

1. The relationship between AI exposure and employment growth is examined using cross-sectional OLS regressions at the occupation level for each period separately. This provides a baseline measure of association between exposure and growth.

2. Robustness checks are performed using trimmed samples that remove extreme growth outliers (1st-99th percentile range) to ensure that results are not driven by a small number of volatile occupations.

3. Distributional analysis is conducted by grouping occupations into AI exposure quartiles and comparing mean growth across groups. This allows identification of non-linear or non-monotonic patterns that are not captured by a single regression slope.

4. Heterogeneity is examined between SOC major occupation groups using z-scores for AI exposure and employment growth, allowing relative comparisons across broad occupational categories.

3.2 Data Construction and Variable Definition

This section describes how the analytical dataset was constructed and how the main variables used in the empirical analysis are defined.

The unit of observation is the occupation, identified by the SOC code. The final dataset is built by merging occupational AI exposure scores with employment levels from U.S. Bureau of Labor Statistics occupational employment files for two separate periods: 2009-2014 and 2019-2024.

3.2.1 AI Exposure Measure

AI exposure is measured using the AI Occupational Exposure score (AIOE), which assigns each SOC occupation a continuous score reflecting the degree to which its task content overlaps with the capabilities of current AI systems. The index is task-based and derived from task descriptions rather than realized technology adoption. Therefore, it should be interpreted as potential exposure, not realized automation.

The exposure variable is treated as continuous in the regressions and is also used to construct exposure quartiles. Because the index is already approximately standardized in its source construction, no further scaling is applied beyond numeric cleaning.

A methodological note on temporal application: the AIOE index was constructed based on AI capabilities as of approximately 2018-2021. It is applied to both the 2009-2014 and 2019-2024 periods in this analysis since the index captures the structural task content of occupations with little to no variance over time. The implicit assumption is that occupations are comprised of more or less the same tasks between the two time frames, even though AI capabilities have evolved at an astounding rate. This assumption is reasonable given the stability of occupational task structures, but it warrants caution: the index does not measure actual AI exposure in the earlier period; rather, it identifies occupations whose task profiles became relevant to AI capabilities.

3.2.2 Employment Measures and Growth Definition

Occupational employment levels are taken from national occupational employment datasets for the years 2009, 2014, 2019, and 2024. Only occupations with valid positive employment in both endpoints of each period are retained because employment growth is measured in logarithmic form.

Employment growth is defined as:

$$Growth\ = log(Employment_{end}) - log(Employment_{start})$$

The logarithmic difference was used because of three main advantages: it approximates percentage growth, it reduces scale distortions across large versus small occupations, and it makes coefficients directly interpretable. Since the dependent variable is measured in log points, the regression coefficient β represents the change in log employment growth associated with a one-unit increase in the AI exposure index. A coefficient of 0.04 means approximately 4 log points.

Observations with zero or invalid employment values are excluded before log transformation to avoid undefined values. This implies that occupations that disappear entirely between endpoints are not included. The sample includes 637 occupations for 2019-2024 and 688 for 2009-2014. The difference of 51 occupations reflects changes in SOC classification, data availability, and occupational definitions across periods. This difference should be kept in mind when interpreting cross-period comparisons, as the later sample excludes some occupations present in the earlier period.

3.2.3 Sample Cleaning and Trimming

After merging exposure and employment data, rows with missing data are dropped along with negative employment values and duplicates, and numeric formatting inconsistencies are corrected.

3.3 Baseline Regression Results

This section presents the baseline regression results linking occupational AI exposure to employment growth. The regression is done for both periods in order to assess whether the relationship is stable across time and across different macroeconomic environments.

The estimated specification relates occupational log employment growth to the AI exposure score at the SOC occupation level using ordinary least squares:

$$Growth_i = \alpha + \beta \times AI_{Exposure_i} + \varepsilon_i$$

where $Growth_i$ is the log employment change for occupation i, $AI_{Exposure_i}$ is the occupational exposure score, α is the constant term, β is the coefficient of interest measuring the association between exposure and growth, and ε_i is the error term.

3.3.1 Period 2019-2024

For the 2019-2024 period, the regression is estimated on 637 occupations after cleaning and filtering. The estimated coefficient on AI exposure is positive and statistically significant.

Baseline regression results:

Statistic	Value
β (AI exposure)	0.042
p-value	< 0.001
R^2	0.031
N observations	637

After trimming the top and bottom 1% of the employment growth distribution to reduce the influence of extreme outliers, the relationship becomes stronger:

Trimmed regression results (1%-99%):

Statistic	Value
β (AI exposure)	0.043
p-value	< 0.001
R^2	0.042
N observations	623

A one-unit increase in the AI exposure index is associated with roughly 4 log points higher employment growth over the period. However, explanatory power remains low. Even in the trimmed sample, AI exposure explains less than 5% of the cross-occupation variation in employment growth. Low R^2 is expected in cross-occupation regressions because employment growth is determined by a multitude of factors such as demand, business cycle, demographic shifts, macroeconomic shocks, and AI exposure is only a factor among many others. The result indicates that while the relationship is statistically detectable, it is economically modest and far from dominant.

Looking at the scatter plots for both periods, we can observe a remarkably high dispersion that indicates substantial heterogeneity: many occupations experience both positive and negative employment growth regardless of exposure. The trend is slightly positive, but the cloud of observations is wide, highlighting again that AI exposure is one of many factors impacting employment growth.

