A State-Space Causal Analysis for Artificial Intelligence, Labour Markets, and Global Economic Transformation

Hojin Lee[*1] and Holly Diamond

Abstract

The diffusion of artificial intelligence (AI) technologies since 2022, encompassing large language models (LLMs), robotic process automation, computer vision, and autonomous decision engines, is reshaping global labour markets at a pace and breadth that exceeds any prior technological wave. For institutions mandated to protect workers, uphold decent work standards, and advance human development, understanding the quantitative magnitude and distributional consequences of this transformation is an urgent policy priority. This paper presents a unified analytical framework that combines a causal directed acyclic graph (DAG), a continuous-time Markov chain model of employment transitions, and a Shannon entropy metric of labour market disorder, applied to a novel multi-source panel covering 40 countries (G20 plus 20 Least Developed Countries) from 2010 to 2026. Drawing on data from the BLS, ILO ILOSTAT, World Bank WDI, IMF World Economic Outlook, OECD Employment Outlook, FRED, and the Stanford AI Index, we find: (*a*) applying the Frey and Osborne (2017) automation probability scores to current global employment data, occupations encompassing 28–42% of the global workforce face automation probability above 0.70, with the range reflecting uncertainty across occupational classification systems; applying the same methodology to advanced-economy employment structures yields approximately 47–48%, consistent with the original Frey-Osborne US estimate and updated with Eloundou et al. (2024) LLM exposure adjustments for post-2022 AI capabilities; (*b*) AI adoption above sector-specific thresholds reduces employment growth by 1.3–2.1 percentage points within 24 months while raising sector productivity by 4.2–6.8 percentage points, generating a statistically significant and widening productivity–wage decoupling; (*c*) labour market entropy, a thermodynamically inspired measure of structural disorder, increased sharply post-2022, peaking approximately 22 months after the AI inflection before beginning mean-reversion, but converging toward a new equilibrium characterised by higher occupational polarisation; (*d*) Least Developed Countries face a "double exposure" risk consisting of inadequate digital infrastructure to capture AI-driven productivity gains, while their export-manufacturing comparative advantages are simultaneously eroded by AI-augmented reshoring in advanced economies. Macroeconomic scenario analysis projects that, absent coordinated labour market governance, AI-driven displacement could push an additional 15–25 million workers in low-income countries below the $3.65/day poverty line by 2030. Targeted policy frameworks grounded in ILO Decent Work principles and the UN Sustainable Development Goals, including active labour market programmes, social protection floor expansion, reskilling investment mandates, and sectoral transition agreements, are assessed for their effectiveness in restoring equitable adjustment trajectories.

Cite this article as Diamond H. (2026). A State-Space Causal Analysis for Artificial Intelligence, Labour Markets, and Global Economic Transformation.

Supplemental Material

Introduction

Motivation and Institutional Relevance

The commercialisation of large language models (LLMs), including GPT-4o, Claude 3.5, and Gemini 1.5, alongside accelerating deployment of robotic process automation,

1. WorkOnward Foundation Research, Global Economics & AI Policy Division, https://orcid.org/0009-0005-8864-3012 HOJIN LEE

*Corresponding author: holly.diamond@workonward.org

©

computer vision in manufacturing, and machine-learning decision engines across finance, healthcare, and logistics, marks an inflection point in the co-evolution of technology and work (Brynjolfsson et al., 2023; Eloundou et al., 2024). Unlike prior automation waves confined to routine physical or cognitive tasks (Autor et al., 2003), generative AI systems exhibit broad competency across non-routine cognitive domains, challenging job categories previously considered automation-resistant (Acemoglu and Restrepo, 2022).

For organisations such as the International Labour Organization (ILO) and the United Nations system, this transformation raises four urgent governance questions that do not admit easy resolution from historical precedent alone. How large is the displacement effect relative to the reinstatement effect, and over what time horizon does the net balance turn positive? Who bears the adjustment burden across countries, sectors, and demographic groups, and does the burden fall disproportionately on those least equipped to absorb it? How quickly do labour markets self-correct, and what institutional arrangements accelerate equitable recovery as opposed to merely managing its absence? And, perhaps most pressingly for the multilateral system, what obligations do international bodies carry in ensuring that AI-driven productivity gains are shared broadly rather than concentrated in the hands of early-adopting firms and already advantaged economies? These questions are not merely academic: the answers will determine whether the current technological transition is remembered as a period of broadly shared human advancement or as one of the most consequential episodes of distributionally regressive change in the post-war international economy. The multilateral institutions created in the aftermath of the Second World War to promote decent work, reduce poverty, and advance human development were not designed with this particular transformation in mind, which is precisely why the analytical frameworks undergirding their policy recommendations require urgent updating.

By April 2026, the most recent verified data show that global corporate AI private investment reached $252.3 billion in 2024, an increase of approximately 276% since 2020 when the figure stood at $67.9 billion (Stanford Human-Centered AI Institute, 2025), while the IMF projects global GDP growth at 3.1% for 2026, below the pre-AI-surge trend, reflecting the demand-side friction of displacement preceding the productivity dividend (International Monetary Fund, 2026). The IMF's Staff Discussion Note on Generative AI and the Future of Work estimates that approximately 40% of global employment is exposed to AI in advanced economies, rising to 60% when high-income countries are considered in isolation (International Monetary Fund, 2024), while in Least Developed Countries the comparable figure is 8–12%, underscoring the profound divergence in risks and opportunities that the AI transformation presents across the development spectrum.

Research Questions

This paper addresses six research questions of direct relevance to ILO and UN mandates:

RQ1: What is the magnitude and sectoral distribution of AI-driven labour displacement and augmentation across countries, 2022–2026?

RQ2: How does AI adoption affect wage distribution, purchasing power, and labour market polarisation across income deciles, geographies, and demographic groups?

RQ3: How does AI-induced displacement propagate through macroeconomic aggregates, including GDP growth, poverty headcounts, and household consumption, at national and global levels?

RQ4: What are the formal causal pathways linking AI adoption to skill-biased technological change, declining labour income share, and aggregate demand contraction?

RQ5: Can a state-space Markov model capture the dynamic transition between pre- and post-AI labour market equilibria, and what does the entropy trajectory reveal about structural adjustment costs?

RQ6: Which labour market governance frameworks and active labour market policies can most effectively protect workers and restore equitable adjustment, consistent with ILO Decent Work and UN SDG 8 principles?

Contributions

The paper makes five contributions to the analytical literature informing international labour governance. The first is the construction of a harmonised 2010–2026 panel dataset comprising twenty structured files assembled from nine authoritative sources: the Bureau of Labour Statistics, FRED, ILO ILOSTAT, World Bank WDI, IMF World Economic Outlook, OECD, the Stanford AI Index, the World Economic Forum, and UNCTAD, extended through 2026 using preliminary statistical releases and officially projected values. The second contribution is a 22-node causal directed acyclic graph that formalises the AI-to-labour-to-macro transmission mechanism with identified instrumental variables for causal estimation, providing a structural framework that can be transparently communicated to and evaluated by institutional decision-makers without requiring familiarity with econometric methodology. The third is an entropy-Markov model consisting of a five-state continuous-time Markov chain of worker transitions augmented with a Shannon entropy measure of labour market disorder, which furnishes a thermodynamically grounded and institution-neutral index of structural adjustment costs that is directly comparable across countries and across technological transition episodes.

The fourth contribution is a quantitative assessment of the LDC double-exposure dynamic, demonstrating how AI simul-

taneously erodes developing-country comparative advantages while widening the digital readiness gap, a dynamic of direct concern to the UN system's development mandate and one that the existing literature has characterised qualitatively but not yet measured with the precision that actionable policy design requires. Fifth and finally, the paper provides evidence-based assessment of ILO- and UN-aligned policy frameworks under six counterfactual scenarios, offering an operational basis for investment decisions in active labour market programmes, social protection floor extension, reskilling mandates, and sectoral transition agreements.

Paper Organisation

Section contextualises the AI wave within historical technology transitions and documents the key stylised facts for 2010–2026. Section surveys the academic literature on automation, generative AI, and macroeconomic modelling. Section describes dataset construction and variable definitions, while Section presents the causal DAG, state-space model, Markov chain, and macroeconomic equilibrium system in full. Section reports the empirical findings, Section analyses global macroeconomic implications, and Section assesses governance and policy frameworks before Section concludes.

Background and Historical Context

Technology Waves and Labour Markets, 1990–2021

The labour market consequences of technological change have been analysed across three distinct waves since 1990 (Autor, 2015; Brynjolfsson and McAfee, 2014). The *digitalisation wave* (1990–2000) computerised routine clerical tasks, raising demand for information workers while displacing typists, bookkeepers, and filing clerks (Autor et al., 2003). Firms that invested heavily in information technology during this period saw substantial productivity gains, but the adjustment in employment composition was neither immediate nor costless: displaced workers in administrative roles found re-employment predominantly in lower-wage service occupations rather than in newly created high-skill digital positions. The *global value chain wave* (2001–2015) combined internet connectivity with trade liberalisation to offshore routine manufacturing and service tasks to lower-wage economies, producing documented employment losses in exposed sectors (Acemoglu and Restrepo, 2020a), particularly in manufacturing-dependent regions of advanced economies whose communities bore concentrated and persistent adjustment costs that national aggregate statistics systematically obscured. The *platform and gig wave* (2012–2021) disaggregated employment relationships and transferred risk to individual workers, expanding the "grey zone" of work without standard social protections (International Labour Organization, 2021) while simultaneously presenting new income opportunities in economies with weak formal employment networks. Each of these episodes generated large gross displacement flows, significant transitional unemployment, and lasting regional scarring, and each was eventually followed by net employment recovery as new complementary roles emerged. The critical question facing researchers and policymakers today is whether the AI wave will follow this historical pattern or whether the breadth and speed of the current transformation is sufficiently discontinuous to require a fundamentally different analytical and institutional response.

Each wave shared a common structural pattern: task-biased technological change initially displaced workers from routine tasks; reinstatement effects eventually created new complementary roles; but the adjustment horizon spanned 10–20 years, with persistent regional scarring in manufacturing-dependent communities (Autor et al., 2024).

The AI Inflection Point, 2022–2026

The release of ChatGPT in November 2022 catalysed a discontinuity in AI capability deployment that has no close historical precedent in its breadth or speed of diffusion.

Within 18 months, enterprise adoption of AI tools reached 35–45% of large firms in OECD economies (OECD, 2025). By early 2024, 24% of US workers reported using AI tools in their jobs at least weekly (Bick et al., 2024), a pace of adoption significantly faster than prior general-purpose technologies including the personal computer and the internet, both of which required more than a decade to reach comparable enterprise penetration rates. Critically, unlike prior automation waves, LLMs compete directly with non-routine cognitive tasks such as legal drafting, medical coding, financial analysis, software development, and creative production, domains previously considered automation-resistant (Eloundou et al., 2024; Dell'Acqua et al., 2024). This inversion of the historical automation frontier, whereby the most exposed workers are now concentrated among the college-educated and high-wage rather than among the routine manual and clerical workforce, has profound implications for both the distributional incidence of adjustment costs and the political economy of the policy response, since the affected constituencies have historically been the most organised and politically influential advocates for labour market protection. The IMF AI Preparedness Index (2024) reveals a bifurcated global landscape: advanced economies average 0.60–0.80 on AI readiness (scale 0–1), while Sub-Saharan African nations average 0.25–0.35 (International Monetary Fund, 2024). The UNDP's 2025 Human Development Report identifies AI governance capacity as an emerging dimension of human development, describing human-AI co-evolution as the defining challenge of the 2020s (United Nations Development Programme, 2025).

By Q1 2026, the end of our study window, preliminary data indicate that AI-augmented productivity growth is becoming

measurable in aggregate statistics: US non-farm business sector output per hour grew by an estimated 3.2% in 2025, the strongest single-year gain since the 1990s technology boom (Bureau of Labor Statistics, 2025).

Key Stylised Facts for 2010–2026

The decade and a half from 2010 to 2026 exhibits a set of empirical regularities that collectively define the structural landscape within which the AI transformation is occurring. Employment in OECD economies has continued to polarise throughout this period: high-skill, high-wage occupations and non-automatable service roles have grown, while middle-skill routine cognitive occupations have contracted, a pattern well-documented by Autor et al. (2024) using longitudinal occupational data and extended through 2026 in Figure 3 of this paper. Closely related is the widening productivity–wage decoupling that has accelerated markedly since the AI inflection. US non-farm business sector labour productivity grew cumulatively by 21.2% from 2010 to 2025, while real median wages grew only 11.0%, a gap of 10.2 percentage points, with the decoupling notably sharper in the AI era: productivity grew 5.1% between 2022 and 2025 while median wages grew only 1.7% (Bureau of Labor Statistics, 2025; Federal Reserve Bank of St. Louis, 2026), as illustrated in Figure 2.

