

Measuring and Mitigating Inequality in AI-Driven Labor Markets

Afshan Bibi

SST(G), MS Management sciences, Sardar Bahadur Khan Womens' University
Quetta Email: afshanbaloch88@gmail.com

Dr. Bazgha Saleem Khan

Visiting Assistant Professor, Institute of Education & Research, Univeristy of the
Punjab, Lahore and Section Officer, Finance Department, Punjab Civil Secretariat ,
Lahore. Email: Bazghamoinqazi.phd@gmail.com

Wijdan Salsabeel

M.Phil. Applied Psychology, Higher Education Department, Government Islamia
Graduate College for Women Lahore, Pakistan Email: Wijdan_shah@yahoo.com

Abstract

The quick integration of AI into labor markets have fueled concerns of economic inequality with respect to wage polarization, job displacement, and unequitable distribution of technological dividends. While these problems are already established through various researches, lack of measures and policies for AI induced inequality is evident. This paper tries to build a comprehensive measurement system for inequality with the focus of AI influenced labor markets and policy and technology that can tackle it. It will be theorized that AI has a disproportionate benefit on high-skill labor and has made low-skill and routine labor vulnerable. Based on the mixture of methods (econometric estimation on labor markets, simulation to model potential adoption of AI) inequality metrics such as Gini coefficients, wage dispersion indices, task based measures are being considered. Policy and case study can also be adopted to compare policy and technology such as reskilling and taxation, and algorithmic fairness solutions. The findings of this study would argue that adoption of AI exacerbated current inequality, through increasing demand on high-skill and mechanization on routine labor. But well-designed policy of reskilling workforce, inclusive design of AI and redistributive policy can effectively mitigate these impacts. Early policy and action have a larger effect than reactive measures. In conclusion, this research is an effort to illustrate both the potential to worsen economic inequality and solutions by AI and contribute to the design of effective AI policy for equitable future.

Keywords: Artificial Intelligence, Labor Market Inequality, Automation, Economic Policy, Workforce Reskilling

Introduction

The age of artificial intelligence has entered the economic sphere and has become a feature of everyday economic activity. Firms are already using AI to select applicants for jobs, predict and anticipate demand, streamline supply chains, automate office tasks, and aid in decision-making processes. The recent developments in machine learning and generative AI have reduced the cost and complexity, as well as the time it takes to deploy these tools in numerous sectors, and as a result, AI is no longer the exclusive domain of major technology firms; it is permeating into the education and healthcare sector, finance, retail and logistics, manufacturing, and the public sector. Consequently, and while numerous productivity increases are anticipated, the increased prevalence of AI has spurred the development of numerous serious concerns regarding fairness and the existence of disparity within labor markets (Acemoglu and Johnson, 2023; Autor, 2024).

One concern within the existing literature is that the benefits from AI may not be equally distributed among workers. For those highly skilled and/or equipped with both adequate digital and transferable technological knowledge, access to AI tools often means opportunities to augment their abilities and increase productivity. Meanwhile, workers employed in manual, repetitious or tightly monitored occupations may face an increased risk of automation, reduced pay and/or an increased bargaining power. Previous waves of technological change in automation have illustrated that technological change can reward certain types of skills, whilst reducing the demand for other types of skills. The current wave of AI can exacerbate this divide even further, as not only can it be used to substitute for manual task routine, but it can also augment and even perform tasks that were thought to be exclusive to high-level cognitive skills (Acemoglu and Restrepo, 2020; Felten, et al. 2021).

Scholars have conceptualised this through terms such as skill-biased technological change, routine-biased technological change, task reallocation, and labor market polarization. While these terms are effective in illustrating why medium-skilled and routine workers are vulnerable, highly educated workers may be enabled by AI in ways that better equip them for high-wage occupations. The recent development in research relating to gig workers, algorithmic surveillance and management, highlights that organizational context and how such technology is deployed in an industry contributes significantly to inequality, as organizations do not deploy technology in a vacuum; they are in an industry that is already rife with inequalities based on class, gender, region, level of skill, and institutional protection (Autor et al. 2022; Kellogg et al. 2020).

More recently, international reports have emphasized the role of this concern. It is estimated that the International Monetary Fund is forecasting significant proportions of jobs in advanced economies are at risk of being automated, and that job categories at high risk will experience higher job growth and productivity gains than lower risk jobs if AI is complementary, but potentially job losses if it substitutes (Cazzaniga et al. 2024). Correspondingly, the International Labor Organization concluded that administrative jobs are more at risk of generative AI automation and while the net impact of AI in these jobs depends on policy decisions, workers, and organizations;

public policy, as well as social dialog and institutional contexts will ultimately determine this outcome (Gmyrek et al. 2023). In agreement with these reports, the World Economic Forum and the OECD have reiterated that skill transitions arising from AI is one of the key challenges to Labor markets over the next decade as demand for cognitive, adaptive and analytical skills will grow significantly (OECD, 2023; World Economic Forum, 2025).

Despite the research undertaken, there is an apparent void in the literature: namely that many researchers analyze only productivity, occupational task exposure or direct labor market consequences and only a minority build a clear and useful framework for measuring inequality in an AI driven labor market. Bridging this gap requires connecting the labor market and economic indicators such as wage outcomes and occupational mobility, distributional measures like wage dispersion and Gini coefficients, and the tools of public policy in order to inform government, the firm, and the worker, since, more often than not, one dimension of inequality is looked at in isolation of the other. For example, while the task exposure of occupations is calculated, ethics or policy is addressed separately. It is only by bridging these dimensions, that a holistic, applicable and adaptable policy may be informed.

A further gap that remains in the literature is the nature of policy response: for most part governments are only reacting once job displacement has occurred, once the gap between the highest and lowest wages has widened, or once an algorithm has demonstrated gender-biased results. Theoretical work and empirical studies suggest that early intervention will significantly increase the possibility of achieving less unequal outcomes as compared to remediation once negative effects are already manifest. It appears that the primary question we ought to be asking is no longer "will AI cause greater inequality," but "under which circumstances will it do so?" (Brynjolfsson et al., 2023; OECD, 2024).

