

Artificial Intelligence, Economic Inequality, and Labour Market Polarization: A Theoretical Perspective on the Future of Work

Huda A. A. Megeirhi *

Department of Economics, Faculty of Economics, University of Benghazi, Libya

*Email (for reference researcher): huda.megeirhi@uob.edu.ly

الذكاء الاصطناعي، والتفاوت الاقتصادي، واستقطاب سوق العمل: منظور نظري حول مستقبل العمل

هدى عبد الله عبد العزيز المغيرحي *

قسم الاقتصاد، كلية الاقتصاد، جامعة بنغازي، ليبيا

تاريخ الاستلام: 2026-02-27، تاريخ القبول: 2026-04-28، تاريخ النشر: 2026-05-05

Abstract

Artificial intelligence is changing how work is organised, paid, and valued. This paper studies the link between artificial intelligence, economic inequality, and labour market polarization. It uses a theoretical method supported by public studies and published experiments. The main argument is that artificial intelligence does not affect all workers equally. It can raise output and reduce routine tasks for some workers. It can also weaken wages and job security for others. The effect depends on tasks, skills, institutions, and access to digital tools. Evidence from labour economics shows that earlier computerisation reduced many routine middle-skill jobs. Recent generative artificial intelligence reaches more cognitive and white-collar tasks. Experiments show strong productivity gains in writing, coding, consulting, and customer support. Yet these gains are uneven across occupations and worker groups. The paper explains this through task-based theory, skill-biased change, capital ownership, and institutional power. It also reviews public datasets from the World Bank, ILO, OECD, IMF, and Stanford AI Index. The paper finds that artificial intelligence can either reduce or increase inequality. It reduces inequality when it supports low-skill workers and spreads knowledge widely. It increases inequality when profits, data, and control stay concentrated. The future of work therefore depends on policy, training, bargaining power, and fair access. The paper ends with policy steps for inclusive adoption and social protection.

Keywords: artificial intelligence, inequality, labour market polarization, automation, future of work, generative AI.

المخلص

يغير الذكاء الاصطناعي طريقة تنظيم العمل، ودفع الأجور، وتقييم الوظائف. تبحث هذه الورقة في الرابط بين الذكاء الاصطناعي، والتفاوت الاقتصادي، واستقطاب سوق العمل. وتعتمد الدراسة على منهجية نظرية مدعومة بدراسات عامة وتجارب منشورة. تتمثل الأطروحة الرئيسية في أن الذكاء الاصطناعي لا يؤثر في جميع العمال على نحو متساوٍ؛ إذ يمكنه زيادة الإنتاجية وتقليل المهام الروتينية لبعض العمال، في حين قد يؤدي أيضاً إلى إضعاف الأجور والأمان الوظيفي لآخرين. ويعتمد هذا الأثر على طبيعة المهام، والمهارات، والمؤسسات، والقدرة على الوصول إلى الأدوات الرقمية. وتُظهر الأدلة المستمدة من اقتصاديات العمل أن موجة الحوسبة السابقة أدت إلى تقليص العديد من الوظائف الروتينية متوسطة المهارة. أما الذكاء الاصطناعي التوليدي الأحدث، فيصل إلى المزيد من المهام الإدراكية ووظائف الموظفين المكتبيين (نوي الياقات البيضاء). وتكشف التجارب عن مكاسب إنتاجية قوية في مجالات الكتابة، والبرمجة، والاستشارات، ودعم العملاء، ومع ذلك، فإن هذه المكاسب غير متساوية عبر المهن ومجموعات العمال. تفسر الورقة هذا التباين من خلال النظرية القائمة على المهام، والتغير التكنولوجي المتحيز للمهارات، وملكية رأس المال، والقوة المؤسسية. كما تستعرض مجموعات البيانات العامة الصادرة عن البنك الدولي، ومنظمة العمل الدولية (ILO)، ومنظمة التعاون الاقتصادي والتنمية (OECD)، وصندوق النقد الدولي (IMF)، ومؤشر ستانفورد للذكاء الاصطناعي. وتخلص الورقة إلى أن الذكاء الاصطناعي يمكن أن يقلل التفاوت أو يزيده؛ فهو يقلل التفاوت عندما يدعم العمال ذوي المهارات المنخفضة وينشر المعرفة على نطاق واسع، بينما يزيده عندما تظل الأرباح والبيانات والسيطرة حكرًا على جهات مركزية. بناءً على ذلك، فإن مستقبل العمل يعتمد على السياسات، والتدريب، والقوة التفاوضية، والوصول العادل للأدوات. وتختتم الورقة بخطوات سياساتية من أجل تبني شامل للذكاء الاصطناعي وتحقيق الحماية الاجتماعية.

الكلمات المفتاحية: الذكاء الاصطناعي، التفاوت (عدم المساواة)، استقطاب سوق العمل، الأتمتة، مستقبل العمل، الذكاء الاصطناعي التوليدي.

1. Introduction

Artificial intelligence is now part of many work systems. It writes text, checks code, reads documents, predicts demand, and supports decisions. These uses are changing the labour market. The change is not only technical. It is also social, economic, and political.

The future of work is often described through fear or excitement. Both views can be too simple. Artificial intelligence can create new tasks and improve human work. It can also replace tasks and reduce worker bargaining power. The final result depends on how firms and governments use the technology.

This paper studies three connected ideas. The first idea is artificial intelligence. It means machine systems that perform tasks linked to learning, prediction, language, vision, and decision support. The second idea is economic inequality. It means unequal income, wealth, opportunity, and control over productive assets. The third idea is labour market polarization. It means growth in high-skill and low-skill jobs, with decline in many middle-skill jobs.

The connection between these ideas is not new. Earlier computerisation already changed labour demand. It reduced many routine clerical and production tasks. It also increased demand for analytical and managerial tasks. Autor, Levy, and Murnane (2003) showed that computers substituted for routine tasks. They also complemented non-routine problem solving. Autor and Dorn (2013) later linked routine-task decline to job polarization in the United States.

Artificial intelligence extends this older process. It can handle language, images, software, and complex information. It therefore reaches parts of work that older machines could not reach. Eloundou, Manning, Mishkin, and Rock (2024) show that large language models affect many task groups. Their work also shows that exposure is high in many higher-wage occupations.