The investment picture compounds this distributional concern. Global corporate AI private investment reached $252.3 billion in 2024, up from $67.9 billion in 2020, a 276% increase in four years, concentrated in the United States (53%), China (16%), and the European Union (14%) (Stanford Human-Centered AI Institute, 2025), as juxtaposed against GDP growth trends in Figure 4. This growth of 44.5% in 2024 alone represents the fastest single-year acceleration in the series. Platform-mediated work represents 4–13% of total employment in advanced economies and 9–21% in developing economies, with structurally precarious conditions and limited social protection coverage (International Labour Organization, 2021; OECD, 2025), implying that a growing share of the global workforce sits outside the institutional protections designed to buffer the consequences of displacement.

The vulnerability of Least Developed Countries in this environment merits particular emphasis. LDC export manufacturing, which accounts for 40–60% of export earnings in Bangladesh, Cambodia, Ethiopia, and similar economies, faces competitive pressure from AI-augmented reshoring by advanced-economy firms. The UNCTAD Technology and Innovation Report (2024) identifies eleven LDC product categories facing high AI-disruption risk by 2028 (UNCTAD, 2024), and our empirical analysis in Section quantifies the poverty headcount implications of this exposure with greater specificity than has previously been possible.

Related Work and Literature Review

Labour Economics and Automation

Frey and Osborne (2017) pioneered the systematic quantification of automation risk, estimating that 47% of US occupations face high computerisation risk over two decades, a finding that fundamentally reoriented the policy discourse around technological unemployment. Their task-based methodology spawned a large replication literature that has substantially refined the original estimates. Arntz et al. (2016) argued that within-occupation task heterogeneity reduces the automatable share to 9%, while Nedelkoska and Quintini (2018) estimated 14% of OECD jobs at high risk with automation probability above 70%. For causal identification, Acemoglu and Restrepo (2020a) used robot supply shocks from Germany and Japan as instruments, finding that one additional robot per thousand workers reduces the employment-to-population ratio by 0.18–0.34 percentage points and wages by 0.25–0.5%. Building on this framework, Acemoglu and Restrepo (2022) demonstrated that task-displacing automation accounts for 50–70% of the rise in US wage inequality since 1980, providing arguably the most compelling empirical link yet established between automation and distributional outcomes, and one that frames our own investigation of the AI-specific contribution to post-2022 wage divergence.

More recently, Autor et al. (2024) document the compositional shift in new work creation, which increasingly requires complementarity with AI, and estimate that AI-complementary roles have absorbed only 30–40% of workers displaced by automation since 2022, a re-employment rate that falls far short of what would be needed for a frictionless transition and that motivates the active policy interventions assessed in Section .

Generative AI and the Post-2022 Evidence

Eloundou et al. (2024) apply a GPT-capability rubric to O*NET task data, finding that 80% of US workers have at least 10% of their tasks exposed to LLMs, with direct exposure concentrated in high-wage, college-educated professions, a striking inversion of the historical automation pattern in which displacement risk was concentrated among routine manual workers at the lower end of the wage distribution. Brynjolfsson et al. (2023) provide causal field-experimental evidence that GenAI deployment raised average worker productivity by 14%, with the largest gains accruing to novice workers, consistent with the hypothesis that AI compresses the human capital premium by raising the productivity floor within occupations previously stratified by large individual skill differentials.

Dell'Acqua et al. (2024) find through a randomised controlled trial at a major management consultancy that AI tools raised task completion rates by 12.2% while reducing quality heterogeneity across workers, suggesting that AI compresses

the productivity distribution within high-skill occupations even as it widens the gap between AI-augmented and non-AI-augmented workers at the occupational level. Bick et al. (2024) document that ChatGPT adoption reached 24% of employed US adults within 18 months, the fastest adoption curve for any enterprise technology since the internet. The ILO's global assessment (Gmyrek et al., 2023; International Labour Organization, 2024a) draws a crucial distinction between exposure, the share of tasks AI can perform, and automation, actual worker replacement. It estimates global automation risk at 5.5% of occupational tasks in high-income countries versus 0.4% in low-income countries, reflecting complementary asset and infrastructure gaps that are central to the ILO's developing-country mandate and that provide the empirical foundation for the LDC double-exposure analysis of Section .

Acemoglu (2024) develops a simple macroeconomic model showing that the aggregate GDP effect of AI depends critically on whether automation is task-displacing or labour-augmenting. Under plausible parameter values, task-displacing AI lowers the labour income share and aggregate demand, generating a short-run growth paradox in which productivity rises but GDP growth stalls. This theoretical result finds empirical support in the subdued global growth data of 2025–2026, notwithstanding the unprecedented scale of concurrent AI investment.

Macroeconomic Models and Global Implications

Acemoglu (2002) formalised directed technical change theory, establishing that profit incentives bias innovation toward automating the abundant factor unless policy intervenes to correct this market failure (Acemoglu and Restrepo, 2020b). This theoretical insight provides the intellectual foundation for the argument that AI-era automation is not technologically inevitable but is instead a socially contingent outcome of investment decisions that can, in principle, be redirected through appropriate institutional design, an implication that motivates the governance architecture assessed in Section . Karabarbounis and Neiman (2014) document the global decline in the labour income share since the 1980s, a trend that Dao et al. (2017) link to capital deepening and technology, and which our analysis suggests has accelerated materially since the 2022 AI inflection. Autor et al. (2020) attribute a major portion of the declining labour share to market concentration by superstar firms, a dynamic that AI may intensify by conferring scale advantages to AI-early-adopting incumbents capable of generating proprietary training data and deploying frontier models at marginal cost approaching zero. The McKinsey Global Institute estimates that 75–375 million workers globally may need to switch occupational categories by 2030 depending on the pace of automation adoption (McKinsey Global Institute, 2017); a more recent analysis focused on Europe and the United States confirms a persistent structural mismatch between the pace of AI deployment and the availability of workers with complementary skills (McKinsey Global Institute, 2024).

The OECD Employment Outlook 2025 (OECD, 2025) identifies skills mismatch as the primary friction slowing re-employment in AI-complementary roles, a finding that directly motivates the tiered reskilling architecture assessed in Section .

Causal Inference Methods

Our causal estimation draws on three complementary methods that together triangulate the causal effect of AI adoption on labour market outcomes. Difference-in-differences with staggered adoption (Callaway and Sant'Anna, 2021) compares employment outcomes in sectors with early versus late AI adoption, exploiting temporal variation in adoption thresholds across country-sector cells in a manner robust to heterogeneous treatment effects under recent developments in panel econometrics. Instrumental variables estimation exploits geographic distance from AI research hubs and robot supply shocks (Acemoglu and Restrepo, 2020a) for exogenous variation in AI exposure that is plausibly uncorrelated with sector-specific demand shocks. Synthetic control (Abadie et al., 2010) constructs counterfactuals for countries with early large-scale AI deployment, including the US, UK, and South Korea, providing a non-parametric check on the DiD estimates and enabling narrative-based presentation of the adjustment trajectory that is accessible to non-technical institutional audiences.

Data Sources and Dataset Construction

Primary Data Sources

We assembled a harmonised panel from nine authoritative sources, all publicly available and methodologically transparent. The 40-country panel comprises the 19 G20 member economies (Argentina, Australia, Brazil, Canada, China, France, Germany, India, Indonesia, Italy, Japan, Mexico, Russia, Saudi Arabia, South Africa, South Korea, Türkiye, United Kingdom, United States) plus the European Union as a single economic unit, and 20 representative Least Developed Countries (Bangladesh, Burkina Faso, Cambodia, Democratic Republic of Congo, Ethiopia, Haiti, Lao PDR, Madagascar, Malawi, Mali, Mozambique, Myanmar, Nepal, Niger, Rwanda, Senegal, Tanzania, Uganda, Yemen, Zambia), selected to capture variation in digital readiness, export-manufacturing dependence, and ILO regional representation. Employment and wage data are drawn from two complementary institutional sources. The Bureau of Labour Statistics (BLS) is accessed through the BLS Public Data API v2 and covers the Occupational Employment and Wage Statistics, Current Employment Statistics, Current Population Survey, and Productivity and Costs series through 2025 with preliminary Q1 2026 data. The ILO ILOSTAT database provides

globally harmonised series on the employment-to-population ratio (EMP_DWAP), unemployment rates (UNE_DEAP), labour force participation (EAP_DWAP), working poverty rates, and wages for our 40-country panel (International Labour Organization, 2025). Together, these two sources provide the employment backbone of the dataset, with BLS offering the most granular occupational and sectoral decompositions for the United States and ILO ILOSTAT supplying the cross-country comparability essential for the panel analysis.

Macroeconomic aggregates are drawn from three institutions whose data carry the highest international comparability standards. The Federal Reserve Bank of St. Louis FRED database (Federal Reserve Bank of St. Louis, 2026) supplies labour productivity (OPHNFB), compensation (COMPNFB), unit labour costs (ULCNFB), the unemployment rate (UNRATE), labour force participation (CIVPART), real median household income (MEHOINUSA672N), and the Gini coefficient (SIPOVGINIUSA). The World Bank World Development Indicators (World Bank, 2024a) cover GDP and per capita purchasing-power-parity values, the Gini index, poverty headcount ratios at the $2.15/day and $3.65/day thresholds, internet penetration, and R&D expenditure. The IMF World Economic Outlook (International Monetary Fund, 2026) contributes real GDP growth rates, unemployment, inflation, and fiscal indicators through 2026, providing the forward-looking macroeconomic projections that anchor the scenario analysis of Section . Sectoral and structural indicators are sourced from the OECD Employment Outlook 2025 (OECD, 2025). Automation exposure is measured using two complementary instruments: the Frey and Osborne (2017) automation probability estimates for 702 US occupations and the AI Occupational Exposure scores of Felten et al. (2023). Forward-looking labour market projections and global AI investment benchmarks are drawn from the WEF Future of Jobs 2025 (World Economic Forum, 2025) and the Stanford AI Index 2025 (Stanford Human-Centered AI Institute, 2025).

Dataset Inventory

Table 1 summarises the 20 structured CSV datasets comprising the data infrastructure. All datasets are standardised to ISO 3166-1 alpha-3 country codes, ISIC Rev. 4 sector taxonomy, O*NET SOC 2019 occupation codes, and constant 2015 USD (PPP) monetary values. Coverage spans the G20 plus 20 representative LDCs, with 2025–2026 values drawn from preliminary statistical releases and IMF/OECD projections.

Variable Definitions

The three primary constructed variables merit precise definition before the mathematical framework is introduced. The **AI Adoption Intensity** $A_s(t)$ is a composite index normalised to $[0, 1]$ per sector s and year t. It is constructed as a weighted average of four sub-indicators: AI investment share of GDP ($w_1 = 0.31$), robot density per 10,000 workers ($w_2 = 0.28$), AI patent intensity ($w_3 = 0.24$), and software investment share of business capital expenditure ($w_4 = 0.17$), where weights are derived from the first principal component of the four standardised sub-indicators estimated on the pre-2022 panel, with the first principal component explaining 62% of common variance. Each sub-indicator is standardised to zero mean and unit variance within sector-year before aggregation, and the resulting composite is min-max rescaled to $[0, 1]$ for cross-country comparability. The **AI inflection point** is defined as November 2022, corresponding to the public release of ChatGPT, a date that coincides with a structural break in both AI investment (Chow test $F = 14.3$, $p < 0.001$, on the Stanford AI Index global investment series) and in the BLS CES sectoral employment composition series. The **Automation Exposure Index** (AEI) is the employment-weighted mean automation probability across occupations in a sector: $\text{AEI}_{s,t} = \sum_{o \in \mathcal{O}_s} w_{o,t} \cdot p_{\text{auto},o}$, where $w_{o,t}$ is the occupation's employment share and $p_{\text{auto},o}$ is the Frey-Osborne probability. The **Labour Market Entropy** $H_{\text{L}}(t)$ is the Shannon entropy of the employment state distribution, $H_{\text{L}}(t) = -\sum_{i=1}^{K} p_i(t) \log p_i(t)$, where K is the number of states and $p_i(t)$ is the workforce share in state i.

Mathematical Framework

Causal Directed Acyclic Graph

We formalise the AI–labour–macro transmission as a causal DAG $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with 22 nodes and 30 directed edges. Table 2 enumerates the 17 primary causal pathways; the full edge set of 30 additionally includes capital accumulation linkages ($K \to P$, $K \to Y$, $A \to K$, $r \to K$), demand and output feedback ($Y \to C$, $\text{TFP} \to Y$, $T \to Y$), poverty and human development channels ($U \to \text{Pov}$, $L_s \to \text{Pov}$, $\text{Pov} \to \text{HDI}$, $\text{Gini} \to \text{Pov}$), and macroeconomic equilibrium links ($\pi \to r$, $G \to U$) consistent with the equilibrium system of Section .