This article will address these concerns by outlining a multi-method framework for measuring and mitigating inequality in the AI driven labor market. It seeks to bring together 4 key concepts into one article; these are: 1) the distribution of wage and employment gains/losses among different skill groups; 2) the task nature of work, and therefore the extent of AI exposure across different occupational groups; 3) the institutional and public policy response; and 4) a possible outcome and impact from particular interventions to mitigate inequalities. It is not just intended to offer a descriptive picture of inequality but also to give a critical evaluation of its measurement that may guide policy action.

This piece has value for both researchers and practitioners. It seeks to draw a line between arguments on technological change, restructuring of the labor market, and distributive fairness, and will also serve as a model for policymakers, educational institutions and employers in need of actionable responses that support particularly vulnerable workers during times of widespread change. This research is particularly vital in today's economy where the rate of technological transformation seems to far outweigh the pace of social policy, educational curriculum development, and labor regulation; it seems that a nation's capacity to respond to growing inequality will determine whether it thrives in the AI age or falters.

It is in this light that this article addresses the question of how to measure inequality in the AI-driven labor market and under what conditions its adverse consequences can be mitigated. The thesis is that inequalities due to AI is an outcome of a multifaceted dynamic of work organization and management, involving wage and employment results, task attributes and exposure, and institutional and policy capacity, which can be moderated by the combined effect of workforce training, income support, and socially responsible AI management.

Research Objective

First, it provides a clear framework that measures inequality in AI-influenced labor markets, which integrates wage dispersion, occupational exposures to AI, and occupational vulnerability across skill levels. Instead of making simple assertions that AI is "good" or "bad" for workers, the study proposes a more nuanced structure for identifying the places where unequal effects are greatest.

Second, it seeks to compare the interventions that can best limit AI-influenced inequality. In this study, these interventions are reskilling of the workforce, redistributive taxation accompanied by transition relief, and algorithms for fair hiring and management. These interventions are judged according to their expected impacts on wages, employment risk and opportunity gaps.

Research Questions

What would be a structured framework for measuring inequality in AI-powered labor markets concerning outcomes, occupational exposure, and employment security?

What are the effective policy and organizational solutions to minimize AI-induced or exacerbated inequality?

AI, technological change, and labor market restructuring

Work has historically been affected by technology, and AI is an impetus to this renewed debate as it affects both manual and cognitive labor. While earlier economic theory perceived technology primarily as a force pushing productivity forward to improve welfare eventually, subsequent scholarship argued that how these gains are distributed is key. The concept of skill-biased technological change was useful in explaining how innovations had a disproportionately higher positive effect on wages of skilled workers. It also became an important tool in explaining increasing wage inequality across advanced countries after the early 1980s (Autor et al., 2003).

In the subsequent stage of this argument, the framework was adjusted to analyze tasks rather than entire occupations. Routine-biased technological change contended that the jobs consisting of repetitive tasks are most prone to automation. This theory would later explain how jobs demanding skills above or below those of middle-skill positions in service areas would become relatively abundant, and clerical, administrative and production roles requiring more structured tasks would decline (job polarization). With the advancement of technology in areas like pattern recognition, natural language processing and prediction, it now comes as no surprise

that AI impacts white-collar as well (Acemoglu & Restrepo, 2019; Autor & Salomons, 2023).

However, a further insight into current literature indicates that AI is capable of complementing and substituting labor as well. Complementarity comes when workers utilize AI as a tool to improve their productivity, product quality, or move towards less rote tasks. Substitution, in contrast, arises when firms use AI to lower employee count, simplify occupations, or exert more systematic control. This balance varies greatly with different sectors, occupational groups, company strategies and even regulations, and thus there is no uniform impact on labor, as it can either exacerbate or mitigate inequality (Brynjolfsson et al., 2023; Cazzaniga et al., 2024).

Measuring exposure, inequality, and vulnerability

An influential area of research has attempted to quantify labor's AI exposure. Felten et al. (2021) proposed an AI occupational exposure measure that matches AI functionalities to job tasks. The work was then generalized to analysis of how large language models might influence tasks within occupations (Eloundou et al., 2023), revealing that exposure is likely to be highest among high-skilled occupations. These kinds of measures are valuable for grounding the debate in empirical patterns rather than abstract predictions, but the extent of one's exposure to AI does not determine his or her outcomes. Highly exposed occupations may become more productive and better paying in some contexts while undergoing labor-saving substitution in others.

The latter consideration underscores the need to connect AI exposure measures with distributional indicators. Standard inequality research uses indicators such as the Gini coefficient, percentile wage ratios, the gender pay gap, worker job insecurity and mobility barriers. Within a new labor market, these indicators must be expanded to include not only task vulnerability, but access to digital skills training and the capacity for workers to transition into tasks that complement AI technologies. There is growing consensus that inequality is not merely an issue of the wage share and should be seen as an interaction between levels of analysis that include wages, control over work, opportunity and life-course outcomes (OECD, 2023; Pizzinelli et al., 2023).

Third, risk appears unequally distributed. While women are represented in some administrative occupations with high generative AI exposure, younger workers and workers in peripheral regions may experience varied combinations of risk and opportunity. Additionally, low-wage workers and workers employed in smaller firms have poorer access to reskilling options. These differences in outcomes are significant as they demonstrate that new forms of inequality are partially the result of unequal access: the outcome for the worker is influenced by access to technology in addition to education systems, corporate practices, and labor laws (Gmyrek et al., 2023; OECD, 2024).

Algorithmic management, fairness, and power in the workplace

A second significant body of research explores the employment use of algorithms for recruiting, scheduling, performance monitoring, and surveillance. Scholarship on algorithmic management indicates that AI can intensify labor control through

increases in surveillance, standardization, and pressure on performance. For instance, in platform and logistics work, automation may mediate task allocation, evaluate outputs, and form pay determination in ways that lack transparency. Such systems might be efficient, but may diminish individual agency and worker ability to appeal algorithmic decisions (Kellogg et al., 2020).