Figure 2 AI Exposure vs Employment Growth Scatterplot (2019-2024)

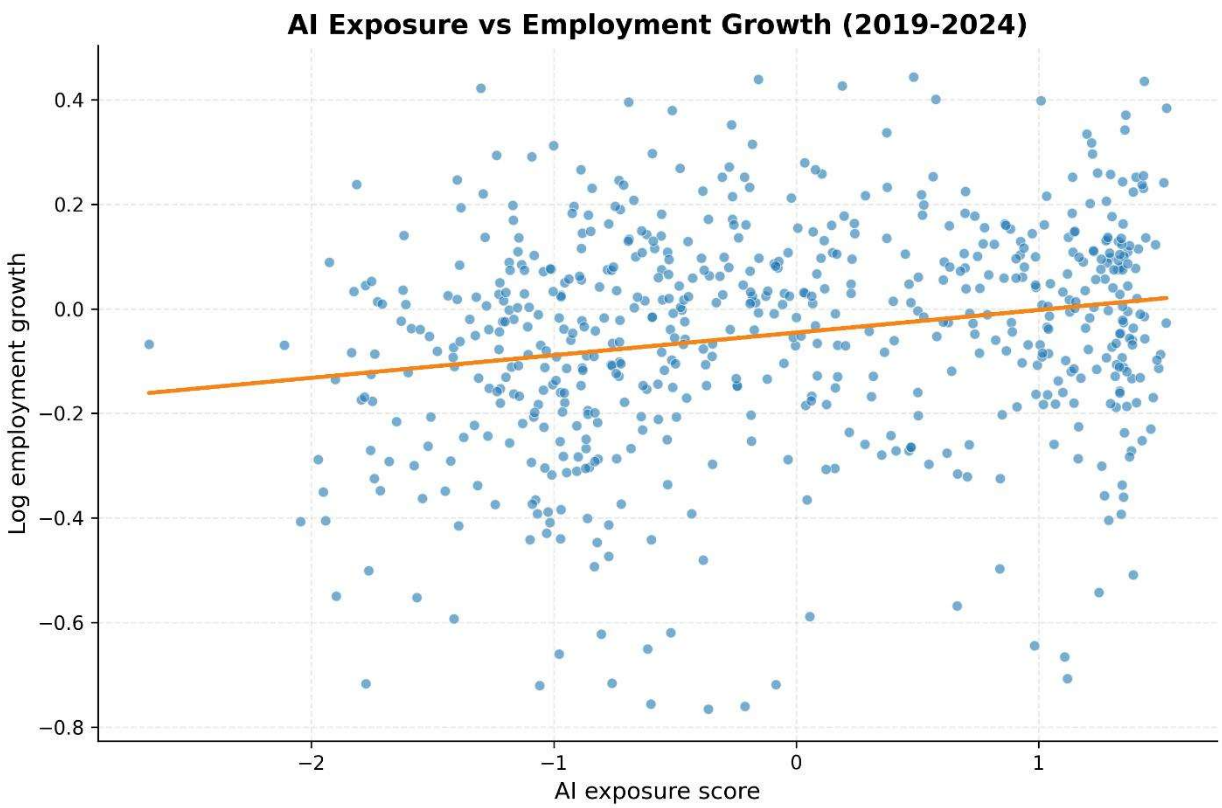

Note: Each point represents one SOC occupation. The fitted line shows the OLS regression relationship. Extreme outliers trimmed at the 1st and 99th percentiles of employment growth.

Source: Author's calculations based on BLS Occupational Employment Statistics and Felten et al. (2021).

3.3.2 Period 2009-2014

For the earlier 2009-2014 period, the regression is estimated on 688 occupations. The estimated coefficient on AI exposure is again positive and statistically significant.

Baseline regression results:

Statistic	**Value**
β (AI exposure)	0.035
p-value	< 0.001
R^2	0.029
N observations	688

Trimmed regression results (1%-99%):

Statistic	**Value**
β (AI exposure)	0.038
p-value	< 0.001
R^2	0.046
N observations	674

The magnitude of the coefficient is similar to the later period (β ≈ 0.035-0.043).

Figure 3 AI Exposure vs Employment Growth Scatterplot (2009-2014)

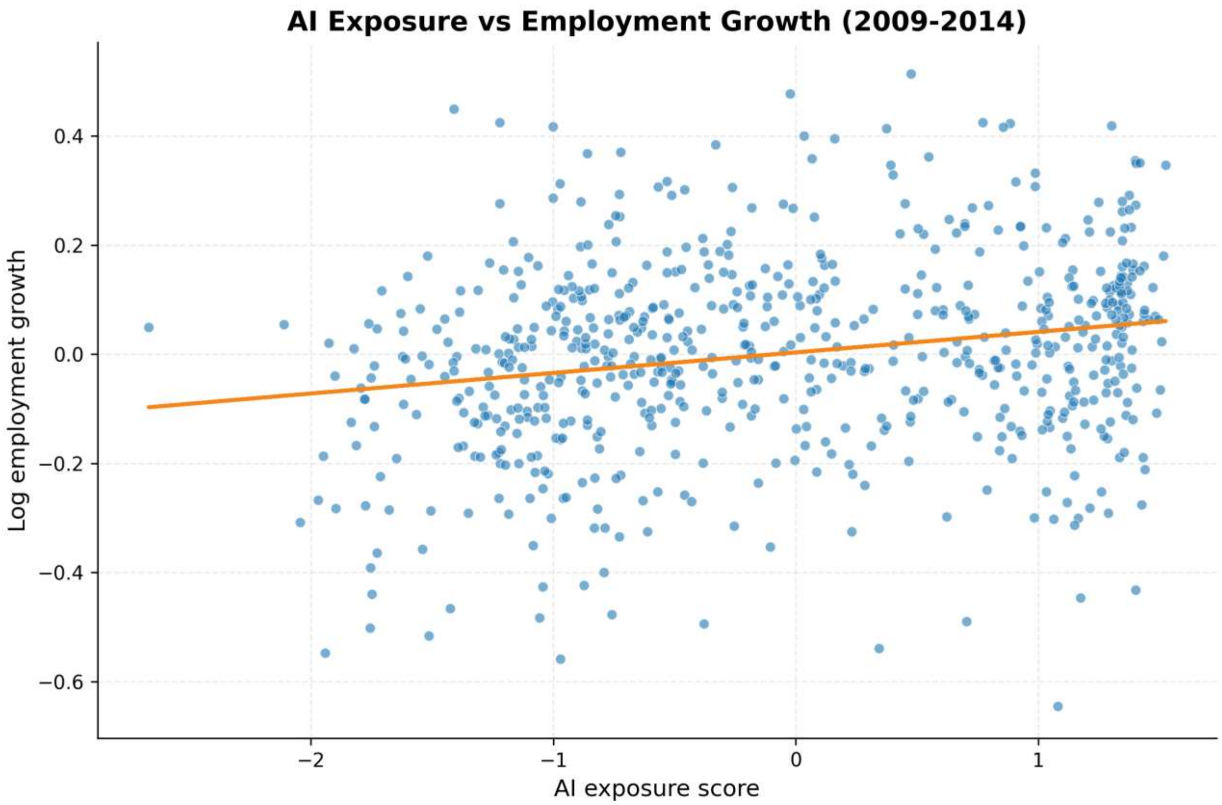

Note: Each point represents one SOC occupation. The fitted line shows the OLS regression relationship. Extreme outliers trimmed at the 1st and 99th percentiles of employment growth.