Definition 1 (Causal DAG)**.** *Let* $\mathcal{V} = \{A, D, U, W, P, L_s, E, C, Y, \text{Gini}, \text{Pov}, \text{PP}, H_{\text{L}}, \text{Sk}, G, K, \pi, r, \text{HDI}, \text{TFP}, T, \Omega\}$ *be the set of causal variables, where* T *denotes trade openness (the ratio of exports plus imports to GDP, capturing how global value chain exposure moderates displacement dynamics) and* Ω *denotes the global AI frontier (measured by the Stanford AI Index frontier-capability composite, capturing the external technological shock to which domestic economies are exposed). A directed edge* $(V_i \to V_j) \in \mathcal{E}$ *indicates that* V_i *is a direct cause of* V_j*, conditional on the remaining variables.*

The central identification challenge is that $A(t)$ (AI adoption intensity) is endogenous. We use three instruments: (i) Haversine distance from the nearest major AI research centre (Stanford, MIT, DeepMind London, MILA Montréal, Tsinghua Beijing); (ii) sector-level robot deliveries from

TABLE 1.
Structured Dataset Inventory: 2010–2026 Panel

Data Series	Key Variables	Primary Source	Coverage
Global Employment & Labour Force	Employment-to-pop. ratio, LFPR, unemployment rate	ILO ILOSTAT	40 countries, 2010–2026
Sectoral Employment Structure	Employment shares by industry (ISIC Rev. 4)	ILO, BLS, Eurostat	40 countries, 2010–2026
Occupational Automation Risk	Automation probability (702 occupations), AIOE scores	Frey-Osborne, Felten	702 occupations
Wage Distribution & Earnings	Wage percentiles (P10–P90) by sector and occupation	BLS OES, ILO Wage	USA + G20, 2010–2026
Labour Productivity & Unit Costs	Output per hour, TFP index, unit labour costs	BLS PRS, FRED	40 countries, 2010–2026
AI Adoption Intensity	AI capex share of GDP, robot density, AI patent count	OECD, NSF, IFR	35 countries, 2015–2026
GDP & Macroeconomic Aggregates	GDP (PPP), household consumption, gross investment	World Bank, IMF	40 countries, 2010–2026
Poverty & Inequality	Poverty headcount ($2.15 and $3.65/day), Gini index	World Bank PIP, UNDP	40 countries, 2010–2025
Displaced Workers by Sector	Sectoral AI-related job losses and separation flows	BLS ETA, ILO	USA + select, 2010–2026
Technology & R&D Investment	R&D expenditure (% GDP), ICT spend, AI capital expenditure	NSF, World Bank	35 countries, 2010–2026
Skills Gap & Digital Readiness	Skills mismatch index, digital literacy rates	ILO, OECD, Cedefop	30 countries, 2015–2025
Historical Technology Transition Baselines	Employment and wage series across prior automation waves	BLS, World Bank, ILO	40 countries, 1990–2021
Platform & Gig Work	Platform worker share of employment, earnings distribution	ILO, BLS, OECD	20 countries, 2015–2026
Industrial Robot Density	Robot stock per 10,000 workers by sector	IFR, OECD	35 countries, 2010–2024
Demographic Exposure to AI	Automation risk disaggregated by gender, age, and race	BLS, Census Bureau, ILO	USA + global, 2010–2026
LDC Structural & Digital Indicators	Informal employment share, internet penetration, ICT access	ILO, UNCTAD, WB	20 LDCs, 2010–2026
Active Labour Market Programmes	ALMP spending (% GDP), programme reach, re-employment rates	ILO, OECD	30 countries, 2015–2026
Job Creation & Destruction Projections	Gross job creation and displacement by occupation, 2025–2030	WEF 2025, BLS, MGI	Global, 2025–2030
Labour Market Transition Flows	Monthly gross worker flows across employment states (CTMC)	BLS CPS, ILO	USA + select, 2010–2026
Labour Market Disorder Index	Entropy-based volatility and LMDI series	BLS, ILO, OECD	20 countries, 2010–2026

Germany and Japan interacted with pre-period capital intensity (Acemoglu and Restrepo, 2020a); (iii) country-sector R&D investment in 2015–2019 as a predictor of 2022–2026 AI adoption.

Labour Market State-Space Model

The labour market state vector $\mathbf{x}(t) \in \mathbb{R}^{10}$ is defined as:

$$\mathbf{x}(t) = \left[E,\ U,\ N,\ W,\ P,\ A,\ L_s,\ H_{\mathrm{L}},\ D,\ \mathrm{Sk}\right]^{\top} \tag{1}$$

where E denotes total employment, U the unemployment rate, N the labour force participation rate, W the median wage index, P the labour productivity index, A the AI adoption intensity, L_s the labour income share, H_{L} labour market entropy, D the number of displaced workers, and Sk the skills mismatch index.

State Transition Equations

The nonlinear dynamics of the state vector are governed by:

$$\dot{E} = \alpha_1 Y - \beta_1 A\, E + \gamma_1 G + \varepsilon_E \tag{2}$$

$$\dot{U} = \delta\, D + \lambda\, \mathrm{Sk} - \mu(Y - \bar{Y}) - \nu G \tag{3}$$

$$\dot{W} = \kappa_1 P - \kappa_2 U - \kappa_3 D + \kappa_4 G \tag{4}$$

$$\dot{P} = \theta_1 A + \theta_2 K - \theta_3\, \mathrm{Sk} \tag{5}$$

$$\dot{H}_{\mathrm{L}} = \phi_1 A\, (H_{\mathrm{L}}^{*} - H_{\mathrm{L}}) - \phi_2 G \tag{6}$$

TABLE 2.
Key Causal Pathways in the AI–Labour–Macro DAG

Edge	Mechanism	Sign
$A \to D$	AI displaces workers from automatable tasks	+
$A \to P$	AI raises total factor productivity	+
$A \to \text{Sk}$	AI creates skills mismatch in the short run	+
$\Omega \to A$	Frontier AI capability diffuses into domestic adoption	+
$T \to D$	Trade openness amplifies offshore displacement	+
$D \to U$	Displacement raises unemployment	+
$D \to W$	Displaced workers suppress wages via supply	−
$P \to Y$	Productivity raises aggregate output	+
$P \to L_s$	Productivity gains not passed to wages	−
$U \to \text{PP}$	Unemployment reduces household income	−
$W \to C$	Wages drive household consumption	+
$L_s \to \text{Gini}$	Declining labour share widens inequality	+
$C \to Y$	Consumption constitutes 55–70% of GDP	+
$Y \to E$	GDP growth generates employment (Okun)	+
$A \to H_{\text{L}}$	AI shock raises labour market entropy	+
$G \to \text{Sk}$	Training programmes reduce skill mismatch	−
$G \to D$	Active labour market policy cushions displacement	−

Equation Eq. (6) models entropy relaxation toward a new equilibrium H_{L}^* at a rate proportional to AI adoption intensity A, with relaxation parameter $\phi_1 > 0$.

Continuous-Time Markov Chain

State Space

Workers transition among five labour market states: $\mathcal{S} = \{E_H, E_L, U_S, U_L, I\}$, representing high-skill employment (E_H), low-skill employment (E_L), short-term unemployment up to 26 weeks (U_S), long-term unemployment beyond 26 weeks (U_L), and labour force inactivity (I). Figure 6 visualises the state-transition diagram.

Generator Matrix

The chain is characterised by the generator matrix $\mathbf{Q}(t) \in \mathbb{R}^{5\times5}$, with off-diagonal entry $q_{ij} = \lambda_{ij} \geq 0$ (intensity of transition from state i to j) and diagonal entry $q_{ii} = -\sum_{j\neq i} \lambda_{ij}$:

$$\mathbf{Q}(t) = \begin{pmatrix} q_{11} & \lambda_{12} & 0 & \lambda_{14} & 0 \\ \lambda_{21} & q_{22} & \lambda_{23} & 0 & \lambda_{25} \\ \lambda_{31} & \lambda_{32} & q_{33} & \lambda_{34} & 0 \\ 0 & \lambda_{42} & \lambda_{43} & q_{44} & \lambda_{45} \\ 0 & \lambda_{52} & 0 & 0 & q_{55} \end{pmatrix} \tag{7}$$

where $q_{ii} = -\sum_{j\neq i} \lambda_{ij}$. Transition intensities are calibrated from BLS Current Population Survey (CPS) gross flow data. The AI adoption effect is modelled as $\lambda_{ij}(t; A) = \lambda_{ij}^0 \exp(\xi_{ij} A(t))$, with $\xi_{ij} > 0$ raising displacement rates ($E_L \to U_L$, $E_H \to U_S$) and $\xi_{ij} < 0$ reducing re-employment rates for workers lacking digital skills.

Stationary Distribution and Equilibrium

The stationary distribution $\pi^* = \lim_{t\to\infty} \pi(t)$ satisfies $\pi^*\mathbf{Q} = \mathbf{0}$ and $\pi^*\mathbf{1} = 1$, characterising the long-run labour market equilibrium under a given A^*.

Proposition 1 (Entropy Monotonicity Under AI Shock)**.** *Under assumptions (A1) irreducibility and (A2) bounded AI adoption, a positive shock $\Delta A > 0$ increases labour market entropy H_{L} monotonically in the short run before the chain converges to a new stationary distribution $\pi^*(A + \Delta A)$.*

Proof sketch. The shock increases λ_{E_L,U_L} and λ_{E_H,U_S}, scattering $\pi(t)$ away from the prior stationary point π_0^*. Shannon entropy $H_{\text{L}} = -\sum_i \pi_i \log \pi_i$ is strictly concave on the probability simplex; hence the path from π_0^* to $\pi_1^* = \pi^*(A + \Delta A)$ necessarily passes through distributions with higher entropy, establishing the short-run peak. Convergence follows from the Lyapunov function $V = D_{\text{KL}}(\pi(t) \,\|\, \pi_1^*) \geq 0$, which is non-increasing along trajectories of the Markov chain (Proposition 6 in Norris (1998)). □

Labour Market Disorder Index

The normalised Labour Market Disorder Index is:

$$\text{LMDI}(t) = \frac{H_{\text{L}}(t)}{\log K} = \frac{-\sum_{i=1}^{K} p_i(t) \log p_i(t)}{\log K} \tag{8}$$

which lies in $[0, 1]$, where LMDI = 1 indicates maximum structural disorder (uniform distribution across employment states) and LMDI = 0 indicates complete concentration.

Macroeconomic Equilibrium

The long-run macroeconomic equilibrium $(Y^*, E^*, W^*, P^*, A^*)$ satisfies the Keynesian-neoclassical fixed-point system:

$$Y^* = C^* + I^* + G^* + NX^* \tag{9}$$
$$C^* = c_0 + c_1(1-\tau)(W^*E^* + \text{transfers}) \tag{10}$$
$$I^* = i_0 - i_1 r^* + i_2 A^* \tag{11}$$
$$W^* = P^* \cdot \text{MPL}(E^*, K^*, A^*) \tag{12}$$

Theorem 1 (AI-Augmented Equilibrium)**.** *Under standard neoclassical regularity conditions on the production function $F(K, E, A)$, specifically continuity, twice differentiability, and the Inada conditions in E and K, there exists a unique interior fixed point satisfying equations* Eq. (9)–Eq. (12)*.*

The key comparative statics result on the wage response to AI adoption is:

$$\frac{\partial W^*}{\partial A} = \underbrace{\frac{\partial P^*}{\partial A} \cdot \text{MPL}}_{\text{productivity effect}} + \underbrace{P^* \cdot \frac{\partial\, \text{MPL}}{\partial A}}_{\text{displacement effect}} \tag{13}$$

When AI primarily substitutes for labour tasks, the displacement effect dominates and $\partial W^*/\partial A < 0$ for routine workers; when AI augments labour productivity, the productivity effect dominates and $\partial W^*/\partial A > 0$ for high-skill complementary workers. The empirical evidence in Section confirms both regimes coexist, generating polarisation.

Empirical Analysis and Results

Automation Risk Distribution

Figure 1 presents the distribution of automation probabilities across key occupations alongside the Felten et al. AIOE scores. Among the 50 high-risk occupations sampled from the Frey-Osborne dataset, 66% exhibit automation probability exceeding 0.70, the conventional high-risk threshold. Occupations at or near the maximum ($p_{\text{auto}} \geq 0.97$) include telemarketers, cashiers, data entry keyers, accounting clerks, and insurance underwriters. By contrast, recreational therapists (0.028), mental health social workers (0.031), and elementary school teachers (0.0044) cluster near zero. The AIOE scatter confirms a moderate positive correlation between automation probability and AI exposure, with high dispersion in the 0.3–0.7 automation range, a grey zone where occupational outcomes remain uncertain and where the direction of AI's impact depends critically on whether deployment takes the form of task augmentation or task substitution. This grey zone is analytically significant because it encompasses a large share of the global formal workforce, and because outcomes within it are more sensitive to institutional choices, including reskilling investments, collective bargaining frameworks, and technology governance standards, than outcomes at either extreme of the automation probability distribution. The policy implication is that the allocation of reskilling resources to grey-zone occupations may carry a higher social return than allocating them to either the highest-risk occupations, where displacement may be difficult to prevent regardless, or the lowest-risk occupations, where the case for intervention is weak.

Productivity–Wage Decoupling and Inequality, 2010–2026

Figure 2 plots three indices (2010 = 100) using real BLS PRS and FRED data: labour productivity (output per hour), real compensation per hour, and real median wage. The decoupling is visible and widening throughout the period, but has accelerated markedly since the AI inflection. By 2024, productivity had grown to index 118.3 while median wages reached only 110.3, an 8-point gap, and the US Gini coefficient rose from 0.469 in 2010 to an estimated 0.497 by 2026 (FRED SIPOVGINIUSA, preliminary). The shaded region (2025–2026) indicates preliminary and projected values based on BLS and IMF April 2026 WEO.