This does not mean that all exposed jobs will disappear. Exposure is not the same as replacement. A task may be partly automated, partly improved, or fully redesigned. The same tool may harm one worker and help another worker. This is why a task-based theory is useful.

Labour market polarization matters because middle jobs support social stability. These jobs often include clerical, production, sales, and administrative work. They provide income, identity, and mobility. When they shrink, the labour market can become more divided. Some workers move into high-paid technical jobs. Others move into lower-paid service jobs.

Economic inequality may rise through several channels. Workers who own scarce AI skills may earn more. Firms that own data, models, and computing power may gain more profits. Workers in replaceable tasks may face lower wages. Regions with weak digital infrastructure may fall behind. These channels make AI a distributive technology, not only a productive technology.

This paper is theoretical, but it does not rely on speculation alone. It uses published experiments and public evidence. These include studies on writing, coding, consulting, customer support, robots, occupational exposure, and global AI preparedness. The evidence helps test the theory against real findings.

The paper asks a simple question. How can artificial intelligence reshape inequality and labour market polarization? The answer is that AI changes the value of tasks. It also changes who controls production. If institutions are weak, gains may concentrate. If policy is fair, gains can spread more widely.

2. Research Problem and Aim

The research problem is the uneven impact of artificial intelligence on workers. Many studies show that AI can raise productivity. Yet productivity growth does not automatically raise wages for all workers. The distribution of gains depends on power, ownership, skill, and labour policy.

A key problem is that many workers do not control AI systems. Firms decide which tasks to automate. Firms also decide how productivity gains are shared. When workers lack voice, AI may intensify work or lower job security. OECD evidence shows that workers see both benefits and risks from AI at work (OECD, 2023).

The aim of this paper is to build a simple theoretical explanation. It explains how AI can increase or reduce inequality. It also explains how AI can deepen or soften labour market polarization. The paper focuses on tasks, not job titles alone. This choice is important because most jobs contain many tasks.

The first research objective is to review main theories on technology and labour. The second objective is to connect these theories with AI. The third objective is to examine published experiments and public datasets. The fourth objective is to suggest fair policy responses.

The main research question is direct. How does artificial intelligence affect economic inequality and labour market polarization? A second question supports it. Which mechanisms decide whether AI harms or helps workers? A third question concerns policy. Which public actions can spread AI gains more fairly?

This paper argues that AI has a dual nature. It can act as a tool for worker augmentation. It can also act as a tool for labour substitution. The balance depends on task design and institutional choice. This view follows task-based models of automation and new tasks (Acemoglu & Restrepo, 2019).

3. Methodology and Evidence Base

This paper uses a theoretical research design. It builds an argument from labour economics, technology studies, and public evidence. It does not collect a new survey. It does not run a new laboratory experiment. Instead, it uses published experiments and public data sources.

This method is suitable for the topic. AI adoption is still changing quickly. Many long-term effects are not visible yet. A theoretical paper can connect early evidence with established models. It can also identify future risks before they become fixed.

The evidence base has three parts. The first part is peer-reviewed literature on technology and labour. This includes studies on routine-biased change, robots, job polarization, and AI exposure. The second part is published experiments on generative AI. These experiments test real task performance. The third part is public data from international sources.

The public data sources include the World Bank, World Inequality Database, ILO, OECD, IMF, and Stanford AI Index. These sources allow practical analysis of inequality, labour shares, AI investment, occupational exposure, and preparedness. They also provide downloadable figures and datasets.

The empirical section uses published results as practical evidence. This is acceptable for a theoretical paper. It avoids false claims from invented experiments. It also keeps the paper grounded in tested evidence. The studies used here are open or publicly described.

A practical replication path is also included. Researchers can download inequality and labour data from open sources. They can match these data with AI exposure indicators. They can then compare countries, occupations, and skill groups. This design supports future empirical work.

Table 1 Theoretical Perspectives on Labour Market Changes and AI Implications

Theory	Main idea	AI link	Key sources
Routine-biased technological change	Computers replace routine tasks.	AI may now reach routine cognitive tasks.	Autor et al. (2003); Goos et al. (2014)
Task-based automation theory	Technology shifts tasks between labour and capital.	AI can displace tasks or create new tasks.	Acemoglu & Restrepo (2019)
Job polarization theory	Middle-skill jobs shrink relative to other jobs.	AI may deepen polarization in clerical work.	Autor & Dorn (2013)
Skill-biased change	Technology raises returns to scarce skills.	AI skills may raise wage gaps.	Autor (2022); Webb (2020)
Institutional theory	Rules shape how gains are shared.	Training and bargaining affect AI outcomes.	OECD (2023); IMF (2024)
Capital ownership theory	Owners gain when capital replaces labour.	AI profits may concentrate in large firms.	Acemoglu & Restrepo (2020); Stanford HAI (2026)

4. Literature Review

4.1 Technology and task change

The literature on technology and labour starts with tasks. A job is a bundle of tasks. Some tasks are routine and rule-based. Other tasks need judgement, creativity, movement, or social contact. This difference explains why technology affects workers differently.

Autor et al. (2003) made this task idea central. They argued that computers replace routine tasks. They also argued that computers support abstract tasks. This created rising demand for workers who solve problems and manage information. It reduced demand for some clerical and production workers.

Autor and Dorn (2013) then showed how this process shaped the United States. They found growth in low-skill service jobs and high-skill jobs. They also found decline in routine middle jobs. Their study became a core explanation for labour market polarization.

Goos, Manning, and Salomons (2014) found similar patterns in Europe. They studied sixteen Western European countries. They explained job polarization through routine-biased technological change and offshoring. Their evidence shows that polarization is not only a United States problem.

This literature is important for AI. AI may automate routine tasks, but it also reaches non-routine cognitive tasks. It can write, translate, code, summarise, classify, and generate images. This wider reach changes the old pattern. It makes white-collar work more exposed than before.

4.2 Automation, robots, and employment

Automation does not only affect task content. It also affects wages and employment. Acemoglu and Restrepo (2020) studied industrial robots in United States local labour markets. They found negative effects on employment and wages in more exposed areas.