Source: Author's calculations based on BLS Occupational Employment Statistics and Felten et al. (2021).

3.3.3 Cross-Period Comparison and Interpretation

The similarity of coefficients across a low-AI-diffusion period (2009-2014) and a high-AI-diffusion period (2019-2024) is a central result. It suggests that the exposure index primarily captures structural occupational characteristics, cognitive task intensity, digital information processing, and education-intensive skill profiles, rather than recent AI adoption shocks. In particular, the AI exposure index is correlated with other occupational characteristics, including average education levels and mean wages, and the coefficient β cannot be interpreted as the marginal effect of AI exposure without controlling for these factors. The results should therefore be interpreted as associational, not causal.

3.4 Distributional Patterns by AI Exposure Quartiles

While the baseline regressions estimate an average linear relationship between AI exposure and employment growth, they do not reveal whether the relationship is monotonic across the exposure distribution or whether certain ranges of exposure are associated with different employment dynamics. To examine potential non-linearities, occupations are grouped into quartiles based on their AI exposure scores (Low, Mid-Low, Mid-High, High), and average employment growth is compared across these groups for each period.

3.4.1 Quartile Analysis: 2019-2024

For the 2019-2024 period, mean log employment growth by exposure quartile is shown in Table 1.

Table 1 Employment Growth by AI Exposure Quartile (2019-2024)

Quartile	Mean Growth	N Occupations	Z-Score
Low	-0.1148	160	-0.28
Mid-Low	-0.0536	159	-0.02
Mid-High	-0.0107	159	0.16
High	-0.0132	159	0.14

Note: Quartile boundaries defined by the distribution of AI exposure scores across all occupations in the 2019-2024 sample (N = 637). Mean growth is the unweighted average log employment growth within each quartile. Z-scores are standardized relative to the overall occupation-level mean (−0.049) and standard deviation (0.412). All quartile means are negative, reflecting overall weak occupation-level employment performance in the period; z-scores capture relative performance within the distribution.

Source: Author's calculations based on BLS OEWS (2019 and 2024) and Felten et al. (2021) AI Occupational Exposure Index.

Two features stand out.

First, there is a strong improvement moving from the Low exposure group to the middle of the distribution. The Low quartile has substantially more negative employment growth than all other quartiles, Low exposure occupations show a z-score of −0.28, while Mid-High exposure occupations show a z-score of +0.16.

Second, the pattern is not strictly monotonic at the top of the distribution. The Mid-High quartile performs better than the High quartile, even though exposure levels are higher in the High quartile by construction. This suggests that the most exposed occupations do not necessarily experience the strongest employment performance. Instead, the best-performing group is “mid-high exposure,” not “highest exposure.”

A cautious interpretation consistent with this pattern is that the relationship between exposure and employment outcomes may reflect two opposing forces:

Complementarity: occupations where a substantial portion of tasks can be replicated by AI systems but where some bottlenecks still remain, here AI may function as a productivity-enhancer that supports labor demand. This task-level complementarity allows AI systems to handle the “boring”, repetitive, and information-intensive tasks, leaving workers more time to focus on more important aspects of the job, like decision-making and other activities that require judgment and creativity (Agrawal et al., 2019; Fossen & Sorgner, 2019). The mid-high quartile likely contains occupations where this augmentation dynamic is strongest.

Substitution risk: occupations with very high exposure, where most tasks fall under AI’s current capabilities, the displacement effect could offset the complementarity benefits. Occupations comprised of predominantly codifiable, rule-based tasks, where little to no human inputs are required, such as social reasoning, are rather threatened by substitution than aided by complementarity.

This is not a causal conclusion, but it provides a coherent way to interpret why the exposure-growth relationship is positive while not fully monotonic across groups.

3.4.2 Quartile Analysis: 2009-2014

For the 2009-2014 period, mean log employment growth by exposure quartile is shown in Table 2.

Table 2 Employment Growth by AI Exposure Quartile (2009-2014)

Quartile	Mean Growth	N Occupations	Z-Score
Low	-0.0653	172	-0.32
Mid-Low	0.0196	172	0.09
Mid-High	0.0258	172	0.12
High	0.0232	172	0.11

Note: Quartile boundaries defined by the distribution of AI exposure scores across all occupations in the 2009-2014 sample (N = 688). Mean growth is the unweighted average log employment growth within each quartile. Z-scores are standardized relative to the overall occupation-level mean (0.001) and standard deviation (0.206). Unlike the 2019-2024 period, the overall mean growth is near zero, so z-scores directly reflect whether quartile growth is above or below the period average

Source: Author's calculations based on BLS OEWS (2009 and 2014) and Felten et al. (2021) AI Occupational Exposure Index.

The pattern in this earlier period is again broadly increasing from Low to higher exposure groups, with the same non-monotonic feature: Mid-High slightly exceeds High.

These results imply that the "bottom quartile disadvantage" is not specific to the later period, and the exposure-growth relationship may not be purely linear even outside the recent AI diffusion phase.

3.4.3 Interpretation of Non-Linear Patterns

Taken together, the quartile analysis refines the regression results in two ways.

First, the positive regression coefficient does not mean "more exposure always implies higher growth." The main difference is between low-exposure occupations and the rest of the distribution. In both periods, the low exposure quartile performs substantially worse than the middle and higher exposure quartiles.