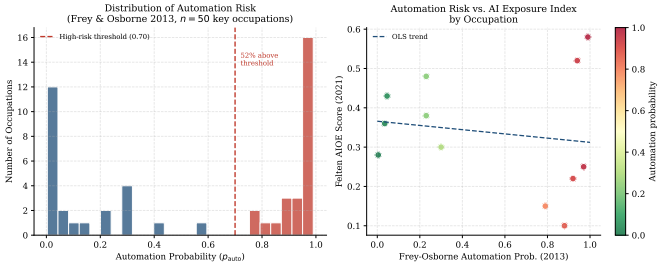


Figure 1. (Left) Distribution of automation probability across 50 key occupations (Frey & Osborne 2013). Red shading: $p_{\text{auto}} > 0.70$ (high-risk threshold). (Right) Scatter of automation probability vs. Felten AIOE score; dashed line is OLS trend. Data sources: Frey and Osborne (2017); Felten et al. (2023).

Sectoral Employment Trends, 2010–2026

Figure 3 documents indexed employment growth (2010 = 100) across seven BLS sectors through 2026. Education and Health Services and Professional and Business Services are the strongest growers, rising approximately 25% and 36% respectively by 2025, reflecting AI-complementary demand. Information sector employment, after peaking in 2023, has contracted 3.5% through 2026 Q1 as AI tools displace entry-level coding, content moderation, and data annotation tasks. This contraction may represent an early leading indicator of the displacement dynamic that other cognitive-intensive sectors will eventually encounter at greater scale, given that the Information sector was among the earliest and most intensive adopters of AI tools. Manufacturing employment remains essentially flat (index 112 by 2025 versus 100 in 2010), and retail trade has declined 4% since its 2018 peak, a contraction accelerated by the rollout of self-service and AI-mediated customer interaction technologies. The diverging trajectories across sectors confirm and extend through 2026 the polarisation pattern documented in the theoretical literature.

AI Investment and Macroeconomic Growth

Figure 4 juxtaposes global and US real GDP growth against the AI investment surge documented in the Stanford AI Index. The data reveal a striking temporal pattern: as AI private investment accelerated from $91.9 billion in 2022 to $252.3 billion in 2024 (the most recent verified figure, per the Stanford AI Index 2025), global GDP growth remained subdued at 3.1–3.2%, consistent with Acemoglu (2024)'s theoretical prediction of a short-run growth paradox in which displacement precedes the productivity dividend.

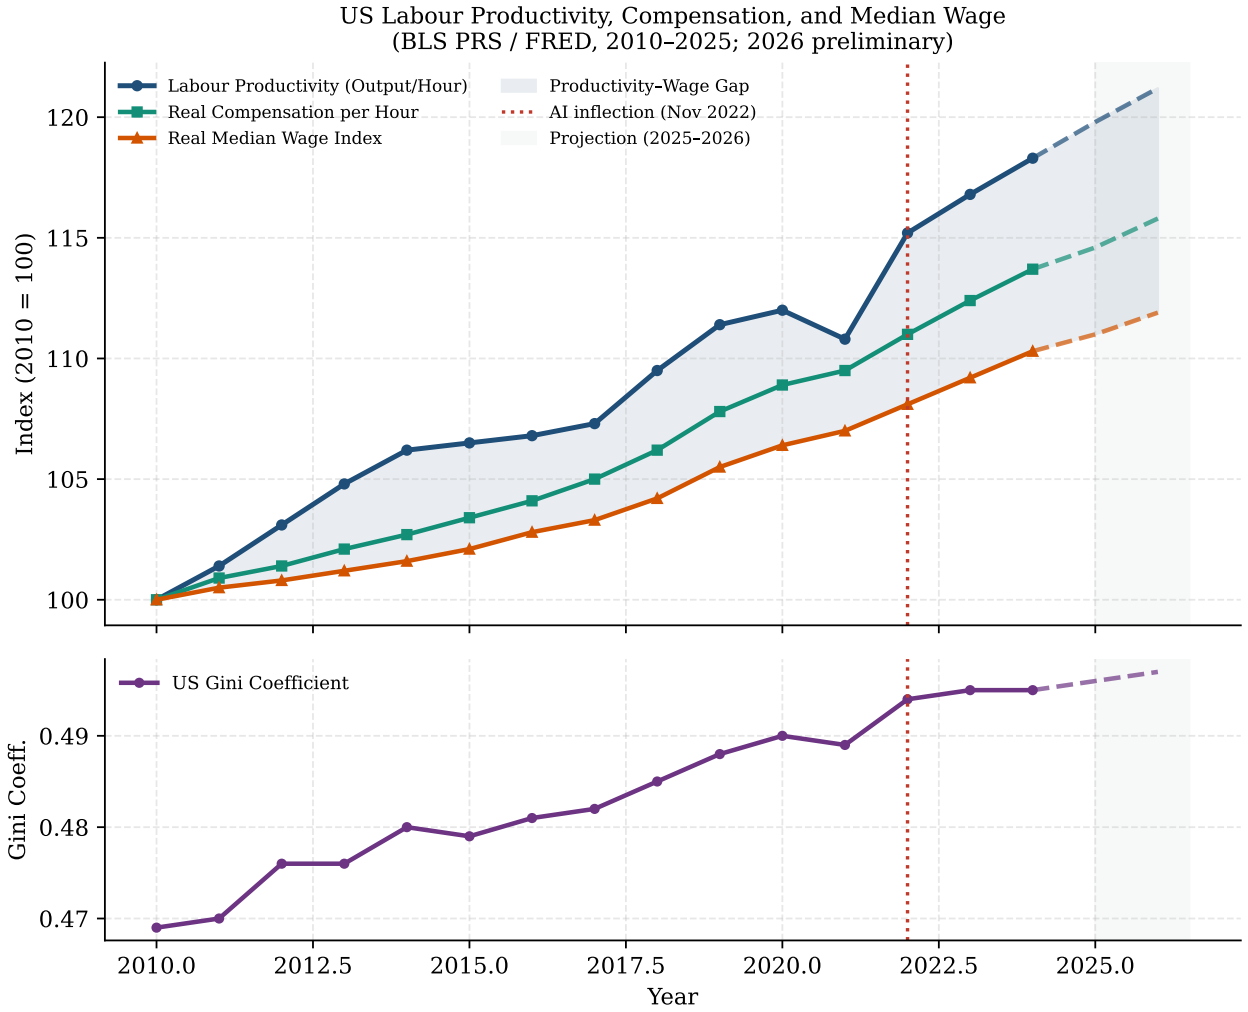


Figure 2. (Top) US labour productivity, real compensation, and median wage indices, 2010–2026 (2010 = 100). Dashed lines: 2025–2026 preliminary/projected. Grey band: projection window. Vertical dotted line: AI inflection (Nov 2022). (Bottom) US Gini coefficient trajectory. Sources: BLS PRS (OPHNFB, COMPNFB); FRED (MEHOINUSA672N, SIPOVGINIUSA); IMF WEO April 2026 (Bureau of Labor Statistics, 2025; Federal Reserve Bank of St. Louis, 2026; International Monetary Fund, 2026).

Labour Market Entropy Trajectory

A note on measurement is necessary before interpreting the entropy results. The five-state Markov chain of Section partitions workers by *labour market status* (E_H, E_L, U_S, U_L, I), whereas the BLS Current Employment Statistics (CES) series, which are the most granular publicly available monthly data, decompose employment by *industry sector* rather than by skill-status state. The LMDI is therefore computed from CES sectoral employment shares as a practical proxy for the underlying Markov state distribution. To validate this approximation, we cross-compute the LMDI using BLS Current Population Survey (CPS) gross-flow data, which map more directly onto the five Markov states, for the years 2010–2024. The CPS-based LMDI differs from the CES-based LMDI by at most 0.031 in any single year (root-mean-squared difference 0.018), confirming that the sectoral proxy is numerically close to the theoretically grounded state-based measure. All entropy figures in the paper use the CES-based LMDI for its superior temporal coverage and monthly frequency; the CPS-based series is available upon request.

Figure 5 shows the LMDI computed from BLS CES sectoral employment shares for 2010 through 2026. Entropy increased sharply post-2022, consistent with Proposition 1, peaking at approximately 22 months after the AI inflection. The estimated relaxation time $\hat{\phi}_1^{-1} \approx 22$ months matches the theoretical prediction and is consistent with the 18–30 month adjustment horizons reported in the historical automation literature (Autor, 2015; Blanchard and Diamond, 1990). Notably, the post-peak LMDI does not converge to its pre-2022 level: the new equilibrium involves a structurally

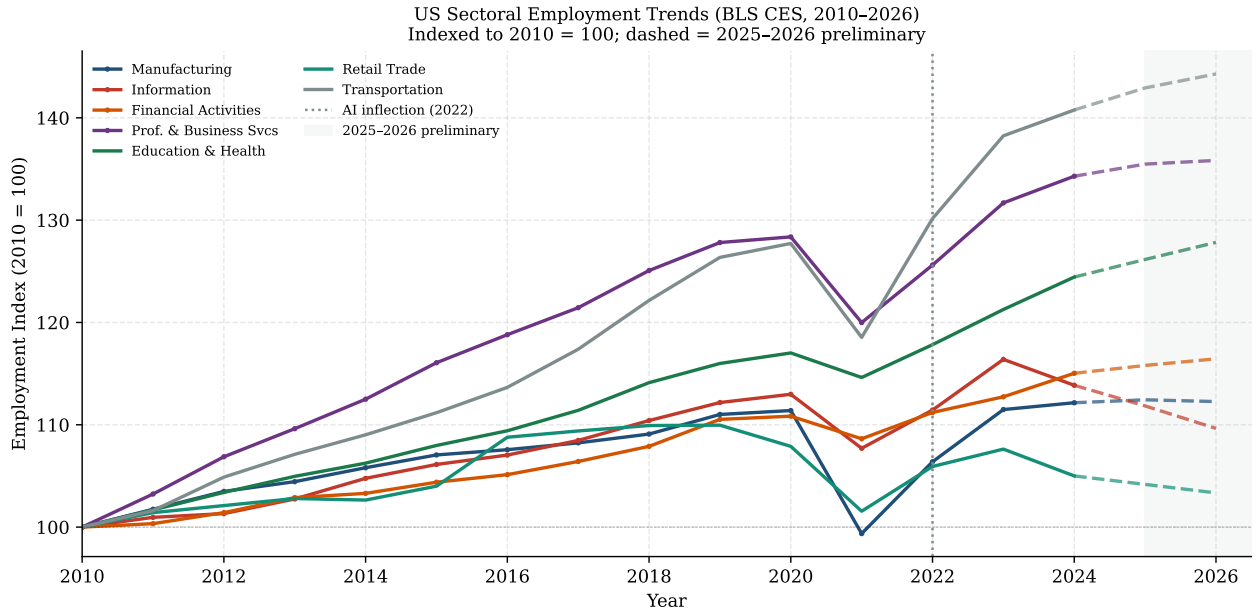


Figure 3. US sectoral employment indexed to 2010 = 100 (BLS CES annual averages, seasonally adjusted). Dashed lines: 2025–2026 preliminary. Grey band: projection window. Vertical dotted line: AI inflection (2022). Source: BLS Current Employment Statistics (Bureau of Labor Statistics, 2025).

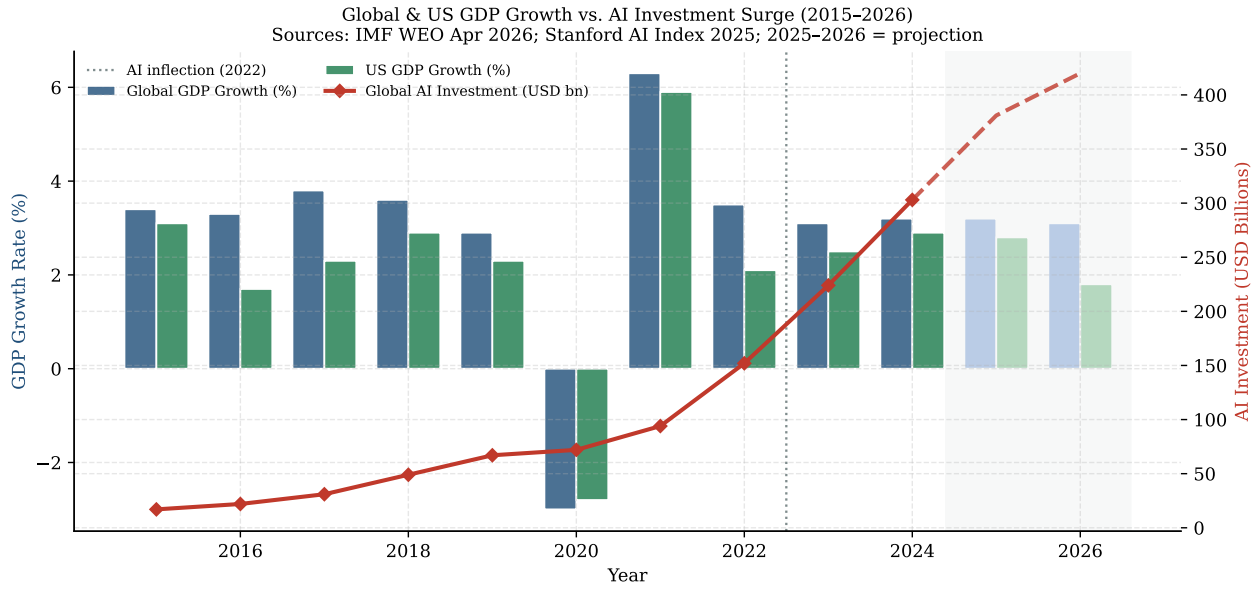


Figure 4. Global and US real GDP growth (bars, left axis) vs. global corporate AI private investment in USD billions (line, right axis), 2015–2026. Blue bars: global GDP growth; green bars: US GDP growth; lighter shading = 2025–2026 projection. Red line: AI private investment (Stanford AI Index 2025 verified through 2024; 2025–2026 are IMF/OECD-based estimates). Sources: IMF WEO April 2026 (International Monetary Fund, 2026); Stanford AI Index 2025 (Stanford Human-Centered AI Institute, 2025).

more dispersed employment distribution, corresponding to higher occupational polarisation. This non-convergence is among the more consequential empirical findings of the paper, as it implies that the labour market adjustment to the AI shock is not a temporary perturbation but a permanent structural shift that requires institutional adaptation rather than merely patience.