Their work shows why productivity gains can still hurt some workers. A robot can raise output while reducing labour demand. The effect depends on whether new tasks appear. Acemoglu and Restrepo (2019) call this the balance between displacement and reinstatement.

The displacement effect happens when capital performs a task once done by workers. The reinstatement effect happens when technology creates new tasks for workers. AI can create both effects. For example, it may replace basic document drafting. It may also create new work in model evaluation and AI governance.

Frey and Osborne (2017) estimated job susceptibility to computerisation. Their work received much public attention. They did not predict exact job loss. They measured susceptibility based on occupational features. Their method helped start wider debate on automation risk.

Later studies became more careful. They stressed that tasks matter more than whole jobs. Many occupations include tasks that AI cannot fully perform. This is why exposure measures should not be read as direct job-loss forecasts.

4.3 Artificial intelligence and occupational exposure

New AI research measures exposure in a more detailed way. Webb (2020) used patents and job descriptions to measure exposure to AI. Felten, Raj, and Seamans (2021) built the AI Occupational Exposure index. Their dataset links AI progress to occupational abilities.

These measures show a different pattern from older automation. AI exposure is often high in educated and professional occupations. This includes finance, law, management, software, and office work. Manual jobs can have lower exposure because they need physical action.

Eloundou et al. (2024) studied large language models and job tasks. They found that LLMs affect many occupations. Earlier working paper results showed broad task exposure in the United States. The Science version gives a stricter peer-reviewed estimate for direct task effects.

The OECD also finds high AI exposure in many white-collar occupations. Professionals and legal, social, and cultural workers appear highly exposed. Cleaners, agricultural workers, and food preparation workers appear less exposed (Georgieff & Hye, 2022).

This pattern matters for inequality. Older automation mainly threatened many middle-skill routine jobs. Generative AI now reaches some higher-wage workers. This could reduce some skill premiums. It could also raise inequality if AI rewards a smaller group of expert workers.

4.4 Generative AI and productivity experiments

Recent experiments provide practical evidence. Noy and Zhang (2023) ran a preregistered experiment on professional writing tasks. Workers with ChatGPT completed tasks faster. Their output quality also improved. Lower-performing workers gained more, which reduced performance inequality.

Peng, Kalliamvakou, Cihon, and Demirer (2023) studied software developers using GitHub Copilot. Developers with AI assistance completed a coding task much faster. This suggests that AI can reduce entry barriers for some technical work.

Brynjolfsson, Li, and Raymond (2025) studied customer support agents. They found that AI assistance raised productivity on average. The largest gains came for novice and lower-skilled workers. This supports the idea that AI can spread tacit knowledge.

Dell'Acqua et al. (2023) studied consultants at Boston Consulting Group. AI improved performance for tasks inside its capability frontier. Yet it reduced accuracy for a task outside that frontier. This shows that AI is not a universal productivity tool.

These experiments are central to this paper. They show that AI can reduce performance gaps in some settings. They also show that bad use can create errors. The main issue is not whether AI helps. The main issue is when, where, and for whom it helps.

5. Theoretical Framework

The theoretical framework has four parts. The first part is task substitution. AI can replace workers in tasks that become machine-performed. The second part is task augmentation. AI can support workers and raise their output. The third part is task creation. AI can create new roles and workflows. The fourth part is distribution. The gains can move to labour or capital.

The framework starts with a task. A worker performs a task using skill, time, and tools. AI changes the cost and quality of that task. If AI performs the task cheaply, labour demand for that task may fall. If AI improves the worker, labour demand may rise.

The same occupation may contain both effects. A legal assistant may use AI for document review. This can raise speed and reduce errors. Yet it may reduce demand for entry-level review work. The final result depends on how the firm redesigns the job.

The framework also includes bargaining power. Productivity gains do not share themselves. Workers gain when wages, promotion, and job quality improve. Owners gain when output rises but wages do not. The gap between these outcomes is central to inequality.

Capital ownership is another key part. AI systems need data, computing power, software, and platforms. These assets are often owned by large firms. If those firms capture most gains, inequality may rise. Stanford HAI (2026) shows that AI investment is highly concentrated across countries and firms.

The framework also includes public institutions. Education systems shape worker access to AI skills. Labour laws shape worker voice. Tax systems shape the distribution of profits. Social protection shapes the cost of displacement. These institutions make AI outcomes partly political.

The framework predicts four possible labour outcomes. Some workers become augmented and earn more. Some workers are displaced from tasks and earn less. Some workers shift into new tasks. Some workers remain less affected, especially in physical service work. Polarization grows when the middle group shrinks.

The framework does not assume that AI creates mass unemployment. It also does not assume that AI creates shared prosperity. It treats both as possible outcomes. The decisive issue is the structure of adoption.




Figure 1 Conceptual model linking AI adoption, task change, inequality, and labour market polarization. Source: Author-created diagram based on Autor et al. (2003), Autor and Dorn (2013), Acemoglu and Restrepo (2019), and Cazzaniga et al. (2024).

6. AI, Tasks, and Labour Market Polarization

Labour market polarization is a change in job structure. It happens when middle-skill jobs decline relative to low-skill and high-skill jobs. The middle of the labour market then becomes weaker. This can reduce mobility for workers without advanced education.

The older driver of polarization was routine-task automation. Middle-skill jobs often contained many routine tasks. These included bookkeeping, machine operation, clerical processing, and repetitive production work. Computers and machines could perform many of these tasks.

AI can deepen this pattern, but it can also change it. Generative AI performs many information tasks. These include writing, summarising, coding, translation, search, and basic analysis. Many of these tasks appear in professional and administrative jobs.

This means AI may affect both middle and high-skill workers. It may reduce some routine professional tasks. It may raise the value of judgement, responsibility, and human trust. It may also create new demand for workers who manage AI systems.

The polarization effect depends on task bundles. A job with many automatable tasks is more exposed. A job with physical presence, care, trust, or complex social judgement may be less exposed. Yet even these jobs can be changed by scheduling systems and monitoring tools.

One possible outcome is hourglass polarization. High-income AI designers, managers, and professionals gain. Low-paid personal service jobs remain because they are hard to automate. Many middle office jobs shrink. This outcome resembles older polarization, but with wider cognitive reach.