Second, the repeated observation that Mid-High exposure performs as well as or better than High exposure suggests that the strongest association is concentrated in the middle-to-upper part of the exposure distribution rather than at the extreme top. This is consistent with the task-based theoretical framework reviewed in Chapter 2: mid-high exposure occupations may retain sufficient task complexity to benefit from AI complementarity, while the very highest exposure occupations may include more tasks vulnerable to full automation. The GPT framework (Felten et al., 2021) predicts exactly this distinction, exposure measures technical feasibility, but actual labor market outcomes depend on task complexity, interdependence with human judgment, and organizational implementation choices.

3.5 Heterogeneity Across Major Occupational Groups

To complement the occupation-level regressions and quartile analysis, results are further divided into SOC major occupational groups.

For each SOC major group, two statistics are computed: average AI exposure score and average log employment growth. Both measures are standardized into z-scores relative to the overall occupation-level mean and standard deviation, enabling direct comparison of how far each occupational family lies above or below the overall average. Figure 4 presents these results in a two-panel visualization, with AI exposure shown in the left panel and employment growth for both periods shown side-by-side in the right panel.

Figure 4 AI Exposure and Employment Growth by SOC Major Group (2009-2014 vs 2019-2024)

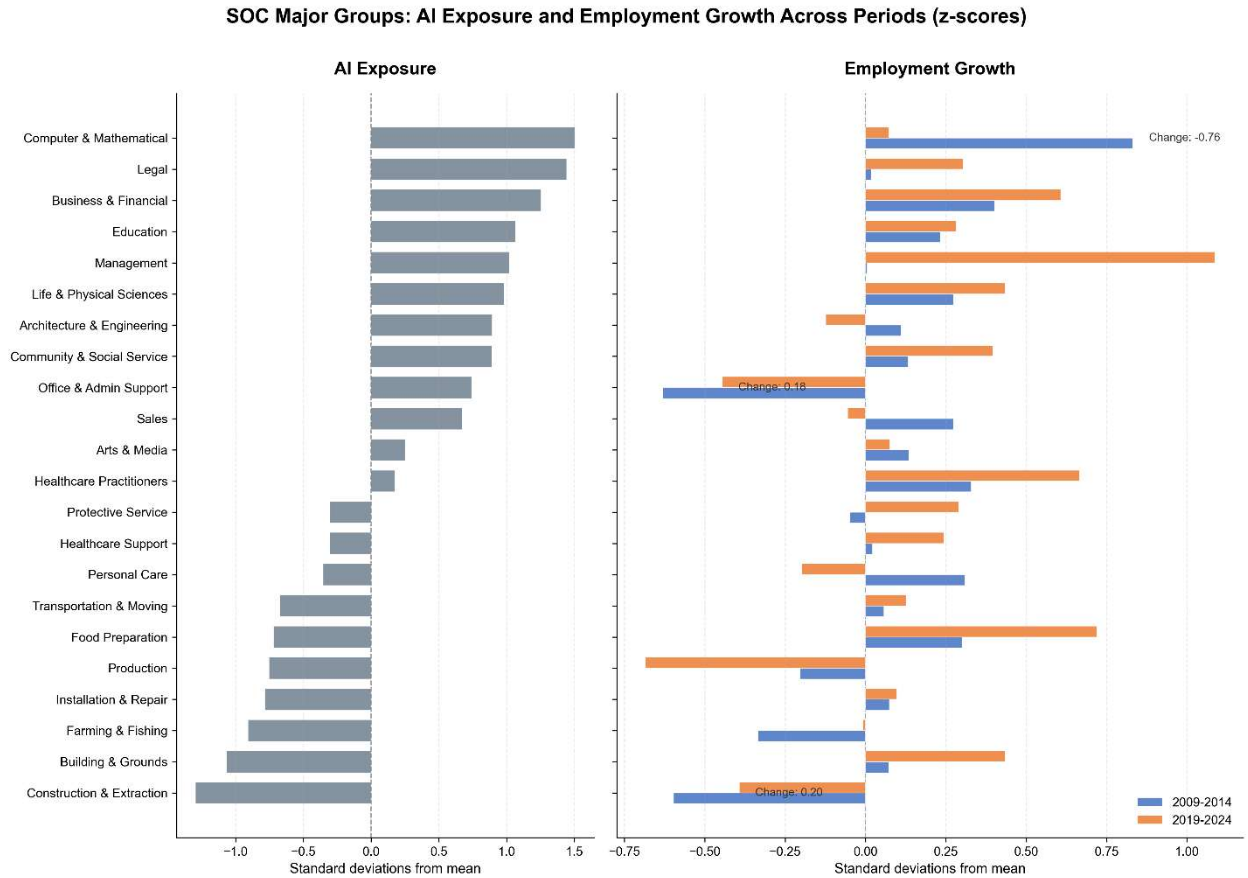

Note: Left panel shows AI exposure z-scores (standardized relative to occupation-level mean and standard deviation). Right panel shows employment growth z-scores for 2009–2014 (blue) and 2019–2024 (orange). Groups sorted by AI exposure level. Annotated changes highlight the largest cross-period divergences.

Source: Author's calculations based on BLS Occupational Employment Statistics and Felten et al. (2021).

Table 3 Selected SOC Major Groups: AI Exposure and Employment Growth (2019-2024)

Occupational Group	**AI Exposure (z)**	**Growth (z)**
Computer & Mathematical	1.51	0.07
Legal	1.44	0.30
Business & Financial	1.26	0.61

Occupational Group	AI Exposure (z)	Growth (z)
Education	1.07	0.28
Management	1.02	1.09
Life & Physical Sciences	0.98	0.43
Architecture & Engineering	0.90	-0.12
Community & Social Service	0.89	0.40
Food Preparation	-0.72	0.72
Installation & Repair	-0.78	0.10
Farming & Fishing	-0.91	-0.01
Building & Grounds	-1.07	0.43
Construction & Extraction	-1.30	-0.39

Note: Table shows a selected subset of SOC major groups, ranked by AI exposure. AI Exposure (z) and Growth (z) are both standardized relative to the occupation-level mean and standard deviation. Full cross-period comparison for all 22 SOC major groups is shown in Figure 4. Growth values reflect the 2019-2024 period only.