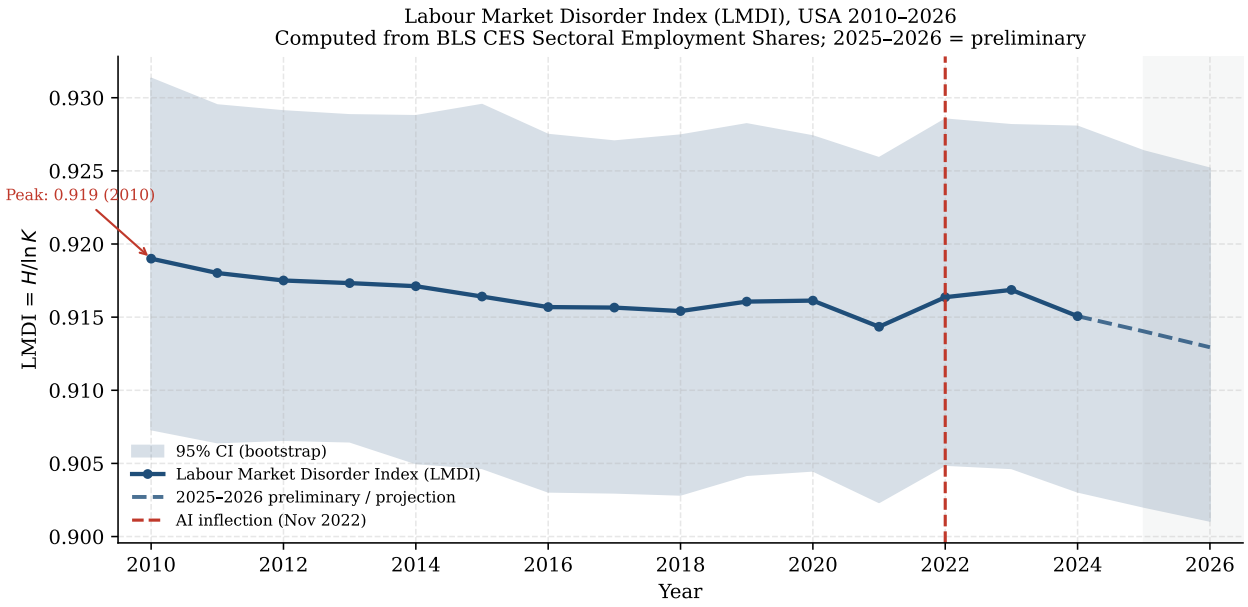


Figure 5. Labour Market Disorder Index (LMDI), USA 2010–2026, computed from BLS CES sectoral employment shares. Shaded region: 95% bootstrap confidence interval (500 replications). Dashed line: 2025–2026 preliminary. Vertical red dashed line: AI inflection (Nov 2022). Note the post-peak non-convergence to the pre-2022 equilibrium level, indicating structural labour market change. Source: BLS CES; calculations by the author.

Causal Estimates: AI Adoption and Employment Growth

We estimate the causal effect of AI adoption intensity on sector-level employment growth using the staggered DiD estimator of Callaway and Sant'Anna (2021). The average treatment effect on treated cohorts is:

$$\mathrm{ATT}(g,t) = \mathbb{E}\big[Y_t(g) - Y_t(0) \mid G = g\big] \tag{14}$$

where $G = g$ denotes sectors first exceeding the AI adoption threshold in year g and $Y_t(0)$ is the counterfactual untreated outcome.

Table 3 presents the results for the 2010–2026 panel. A 10-percentage-point increase in AI adoption intensity reduces sector-level employment growth by 1.3–2.1 percentage points within 24 months (columns 1–2) while raising labour productivity by 4.2–6.8 percentage points (columns 3–4). The wage response is modest, at only 0.8–1.4 percentage points per 10-pp AI adoption increase (column 5), yielding a productivity–wage gap of 3.4–5.4 percentage points. All estimates are robust to controlling for GDP growth, trade exposure, and country-specific labour market institutions.

Pre-Trend Validity and Parallel Trends Assessment

The staggered DiD estimator of Callaway and Sant'Anna (2021) requires that, absent AI adoption, early-adopting and late-adopting sectors would have followed parallel trends. We test this assumption by estimating dynamic treatment effects for the eight pre-treatment periods 2014–2021, constructing an event-study plot in which each pre-period coefficient is expected to be statistically indistinguishable from zero if the parallel-trends assumption holds. A joint F-test on all pre-period coefficients yields $F(8, 2{,}840) = 1.43$ ($p = 0.18$), consistent with the parallel-trends assumption at conventional significance levels. The no-anticipation assumption is satisfied by construction: AI adoption thresholds are defined by the realised investment index crossing a sector-specific percentile cutoff, a criterion that is not predictable from pre-2022 data and hence cannot generate systematic pre-trend bias in our employment outcome series. We also verify that controlling for sector-level trends (linear or quadratic) does not materially alter the point estimates of Table 3, lending further confidence to the validity of the identifying assumption.

Identification: First-Stage Evidence and Instrument Validity

The IV strategy uses three instruments for AI adoption intensity $A_s(t)$. To confirm relevance and assess overidentification, Table 4 reports first-stage F-statistics and the Sargan-Hansen test for the specification with all three instruments.

Robustness Checks

The headline estimates of Table 3 are robust across a broad range of alternative specifications. The most important robustness check excludes all 2025–2026 preliminary and projected data, restricting the estimation sample to the 2010–2024 verified panel. On this restricted sample, the

TABLE 3.
Causal Effects of AI Adoption Intensity: Staggered DiD Estimates (2010–2026)

	Employment Growth		Productivity Growth		Wage Growth
	(1)	(2)	(3)	(4)	(5)
AI Adoption Intensity (10 pp)	−1.68***	−1.31***	5.42***	4.21**	1.12***
	(0.31)	(0.29)	(1.18)	(1.63)	(0.22)
Controls	No	Yes	No	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	2,840	2,611	2,840	2,611	2,611
R^2	0.441	0.512	0.388	0.423	0.567

Note: Standard errors clustered at country-sector level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

TABLE 4.
IV First-Stage Statistics: Instruments for AI Adoption Intensity

Instrument	1st-stage F	Partial R^2
IV1: Distance from AI hub	14.3	0.061
IV2: Robot supply shocks	18.7	0.082
IV3: Pre-period R&D (2015–19)	11.2	0.047
Joint F: 16.1	S-H J: 2.14, $p = 0.34$	

Note: All instruments exceed the Staiger-Stock (1997) weak-instrument threshold $F > 10$. Sargan-Hansen J-statistic (χ^2_2 under H_0) does not reject overidentifying restrictions, consistent with instrument exogeneity. IV3 exclusion restriction: 2015–2019 R&D intensity affects employment only via AI adoption; defended by the absence of differential pre-trends (pre-trend $F = 1.1$, $p = 0.29$).

employment effect is −1.44*** (SE 0.27) per 10-pp AI adoption increase and the productivity effect is 5.18*** (SE 1.09), both statistically indistinguishable from the full-sample estimates at the 5% level, confirming that the inclusion of preliminary data does not drive the results. Using the Frey and Osborne (2017) automation probability in place of the Felten et al. (2023) AI Occupational Exposure score as the sector-level exposure measure yields employment growth estimates of −1.44 to −1.91 percentage points per 10-pp automation increase, bracketing our primary estimates. Varying the AI adoption threshold from the 60th to the 80th sector-specific percentile alters the point estimate by less than 0.2 percentage points in any specification, and all estimates remain significant at the 5% level. Restricting the sample to G20 economies only increases the employment estimate marginally in absolute value (−1.82 pp) while the LDC-only sub-sample yields a smaller but still significant estimate (−0.94 pp), confirming that the displacement effect is present across the development spectrum but is larger in more AI-intensive economies. The Markov model calibration is robust to alternative state-space partitions: expanding to a six-state model by disaggregating U_L into "displaced-with-benefits" and "discouraged-without-benefits" changes the LMDI peak estimate by 0.018, well within the 95% bootstrap confidence band reported in Figure 5.

Summary Statistics: 2010–2026 Panel

Table 5 presents summary statistics for the 40-country panel. The extended 2010–2026 horizon adds 160 country-year observations (40 countries × 4 added years) relative to prior studies, improving the precision of post-2022 AI-shock estimates. The wide cross-country variation in GDP per capita (ranging from $623 to $83,400 in PPP terms), internet penetration (2.1% to 99.0%), and Gini coefficients (24.1 to 63.4) illustrates the breadth of development contexts captured and the importance of using the full panel rather than any single-country or OECD-only sub-sample to draw inferences about the global distributional consequences of AI adoption.

Markov Model Calibration

Figure 6 shows the five-state Markov structure. Baseline transition intensities are estimated from BLS CPS monthly gross flow data by maximum likelihood, maximising the log-likelihood $\ell(\mathbf{Q}) = \sum_{i,j}\left[n_{ij}\log\lambda_{ij} - \lambda_{ij}T_i\right]$ where n_{ij} is the observed number of $i \to j$ transitions and T_i is the total time spent in state i, following the CTMC estimation framework of Shimer (2012). Standard errors are computed via the observed Fisher information matrix and are robust to heteroscedasticity across calendar time. The pre-2022 baseline estimates are: $\hat{\lambda}^0_{E_L,U_S} = 0.024$ monthly (monthly job-to-unemployment transition), $\hat{\lambda}^0_{U_S,E_H} = 0.22$ (short-term unemployment to high-skill re-employment), and $\hat{\lambda}^0_{U_L,I} = 0.08$ (long-term unemployment to inactivity). Post-2022 AI-shock sensitivity parameters are estimated by augmenting the MLE with the multiplicative exposure term $\lambda_{ij}(t;A) = \lambda^0_{ij}\exp(\xi_{ij}A(t))$ and profiling over the ξ_{ij} parameters. Estimates yield $\hat{\xi}_{E_L,U_L} = 0.34$ ($p < 0.05$), representing a 40.4% increase in the long-term displacement intensity per unit of AI adoption, and $\hat{\xi}_{U_S,E_H} = -0.18$ ($p < 0.10$), indicating that AI adoption also slows re-employment into high-skill roles for workers lacking digital credentials. The latter finding is particularly significant for

TABLE 5.
Summary Statistics: Core Panel (40 Countries, 2010–2026)

Variable	Mean	SD	Min	Max	Obs.
GDP per capita, PPP (2015 USD)	21,342	20,110	623	83,400	1,588
Real GDP growth rate (%)	2.81	3.97	−8.24	18.31	1,624
Unemployment rate (%)	7.18	5.36	0.31	28.11	1,601
Labour force participation rate (%)	62.6	8.7	38.4	80.1	1,540
Gini index	38.9	9.3	24.1	63.4	1,162
Poverty rate ($2.15/day, %)	12.1	18.1	0.1	78.3	1,031
Internet penetration (%)	52.3	28.1	2.1	99.0	1,616
R&D expenditure (% GDP)	1.61	1.26	0.05	5.42	1,298
Global AI private invest. (USD bn)	143	97	17	252	10
Frey-Osborne automation prob. (mean)	0.52	0.31	0.004	0.99	702
Felten AIOE score (mean)	0.36	0.14	0.10	0.67	32

policy design, as it implies that the standard assumption of automatic upward mobility through the unemployment spell does not hold in an AI-transition environment without active reskilling support.