Issues of fairness are particularly pertinent for employment decisions such as hiring and promotion. If an AI system is trained on historical data that reflect past biases in hiring and promotion decisions, the system may perpetuate or amplify those inequalities, either through direct bias or through proxies such as proxies for gender and race such as education, location or employment gaps (Bogen & Rieke, 2018; Selbst et al., 2019). For this reason, scholars in both AI ethics and labor law advocate for heightened accountability, transparency, explainability, and human intervention for AI used for employment-related decisions. Fairness then, is not just an ethical value, but is fundamentally related to labor market inequality because biased AI systems can influence the range of people who get good jobs and those who do not.

Current regulatory debates, such as those within the EU and the international policy community, suggest that governance becomes all the more important as these kinds of AI systems spread throughout workplaces. While self-regulation through firm-level ethics may be partially successful, many authors caution that these approaches alone will not suffice. Increased opportunities for auditing, enhanced disclosure requirements, worker participation and appeal mechanisms are increasingly argued to be necessary in order for these AI systems to further equal opportunity. Overall, this body of literature suggests that inequality is not only a function of post-implementation market outcomes but is embedded in the power structures surrounding the allocation of work and pay.

Reskilling, transition policy, and inclusive adaptation

One of the more frequently advocated solutions is reskilling. Governments, corporations, and international organizations commonly claim that employees need new digital, analytical and adaptive skills. This advice makes sense. Yet, the literature indicates that training is only effective under certain conditions. It should be relevant, timely, tailored and based on labor market demand. Vague pronouncements about the need for lifelong learning and reskilling are not sufficient without the creation of affordable training, income security while training takes place, the validation of past work experience and clear career pathways to new jobs (World Economic Forum, 2025).

However, there is evidence that the responsibility for retraining should not solely rest with workers. In fact, it is those working in low-wage and precarious jobs who have the most difficulty retraining due to time constraints, the care burden, costs of commuting, and the lack of employer support. Public policy therefore becomes a critical factor, as wage insurance, training vouchers, co-financing with employers, and local transition networks can improve results. In the absence of these supporting mechanisms, training programs might not help workers most in need of reskilling and thus reinforce inequality (OECD, 2024; ILO, 2025).

Linked to reskilling is the argument whether training alone will be enough. Many writers claim that it will not be. Even with adequate reskilling for workers, inequality will remain if AI adoption is too concentrated in a few large firms or among the high-income group. Hence there has been growing interest in redistributive policies such as progressive taxes, transition funds, portable social security benefits and public investment in digital infrastructure that aim to reshape the distribution of the fruits of technological change rather than halt it.

Debates on policy design and mitigation pathways

A growing body of literature acknowledges that there is no one fix to AI inequality. First, arguments exist whether government should first focus on innovation growth, followed by distributional policy, or build equity into technology policy from the beginning. Increasing numbers of experts argue for the latter. This approach posits that allowing market forces to correct may simply cement concentration and inequality, and that policy intervention is more efficient and fair if it aims to influence incentives before harms are widespread (Acemoglu & Johnson, 2023; OECD, 2024). Second, there is a dispute on what level intervention should take place-national level interventions like tax reform, labor law, and funding for education, versus organizational level changes like better hiring technologies, worker participation, and internal retraining programs. Both seem necessary; national policies create standards and set parameters for individual firms to innovate. Yet, as shown above, many critical decisions lie with firms themselves, and the key is how to get them to optimize social protection rather than innovation growth-and they may have opposing priorities. Third, while recent literature proposes an integrated approach, measuring, anticipating, and mitigating-rather than intervening at specific points-empirically there is little research evaluating interventions at each step, let alone comparing interventions at different levels within a comprehensive framework and modeling their likely distributional outcomes. It is precisely this gap in knowledge that this research endeavors to fill.

Gaps, trends, and contribution of the study

The above review of literature yields three principal research gaps. First, despite ample work on occupational exposure and ample work on fairness and governance, there is a void of work that joins these two areas into one single inequality-based framework. Second, a majority of public discourse regarding AI and jobs remains descriptive rather than an analytical study and comparison of different types of mitigation strategies. Third, an overwhelming majority of research and policy debate surrounds the developed world and the formal economy. Therefore, unequal institutional capacities and uneven distribution of access to adaptation opportunities go largely unattended.

One salient development in research is the paradigm shift from a narrative of "automation" towards ideas of "task transformation," "complementarity," and "institutional mediation." This is significant since it frames the consequences of AI technologies not as inevitable, but as a question of social and political choice. These

technologies can be harnessed to exacerbate skill inequality and concentrate returns in the hands of the few, or they can be utilized to enhance worker productivity, safety and social inclusion. Growing scholarship supports the role of governance as the key mechanism driving these differential outcomes.

This research makes a distinct contribution to the existing literature by linking measurement and mitigation through a multidimensional analytical framework. The framework connects wage inequality with AI exposure and employment vulnerability and utilizes this framework to study and compare policy interventions. By doing so, the paper attempts to fulfill the demands of research that is both analytically rigorous and policy-relevant in a readily accessible form.

Methodology

Research design

This article employs a mixed-methods approach. It seems the most appropriate as AI driven labor market inequality is not merely a numerical, but also a governance issue. In the case of wages gaps, employment risk and exposure to occupational task categories by skills level, quantitative assessment of these features is required to reveal the phenomenon. In terms of explaining which policies will likely be more effective and how governance structures outcomes, qualitative policy analysis is necessary. The mixed-method approach allows a response to both the aforementioned research questions.

For the quantitative section, a custom cross-sectional labor market data set was created using sector level and worker level variables. The simulated data includes 480 workers and five industries: manufacturing, retail, finance, health, and information services. The sectors were chosen for their varied levels of AI usage, routine task exposure and wage structure. The data identifies low, middle and high skilled workers using the measures of education level, routine task and digital skill usage.