Source: Author's calculations based on BLS OEWS (2019 and 2024) and Felten et al. (2021) AI Occupational Exposure Index.

3.5.1 Dramatic Divergence in Employment Growth

Employment growth patterns show dramatic variation across the two periods. The right panel of Figure 4 reveals several striking patterns:

The Computer & Mathematical Decline: In the earlier period, this group showed extraordinary growth of 0.83 SD above the mean in sharp contrast with the more recent period, where growth fell to barely above average (0.07 SD above the mean). This is by far the largest spread in employment growth between the two periods (-0.76 SD). This finding is especially interesting because Computer & Mathematical is the most exposed occupational category, reinforcing the substitution theory. However, it is important to acknowledge that this major group contains a small number of observations in both periods (5 in 2009-2014 and 4 in 2019-2024), so the results may be driven by the trends of individual occupations rather than reflecting a broad-based decline across all computing-related roles. Detailed occupation-level decomposition would be needed to confirm the robustness of this finding.

Widespread Decline in High-Exposure Technical Groups: Among other occupations that experienced strong declines: Architecture & Engineering declined by −0.23 (from +0.11 to −0.12), and Sales by −0.22 (from +0.27 to +0.05). Looking closely, we find that most occupations that experienced employment growth declines are technical in nature, suggesting again that the substitution effect could be more prevalent in those occupations where tasks are repetitive, routine, and information-intensive.

Mixed Performance Among High-Exposure Professional Groups: In sharp contrast, Legal occupations improved by +0.28 (from +0.02 to +0.30), Management by +1.08 (from 0.00 to +1.09), and notably, Office & Admin Support, despite moderate AI exposure, showed one of the largest improvements (+0.42). This heterogeneity within high and mid-high exposure groups reinforces the idea that exposure is not the single determinant of employment trajectory.

Improvement in Low-Exposure Manual Occupations: Conversely, several low-exposure occupational groups showed substantial improvement. Construction & Extraction improved by +0.21 (from −0.60 to −0.39) and Building & Grounds by +0.36 (from +0.07 to +0.43). Even if these groups show below-average employment growth, their relative performance improved significantly between the two periods. One explanation of this phenomenon is Moravec's Paradox, already discussed in earlier chapters. Occupations within these groups require physical presence, sensorimotor coordination, and environmental adaptability, all skills that may seem easy for humans but require enormous computational resources for them to be replicated by AI systems. This provides a structural guardrail for labor demand in these sectors while cognitive tasks face increasing automation pressure.

Taken together, these patterns carry several implications. The Computer & Mathematical decline resonates with the theoretical prediction discussed in Section 2.2 that AI extends automation "up the skill ladder" to affect high-skilled occupations previously considered immune (Aghion et al., 2017; Nedelkoska & Quintini, 2018). The heterogeneity within high-exposure groups is equally telling: Legal and Management improved despite high exposure, consistent with the strong complementarity potential identified for these professions in Section 3.4 (Cazzaniga et al., 2024), while Architecture & Engineering and Sales declined. This confirms that task composition, the balance between complementable and substitutable tasks, and sector-specific demand dynamics all mediate outcomes.

In summary, AI exposure is structurally stable across periods while employment growth is highly variable, confirming that the exposure-growth relationship is context-dependent rather than deterministic.

3.6 Industry-Level AI Exposure Mapping

In addition to occupation-level exposure measures, the analysis incorporates industry-level AI exposure scores based on the NAICS classification (AIIE index). This section is descriptive and structural in purpose: it is not used in the regression model but serves to map where AI exposure is concentrated across industries and to support interpretation of the occupation-based results. Specifically, the industry mapping provides cross-validation of the occupation-level findings by examining whether the same task-structure patterns that drive occupational exposure rankings are also reflected at the industry level.

The industry dataset contains standardized exposure scores (AIIE), already centered approximately at zero and scaled in standard deviation units. Industries are ranked and divided into quartiles.

3.6.1 High-Exposure Industries

The highest exposure industries are concentrated in:

- Securities and financial investment activities
- Accounting, tax, and bookkeeping services
- Legal services

- Insurance and investment funds
- Software publishing
- Financial intermediation and related services

Top AIIE scores are around +2 standard deviations, indicating strong alignment between industry task structure and AI-replicable capabilities. These industries share common characteristics: high cognitive task intensity, strong information processing components, formalized rule-based analytical workflows, and digital data environments. The concentration of high exposure in finance, legal, and professional services aligns closely with the KIS analysis presented in Section 2.1, where these sectors were estimated to have 50-80% of tasks technically exposed to AI capabilities (Eloundou et al., 2024; Filippucci et al., 2025). This convergence of occupation-level and industry-level rankings strengthens confidence that the AIOE index captures genuine structural differences in task content.

3.6.2 Low-Exposure Industries

The lowest exposure industries include:

- Support activities for crop production
- Building and finishing contractors
- Animal processing and slaughtering
- Warehousing and storage
- Rail and water transport support
- Sawmills and wood processing

These industries show AIIE scores near −2 standard deviations. They are characterized by physical task intensity, manual and on-site activity, low digital task share, and high embodiment requirements. This mirrors the occupation-level findings and confirms the theoretical distinction between routine cognitive and routine manual tasks: AI technologies primarily target cognitive and information-processing tasks, while manual tasks requiring physical presence and environmental adaptability remain difficult to automate with current AI capabilities.