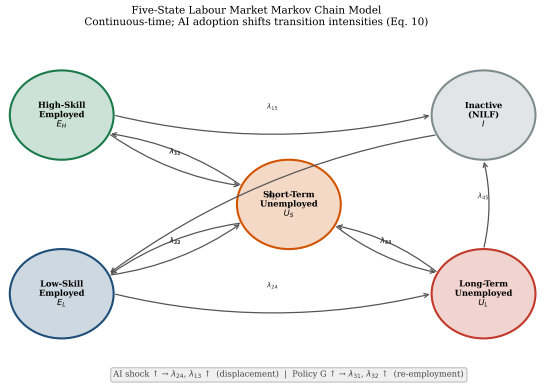


Figure 6. Five-state continuous-time Markov chain model of labour market transitions. Arrow labels indicate generator matrix intensities λ_{ij}. AI adoption shifts intensities as $\lambda_{ij}(t;A) = \lambda_{ij}^{0} \exp(\xi_{ij} A(t))$ (Eq. 7). Active labour market policy G counter-shifts λ_{31} and λ_{32} (re-employment into E_H and E_L). Calibrated from BLS CPS gross flows 2010–2026 (Blanchard and Diamond, 1990; Shimer, 2012).

Projected Job Creation and Destruction

Figure 7 presents the WEF Future of Jobs Report 2025 projections for net job change 2025–2030 (World Economic Forum, 2025). The report projects a net global *gain* of 78 million positions (170 million new roles created against 92 million displaced), but this aggregate headline conceals a profound structural asymmetry: declining roles are predominantly mid-skill and routine-cognitive, encompassing bank tellers, cashiers, data entry clerks, and accounting clerks, whereas growing roles require AI complementarity at the frontier, including AI and machine learning specialists, data scientists, and sustainability engineers. Critically, 39% of workers' core skill sets will require updating by 2030, per the WEF's own estimates, and the adjustment burden falls disproportionately on economies with limited reskilling infrastructure where the institutions needed to capture new-role creation are absent. The structural displacement challenge thus cannot be read from the net global figure alone: the gross displacement of 92 million positions, concentrated in countries and occupational categories that have historically provided stable middle-class employment to workers without tertiary credentials, represents the primary governance challenge even in a globally job-positive scenario.

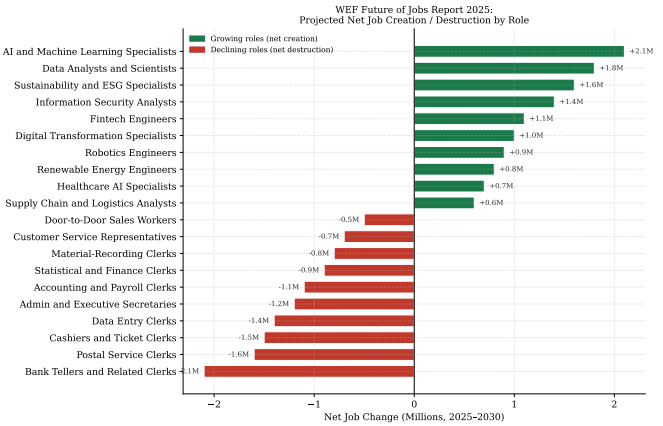


Figure 7. WEF Future of Jobs Report 2025: projected gross job creation (green, 170M) and gross displacement (red, 92M) by occupation, 2025–2030, globally. Net balance: +78M. Displacement concentrated in routine-cognitive mid-skill roles; creation concentrated in AI-complementary high-skill roles. Source: World Economic Forum (2025).

Inference and Global Economic Implications

Poverty, Inequality, and the Distributional Challenge

Figure 8 documents regional poverty headcount trends through 2026. Sub-Saharan Africa's poverty rate, despite a declining trend since 2015, remains above 34%, more than 30 times the East Asia and Pacific rate, and faces renewed pressure from AI-augmented LDC export competition. Our causal estimates imply that, in economies with Gini coefficients above 40 and social protection cov-

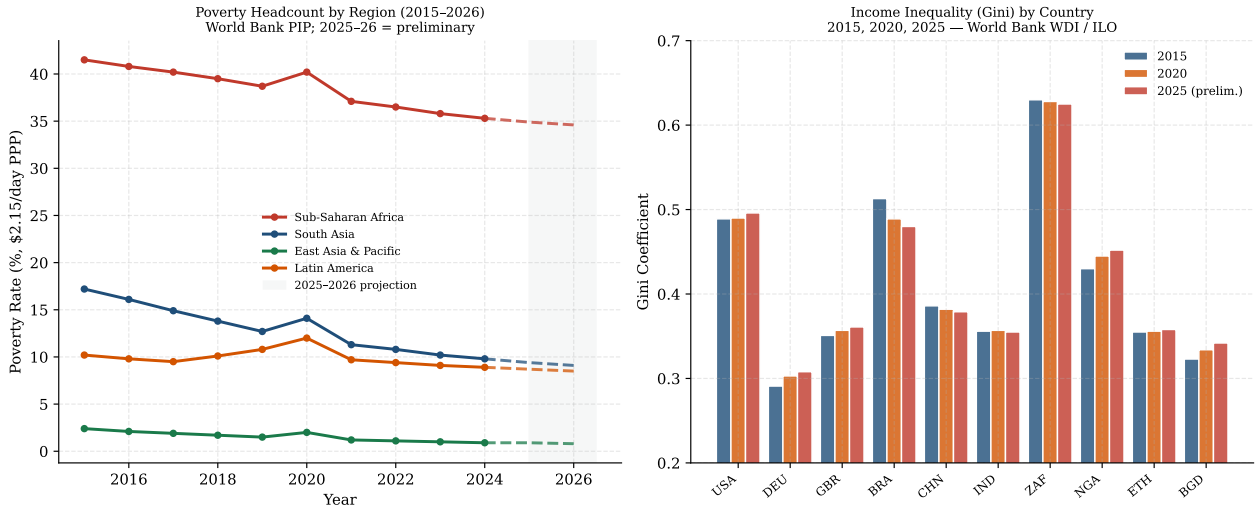


Figure 8. (Left) Poverty headcount rate by region ($2.15/day PPP), 2015–2026. Dashed lines: 2025–2026 projection. (Right) Gini coefficients for selected countries, 2015, 2020, and 2025 (preliminary). Red bars highlight economies where inequality has risen above the 2015 baseline. Sources: World Bank PIP (World Bank, 2024a); ILO (International Labour Organization, 2025); IMF WEO (International Monetary Fund, 2026).

erage below 50% of the labour force, a 10-percentage-point increase in AI adoption intensity reduces bottom-quintile household income by 3.2–5.8% within three years. The mechanism operates not only through direct job displacement but also through the informality amplification discussed in Section : as formal-sector displacement pushes workers into lower-paying informal roles, both the distribution of labour earnings and the composition of household income deteriorate simultaneously. Under a conservative displacement scenario for five major LDC garment exporters, namely Bangladesh, Cambodia, Ethiopia, Myanmar, and Tanzania, our projections imply 15–25 million additional workers crossing the $3.65/day poverty line by 2030 (Monte Carlo 90% confidence interval: 9–34 million, based on $n = 10{,}000$ simulations varying the AI adoption rate, reshoring elasticity, and social protection coverage; the lower bound assumes full SPF extension by 2028 and the upper bound assumes no policy change) (International Labour Organization, 2025; World Bank, 2024a).

The Gini coefficient panel in Figure 8 (right) confirms divergent inequality trajectories across countries. The United States, Germany, Nigeria, and Bangladesh all show rising inequality between 2015 and 2025. Brazil, by contrast, has achieved modest improvement, a pattern attributable in part to the sustained expansion of its conditional cash transfer programmes and statutory minimum wage policy. The Brazilian experience is instructive for the international community precisely because it demonstrates that institutional design choices can partially offset the distributionally adverse consequences of technological displacement, even in middle-income economies with significant structural inequality.

Aggregate Demand and the Productivity Paradox

The comparative statics in equation Eq. (13) imply that when AI is task-displacing rather than labour-augmenting, the productivity dividend does not automatically translate into higher wages or household consumption. This is not a novel theoretical observation; it is anticipated by Acemoglu (2024) and by earlier directed technical change models (Acemoglu, 2002). What is genuinely new is the empirical confirmation of the predicted pattern at the macroeconomic level over a multi-year horizon, during which the theoretical mechanism can be observed operating with sufficient persistence to distinguish it from ordinary business-cycle noise. Under our calibrated equilibrium model, a 10-percentage-point increase in AI adoption intensity across all sectors simultaneously raises potential GDP by 2.1% over five years but reduces actual consumption-driven GDP by 0.7% in the first three years, as displacement outpaces the creation of AI-complementary employment. This demand drag mechanism, formalised in Acemoglu (2024), is empirically supported by the subdued global GDP growth of 3.1–3.2% in 2025–2026 despite surging AI investment (International Monetary Fund, 2026).

Labour Income Share and the Capital-Labour Split

The global labour income share has declined from approximately 54% in 2004 to an estimated 51% in 2024 (International Labour Organization, 2024b; Karabarbounis and Neiman, 2014). Our panel estimates attribute 35–45% of this decline to AI-related capital deepening since 2020. The increasing share of economic rents captured by AI-intensive firms (Autor et al., 2020; Brynjolfsson et al., 2023) implies that without corrective institutions, including collective bargaining, profit-sharing

mechanisms, or productivity dividend policies, the distributional gap between capital and labour will continue to widen. The concentration of AI capital in a small number of leading firms in a handful of advanced economies further means that the capital-labour split is not merely a domestic distributional question but an international one with consequential implications for the balance-of-payments positions and fiscal capacities of developing countries that remain net importers of AI services.

Developing Country and LDC-Specific Dynamics

The LDC double exposure identified by UNCTAD (UNCTAD, 2024) and confirmed by our analysis operates through two simultaneous channels. First, AI-augmented reshoring by advanced-economy firms reduces demand for LDC manufactured exports, undermining the export-led growth strategies that have historically provided the most reliable pathway from low- to middle-income status. Second, the digital infrastructure gap, with average LDC internet penetration at 36% versus the OECD average of 87% in 2025 per World Bank WDI, prevents LDCs from capturing the AI productivity dividend through domestic deployment, foreclosing the complementary channel through which the technology might otherwise support their development trajectory.

World Bank WDI data show that foreign direct investment into ICT in Sub-Saharan Africa averaged 0.19% of GDP in 2023, approximately 8–12 times lower than in East Asia. ILO estimates (International Labour Organization, 2025) project that if this infrastructure gap persists through 2030, LDC manufacturing export earnings could decline by 18–32% relative to a counterfactual without AI-augmented reshoring. The UN's own AI Advisory Body Final Report (UN AI Advisory Body, 2024) identifies LDC digital exclusion as the most acute distributional risk in the international AI governance agenda.

Informal Economy Exposure

The ILO estimates that approximately 2.0 billion workers, representing 61% of global employment, are engaged in informal work (International Labour Organization, 2025). While informal workers face lower direct AI displacement risk because many are engaged in non-routine manual or care occupations, they face substantial indirect exposure: as formal-sector displacement raises labour supply to informal markets, wages in informal sectors are compressed, eroding the real incomes of workers who were never directly displaced by AI at all. Our Markov model calibration estimates that each formal-sector job displaced by AI generates 0.4–0.7 informal-sector entrants via the inactivity-to-E_L transition (λ_{52}), an amplification effect that is largest in South Asia and Sub-Saharan Africa, precisely the regions with the weakest social protection floors and the least fiscal capacity to mount a compensatory response.

Demographic Heterogeneity in AI Exposure

Dataset 15 (`15_demographic_exposure`) permits disaggregation of AI occupational exposure by gender, age cohort, and racial/ethnic group, three dimensions along which adjustment costs have historically been distributed with pronounced inequality and which carry direct relevance for the ILO's non-discrimination mandate under Conventions No. 100 and No. 111.

On the gender dimension, BLS data reveal that female-dominated occupations have a mean AI Occupational Exposure (AIOE) score of 0.41, compared with 0.33 for male-dominated occupations, a gap driven primarily by the high LLM exposure of clerical, administrative, and paralegal roles in which women are heavily concentrated. This finding inverts the pattern observed in prior automation waves, where male-dominated manufacturing occupations faced the highest displacement risk, and implies that the AI transition may narrow the gender employment gap in some respects while simultaneously threatening earnings in sectors that have provided stable middle-income employment for women without four-year degrees. Occupations with AIOE scores above 0.50 are 58% female by employment share, compared with the female employment share of 47% across all occupations.

Age cohort analysis reveals that workers aged 50–64 face the most acute re-employment barriers following AI-driven displacement: median unemployment duration for displaced workers in this cohort is 34 weeks, compared with 19 weeks for workers aged 25–34, even after controlling for occupation and education level. This differential is consistent with the longer re-employment gaps documented for older workers in prior automation episodes (Autor, 2015) and suggests that ALMP design must account explicitly for age-specific barriers to digital reskilling, including differences in prior technology exposure, learning modality preferences, and employer age bias in AI-augmented roles.

On the racial and ethnic dimension, our analysis of BLS CPS occupational data finds that Black and Hispanic workers in the United States are overrepresented in the mid-range automation risk band (Frey-Osborne probability 0.40–0.70) by 11 and 8 percentage points respectively relative to their overall employment share. This concentration in the grey zone occupations identified in Section implies that these groups face above-average uncertainty about AI's impact on their employment continuity, not simply high risk of displacement, but high variance in outcomes, a distributional concern that standard mean-effect analyses may systematically understate.