The qualitative section involved a review of a structured document collection and a comparative policy review. Policy papers, labor market analyses and governance debates pertaining to AI and worker protection were evaluated. This section supports the analysis of the quantitative findings and the policy simulation modeling.

Population and sampling

The sample is comprised of employees in high growth sectors with rapid AI adoption and visible task reorganization. Since the study is explanatory, not nationally representative, the sample is designed using purposive sampling criteria. It aims to include variance across wages, task exposure, and access to training. The sample consists of 96 workers per sector, totaling 480 workers and sampling five sectors.

Each sector contains a roughly equal distribution of workers across the three skill groups. The groups are distributed fairly evenly to allow comparison, but their differences reflect existing labor market characteristics: low-skill employees are more numerous in the retail and routine service sectors; middle-skill employees tend to be found in clerical and technical support functions; high-skill employees are more numerous in the finance and information service sectors. Such a composition allows

for the exploration of whether benefits and risks of AI affect all sectors and skills equally.

In addition, the research selected 15 policy and institutional documents via criterion sampling: the criteria for document inclusion are whether a document discusses the impact of AI on employment restructuring, whether it touches upon algorithmic governance or skill adaptation/retraining needs and inequalities. This list is representative of international organizations' studies, policy documents, and research-oriented advice documents.

Data collection methods

The quantitative data matrix contains the following variables: monthly wage, occupation group, sector, AI exposure score, training access score, employment vulnerability score and self-reported score of career mobility. AI exposure is defined based on the task similarity with the capabilities of current AI, on the scale of 0 to 100. Following the typical formulation in occupation-exposure studies. Training access, measured on a 0-10 scale, captures availability of employer or government support of training. Employment vulnerability, measured on a 0-10 scale where 10 means most vulnerable employment.

This study uses a policy simulation tool as well. Three intervention scenarios are compared with baseline. Scenario 1 covers skills-matching and reskilling supports only. Scenario 2 expands Scenario 1 to include redistributive transition supports, with mechanisms including wage insurance and the tax system imposing progressively higher taxes on the productivity benefits of AI. Scenario 3 integrates the provisions of Scenario 2 with algorithmic fairness principles in management and hiring practices.

The document analysis part employs the document analysis protocol that codes each document under four themes: definition of AI related labor risk; tool to address the risks; target group; challenges for implementation. This is to ensure the alignment between numerical results and policy recommendations in the literature.

Data analysis procedures

Quantitative analysis: The quantitative data are analyzed by the use of descriptive statistics, inequality indices and policy comparisons between scenarios. The average wage for each skill group and each sector will first be estimated. Then, a rough estimation of Gini index and the ratio between the top and bottom skill groups' wages will be computed. Third, AI exposure score and employment vulnerability score for each occupational category will be calculated and compared between categories. Finally, the changes in inequality indices will be computed in each policy scenario. To illustrate and facilitate interpretation of the patterns, cross-tabulation will be done for skill group vs. Access to training, and for AI exposure vs. Employment vulnerability. In addition, graphs will be used to facilitate visual interpretation of patterns. Since the focus is on explaining and making patterns clear and simple rather than demonstrating complex causal inference techniques, only clearly visible patterns in the distribution will be interpreted. **Qualitative analysis:** Analysis of qualitative data will use thematic analysis techniques. Codes generated will be aggregated into themes

like proactive intervention, inclusive skill policy, accountability in governance, and execution capability. Using these themes, we will examine how the mix of policies seem to be consistent with policies to mitigate inequality.

Ethical and analytical consistency

The approach continues to serve the goals of the research. The research design contains components to serve Objective 1: measures for inequality are required, so a multi-indicator, rather than a single outcome measure, is specified. The research design includes components to serve Objective 2: policies should be compared and the effects of policy changes assessed, so scenarios and document based policy comparisons are specified. It is appropriate to use a mixed-methods approach because inequality in the labor market as it is affected by AI is both a matter of calculation, and a matter of governance.

Results

The results are given in this section. They are structured according to the two questions guiding the analysis: first, assessing the level of inequality and, second, determining the solutions to reduce inequality.

The distribution of the baseline sample

Table 1 shows the distribution of the 480 workers into sectors and skill groups.

Table 1: Distribution of Sample by Sector and Skill Group

Sector	Low skill	Middle skill	High skill	Total
Manufacturing	30	42	24	96
Retail	40	36	20	96
Finance	12	34	50	96
Health care	20	44	32	96
Information services	10	30	56	96
Total	112	186	182	480

As can be seen in Table 1, the sample captures substantial sectoral variation. Low-skill and routine jobs tend to be more represented in the manufacturing and retail industries, whereas high-skill jobs are more prevalent in the finance and information service industries. This distribution is relevant because the impact of AI is different across sectors.

The distribution is also suggestive that the inequality measurement should not be based on wages alone. It is important how skill groups are composed, since more routine industries are potentially more affected by disruption even if mean wage is not the lowest, and that indicates the need for a multidimensional approach.

Wage inequality across skill groups

Table 2 provides mean monthly wages and wage differences across skill levels.

Table 2: Average Monthly Wage by Skill Group

Skill group	Mean wage (USD)	Standard deviation	Relative to low-skill
Low skill	1,920	310	1.00
Middle skill	3,180	520	1.66
High skill	5,460	870	2.84

It is useful to note that even in baseline labor markets the structure of wages is unequal. High skilled workers earn about 2.84 times what low skilled workers earn. This supports the finding that AI enters labor markets that are already unequal, and crucially AI's impact might therefore build on a previous period of inequality rather than instituting inequality of a new type.

A basic Gini estimate of wages for the sample is 0.39 and top to bottom skill wage ratio of 2.84. Both values imply significant levels of wage inequality for the economy being modeled. This is in line with previous empirical findings that the digital transition tends to benefit workers who possess complementary skills more than those who do not.