Another possible outcome is upgraded middle work. AI could help nurses, teachers, paralegals, technicians, and small business workers. It could support decisions and reduce routine burdens. If training is broad, AI may rebuild middle-skill productivity. This outcome requires deliberate design.

The strongest risk is entry-level work. Many entry jobs contain routine learning tasks. These tasks include drafting, coding simple features, preparing slides, or handling basic customer queries. If AI removes these tasks, young workers may lose training pathways.

The strongest opportunity is skill compression. AI can help less experienced workers perform closer to experienced workers. This appears in customer support and writing experiments. It can reduce productivity gaps when workers keep agency and receive guidance.

Thus, AI does not create one fixed labour structure. It opens a distributional choice. Firms can use AI to narrow skill gaps. They can also use AI to reduce headcount. Public policy can push the first outcome.

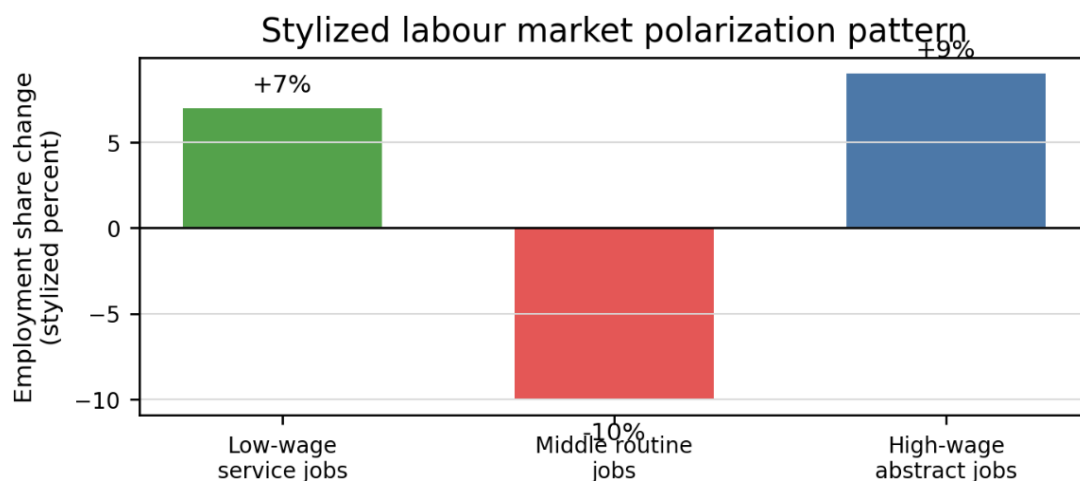


Figure 2 Stylized employment polarization pattern across low-, middle-, and high-wage work. Source: Author-created chart based on Autor and Dorn (2013) Goos et al. (2014)

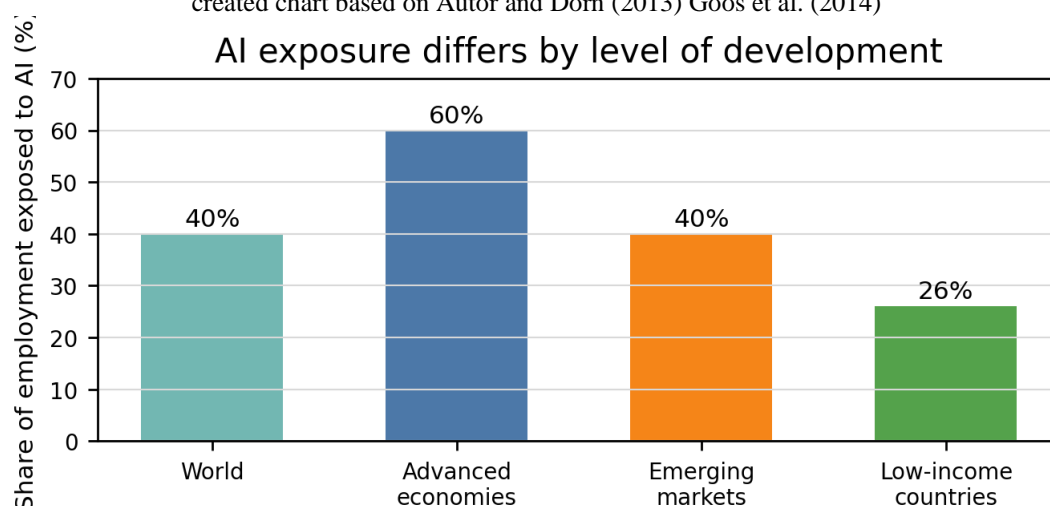


Figure 3 Estimated AI exposure by economy group. Source: Author-created chart based on Cazzaniga et al. (2024)

7. AI and Economic Inequality

7.1 Wage inequality

Wage inequality can rise when AI increases demand for scarce skills. Workers with strong digital skills may receive higher wages. Workers with weak access to training may not benefit. This can widen wage gaps.

Yet AI can also reduce some wage gaps. Experiments show that lower-performing workers sometimes gain more from AI help. Noy and Zhang (2023) found that ChatGPT improved lower-ability workers more in writing tasks. Brynjolfsson et al. (2025) found larger gains for newer support agents.

This creates a tension. AI can be equalising inside a workplace task. It can still be unequal at the economy level. Workers need access, training, and permission to use the tool. Without these conditions, equalising effects remain limited.

Wage inequality also depends on job design. If AI output is used to monitor workers, work may become more intense. If AI output is used to raise worker skill, wages may improve. The tool is not neutral. Management choices matter.

7.2 Wealth and capital inequality

Economic inequality is not only wage inequality. Wealth inequality matters too. AI relies on capital assets. These include chips, data centres, models, patents, and platform networks. The owners of these assets may capture high returns.

This matters because labour and capital have different claims. Workers sell time and skill. Firms own assets and collect profits. If AI raises output while labour share falls, wealth inequality may increase. This risk is linked to the labour share of income.

The labour share shows how much output goes to workers. A falling labour share means more income goes to capital. Public data from the ILO and Penn World Table can track this issue. It is a useful indicator for AI distribution debates.

