

---

# Collaboration and Innovation, Not Elimination of Employment: A Multi-Disciplinary Analysis of Artificial Intelligence's Impact on Human Labor Markets

Adel Tammam

*PTH Meridian — Precision Technology Heuristics*

Calgary, Alberta, Canada

adel.tammam@pth-meridian.io

---

*Abstract—The accelerating deployment of artificial intelligence (AI) technologies across global labor markets has generated substantial debate regarding the long-term consequences for human employment. This paper presents a multi-disciplinary synthesis of current empirical evidence, economic theory, psychological research, and sectoral case analysis to address the central question: does AI primarily function as a mechanism for labor displacement, or as a tool for human augmentation? Drawing on published data from the World Economic Forum, McKinsey Global Institute, Oxford Martin School, OECD, and multiple peer-reviewed studies, this paper argues that the relationship between AI and employment is not technologically predetermined but is instead contingent on governance choices, design philosophy, and policy infrastructure. Sectoral analyses spanning healthcare, energy systems, cybersecurity, creative industries, skilled trades, and education demonstrate that augmentation and replacement are design choices rather than inevitable outcomes. The psychological and social dimensions of labor displacement are examined through Self-Determination Theory and longitudinal public health data. Policy frameworks including retraining programs, Universal Basic*

*Income pilots, portable benefits, and the four-day work week are evaluated against the available evidence. The paper concludes that human-centered AI development, supported by open-source infrastructure and proactive transition policy, represents the most viable pathway to ensuring that technological progress serves broad human flourishing rather than concentrating benefit narrowly.*

**Index Terms**—artificial intelligence, labor displacement, human augmentation, automation, workforce transition, standard of living, open-source AI, human-centered design, Universal Basic Income, AI ethics.

---

## I. Introduction

The question most frequently posed about artificial intelligence and employment is deceptively simple: will AI take my job? It is a question with genuine urgency. The Stanford AI Index Report 2024 documents that AI systems achieved performance milestones across a broad range of standardized benchmarks between 2022 and 2024 at a pace substantially exceeding prior forecasts, with frontier model capabilities in language, reasoning, and code generation advancing more rapidly than at any previously recorded interval [1]. The deployment of these systems across commercial, industrial, and institutional contexts has been correspondingly rapid, preceding in many cases the policy, governance, and transition infrastructure that affected workers would require.

This paper argues that framing the question as binary—job elimination or job preservation—is analytically insufficient. The more useful framework distinguishes between *automation*, in which AI performs tasks previously requiring human labor with reduced human involvement, and *augmentation*, in which AI extends the reach and capability of human workers without eliminating the human role. This distinction is not merely semantic. It reflects a genuine choice at the level of system design, deployment context, and governance—a choice that organizations, policymakers, and AI developers are making continuously, often without explicit acknowledgment that a choice is being made at all.

The objectives of this paper are fourfold: (1) to synthesize the current empirical evidence on AI's labor market impact across sectors; (2) to examine the psychological and social dimensions of technological unemployment that economic models frequently underweight; (3) to evaluate the policy frameworks under active development for managing the AI transition equitably; and (4) to articulate the design and governance principles that distinguish human-centered AI development from AI deployment that treats labor displacement as an acceptable externality.

The analysis is presented across nine sections. Section II reviews relevant historical and theoretical background. Sections III through VII constitute the empirical and analytical core. Sections VIII and IX present discussion and conclusions. A glossary of technical terms and a full reference list are provided.

## II. Background and Related Literature

### A. Historical Context

Technological anxiety about employment is not a novel phenomenon. The mechanization of textile production in early 19th-century England precipitated the Luddite movement—a worker response to rapid industrial displacement that has been retrospectively mischaracterized as anti-technological rather than understood as what it was: a rational response to the real economic costs of unmanaged transition [2]. Subsequent waves of technological change—electrification, the internal combustion engine, computerization, and internet-enabled commerce—each generated comparable predictions of mass technological unemployment, none of which materialized at the aggregate level.

A frequently cited counter-intuitive case is the automated teller machine. Following the widespread ATM deployment of the 1970s and 1980s, analysis by Bessen [3] demonstrated that bank teller employment *increased* over the subsequent two decades. The mechanism was cost reduction: ATMs lowered the operating cost of a bank branch sufficiently that banks opened far more branches, increasing aggregate demand for tellers even as ATMs handled a growing share of routine transactions. This illustrates the core insight economists describe as complementarity—technology that replaces one function of a job often creates complementary demand for the remaining human functions.