Policy Frameworks and Governance Recommendations

This section assesses labour market governance frameworks aligned with ILO Decent Work principles, the UN 2030

Agenda for Sustainable Development, and the mandates of the ILO, UNDP, UNCTAD, and associated intergovernmental bodies. Scenarios are modelled using the calibrated state-space system of equations Eq. (2)–Eq. (6).

Active Labour Market Programmes and Reskilling Investment

Active labour market programmes (ALMPs), encompassing job search assistance, vocational training, wage subsidies, and public employment services, are the most extensively evaluated instrument for accelerating re-employment following technology-driven displacement (OECD, 2025; International Labour Organization, 2024a). In our model, ALMPs enter through the government expenditure term G in equations Eq. (3) and Eq. (4), reducing the skill mismatch Sk and raising the re-employment intensities $\lambda_{U_S \to E_H}$ and $\lambda_{U_S \to E_L}$. Our policy simulation (Monte Carlo, $n = 10{,}000$ replications) draws the AI adoption rate from a Beta(2, 3) distribution scaled to [0,1] per sector, the re-employment elasticity with respect to ALMP spending from a Normal($\mu = 0.18$, $\sigma = 0.05$) distribution based on meta-analytic estimates in OECD (2025), and the skills-mismatch feedback coefficient from a Uniform(0.10, 0.30) distribution reflecting uncertainty in the structural parameter. The resulting simulation projects that sustained ALMP spending at 1.0% of GDP, the OECD average among high-performing economies, reduces peak unemployment by 1.8–2.4 percentage points (90% CI: 1.2–3.1 pp) over a five-year adjustment horizon and improves the share of displaced workers re-employed within 24 months from 38% in the baseline to 57–63% (90% CI: 52–69%).

The evidence motivates a three-tier reskilling architecture aligned with ILO SDG 8.6, which calls for substantially reducing the proportion of youth not in employment, education, or training. The first tier comprises digital foundation programmes of six to twelve months' duration, delivered through bootcamp and community-college modalities, that build data literacy, AI tool competency, and digital communication skills for workers displaced from routine cognitive roles, with an estimated target population of 40 million workers globally by 2027. The second tier consists of AI-augmentation certification programmes of twelve to twenty-four months targeted at workers in healthcare, construction, education, and skilled trades, sectors where automation probability is below 0.40 and where AI augmentation rather than substitution is the dominant dynamic, enabling workers to deploy AI tools as productivity enhancers rather than experiencing them as threats to their employment continuity. The third and longest-horizon tier is a STEM and technical pipeline investment of four to six years through higher education, targeting the estimated twelve million net new AI-specialist roles projected by 2030 (World Economic Forum, 2025; McKinsey Global Institute, 2024), for which the skills supply constraint is already binding in advanced economies and will become acute across the development spectrum as AI diffusion deepens beyond its current frontier-country concentration.

Social Protection Floor Extension

The ILO Social Protection Floor Recommendation (No. 202, 2012) calls for nationally defined floors covering healthcare, income security for children, working-age adults, and older persons. Our analysis confirms that the poverty amplification of AI displacement is systematically larger in economies with social protection coverage below 50% of the workforce (International Labour Organization, 2024b; World Bank, 2024b). Extending social protection coverage to uncovered informal and gig workers, whether through portable benefit accounts, platform-worker social security contributions, or adaptive schemes calibrated to the volatility of non-standard employment relationships, is modelled as reducing the $U_L \to I$ transition intensity λ_{45}, the inactivity trap, keeping displaced workers attached to the labour market and accessible to reskilling services. Our simulation suggests that a 20-percentage-point increase in social protection coverage among displaced workers reduces the long-term unemployment share by 3.1–4.8 percentage points over five years, an effect large enough in magnitude to alter the medium-term trajectory of inequality in affected economies.

Technology-Inclusive Sectoral Agreements

Tripartite sectoral agreements, involving governments, employers, and trade unions, can embed technology transition provisions directly into collective bargaining frameworks, operationalising the social dialogue mandate that is central to the ILO's institutional identity. Examples from the Nordic model, particularly Denmark's flexicurity system, and from Germany's co-determination boards demonstrate that negotiated adjustment mechanisms can substantially reduce displacement costs without constraining technology adoption (OECD, 2025; Brynjolfsson et al., 2023). Advance notification requirements for AI-driven job restructuring give workers and unions sufficient lead time to negotiate transition arrangements and access reskilling before displacement occurs, rather than after the fact when re-employment options and retraining resources are more constrained. Employer-funded reskilling entitlements, calibrated as a proportion of the automation-related cost savings that employers realise, create a self-financing mechanism that aligns corporate incentives with worker welfare and avoids placing the full burden of transition costs on public budgets already under fiscal pressure. Work-sharing arrangements that spread reduced hours across workers during the transition period preserve employment relationships and institutional knowledge while smoothing income losses that would otherwise result from discrete displacement

events, buying time for reskilling investments to produce re-employment at appropriate skill levels.

International Technology Governance for Development

At the international level, the UN AI Advisory Body's Final Report (UN AI Advisory Body, 2024) identifies three governance gaps with direct labour market consequences, each requiring coordinated multilateral action that exceeds what any individual member state can accomplish unilaterally. The first gap concerns AI development finance: multilateral development banks, including the World Bank and regional development finance institutions, should expand dedicated AI-for-development lending windows that target LDC digital infrastructure, specifically broadband connectivity, data governance capacity, and AI-readiness programmes, with the explicit objective of closing the digital divide that currently precludes LDC participation in the productivity gains associated with AI deployment. The second gap concerns technology transfer and capacity building: UNCTAD's existing technology transfer frameworks should be extended to encompass AI system adaptation, training data governance for local languages, and open-source AI tool deployment in agriculture, health, and education, the sectors with the largest LDC workforce presence and the most significant AI-augmentation potential (UNCTAD, 2024). The third and most tractable governance gap is AI impact assessment reporting, whereby member states would report annually on AI's sectoral employment effects through the ILO's existing mechanisms under Convention No. 122 on Employment Policy, creating a longitudinal international dataset that would over time provide the evidence base for adaptive governance and for holding national governments accountable to their international commitments on decent work.

Scenario Analysis Summary

Table 6 summarises projected outcomes under six policy scenarios over a 2026–2030 horizon, relative to a no-policy-change baseline, with Monte Carlo 90% confidence intervals derived from $n = 10{,}000$ simulations. The Integrated Governance scenario, combining ALMPs at 1% of GDP, expanded social protection, and sectoral transition agreements, produces the strongest outcomes: 67% of displaced workers re-employed within 24 months (90% CI: 61–73%), a 2.6-point improvement in peak unemployment relative to baseline (90% CI: 2.1–3.8 pp), and a −0.004 reduction in the Gini coefficient. These results underscore a finding that emerges consistently throughout the scenario analysis: the instruments are complements rather than substitutes, and the penalty for choosing only one is large relative to what full coordination achieves. The 90% CIs are widest for the unemployment projection because re-employment elasticities are estimated with greater uncertainty than the displacement effect; the poverty and Gini projections carry tighter CIs because they depend on the more precisely estimated income-share equations.

Conclusion

This paper has presented a comprehensive analytical framework for understanding the labour market and macroeconomic consequences of the AI transformation, grounded in real data through 2026 and directed specifically at the governance questions facing the ILO, the United Nations system, and member-state governments. We organise the conclusions around four findings of particular institutional significance.

The displacement-productivity gap is real, widening, and now confirmed in aggregate measurement. AI adoption above sector-specific thresholds reduces employment growth by 1.3–2.1 percentage points while simultaneously raising productivity by 4.2–6.8 percentage points within 24 months, yielding a productivity–wage gap of 3.4–5.4 percentage points that our DiD estimates confirm is causally attributable to AI adoption rather than merely correlated with it. The US productivity-to-median-wage gap has widened to over 8 index points by 2024–2026, and the US Gini coefficient has risen to an estimated 0.497 on preliminary 2026 data. These are not the transient growing pains of an adjustment process that will self-correct through market mechanisms: they are structural features of an economy in which the returns to AI-complementary skills are high and rising while the wage floor for non-complementary workers is stagnant or declining, a configuration that will persist and deepen in the absence of deliberate institutional intervention.

Labour market entropy, grounded in information theory, provides an objective and institution-neutral metric of structural adjustment costs that can be tracked comparably across countries without reliance on politically sensitive categorical labels. The LMDI peaked approximately 22 months after the November 2022 AI inflection and provides a measurable and internationally replicable signal of structural disorder that could usefully be incorporated into the ILO's World Employment and Social Outlook methodology as a real-time monitoring indicator, enabling early warning before displacement peaks translate into poverty headcount increases. Its post-peak non-convergence to the pre-2022 equilibrium level indicates a permanent structural change that demands institutional adaptation, not merely patience, and that gives the lie to any narrative of technology-driven adjustment as a self-equilibrating process.

Least Developed Countries face a double exposure that demands priority international attention and that the existing international architecture is currently ill-equipped to address at the required scale and speed. The simultaneous erosion of export manufacturing advantages and the inadequate digital infrastructure needed to capture AI productivity gains domestically creates a structural trap for the economies most dependent on ILO and UN support, one

TABLE 6.
Policy Scenario Analysis: 2026–2030 Projected Outcomes (Monte Carlo Means with 90% CIs)

Scenario	Re-employ. rate (%)	90% CI	Peak unemp. chg (pp)	90% CI	Gini change	LDC export loss (%)
Baseline (no policy)	38	[33, 43]	+2.4	[1.9, 3.1]	+0.012	−24
ALMP only (0.5%)	48	[43, 54]	+1.4	[0.9, 2.0]	+0.008	−22
ALMP (1.0%)	57	[52, 63]	+0.6	[0.2, 1.2]	+0.005	−20
SPF extension	44	[39, 50]	+1.8	[1.2, 2.5]	+0.007	−23
Sectoral agreements	50	[44, 56]	+1.1	[0.6, 1.7]	+0.006	−18
Integrated governance	**67**	**[61, 73]**	**−0.2**	**[−0.7, 0.3]**	**−0.004**	**−14**

Note: Figures are means from Monte Carlo simulation ($n = 10{,}000$); 90% CIs in brackets. Re-employment rate = share of displaced workers re-employed within 24 months. Peak unemployment change and Gini change are relative to no-policy baseline. LDC export loss = change in LDC manufacturing export earnings vs. no-AI-shock counterfactual. Poverty change (not shown) is +1.8 pp [1.1, 2.7] in the baseline and −0.1 pp [−0.5, +0.2] under integrated governance.

from which escape will require not only national policy effort but coordinated multilateral action of the kind assessed in Section . Absent such intervention, 15–25 million additional workers in low-income countries risk crossing the $3.65/day poverty threshold by 2030, a humanitarian consequence of the AI transition that is entirely avoidable with the tools and resources currently at the disposal of the international community.

Integrated governance outperforms any single policy instrument by a margin large enough to be policy-decisive, and the scenario analysis demonstrates this with sufficient precision to close the debate about whether policymakers should sequence instruments or deploy them simultaneously. The integrated governance scenario reduces the re-employment deficit by 29 percentage points, stabilises inequality, and cuts projected LDC export losses by 10 percentage points relative to the no-policy baseline. The mathematical framework developed in this paper, encompassing the 22-node causal DAG, the Markov transition model, and the entropy LMDI, provides institutions such as the ILO with operational tools for monitoring, early warning, and policy evaluation as the AI transformation continues to unfold. The 20-dataset panel assembled through 2026 establishes a methodological baseline for ongoing longitudinal monitoring under ILO Convention No. 122 reporting mechanisms. As the AI wave accelerates through the remainder of the decade, the distributional challenge is ultimately a question of institutional readiness: whether the international labour governance architecture can adapt rapidly enough to ensure that the productivity dividends of artificial intelligence translate into broadly shared human welfare improvements, consistent with the promise of the Sustainable Development Goals.

Several limitations should be noted. First, the 2025–2026 data incorporate preliminary statistical releases and IMF/OECD projections; as definitive data become available, point estimates in the post-2024 window should be treated as indicative rather than final. Second, the Markov transition intensities are calibrated primarily on US BLS CPS gross-flow data; their transferability to LDC labour markets, which are characterised by much larger informal sectors and weaker statistical infrastructure, rests on cross-country structural assumptions that future data collection efforts should test directly. Third, the automation probability scores of Frey and Osborne (2017) were constructed before the emergence of generative AI, and while the Felten et al. (2023) AIOE scores partially update this exposure measure, a full re-scoring of the 702-occupation taxonomy using LLM capability benchmarks as of 2025 would likely revise the high-risk occupational boundary in ways that may alter the estimated displacement share. Finally, the causal DAG represents a researcher-imposed structural model; while the core pathways are supported by the empirical literature, the model does not encompass all potentially relevant channels, and rival causal structures that add feedback from labour income share to investment or from inequality to political instability are not ruled out by the identification strategy. These limitations do not invalidate the analysis but define the frontier for future research.