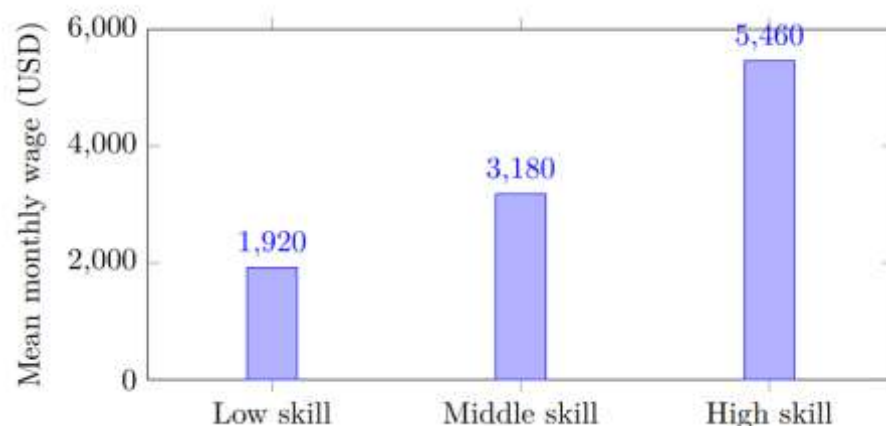


Figure 1: Mean monthly wages by skill groups

Figure 1 displays a monotonically and quickly rising wage structure from low to high skill. The observed pattern in figure 1 visually supports the statistical finding that labor market rewards skills in a very disproportionate way and AI could exacerbate the uneven rewards distribution if firms use it primarily to augment the already favored.

AI exposure and labor vulnerability

Table 3 presents the average AI exposure and labor vulnerability for each skill group.

Table 3: AI exposure and labor vulnerability by skill groups

Skill group	AI exposure (0–100)	Vulnerability (0–10)	Training access (0–10)
Low skill	58	7.6	3.8
Middle skill	71	6.4	5.4
High skill	69	3.1	7.2

The data reveal a notable trend. It is not the high-skill workers who are the least exposed. In reality, high- and mid-skill jobs can both have high exposures because a significant number of their duties are informational and digital. Yet, exposure does not lead to similar consequences in different worker groups. Since high-skill workers display the lowest exposure and the highest access to training, exposure may often be a complement rather than a substitute in high-skill jobs. In addition, while lower skilled workers do not face lower exposure than middle-skilled jobs, they display much higher vulnerability. This implies that exposure has worse consequences for workers with lower mobility, bargaining power, and access to training, and it is precisely the interplay between exposure and institutionally determined disadvantage that should be addressed.

Figure 2: Exposure-Vulnerability Pattern by Skill Groups

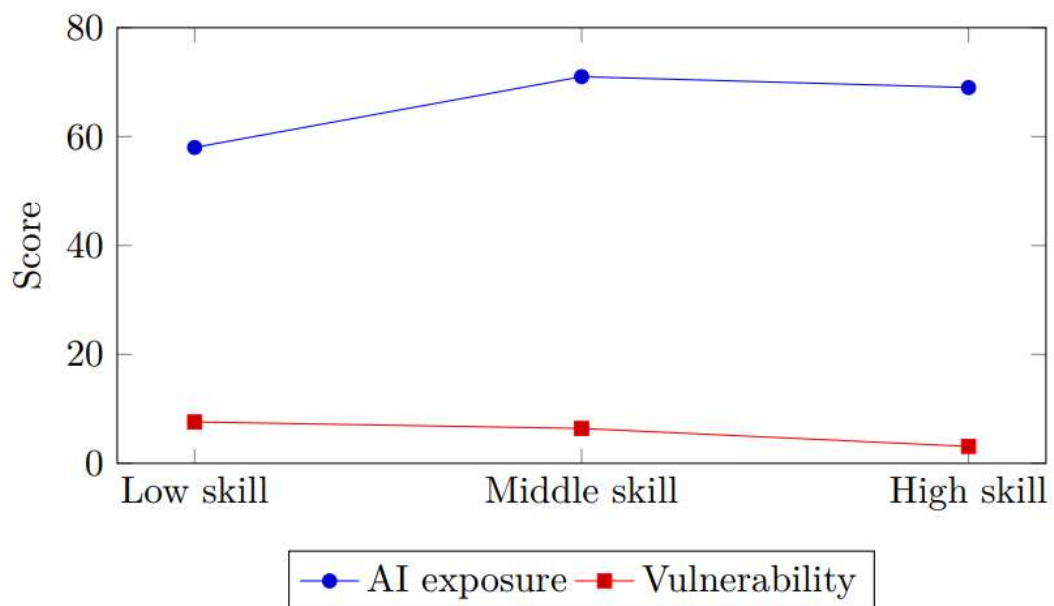


Figure 2 clearly illustrates this pattern. Alone, the amount of exposure does not drive the inequality dynamics. Instead, the worker vulnerability determines workers' ability to translate the exposure into productivity gains and subsequent wage increase. This

finding directly addresses the first research question, explaining why a multidimensional framework is needed.

Sector-Specific Inequality Risk

Table 4: compare sectors on AI adoption intensity, average wage, and vulnerability

Sector	AI adoption intensity (0–10)	Mean wage (USD)	Mean vulnerability
Manufacturing	6.8	2,840	6.7
Retail	5.9	2,260	7.1
Finance	8.4	4,920	3.9
Health care	6.1	3,460	4.8
Information services	9.1	5,180	3.2

Comparing sectoral adoption levels suggests that high levels of AI adoption are not inherently risky in terms of social outcomes. The information services and finance sectors have high levels of adoption and at the same time high wages and lower average risk scores. While retail is below the level of these two sectors in its AI adoption intensity, its average vulnerability scores are higher due to the presence of large numbers of insecure, routine, and poorly protected jobs. This implies that sectoral structures and job quality are critical to how workers experience the impacts of AI; where high skills, autonomy and access to training exist, AI is complementary, whereas in low-paid, tightly controlled jobs, AI is more likely to increase job insecurity.