Figure 5 AI Exposure by Industry (Top & Bottom 10 NAICS Industries)

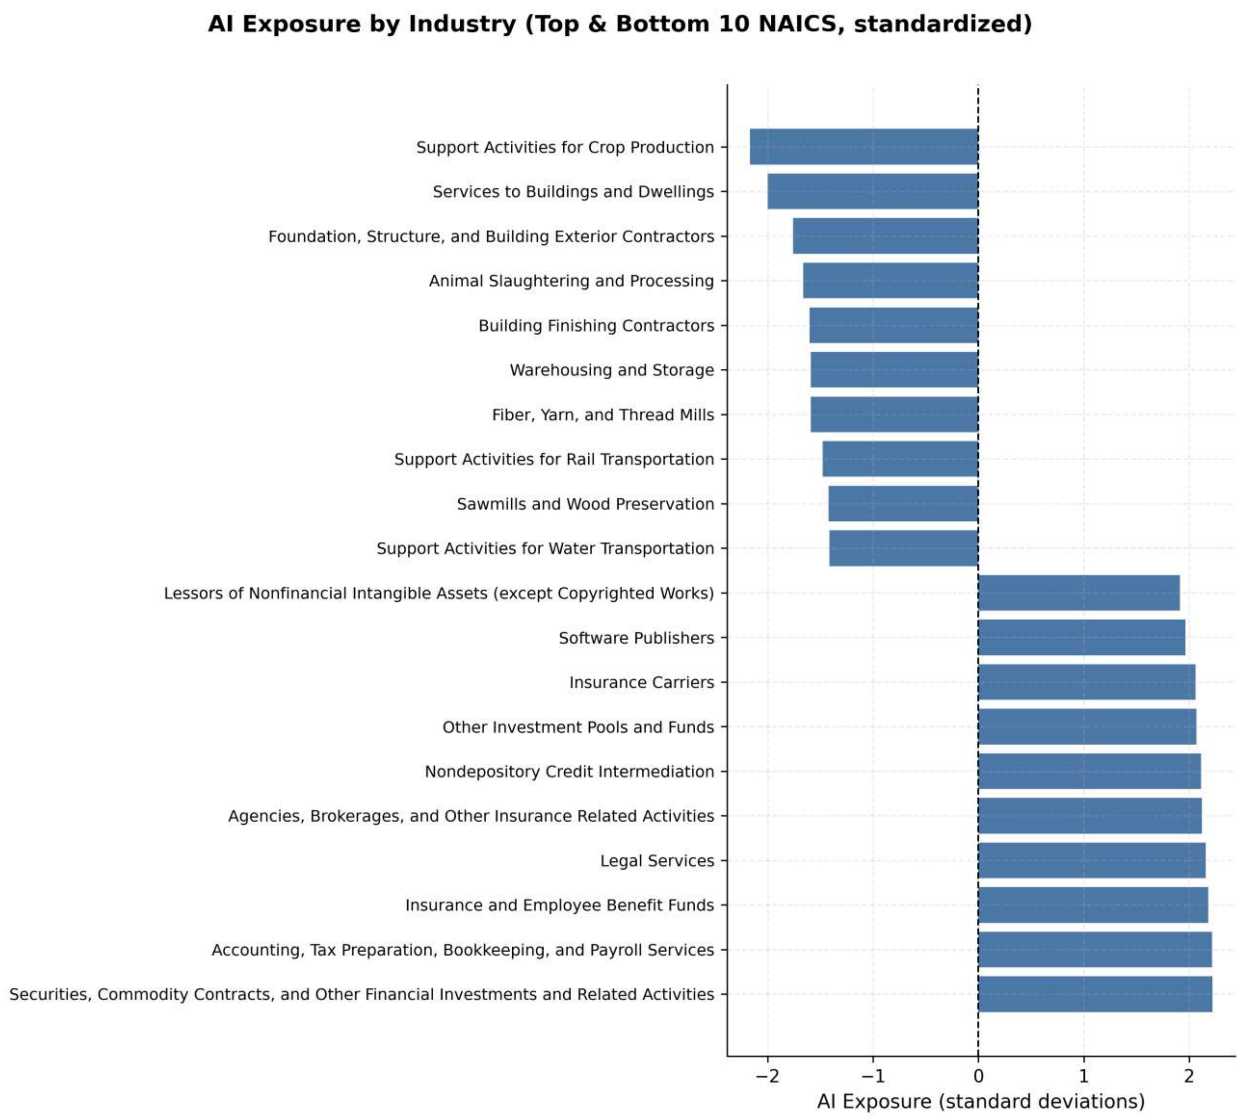

Note: Bars show the AI Industry Exposure Index (AIIE), already standardized in its source construction (mean ≈ 0, standard deviation ≈ 1). Only the 10 highest and 10 lowest exposure industries are shown for clarity.

Source: Felten et al. (2021), NAICS industry-level AIIE scores.

3.6.3 Cross-Level Consistency

The NAICS exposure mapping adds structural depth to the chapter in two ways:

First, cross-level consistency: high exposure appears in both cognitive occupations and knowledge-intensive service industries. This consistency strengthens confidence that the exposure measure captures a real structural dimension rather than noise.

Second, exposure is not confined to "tech sectors": finance, legal, and accounting services rank at the very top, often above software publishing. This supports the argument that AI exposure is task-based, not sector-label-based.

3.7 Integrated Discussion and Interpretation

The empirical analysis yields three major findings: (1) a statistically robust but economically modest positive association between AI exposure and employment growth ($\beta \approx 0.04$ log points) that is stable across both periods; (2) a suggestive non-monotonic pattern where mid-high exposure occupations outperform both low and very high exposure groups; and (3) dramatic heterogeneity across major occupational groups, most notably the −0.76 standard deviation growth decline in Computer & Mathematical occupations.

The stability of the coefficient across a low-AI and high-AI diffusion period (Section 3.3.3) speaks to the distinction between exposure and adoption emphasized in Section 3.2: the AIOE index measures potential technical exposure based on task content, not realized AI deployment. The fact that the coefficient is similar in 2009-2014, when AI adoption was negligible, and 2019-2024, when adoption entered its steep diffusion phase (Section 3.3), suggests that the employment patterns observed here are driven more by structural occupational characteristics than by actual AI implementation.

The Computer & Mathematical Puzzle: Perhaps the most interesting finding is the spread of the employment growth between the two time periods in Computer & Mathematical occupations: the strongest growth in 2009-2014 ($z = +0.83$), declining to barely above average in 2019-2024 ($z = +0.07$). This decline can be attributed to three possible reasons: (1) market saturation following the tech boom, (2) AI systems beginning to automate technical work previously considered immune, or (3) structural changes in the tech sector, including offshoring and consolidation. As stated, the data is based on four observations and should be interpreted with caution.