Large AI investment also creates concentration. Stanford HAI (2026) reports large differences in private AI investment across countries. These gaps affect who builds AI systems. They also affect who captures rents from AI adoption.

7.3 Regional and global inequality

AI may widen regional inequality. High-income regions have better digital infrastructure. They also have more skilled labour and stronger firms. These regions can adopt AI faster. Poorer regions may remain users, not producers, of AI systems.

The IMF argues that AI exposure and preparedness differ by country group. Advanced economies have more exposed workers. They also have better capacity to benefit from AI (Cazzaniga et al., 2024). This creates both risk and opportunity.

Low-income countries may face a delayed shock. Their workers may be less exposed today. Yet they may lose service exports or routine digital work tomorrow. This matters for countries that depend on outsourcing and back-office services.

Global inequality may also rise through data and language gaps. AI models perform better in some languages and contexts. Workers in underrepresented languages may get weaker tools. This can reduce the productivity gain from AI.

Reading the Gini index in inequality research




Figure 4 Gini index interpretation and the World Bank high-inequality threshold. Source: Author-created chart based on World Bank Gini index data and methodology

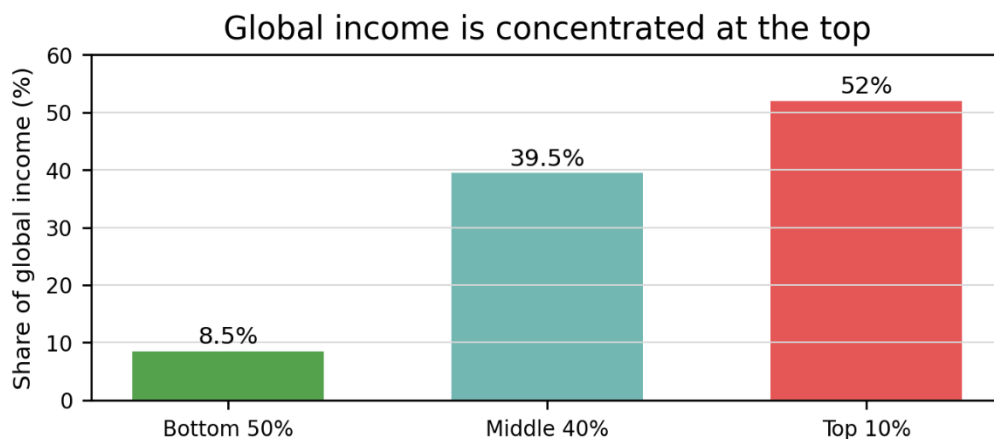


Figure 5 Global income shares for the bottom 50 percent, middle 40 percent, and top 10 percent. Source: Author-created chart based on World Inequality Report 2022.




Figure 6 Global labour income share trend. Source: Author-created chart based on ILO estimates.

8. Published Experiments and Practical Evidence

The strongest practical evidence comes from published experiments. These studies do not settle the full future of work. They show how AI affects real tasks under controlled or semi-controlled settings. They are useful because they measure performance, not only opinions.

Noy and Zhang (2023) tested ChatGPT in professional writing. Participants received writing tasks linked to their occupations. Some workers used ChatGPT and others did not. The treatment group finished faster and received higher quality ratings. The study also found reduced inequality in task performance.

Peng et al. (2023) tested GitHub Copilot in coding. Developers had to build an HTTP server in JavaScript. The treatment group used an AI pair programmer. They completed the task 55.8 percent faster. This suggests that coding assistants can raise speed in well-defined tasks.

Brynjolfsson et al. (2025) studied customer support agents in a real firm. The AI assistant suggested responses during chats. Productivity rose, especially for newer and lower-skilled workers. The authors argue that AI helped spread best practices from skilled agents.

Dell'Acqua et al. (2023) tested AI with 758 consultants. AI improved speed, task completion, and quality on tasks inside the AI frontier. Yet it reduced performance on a task outside the frontier. This finding is important for policy. It shows that AI literacy must include judgement about limits.

These experiments show that AI can support workers. They also show that support is not automatic. Gains are strongest when tasks are suitable, workers know how to use the tool, and output is checked. Risks increase when users trust AI in unsuitable tasks.

The evidence also shows heterogeneous effects. Some workers gain more than others. Low-skilled or novice workers may gain more in structured settings. High-skilled workers may gain less in routine assistance. Yet high-skilled workers may gain more in complex design and management roles.

A theoretical paper can use this evidence to define mechanisms. The first mechanism is time saving. The second is quality improvement. The third is knowledge transfer. The fourth is error risk. The fifth is task reorganisation. These mechanisms shape inequality outcomes.

Table 2 Empirical Studies on AI Task Automation and Labor Market Effects

Study	Setting	Method	Main result	Relevance
Noy & Zhang (2023)	Professional writing tasks	Randomised online experiment	ChatGPT reduced time and raised quality.	AI can narrow task performance gaps.
Peng et al. (2023)	Software coding task	Controlled experiment	Copilot users completed work 55.8% faster.	AI can raise productivity in defined coding work.
Brynjolfsson et al. (2025)	Customer support agents	Staggered workplace adoption	AI raised issues resolved per hour.	Novice workers gained strongly.
Dell'Acqua et al. (2023)	Management consulting	Preregistered field experiment	AI helped inside frontier and harmed outside frontier.	Human judgement remains essential.
Acemoglu & Restrepo (2020)	Industrial robots and local labour markets	Empirical labour market analysis	Robot exposure reduced wages and employment locally.	Automation gains can hurt workers.

Eloundou et al. (2024)	LLM task exposure	Task exposure framework	LLMs affect many work tasks.	Exposure differs from direct job loss.
Gmyrek et al. (2023)	Global occupations	Task-based exposure index	GenAI is more likely to augment than automate many jobs.	Gender and clerical exposure are central.