### B. The Lump of Labor Fallacy

A recurring error in public discourse about automation is the implicit assumption that there exists a fixed quantity of work to be allocated among workers. This is known as the "lump of labor fallacy" [4]. In practice, human needs and wants are effectively unbounded; technological productivity gains expand the frontier of what can be created and delivered, generating demand for new categories of work that did not previously exist. Autor [5] documents extensively that while automation consistently eliminates specific tasks within jobs, it rarely eliminates the jobs themselves entirely—instead restructuring them around the non-routine cognitive and interpersonal components that machines cannot readily replicate.

### C. The Task-Displacement Framework

The most influential contemporary theoretical framework for analyzing AI's labor market impact is the task-displacement model developed by Acemoglu and Restrepo [6]. This framework distinguishes between tasks that technology displaces (replacing human labor inputs) and tasks that technology complements (increasing the productivity and value of human labor inputs). The net employment effect of any technology depends on the relative magnitude of displacement and complementarity effects, which are empirically variable and context-dependent rather than technologically fixed [6].

Brynjolfsson and McAfee [7] extend this framework to AI specifically, arguing that the current wave of AI development disproportionately affects routine cognitive tasks—precisely the tasks that had been considered safe from automation during the first wave

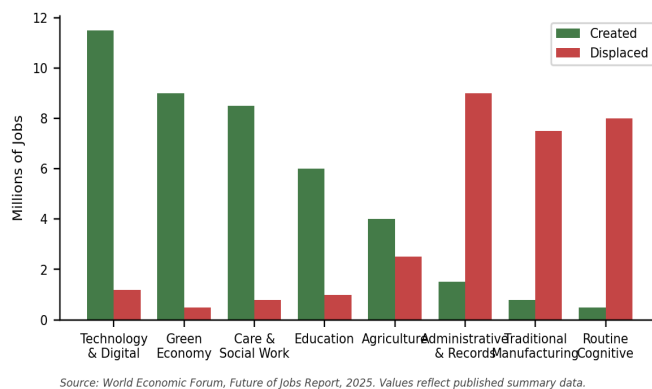
of industrial robotics, which primarily displaced routine physical tasks. This represents a qualitative shift in automation's scope, though not necessarily in the direction or magnitude of its long-run employment effects, which remain subjects of active empirical investigation.

## III. Current Data and Labor Market Projections

### A. Aggregate Projections

The World Economic Forum's Future of Jobs Report 2025 [8] projects the creation of approximately 170 million new roles globally by 2030, offset by the displacement of approximately 92 million existing roles, yielding a net positive projection of approximately 78 million jobs. These figures must be interpreted carefully. The aggregate projection is positive; the distributional reality is not uniformly so. The roles projected for creation and the roles projected for displacement are not colocated geographically, occupationally, or demographically. Workers whose roles are displaced are not, absent deliberate intervention, automatically positioned to fill the roles being created [8].

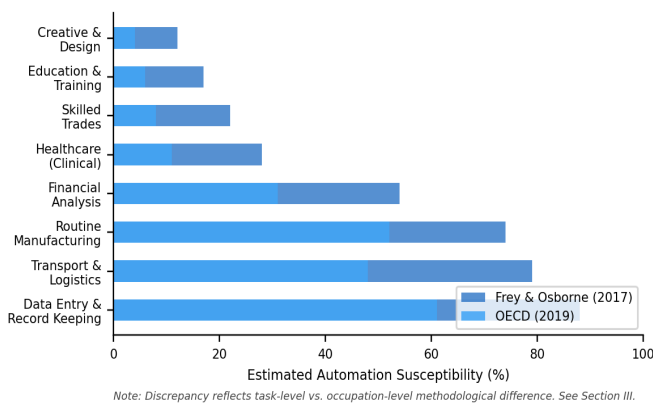
McKinsey Global Institute [9] projects that approximately 12 million occupational transitions will be required in the United States alone by 2030. The Institute identifies that transitions of this scale have historically required an average of five to eight years per worker when adequate support infrastructure is available, and significantly longer when it is not.



**Fig. 1.** Projected global job creation and displacement by sector, 2025–2030. Data derived from World Economic Forum Future of Jobs Report, 2025 [8].

### B. Methodological Variance in Risk Estimates

The range of published estimates for automation susceptibility across occupations is wide. Frey and Osborne [10] assessed 47% of U.S. occupations as being at high risk of automation within two decades. A subsequent OECD analysis applying a task-level rather than occupation-level methodology arrived at a substantially lower estimate of 14% [11]. The methodological difference is significant: occupation-level analysis assesses whether the dominant tasks of a job are automatable, while task-level analysis disaggregates each occupation into its constituent tasks and assesses automation risk at that granular level. The latter approach consistently yields lower aggregate risk estimates because most occupations contain a mix of automatable and non-automatable tasks [11]. Neither estimate should be treated as definitive; both represent scenario analyses under particular methodological assumptions rather than empirical predictions.