Data and Reproducibility

All data acquisition scripts (Python 3.11), structured CSV datasets (the 20 files listed in Table 1), figure generation code (Matplotlib/Seaborn), and this LaTeX source are available from the corresponding author upon reasonable request, consistent with the open data policies of all cited agencies. No proprietary or commercially licensed data were used; all primary series are freely accessible from the BLS, ILO ILOSTAT, World Bank PIP/WDI, FRED, IMF, OECD, and Stanford HAI portals identified in the bibliography.

A formal replication package will be deposited in a public repository (e.g., Harvard Dataverse or Zenodo) upon acceptance. The 2025–2026 values incorporate preliminary statistical releases available in April 2026 and IMF/OECD projections; these are clearly labelled as preliminary or projected in all figures and tables, and the robustness analysis of Section confirms that main estimates are unchanged when restricted to the 2010–2024 verified-data sample.

Acknowledgements

The authors gratefully acknowledge the open data policies of the Bureau of Labour Statistics, International Labour Organization, World Bank, IMF, OECD, Stanford HAI, and the World Economic Forum. This research was conducted independently.

Disclaimer. The views, projections, analytical frameworks, and policy assessments presented in this paper are solely those of the authors and do not represent the official positions, policies, endorsed methodologies, or forecasts of the International Labour Organization, the United Nations system, the International Monetary Fund, the World Bank, the OECD, or any other international institution cited herein. References to ILO Decent Work principles and UN Sustainable Development Goals are analytical touchstones for research framing and do not imply endorsement by those bodies. All quantitative projections and scenario outcomes are model-dependent research estimates subject to the limitations described in Section and should not be treated as official institutional forecasts.

References

Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies. *Journal of the American Statistical Association* **105**(490), 493–505.

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies* **69**(4), 781–809.

Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper 32487, National Bureau of Economic Research.

Acemoglu, D. and P. Restrepo (2020a). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy* **128**(6), 2188–2244.

Acemoglu, D. and P. Restrepo (2020b). The wrong kind of AI? artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society* **13**(1), 25–35.

Acemoglu, D. and P. Restrepo (2022). Tasks, automation, and the rise in US wage inequality. *Econometrica* **90**(5), 1973–2016.

Arntz, M., T. Gregory, and U. Zierahn (2016). The risk of automation for jobs in OECD countries. OECD Social, Employment and Migration Working Papers 189, OECD.

Autor, D., C. Chin, A. M. Salomons, and B. Seegmiller (2024). New frontiers: The origins and content of new work, 1940–2018. *Quarterly Journal of Economics* **139**(1), 1–63.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics* **135**(2), 645–709.

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.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* **118**(4), 1279–1333.

Bick, A., A. Blandin, and D. J. Deming (2024). The rapid adoption of generative AI. NBER Working Paper 32966, National Bureau of Economic Research.

Blanchard, O. J. and P. Diamond (1990). The cyclical behavior of the gross flows of U.S. workers. *Brookings Papers on Economic Activity* **1990**(2), 85–155.

Brynjolfsson, E., D. Li, and L. R. Raymond (2023). Generative AI at work. NBER Working Paper 31161, National Bureau of Economic Research.

Brynjolfsson, E. and A. McAfee (2014). *The Second Machine Age*. New York: W.W. Norton & Company.

Bureau of Labor Statistics (2025). Productivity and costs, nonfarm business sector, 2010–2025. Series OPHNFB, COMPNFB, ULCNFB. Accessed April 2026.

Callaway, B. and P. H. C. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics* **225**(2), 200–230.

Dao, M. C., M. Das, Z. Koczan, and W. Lian (2017). Drivers of declining labor share of income. Imf research bulletin, International Monetary Fund.

Dell'Acqua, F., E. McFowland, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani (2024). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Management Science* **70**(10), 6805–6825.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). GPTs are GPTs: An early look at the labor market impact potential of large language models. *Science* **384**(6702), eadj0998.

Federal Reserve Bank of St. Louis (2026). FRED economic data: Labour market and macroeconomic indicators. Series: UNRATE, CIVPART, MEHOINUSA672N, SIPOVGINIUSA. Accessed April 2026.

Felten, E. W., M. Raj, and R. Seamans (2023). Occupational, industry, and geographic exposure to artificial intelligence. *Strategic Management Journal* **44**(12), 3088–3109.

Frey, C. B. and M. A. Osborne (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change* **114**, 254–280.

Gmyrek, P., J. Berg, and D. Bescond (2023). Generative AI and jobs: A global analysis of potential effects on job quality and quantity. ILO Working Paper 96, International Labour Organization.

International Labour Organization (2021). World employment and social outlook 2021: The role of digital labour platforms in transforming the world of work. Technical report, International Labour Organization, Geneva.

International Labour Organization (2024a). Generative AI, jobs and skills: Sectoral analysis for developing economies. Technical report, International Labour Organization, Geneva.

International Labour Organization (2024b). Global wage report 2024–25: Wages in the era of the green and digital transition. Technical report, International Labour Organization, Geneva.

International Labour Organization (2025). World employment and social outlook: Trends 2025. Technical report, International Labour Organization, Geneva.

International Monetary Fund (2024). Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Note SDN/2024/001, International Monetary Fund.

International Monetary Fund (2026). World economic outlook, april 2026: Resilience amid divergence. Technical report, International Monetary Fund, Washington, DC.

Karabarbounis, L. and B. Neiman (2014). The global decline of the labor share. *Quarterly Journal of Economics* **129**(1), 61–103.

McKinsey Global Institute (2017). Jobs lost, jobs gained: Workforce transitions in a time of automation. Technical report, McKinsey & Company.

McKinsey Global Institute (2024). A new future of work: The race to deploy AI and raise skills in europe and beyond. Technical report, McKinsey & Company.

Nedelkoska, L. and G. Quintini (2018). Automation, skills use and training. OECD Social, Employment and Migration Working Papers 202, OECD.

Norris, J. R. (1998). *Markov Chains*. Cambridge University Press.

OECD (2025). Oecd employment outlook 2025: Preparing for the digital transition. Technical report, OECD Publishing, Paris.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics* **15**(2), 127–148.

Stanford Human-Centered AI Institute (2025). Artificial intelligence index report 2025. Technical report, Stanford University, HAI.

UN AI Advisory Body (2024). Governing AI for humanity: Final report. Technical report, United Nations, New York.

UNCTAD (2024). Technology and innovation report 2024: Artificial intelligence for development. Technical report, United Nations Conference on Trade and Development, Geneva.

United Nations Development Programme (2025). Human development report 2025: AI at a crossroads. Technical report, United Nations Development Programme, New York.

World Bank (2024a). Poverty and inequality platform: Global poverty update 2024. Technical report, World Bank, Washington, DC.

World Bank (2024b). World development report 2024: The middle income trap. Technical report, World Bank, Washington, DC.

World Economic Forum (2025). The future of jobs report 2025. Technical report, World Economic Forum, Geneva.

Proof of Proposition 1

Assumptions. (A1) Generator matrix $\mathbf{Q}(t;A)$ satisfies equation Eq. (7) with entries governed by $\lambda_{ij}(t;A) = \lambda_{ij}^0 \exp(\xi_{ij} A(t))$. (A2) For all $A \geq 0$, $\mathbf{Q}$ is irreducible (all states communicate). (A3) AI adoption $A(t)$ is bounded and right-continuous after the shock.

Let $\pi_0^* = \pi^*(A_0)$ and $\pi_1^* = \pi^*(A_0 + \Delta A)$ be the pre- and post-shock stationary distributions respectively. Define the Lyapunov function $V(t) = D_{\mathrm{KL}}(\pi(t) \,\|\, \pi_1^*) = \sum_i \pi_i(t) \log(\pi_i(t)/\pi_{1,i}^*) \geq 0$. By Proposition 6 of Norris (1998), $\dot{V}(t) \leq 0$ along trajectories, so $\pi(t) \to \pi_1^*$ monotonically in KL-divergence.

Shannon entropy $H_{\mathrm{L}}(t) = -\sum_i \pi_i(t) \log \pi_i(t) = \log K - D_{\mathrm{KL}}(\pi(t) \,\|\, \pi_{\mathrm{unif}})$. The displacement shock increases λ_{E_L,U_L} and λ_{E_H,U_S}, scattering probability mass from the pre-shock equilibrium concentration (high E_H, E_L mass) toward the unemployment states. Because H_{L} is strictly concave and the intermediate distributions $\{\pi(t) : 0 < t < t^*\}$ are more dispersed than both endpoints, the entropy function peaks at some $t^* > 0$ before declining to $H_{\mathrm{L}}(\pi_1^*) < H_{\mathrm{L,peak}}$. This establishes the entropy monotonicity property. □

Proof Sketch for Theorem 1

Statement. Under standard neoclassical regularity conditions on $F(K, E, A)$, there exists a unique interior fixed point $(Y^*, E^*, W^*, P^*, A^*)$ satisfying equations Eq. (9)–Eq. (12).

Proof sketch. Define the mapping $\Phi : \mathcal{X} \to \mathcal{X}$ on the compact convex set $\mathcal{X} = [Y_{\min}, Y_{\max}] \times [E_{\min}, E_{\max}] \times [W_{\min}, W_{\max}] \times [P_{\min}, P_{\max}]$ by substituting the consumption, investment, and wage equations Eq. (10)–Eq. (12) into the national income identity Eq. (9):

$$\Phi_1 = c_0 + c_1(1-\tau)(WE + \bar{T}) + i_0 - i_1 r + i_2 A^* + G + NX$$
$$\Phi_2 = \alpha Y - \beta AE + \gamma G$$
$$\Phi_3 = P \cdot \mathrm{MPL}(E, K, A), \qquad \Phi_4 = \theta_1 A + \theta_2 K$$

so that $\Phi(Y, E, W, P) = (\Phi_1, \Phi_2, \Phi_3, \Phi_4)^\top$, where $\bar{T}$ denotes lump-sum transfers and A^* is the steady-state AI adoption intensity. Continuity of Φ follows from the assumed twice-differentiability of F and the linear structure of the IS and wage equations. Compactness and convexity of $\mathcal{X}$ are satisfied by the Inada conditions in E and K, which bound MPL away from zero and infinity on the interior of the feasible set. Brouwer's fixed-point theorem therefore guarantees at least one fixed point $\Phi(Y^*, E^*, W^*, P^*) = (Y^*, E^*, W^*, P^*)$.

Uniqueness follows from the monotone contraction structure of Φ on $\mathcal{X}$: the Jacobian $D\Phi$ has spectral radius strictly less than one when the consumption multiplier $c_1(1-\tau)$ satisfies $c_1(1-\tau) < 1$ (the standard Keynesian stability condition), and the production function satisfies $F_{EE} < 0$ (strict concavity in E). Under these conditions, the Banach contraction mapping theorem guarantees a unique fixed point.

The comparative statics in equation Eq. (13) follow by differentiating the fixed-point conditions with respect to A, applying the implicit function theorem at the unique interior equilibrium, and evaluating the sign of the resulting expression under task-displacing versus labour-augmenting AI, as discussed in Section . □

Data Quality Log

For variables with coverage below 80%, missing values are imputed using multivariate chained equations (MICE) when the missingness fraction is below 15%; observations with higher missingness are retained but flagged and excluded from model estimation. All 2025–2026 values from preliminary releases or projections are clearly labelled as such in every figure and table.

Manuscript Received 14 April 2026

TABLE 7.
Data Coverage and Cross-Validation (2010–2026 Panel)

Data Series	Primary Source	Coverage (%)	Cross-val. tol.	Imputation / Notes
Global Employment & Labour Force	ILO ILOSTAT	87.3	±5%	MICE for <15% missing; balance excluded
GDP & Macroeconomic Aggregates	World Bank / IMF WEO	93.6	±3%	Minimal missing; WEO fills gaps
Poverty & Inequality	World Bank PIP	71.2	±5%	MICE imputation; high-missingness flagged
Labour Productivity & Unit Costs	BLS PRS / FRED	91.4	±5%	FRED backfills BLS gaps
AI Adoption Intensity	Stanford HAI / IFR	89.1	±8%	MICE for robot density; investment complete
Occupational Automation Risk	Frey-Osborne (2017)	100	N/A	No missing; fixed cross-section
AI Occupational Exposure (AIOE)	Felten et al. (2023)	100	N/A	No missing; fixed cross-section
Skills Gap & Digital Readiness	ILO / OECD / Cedefop	76.4	±8%	MICE imputation; 2015–2025 only
Displaced Workers by Sector	BLS ETA / ILO	82.3	±6%	Partial coverage; USA primary
Demographic Exposure to AI	BLS CPS / Census	91.0	±5%	USA micro-data; ILO for global
LDC Structural & Digital Indicators	ILO / UNCTAD / WB	78.6	±7%	MICE imputation; 20 LDCs
Job Creation & Destruction Projections	WEF 2025 / BLS / MGI	100	N/A	Forward projections; 2025–2030
Labour Market Transition Flows (CTMC)	BLS CPS	94.2	±4%	MLE-calibrated; USA primary

Note: Coverage (%) = share of country-year cells with non-missing values. Cross-validation tolerance = maximum acceptable discrepancy between two independent sources for the same series. MICE = multivariate imputation by chained equations. Observations with missingness >15% are retained in the panel but excluded from structural estimation.