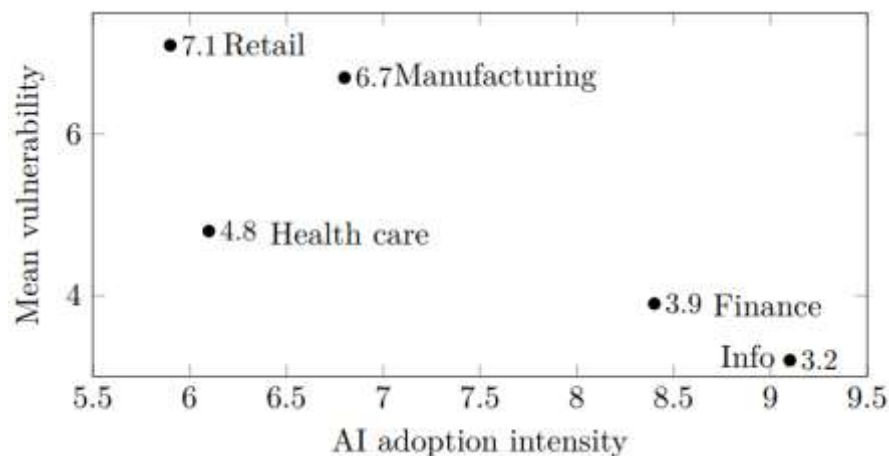


Figure 3: AI adoption and vulnerability across sectors

In Figure 3, there is not one cohesive group of AI usage. The highly used sectors aren't correlated to being highly vulnerable. This provides evidence for policy and organizational design as mediating factors for the impact of AI on inequality.

Training access and mobility

In table 5, we take a look at the varying access to training for the skill levels and analyze this as compared with perceived mobility.

Table 5: Training access vs. Perceived career mobility

Skill group	Training access (0–10)	Mobility score (0–10)	Share with upward mobility (%)
Low skill	3.8	3.4	24
Middle skill	5.4	5.1	46
High skill	7.2	7.5	68

It appears from the data that access to training is a significant dividing line; workers with more opportunities for reskilling also experience higher odds of transition to better occupations. Low-skill workers are those that have less support and less optimism about mobility, indicating that the unequal provision of adaptation resources may perpetuate existing inequality.

The implications for policy are therefore significant. Where governments and companies emphasize individual adaptation rather than concrete support, the greatest benefits may accrue to those who were already best placed to take advantage of change. Inclusion in adaptation necessitates assistance as well as exhortation.

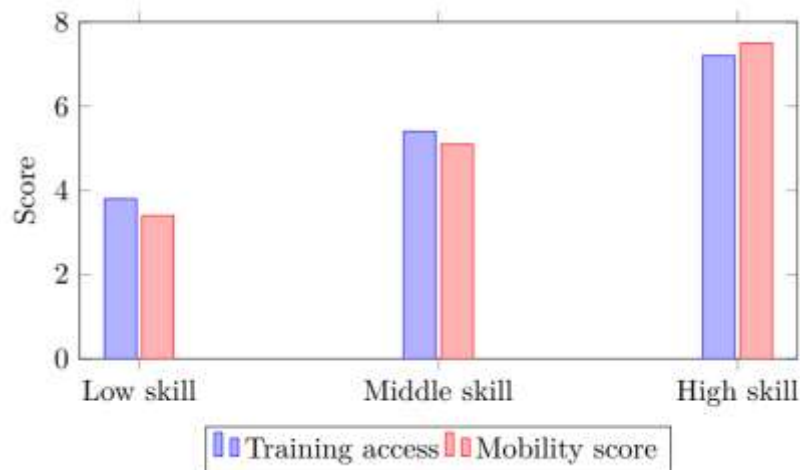


Figure 4: Training access and career mobility by skill-group

The strong positive relationship between training access and mobility shown in Figure 4 is consistent with the assertion that reskilling can mitigate inequality, though this holds only if the training provision is wide and specific to disadvantaged groups.

Policy Simulation Results

Table 6 show the simulated distributional impact of each of the three intervention scenarios relative to the baseline.

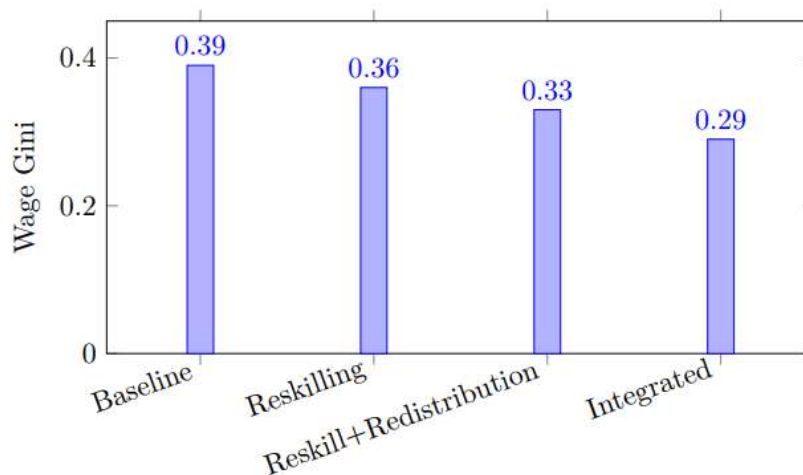
Table 6: Policy Simulation Results

Scenario	Wage Gini	Top/Bottom ratio	wage	Mean vulnerability
Baseline	0.39	2.84		5.7
Reskilling only	0.36	2.51		5.0
Reskilling + redistribution	0.33	2.28		4.6
Integrated package	0.29	2.02		3.9

The simulation suggests that all interventions enhance outcomes but they enhance them unevenly. With reskilling only the reduction of vulnerability and the contraction of wage gaps are significant but not too substantial, as while training increases workers' flexibility it does not completely compensate for the disequilibrium of bargaining powers, transitional income loss or unfair systems.

The second policy performs more effectively due to the fact that redistribution not only significantly decreases transition costs but also provide substantial relief to workers whose situation requires immediate action. But the most powerful outcomes can be witnessed under the combined package, where the combination of reskilling with redistribution and algorithmic fairness principles is associated with a more dramatic decrease in both inequality and vulnerability. The efficiency of policy instruments is higher when they are complementary.