3.7.1 Limitations and Robustness

Several limitations must be acknowledged:

AI exposure is correlated with many things, such as education levels, wages, and occupational skill requirements, suggesting that the positive coefficient may reflect, in part, skill-biased technological change along with the effect of AI. R^2 values below 5% confirm that AI exposure explains only a part of cross-occupation employment growth. Future models adopting controls for education composition or pre-existing employment trends would help isolate the effect of AI exposure.

Second, both periods capture exogenous shocks that may influence results. The 2009-2014 window captures the final part and subsequently the recovery of the 2007-2009 Great Recession, while the 2019-2024 period includes the COVID-19 pandemic (2020-2021). Both are characterized by uneven sectoral growth and employment polarization. For example, during the pandemic, the hospitality sector suffered a great blow while remote work adoption and digitalization thrived. Alternative time windows might reduce these distortions but would probably introduce other trade-offs. The fact that systematic patterns related to the relationship between AI exposure and employment growth are still noticeable suggests that the underlying relationship is robust.

Despite these limitations, several features strengthen confidence in the core findings. The low-exposure disadvantage is evident in the consistently weaker performance of the bottom quartile relative to all other groups in both periods, reflecting genuine structural differences rather than period-specific noise. The positive sign on the regression coefficient is likewise stable across both periods and robust to trimming. The SOC major group analysis provides additional detail to the findings.

3.7.2 Implications for Theory and Policy

The findings carry three implications for the theoretical frameworks reviewed in Chapters 1 and 2: the positive coefficients found in the regression contradict the popular automation narratives, the non-monotonic patterns found in the quartile analysis confirm that employment growth depends highly on task composition rather than exposure, and the cross-period stability suggests the patterns reflect structural occupational characteristics rather than AI-specific shocks.

From a policy perspective, as discussed earlier in Chapter 3, whether AI will complement or substitute workers depends critically on how firms choose to implement this technology, investments in reskilling, digital infrastructure, and social protection from governments will tip the balance between complementarity and substitution. The finding that AI exposure does not necessarily translate into job destruction reinforces the argument that proactive policy frameworks will pave the way towards augmentation rather than displacement in the labor market.

Summarizing this chapter, we can conclude that there is a modest positive association between AI exposure and employment growth in the US labor market that is stable across time; non-linearity across the distribution and heterogeneity across occupational groups suggest that AI exposure, although increasingly related to employment outcomes, is neither a dominant driver of employment growth nor a predictor of job destruction. The findings point to complex relationships characterized by task composition, complementarity dynamics, structural occupational characteristics, and the policy and regulatory frameworks governing AI adoption and implementation.

Conclusion

This thesis analyzed whether occupational AI exposure is associated with varying employment dynamics across the U.S. labor market. The evidence shows a statistically significant but economically modest positive association between exposure and employment growth, stable across a pre-diffusion (2009-2014) and a post-diffusion (2019-2024) period.

Exposure is related to employment outcomes, but it is neither a dominant driver of growth nor a predictor of job destruction.

These findings connect directly to the theoretical framework developed in the earlier chapters. Chapter 1 framed AI as a general-purpose technology with simultaneous substitution and complementarity effects, where net employment outcomes depend on task composition rather than aggregate exposure. The empirical results confirm this: the non-monotonic distributional pattern and the heterogeneity across occupational groups show that task structure, not exposure level alone, shapes employment dynamics. The resilience of low-exposure manual occupations is consistent with the automation bottlenecks discussed in Chapter 1, including Moravec's Paradox and the limits of replicating sensorimotor and tacit-knowledge tasks. Chapter 2 emphasized the distinction between technical exposure and realized adoption, and the stability of the regression coefficient across two very different AI diffusion periods supports this distinction: the exposure index captures structural occupational characteristics rather than short-run adoption shocks.

From a policy perspective, these findings suggest that the focus should be on shaping how AI is deployed rather than on preventing its diffusion. Investments in reskilling, digital infrastructure, and labor market adjustment mechanisms remain central to steering adoption toward augmentation rather than displacement.

Bibliography

Acemoglu, D., & Restrepo, P. (2018). *Artificial Intelligence, Automation and Work*.

Aghion, P., Jones, B. F., & Jones, C. I. (2017). *NBER WORKING PAPER SERIES*.

Agrawal, A., Gans, J., & Goldfarb, A. (2019). *Economic Policy for Artificial Intelligence*.

Agrawal, A., Gans, J. S., & Goldfarb, A. (n.d.). Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction. *2019*.

Alnuaimi, A., & Albaldawi, T. (2024). An overview of machine learning classification techniques. *BIO Web of Conferences*, *97*, 00133. https://doi.org/10.1051/bioconf/20249700133

André, C., Bétin, M., Gal, P., & Peltier, P. (2025). *Developments in Artificial Intelligence markets: New indicators based on model characteristics, prices and providers* (37th ed., OECD Artificial Intelligence Papers) [OECD Artificial Intelligence Papers]. https://doi.org/10.1787/9302bf46-en

Autor, D. H., Levy, F., & Murnane, R. J. (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

Baumol, W. J. (1967). Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis. *The American Economic Review*, *57*(3), 415–426.

Bessen, J. (2018). *AI and Jobs: The role of demand*.

Bessen, J., Impink, S. M., & Seamans, R. (2018). *The Business of AI-producing Startups: Evidence from a Worldwide Survey*.

Brynjolfsson, E., & McAfee, A. (2017, July 18). The Business of Artificial Intelligence. *Harvard Business Review*. https://hbr.org/2017/07/the-business-of-artificial-intelligence

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, *108*, 43–47. https://doi.org/10.1257/pandp.20181019

Brynjolfsson, E., Rock, D., & Syverson, C. (2017). *Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics* (No. W24001; p. w24001). National Bureau of Economic Research. https://doi.org/10.3386/w24001

Brynjolfsson, E., Rock, D., & Syverson, C. (2020). *The Productivity J-Curve: How Intangibles Complement General Purpose Technologies*.