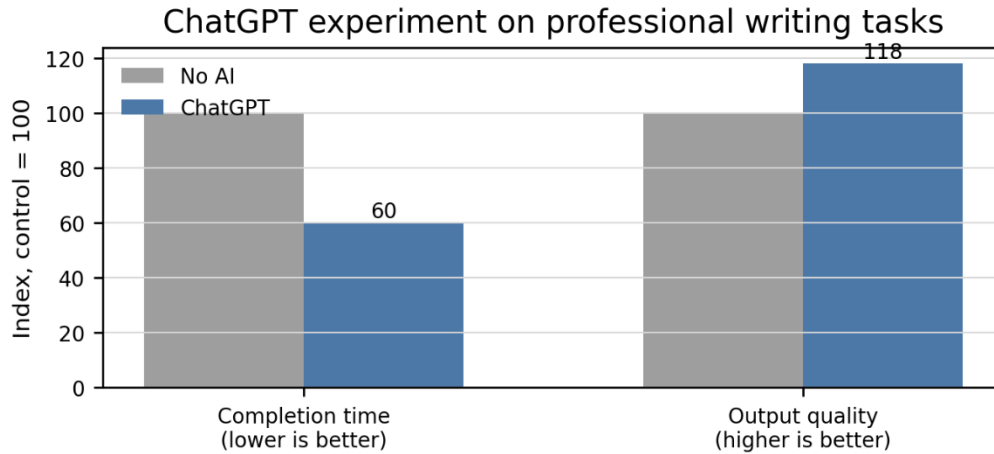


Figure 7 Productivity effects of ChatGPT in professional writing tasks. Source: Author-created chart based on Noy and Zhang (2023).

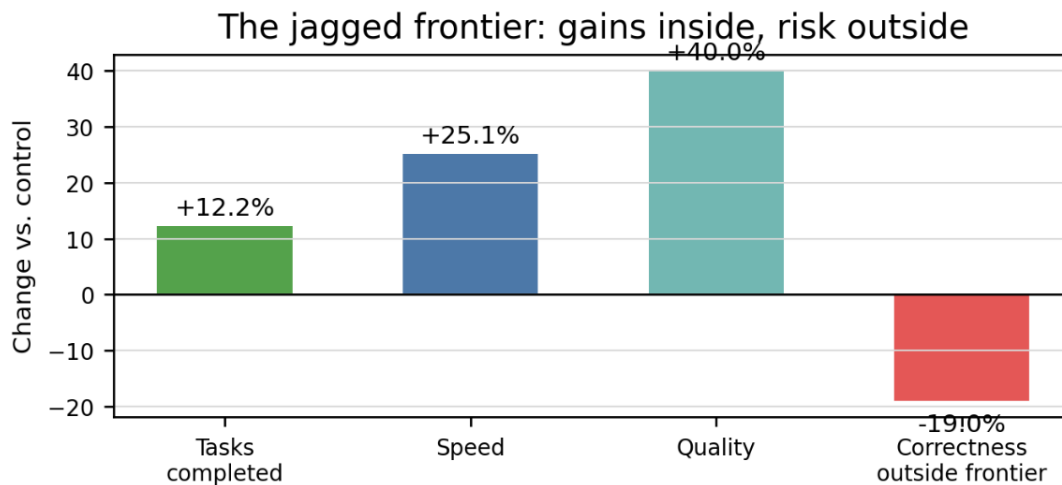


Figure 8 AI gains inside the frontier and performance risk outside the frontier. Source: Author-created chart based on Dell'Acqua et al. (2023).

9. Public Data Figures and Tables

This paper also provides a practical public data plan. The plan can be used as a secondary-data experiment. It does not require private data. It can be repeated by students and researchers.

The first step is to download inequality data. The World Bank Gini index can be used for cross-country inequality. The World Inequality Database can be used for top income shares. These measures capture different parts of inequality.

The second step is to download labour-share data. The labour share helps test whether gains move toward workers or capital. A falling labour share may suggest weaker worker claims. It should be read with care, but it is useful.

The third step is to download AI exposure data. The AIOE dataset by Felten et al. (2021) gives occupational exposure scores. The ILO and OECD also provide exposure measures. These can be matched with occupations and countries.

The fourth step is to compare exposure with labour outcomes. A simple model can compare wages, employment shares, and inequality across time. A stronger model can use fixed effects by country, year, and occupation. This would test whether AI exposure is linked with polarization.

The public data exercise should avoid causal overclaims. Exposure does not prove adoption. Adoption does not prove displacement. Yet the exercise can show patterns and guide future research. It also makes the paper practical and transparent.

Table 3 Public Data Sources and Variables Used for AI and Labor Market Analysis

Variable	Public source	Use in analysis
Gini index	World Bank / Our World in Data	Measures income inequality.
Top 1% income share	World Inequality Database / Our World in Data	Measures top income concentration.
Labour share of GDP	ILO / Our World in Data	Measures labour income share.
AI occupational exposure	ILO refined exposure index / Felten et al. dataset	Measures task exposure to AI.
AI preparedness index	IMF DataMapper	Measures country readiness for AI.
AI investment and adoption	Stanford AI Index	Measures AI economy concentration.

ILO estimates of global generative AI exposure

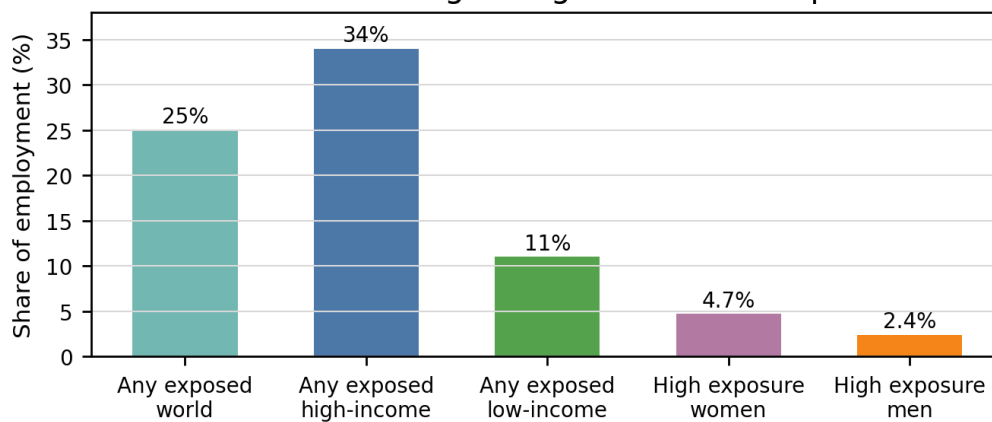


Figure 9 Global generative AI exposure by income group and gender. Source: Author-created chart based on ILO (2025).