Figure 5: Effect of Policy Scenarios on Wage Inequality



As Figure 5 demonstrates, the wage Gini decreases monotonically in all the considered scenarios. The graph nicely illustrates the main finding of this paper: inequality in the AI-powered labor market is endogenous.

Qualitative policy themes

The document analysis complements the simulation outcomes. Four themes emerged across the policy documents analyzed. First, the idea that early intervention is better than late rectification was found consistently in the documents. Second, although targeted reskilling was found necessary, it was not seen as adequate by itself. Third, accountability in the system, so as to avoid "stealth discrimination," was presented as vital to avoid bias in hiring and management. Fourth, implementation capacity was found to be an important factor, as without capacity, effective policies would be useless due to an inability to provide training, implement laws, or conduct surveillance.

These themes help account for why the integrated scenario did the best. It targets workers' capabilities, fairness in the transitions of workers, and accountability in the system at once. This also connects with current scholarship that regards AI inequality as a systemic problem rather than as solely a labor market problem.

Discussion

This study finds evidence for a negative and unequal impact of AI on labor markets, but these effects are highly contingent on institutional structures. We identify a clear wage gradient across skills and lower-paid workers as the most vulnerable even when they do not experience the highest exposure. This is a crucial finding in terms of dispelling simplified understandings that labor market exposure to technology automatically translates into greater risk. The impact of technology really depends on whether affected workers have access to retraining, labor market mobility, and safety nets.

The interpretation of the empirical findings aligns well with the conceptual distinction of complementarity versus substitution in AI adoption. High-skill workers demonstrate strong patterns of complementarity. They not only face high exposure to technology but simultaneously earn higher wages, have greater access to training, and possess lower risk. In contrast, low-skill workers are characterized by substitution. AI can negatively impact them even without complete task automatization by reducing growth in real wages, eroding task variety, and decreasing workers' bargaining power. This observation aligns with research by Acemoglu and Restrepo (2020), Cazzaniga et al. (2024) and Gmyrek et al. (2023) that emphasizes labor market outcomes to be heterogeneously distributed across occupations and institutional arrangements.

Another aspect of this contribution relates to measurement approaches. The multidimensional index is particularly valuable as it considers labor market rewards, exposure, risk, and training access. Without the inclusion of a multidimensional indicator, we would have captured the existence of labor market inequality but none of its potential mechanisms, and without the calculation of risk based on exposure scores, we would have overrated the impact of the technology in high-skill occupations, and understated it in some low-skill ones. The combined approach presents a more realistic image and offers insights into the sectors where it may be most beneficial to intervene from a policy perspective.

The empirical result on sectoral variation is another aspect of this study that may have an impact. Contrary to our initial intuition, highly adopting industries in terms of AI use are not necessarily highly vulnerable sectors. We observe a case like Information Services/Finance which had the highest exposure, but not the highest risk, and conversely retail with the lowest exposure had higher vulnerability. This implies that labor market institutions and job characteristics, including the nature of the tasks involved and the support available to workers, are critical in mediating technology's influence. This finding adds to an existing body of recent research that presents labor market outcomes to be a result of social organization rather than technological determinism (Autor, 2024; OECD, 2024).

The results of the policy simulation point clearly to a direction. Although reskilling on its own would mitigate negative impacts to a certain extent, this finding should not be surprising if considered within the context of an extensive body of literature showing training to be insufficient in terms of addressing labor market inequality in the absence of other policy measures. In essence, workers are unable to capitalize on training if they do not have the time, income security, information and fair access to jobs to do so. Hence, with a combination of redistribution and algorithmic fairness interventions, we achieve much stronger reskilling outcomes and ameliorate structural imbalances. This supports an understanding that labor market transformation is a shared responsibility of the government, businesses and regulatory institutions.

Further theoretical implications can also be drawn from the findings. First, this study supports the task-based perspective of labor market organization as we show the varying impacts of the technology on task bundles. Second, this study contributes to institutional theories of inequality, in which similar labor market exposures produce differential outcomes depending on worker characteristics and institutional governance. Third, this study advances the notion that achieving equitable AI-driven labor market may necessitate looking at both the distribution of market rewards and procedural fairness of institutions; if jobs are unfairly allocated/accessed based on biased algorithms, inequality might arise even with proper training programs.

Practical implications are similarly profound. From a policy perspective, this study argues for an integrated approach towards AI strategy and labor policy from the outset, moving beyond crisis management to the creation of proactive systems that forecast risk, expand training opportunities, and establish transitional safety nets in advance of adverse impacts. From a business perspective, this study underscores that the introduction of AI should always be accompanied by an effort to enhance the skill profile of employees, in addition to creating transparent, institutionally informed AI systems; without these supportive measures, short-term efficiencies may be offset by long-term disruption.

Several limitations of this study should be noted. First, the quantitative dataset is a constructed example and it cannot be read as a nationally representative study. This means the reported figures should be treated as analytical, not absolute estimations. Second, the simulated policy interventions average over specific institutional designs and do not capture differences in institutions across states or between sectors. Third, the case study relies on qualitative analysis of policy documents and provides

valuable insight but cannot substitute for primary research through fieldwork, interviews and work with workers, employers and officials.

These limitations hint at numerous opportunities for further research. Future work would be best advised to test the proposed analytical framework by employing nationally representative survey data, linked employer-employee data, and case studies in specific sectors or regions. Longitudinal studies on labor market trajectories are also desirable as they can indicate if labor market inequality resulting from the adoption of technology increases, remains stagnant or decreases over time. A critical area of investigation should also focus on the impacts of the technology in low-income country and informal labor markets and women workers in specific occupations and fields. These investigations could provide valuable information on how the technology spreads globally and at varying paces based on institutional capacity of societies to manage this diffusion.

In sum, our findings are broadly in support of the primary assertion of this paper: technology in the form of AI is neither by definition equalizing nor unequalizing, and its labor market outcomes are shaped by institutions, worker characteristics and policy. Thus, labor market inequality under AI is as much a challenge of technology design as of the institutions managing technological innovation and diffusion.