Cazzaniga, P. M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., & Tavares, M. M. (2024). *Gen-AI: Artificial Intelligence and the Future of Work*.

Cockburn, I., Henderson, R., & Stern, S. (2018). *The Impact of Artificial Intelligence on Innovation*.

Comin, D., & Mestieri, M. (2018). If Technology Has Arrived Everywhere, Why Has Income Diverged? *American Economic Journal: Macroeconomics*, *10*(3), 137–178. https://doi.org/10.1257/mac.20150175

Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future Healthcare Journal*, *6*(2), 94–98. https://doi.org/10.7861/futurehosp.6-2-94

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, *384*(6702), 1306–1308. https://doi.org/10.1126/science.adj0998

Eurostat. (2025). *Use of artificial intelligence in enterprises*. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, *42*(12), 2195–2217. https://doi.org/10.1002/smj.3286

Filippucci, F., Gal, P., Jona-Lasinio, C., & Nicoletti, G. (2024). *The impact of Artificial Intelligence on productivity, distribution and growth: Key mechanisms, initial evidence and policy challenges* (OECD Artificial Intelligence Papers No. 15; OECD Artificial Intelligence Papers, Vol. 15). https://doi.org/10.1787/8d900037-en

Filippucci, F., Gal, P., Laengle, K., & Schief, M. (2025). *Macroeconomic productivity gains from Artificial Intelligence in G7 economies* (41st ed., OECD Artificial Intelligence Papers) [OECD Artificial Intelligence Papers]. https://doi.org/10.1787/a5319ab5-en

Fossen, F. M., & Sorgner, A. (2019). New Digital Technologies and Heterogeneous Employment and Wage Dynamics in the United States: Evidence from Individual-Level Data. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3390231

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

Geroski, P. A. (2000). Models of technology diffusion. *Research Policy*, *29*(4), 603–625. https://doi.org/10.1016/S0048-7333(99)00092-X

Gries, T., & Naudé, W. (2018). Artificial Intelligence, Jobs, Inequality and Productivity: Does Aggregate Demand Matter? *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3301777

Hall, B. H. (2006). Innovation and Diffusion. In J. Fagerberg & D. C. Mowery (Eds.), *The Oxford Handbook of Innovation* (p. 0). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199286805.003.0017

ifo Institute. (2025). *Companies in Germany Increasingly Relying on Artificial Intelligence | ifo Business Survey | ifo Institute*. https://www.ifo.de/en/facts/2025-06-16/companies-germany-increasingly-relying-artificial-intelligence

INSEE. (2025). *Use of information and communication technologies by companies in 2024—Insee Première—2061*. https://www.insee.fr/en/statistiques/8605307

ISTAT. (2025). *Imprese e Ict – Anno 2025*. https://www.istat.it/comunicato-stampa/imprese-e-ict-anno-2025/

Lane, M., & Saint-Martin, A. (2021). *The impact of Artificial Intelligence on the labor market: What do we know so far?* (OECD Social, Employment and Migration Working Papers No. 256; OECD Social, Employment and Migration Working Papers, Vol. 256). https://doi.org/10.1787/7c895724-en

Lassébie, J., & Quintini, G. (2022). *What skills and abilities can automation technologies replicate and what does it mean for workers?: New evidence* (OECD Social, Employment and Migration Working Papers No. 282; OECD Social, Employment and Migration Working Papers, Vol. 282). https://doi.org/10.1787/646aad77-en

Lipsey, R. G., Carlaw, K. I., & Bekar, C. T. (2005). *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*. OUP Oxford.

Nedelkoska, L., & Quintini, G. (2018). *Automation, skills use and training* (OECD Social, Employment and Migration Working Papers No. 202; OECD Social, Employment and Migration Working Papers, Vol. 202). https://doi.org/10.1787/2e2f4eea-en

OECD. (2019a). *Artificial Intelligence in Society*. OECD Publishing. https://doi.org/10.1787/eedfee77-en

OECD. (2019b). Scoping the OECD AI principles: Deliberations of the Expert Group on Artificial Intelligence at the OECD (AIGO). *OECD Digital Economy Papers*. https://doi.org/10.1787/d62f618a-en

ONS. (2025). *Business insights and impact on the UK economy—Office for National Statistics*. https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/bulletins/businessinsightsandimpactontheukeconomy/2october2025

Polanyi, M. (2009). *The Tacit Dimension* (A. Sen, Ed.). University of Chicago Press. https://press.uchicago.edu/ucp/books/book/chicago/T/bo6035368.html

Rogers, E. M. (1983). *Diffusion of innovations* (3. ed). Free Press [u.a.].

Statistics Canada. (2025, June 16). *Analysis on artificial intelligence use by businesses in Canada, second quarter of 2025*. https://www150.statcan.gc.ca/n1/pub/11-621-m/11-621-m2025008-eng.htm

Stéphane Sorbe, Peter Gal, & Valentine Millot. (2018). *Can productivity still grow in service-based economies?: Literature overview and preliminary evidence from OECD countries* (OECD Economics Department Working Papers No. 1531; OECD Economics Department Working Papers, Vol. 1531). https://doi.org/10.1787/4458ec7b-en

Tankwa, B., Vazquez Bassat, L., Barbrook-Johnson, P., & Farmer, J. D. (2025). *Technological Progress at National Level: Increasing Diffusion Speeds with Ever-Changing Leaders and Followers*. SSRN. https://doi.org/10.2139/ssrn.5099043

Varian, H. (2018). *Artificial Intelligence, Economics, and Industrial Organization*.

Webb, M. (2019). *The Impact of Artificial Intelligence on the Labor Market* (SSRN Scholarly Paper No. 3482150). Social Science Research Network. https://doi.org/10.2139/ssrn.3482150

Zeng, M. (2018, September 1). Alibaba and the Future of Business. *Harvard Business Review*. https://hbr.org/2018/09/alibaba-and-the-future-of-business