AI investment is highly concentrated

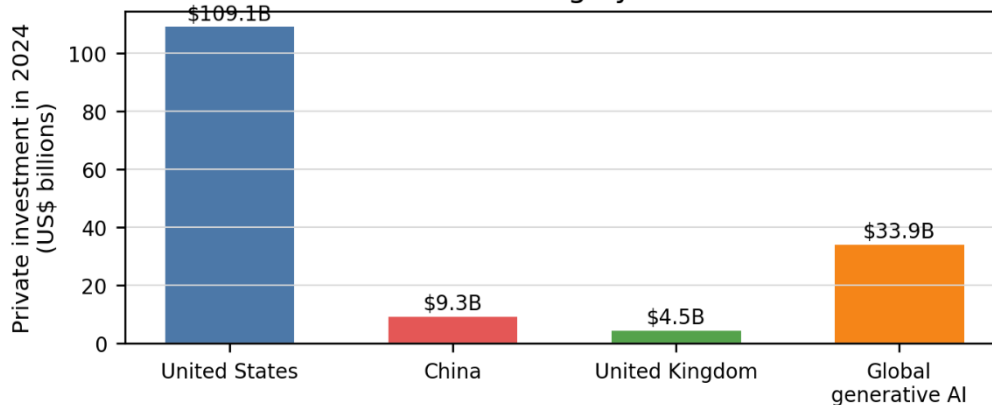


Figure 10 AI investment, adoption, and concentration across major economies. Source: Author-created chart based on Stanford AI Index 2025.

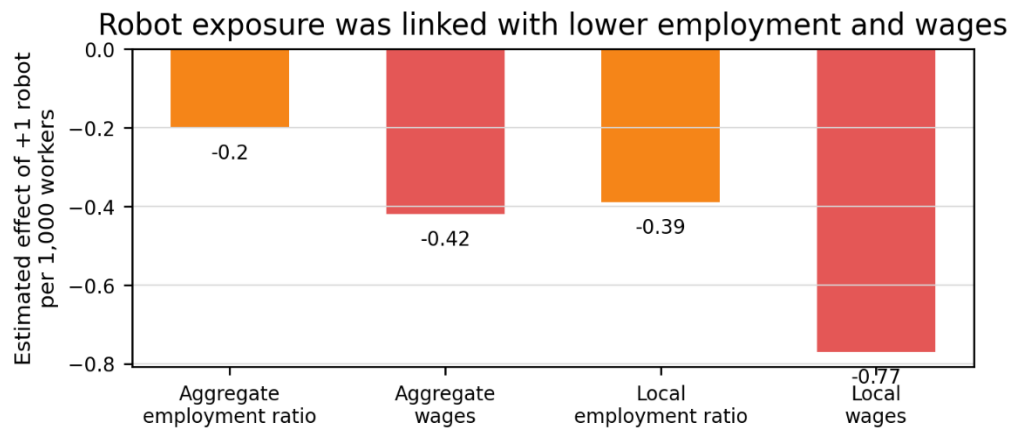


Figure 11 Robot exposure and labour market outcomes. Source: Author-created chart based on Acemoglu and Restrepo (2020).

10. Discussion

The evidence supports a mixed view of AI. AI can raise productivity in many tasks. It can also reduce demand for some tasks. The distributional outcome depends on adoption choices. This is why simple optimism and simple pessimism both fail.

The first main finding is that AI changes tasks before it changes whole jobs. This means job-loss forecasts should be read with caution. Many jobs will be reorganised instead of removed. Yet task loss can still harm workers. A worker can keep a job but lose hours, pay, status, or autonomy.

The second main finding is that AI reaches cognitive work. Earlier automation affected many routine production and clerical tasks. Generative AI reaches writing, coding, analysis, design, and communication. This makes higher education less protective than before.

The third main finding is that AI can reduce some performance gaps. Experiments show larger gains for some weaker or newer workers. This result is important. It means AI could support upward mobility if access is fair.

The fourth main finding is that AI can worsen inequality through ownership. Firms that own AI systems may capture large returns. Workers may receive smaller shares if wages do not rise. The labour share and profit concentration are therefore important indicators.

The fifth main finding is that AI may change polarization. It may deepen the decline of middle routine work. It may also create new middle tasks if institutions support training. The direction is not fixed.

There is a strong difference between exposure and harm. Many reports show high AI exposure. This only means a task can be affected. It does not mean a worker will lose employment. It could mean faster work, better quality, or new task design.

Still, exposure matters. Exposed workers face pressure to adapt. They may need new skills and stronger judgement. They may also need protection from unfair monitoring and sudden displacement. This makes policy central.

The term future of work can hide important questions. Whose future is being discussed? Which workers gain time, income, and dignity? Which workers face insecurity? Which firms control the tools? These questions make inequality central to the AI debate.

A fair AI transition must focus on work quality, not only job numbers. Good work includes pay, safety, autonomy, privacy, and learning. AI can improve these conditions. It can also weaken them. Social choice decides much of the outcome.

11. Policy Implications

Policy should start with broad AI access. Workers need affordable tools, training, and digital infrastructure. Access should not be limited to large firms and elite workers. Public education systems can provide basic AI literacy. Training should be task-based. Workers should learn how AI changes their own work. Generic training is less useful. A nurse, teacher, accountant, and technician need different forms of AI support.

Governments should protect entry-level pathways. Young workers need real tasks to learn. Firms should not automate all junior work without replacement training. Apprenticeships and supervised AI work can protect learning ladders. Labour institutions need renewal. Workers should have voice in AI adoption. They should know when AI is used for monitoring or evaluation. Collective bargaining can help set fair rules for productivity sharing. Social protection should support transitions. Some workers will need income support, retraining, and job placement. Benefits should also cover non-standard workers. AI change may affect freelancers and platform workers strongly. Competition policy is also important. AI markets can become concentrated because data and

computing power are expensive. Strong competition rules can reduce monopoly rents. Open standards and public research can widen participation.