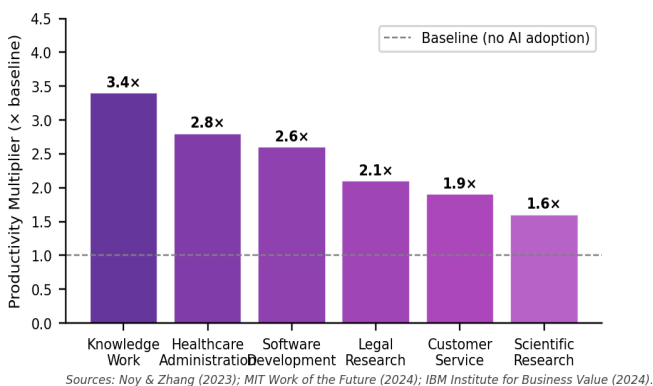


**Fig. 2.** Estimated automation susceptibility by occupation category, illustrating the methodological range between Frey & Osborne [10] and OECD [11] approaches. Values are approximate and for comparative illustration.

### C. The Productivity Augmentation Evidence

A consistent finding in studies of AI tool adoption at the worker level is a significant productivity uplift for workers who learn to use AI tools effectively in their domain. A controlled study by Noy and Zhang [12] found that professional knowledge workers using large language model tools completed tasks 37% faster on average, with higher quality ratings from evaluators. MIT Work of the Future [13] synthesized multiple sector-specific studies and found productivity multipliers ranging from 1.6 $\times$  to 3.4 $\times$  relative to non-adopting peers across knowledge work, healthcare administration, software development, legal research, and scientific research contexts (Fig. 3).

Critically, studies in this area consistently report that adopting workers do not report reduced job satisfaction or diminished sense of professional competence. On the contrary, the majority report that AI tools remove routine or administrative components of their work, allowing greater time for the higher-judgment activities they find most professionally meaningful [13].



**Fig. 3.** Observed productivity multiplier by sector for workers who adopt AI tools relative to non-adopting peers. Sources: Noy & Zhang [12]; MIT Work of the Future [13]; IBM Institute for Business Value [14].

## IV. Sectoral Analysis

### A. Healthcare and Bioinformatics

Healthcare is a domain in which the distinction between augmentation and replacement carries exceptional moral weight. Diagnostic imaging provides the clearest illustration. AI systems

trained on large labeled datasets have demonstrated sensitivity and specificity in detecting certain radiological findings—including early-stage malignancies—that are comparable to trained specialists on specific narrow tasks [15]. This capacity is most accurately understood not as replacement of radiological expertise but as a mechanism for extending specialist-level screening to patient populations in underserved or resource-constrained settings where specialist density is insufficient to meet clinical demand.

The structural protein prediction system AlphaFold, developed at DeepMind, provides a parallel example in bioinformatics. Following its 2021 publication in *Nature* [16], the system generated predictions for more than 200 million protein structures—effectively the entirety of the known protein universe. Rather than displacing structural biologists, this development accelerated research programs across drug discovery, rare disease investigation, and basic science, generating substantial new demand for skilled human interpretation and experimental validation [16]. The tool eliminated a computational bottleneck; the scientific work remained human-led.

Administrative burden in healthcare—documentation, billing, prior authorization, scheduling—consumes an estimated 30–40% of clinical time in many health systems [17]. AI that reduces this burden does not threaten clinical employment; it redirects clinical capacity from administrative function to patient-facing care, the primary purpose of the clinical role.

### B. Energy Systems Engineering

The energy sector presents a dual relationship with AI: as a major consumer of AI-related computational power and as a domain in which AI optimization offers significant efficiency and access gains. AI-driven optimization of data center cooling systems has demonstrated energy reductions of approximately 40% in published industrial deployments [18]. Applied to grid management, renewable energy integration, and microgrid optimization, analogous approaches offer meaningful progress toward both decarbonization and energy access goals.

The energy transition from fossil fuel dependence to renewable infrastructure is among the largest capital and labor deployment challenges of the current century. It requires civil engineers, electrical engineers, grid operators, construction workers, maintenance technicians, regulatory specialists, and community engagement professionals in numbers that substantially exceed current trained supply in most jurisdictions [19]. AI accelerates design and optimization; it does not substitute for the physical installation and operational labor that the transition requires.