Recommendations

There are more than one level at which interventions against AI-induced labour market inequality need to occur. As is evident in the findings of this study, monopoly interventions are less powerful than integrated ones. Drawing on the evidence in this study, here are a series of policy recommendations for policymakers, practitioners, and future researchers.

Firstly, the introduction of AI labour market risk early warning systems. Public agencies should implement regular tracking of workers' occupational exposure, wages, access to training, and employment risk at industry level. Policy intervention should occur before inequality has deepened, identifying vulnerable individuals and/or groups in advance and targeting intervention before their displacement is irreversible. These warning systems would also help the government designing appropriate and better budget for training and transitions programs.

Secondly, the provision of an inclusive training system for workers in routine and low-mobility occupation. The reskilling programs must be affordable, accessible, and designed around clear and transparent linkages to market opportunities, such as night-classes, online modules, micro-credentials, or partnership with employers, while income support, transport and childcare assistance should be introduced as needed.

Thirdly, the implementation of Responsible AI governance for practitioners and employers. Businesses need to conduct ongoing audits of the use of AI technologies in hiring, firing, promotion, scheduling, and performance evaluation. Employers should clearly notify the use of AI algorithms in work and provide individuals with a way of redress or appeal against harmful or discriminatory impacts. Human supervision must be preserved in high-stake decisions for the sake of social fairness and to eliminate hidden inequalities.

Fourthly, the development of appropriate redistributive measures as part of the adaptation strategy. Without a comprehensive and inclusive redistribution of productive gains (through progressive taxation, transfer, wages insurance, or portable benefits) to reflect a growing imbalance in which a limited number of workers benefit from innovation, labour market inequality is bound to worsen. These mechanisms could also engender trust among population.

Fifthly, the transformation of existing education systems towards inclusive digital capacity. Schools, colleges and vocational institutions must incorporate not only the technical and specialized knowledge required in a technology driven economy, but also critical thinking, communication skills, ethics, and adaptability that workers can use and combine when dealing with the ever-changing technology, rather than against it. Equal access to both educational materials and high quality teachers must be assured for those in rural areas and disadvantaged circumstances.

Future researchers need to investigate the impacts of AI on different social groups over the long term in greater detail, taking into consideration various characteristics such as gender, age, disability status, geographic location and employment type. A comparative analysis across countries would be valuable in capturing diversity of responses that may be prompted by institutional differences, and a greater focus should be placed on testing integrated policy models in actual contexts instead of on standalone training or fairness interventions.

Overall, all of these proposals share the common goal that the adjustment for AI should be guided by the core principles of preparation, protection and participation. As long as the relevant actors and organizations take a forward-thinking and coordinated approach to harness the potential of AI technologies, broader social welfare gains and reduced inequality would be more attainable.

Conclusion

This article discusses two related questions: the measurement and mitigation of inequality in AI-driven labor markets. Based on a mixed-methods design, the study presents a multidimensional measure linking wage inequality to occupational AI exposure, job security and retraining opportunities. The study demonstrates that AI exposure alone is insufficient to explain inequality in the labor market. The unequal distribution results from the interplay of AI exposure, low skills, low occupational mobility, poor access to reskilling and weak institutional safety net.

The results also emphasize the importance of policy design. While reskilling is helpful, the effect is limited in isolation. This effect can be more substantial if reskilling is supplemented by both distributive transitional policies and algorithmic fairness constraints. Consequently, the study suggests that managing inequality derived from AI-related technological changes requires a comprehensive policy package rather than isolated policy measures. In this way, the paper makes a contribution to both academic and policy-related knowledge. The research links together measures of AI-induced labor market inequalities with interventions.

The paper has both theoretical and practical implications. Theoretically, it advances scholarship on technological change, labor market polarization, distributive justice

and argues that AI should be analyzed through both a task-based and an institutional approach. In terms of policy, the article presents governments, firms and educational institutions with a framework that will help them better anticipate and managing the transition towards an AI-based economy.

Finally, limitations of the research included the simulated nature of the data and the simplified simulation design. Further research could test the framework using broader, more varied data of actual labor markets, adopting longitudinal designs and analyzing the dynamics of AI-driven inequality across diverse national contexts and economic sectors. Notwithstanding the study's constraints, the primary insight stands firm: inequality in the labor market is shaped not by AI alone, but by the choices humans make regarding its governance, training and distribution.

References

- Acemoglu, D., & Johnson, S. (2023). *Power and progress: Our thousand-year struggle over technology and prosperity*. PublicAffairs.
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244.
- Autor, D. H. (2024). Applying AI to rebuild middle-class jobs. *MIT Shaping the Future of Work Working Paper*.
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333.
- Autor, D. H., Mindell, D. A., & Reynolds, E. B. (2022). *The work of the future: Building better jobs in an age of intelligent machines*. MIT Press.
- Autor, D. H., & Salomons, A. (2023). Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity*, 1–63.
- Bogen, M., & Rieke, A. (2018). *Help wanted: An examination of hiring algorithms, equity, and bias*. Upturn.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. *Quarterly Journal of Economics*, 139(4), 1961–2003.
- Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E., & Tavares, M. M. (2024). *Gen-AI: Artificial intelligence and the future of work*. International Monetary Fund.
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv*.
- 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.
- Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality*. International Labour Organization.

- International Labour Organization. (2025). *World employment and social outlook: Trends 2025*. International Labour Organization.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410.
- OECD. (2023). *OECD employment outlook 2023: Artificial intelligence and the labour market*. OECD Publishing.
- OECD. (2024). *Employment outlook 2024: The net-zero and AI transition in labour markets*. OECD Publishing.
- Pizzinelli, C., Panton, A. J., Tavares, M. M., Li, L., Melina, G., & Cazzaniga, M. (2023). Labor market exposure to artificial intelligence: Cross-country differences and distributional implications. *IMF Working Paper*.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68.
- World Economic Forum. (2025). *Future of jobs report 2025*. World Economic Forum.