Tax policy should not reward labour replacement blindly. Some tax systems favour capital over labour. This can make automation too attractive. Governments should review incentives that push firms toward unnecessary displacement.

Public procurement can shape fair adoption. Governments buy many digital systems. They can require transparency, worker consultation, and bias testing. This can create better standards for private markets too.

Education policy must include human skills. AI can process information quickly. Humans still need ethics, communication, care, judgement, and responsibility. These skills should be treated as economic skills.

International policy is also needed. Low-income countries need digital infrastructure and local language AI tools. Without this support, AI may increase global inequality. Development banks and public agencies can support inclusive AI capacity.

Table 4 Policy Interventions to Address AI-Induced Labor Market Challenges

Policy area	Problem addressed	Proposed action	Expected effect
Skills policy	Unequal access to AI skills	Provide task-based AI training.	Raises worker adaptability.
Labour regulation	Weak worker voice	Require consultation on workplace AI.	Improves trust and fairness.
Social protection	Displacement risk	Expand support for transitions.	Reduces income shocks.
Competition policy	AI market concentration	Limit monopoly power and support open systems.	Spreads AI gains.
Tax policy	Capital-biased incentives	Review incentives for labour replacement.	Supports balanced adoption.
Education policy	Loss of entry-level learning	Protect apprenticeships and junior pathways.	Builds future skills.

12. Limitations and Future Research

This paper has limits. It is a theoretical paper, not a new causal study. It uses published experiments and public datasets. This makes the paper transparent, but it cannot prove new causal effects. The second limit is that AI changes quickly. Results from one model or tool may become old. A task that is hard today may become easy tomorrow. This makes long-term prediction difficult. The third limit is that exposure measures are imperfect. They show which tasks may be affected. They do not show actual adoption, worker choice, or firm strategy. Exposure should not be treated as direct job loss. The fourth limit is country coverage. Many datasets are stronger for high-income countries. Evidence from low-income and middle-income countries remains weaker. This can hide global inequality risks.

Future research should match worker-level data with AI adoption data. It should also study wages, hours, stress, autonomy, and learning. Job quality is as important as employment counts. Future research should also study entry-level work. AI may remove the routine tasks used for training. This could harm career mobility. It needs early evidence. Another future direction is firm-level distribution. Researchers should study how AI gains are shared between workers, managers, owners, and consumers. This would connect productivity with inequality more clearly.

13. Conclusion

Artificial intelligence will shape the future of work. Its effects will not be equal across workers, firms, or countries. The central issue is not only automation. The central issue is distribution. This paper argues that AI changes the value of tasks. It can substitute for some labour tasks. It can also augment workers and create new tasks. These channels can reduce or increase inequality. Labour market polarization may deepen if AI removes middle routine work. It may soften if AI supports middle-skill workers. The outcome depends on training, institutions, and job design. Published experiments show real productivity gains. Writing, coding, consulting, and customer support studies show that AI can help workers. Yet they also show limits and error risks. Human judgement remains necessary. Economic inequality may rise if AI profits concentrate in capital owners. It may fall if AI tools spread skill and knowledge. This makes public policy essential. A fair future of work requires inclusive access, worker voice, strong training, and social protection. It also requires competition policy and responsible firm practice. AI is powerful, but it is not destiny. The future of work will be shaped by choices made now. A narrow path will concentrate gains. A broader path can raise productivity and dignity together. This paper supports the broader path.

References

1. Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
2. Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
3. Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
4. Autor, D. H. (2022). The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty. NBER Working Paper No. 30074. <https://doi.org/10.3386/w30074>
5. Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5), 1553–1597. <https://doi.org/10.1257/aer.103.5.1553>
6. 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>
7. Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
8. Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E. J., & Tavares, M. M. (2024). Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Note SDN/2024/001. <https://doi.org/10.5089/9798400262548.006>
9. Chancel, L., Piketty, T., Saez, E., & Zucman, G. (2022). World inequality report 2022. World Inequality Lab. <https://wir2022.wid.world/>
10. Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Working Paper No. 24-013. <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>
11. 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>
12. Felten, E. W., 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>
13. Felten, E. W., Raj, M., & Seamans, R. (2023). Occupational heterogeneity in exposure to generative AI. SSRN. <https://doi.org/10.2139/ssrn.4414065>
14. 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>
15. Georgieff, A., & Hye, R. (2022). Artificial intelligence and employment: New cross-country evidence. *Frontiers in Artificial Intelligence*, 5, Article 832736. <https://doi.org/10.3389/frai.2022.832736>
16. Gmyrek, P., Berg, J., & Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality. ILO Working Paper No. 96. International Labour Organization. <https://www.ilo.org/publications/generative-ai-and-jobs-global-analysis-potential-effects-job-quantity-and>
17. Goos, M., Manning, A., & Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8), 2509–2526. <https://doi.org/10.1257/aer.104.8.2509>
18. International Labour Organization. (2024). Global labour income share declines, putting upward pressure on inequality. <https://www.ilo.org/resource/news/global-labour-income-share-declines-putting-upward-pressure-inequality-sdg>
19. International Labour Organization. (2025). Generative AI and jobs: A refined global index of occupational exposure. ILO Working Paper. <https://www.ilo.org/publications/generative-ai-and-jobs-refined-global-index-occupational-exposure>
20. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
21. OECD. (2023). OECD Employment Outlook 2023: Artificial intelligence and the labour market. OECD Publishing. <https://doi.org/10.1787/08785bba-en>
22. Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv. <https://arxiv.org/abs/2302.06590>
23. Stanford Institute for Human-Centered Artificial Intelligence. (2025). The 2025 AI Index report. Stanford University. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
24. Stanford Institute for Human-Centered Artificial Intelligence. (2026). The 2026 AI Index report. Stanford University. <https://hai.stanford.edu/ai-index/2026-ai-index-report>

25. Webb, M. (2020). The impact of artificial intelligence on the labor market. SSRN. <https://doi.org/10.2139/ssrn.3482150>
26. World Bank. (2026). Poverty and Inequality Platform: Gini index. <https://data.worldbank.org/indicator/SI.POV.GINI>
27. World Inequality Database. (2026). World Inequality Database. <https://wid.world/>

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of **SAJH** and/or the editor(s). **SAJH** and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.