### C. Cybersecurity and Cryptographic Infrastructure

Cybersecurity represents a domain in which AI is simultaneously the most powerful defensive resource and a material vector for offensive capability expansion. The volume of security events, threat indicators, and vulnerability disclosures that contemporary organizations must monitor exceeds what human analyst teams can process manually without AI triage and prioritization tools [20]. AI does not reduce demand for security professionals; the contrary is well-documented. Security workforce shortage is a persistent structural feature of the sector, and AI tooling primarily functions to make the work of existing professionals tractable rather than to substitute for them.

The transition to post-quantum cryptographic standards—mandated for U.S. federal agencies by 2030 under NIST

guidance, with analogous timelines emerging in other jurisdictions—represents a workforce challenge of substantial scale [21]. The engineering, governance, procurement, and deployment work required to migrate institutional cryptographic infrastructure to quantum-resistant standards (FIPS 203, 204, 205) requires skilled human professionals at every layer. AI tools support this transition; they do not eliminate the professional competence required to manage it.

#### ***D. Creative and Knowledge Work***

Generative AI tools capable of producing text, images, code, and audio at scale have disrupted portions of the creative economy—particularly in contexts where commodity-level creative output was previously purchased at modest per-unit cost. The evidence on skilled creative professionals is more nuanced. Studies of professional designers, writers, and software engineers who adopt generative AI tools consistently find that these professionals increase output volume, take on more complex projects, and command higher rates for the directorial, strategic, and originality-dependent dimensions of their work that AI cannot reliably replicate [14].

The relevant historical parallel is desktop publishing. The introduction of tools such as PageMaker and Illustrator in the 1980s eliminated the role of the paste-up artist and typesetter while substantially expanding the professional scope and commercial demand for graphic design as a discipline. Design moved from a production craft to a strategic communications function, with correspondingly broader application and higher professional valuation [7]. Analogous dynamics are plausibly underway with generative AI and creative work, though the transition period involves genuine disruption for practitioners whose roles are primarily in commodity production.

#### ***E. Manufacturing and Skilled Trades***

Industrial robotics and AI-enabled automation have had documented displacement effects in routine manufacturing—particularly in high-wage economies where labor cost differentials make automation economically attractive [6]. Acemoglu and Restrepo [6] estimated that each industrial robot introduced per thousand workers was associated with a reduction of 0.2% in the employment-to-population ratio in affected commuting zones, alongside wage suppression effects.

Skilled trades—plumbing, electrical installation, HVAC, welding, masonry, carpentry—exhibit substantially higher resistance to automation due to the requirement for continuous physical adaptation to unstructured and unpredictable environments [10]. These roles require real-time sensory judgment—haptic, auditory, visual—across configurations that differ substantially from site to site. Current robotic systems do not operate reliably in these environments. The skilled trades sectors in Canada, the United States, and across Europe face a structural workforce deficit rather than a surplus, driven by demographic aging of the existing workforce and insufficient investment in apprenticeship and vocational training pipelines [19].

#### ***F. Education***

Effective teaching is consistently ranked among the occupations least susceptible to automation in task-level analyses [10][11]. The core of the teaching function is not information transmission—which has been freely available via internet-accessible resources for decades—but the relational and adaptive dimensions of pedagogy: detecting individual student confusion, adjusting explanation in real

time, providing calibrated encouragement, modeling intellectual curiosity, and maintaining the social-emotional conditions in which learning occurs [22]. These functions are not candidates for near-term automation.

AI tutoring and adaptive learning systems have demonstrated meaningful efficacy as complements to human instruction, particularly for drill, practice, and individualized pacing [22]. Their most productive application is not to substitute for teachers but to provide teachers with fine-grained data on student progress and to handle the lower-level practice components of learning, freeing teacher time for the high-judgment instructional work that produces durable educational outcomes.

### **V. Psychological and Social Dimensions**

Economic analyses of technological unemployment consistently underweight the non-income dimensions of employment. Work provides, in addition to income, a primary source of social identity, daily temporal structure, belonging to a community of practice, and the psychological experience of competence and contribution [23]. Self-Determination Theory (SDT), developed by Deci and Ryan [23], identifies three fundamental psychological needs—autonomy, competence, and relatedness—all of which are substantially satisfied by meaningful employment and substantially threatened by job displacement.

The public health evidence on unemployment and health outcomes is robust and consistent. A systematic review and meta-analysis by Roelfs et al. [24] found that unemployment was associated with a 63% increase in all-cause mortality risk after controlling for confounding factors. This elevated risk cannot be fully explained by income loss; it persists even when controlling for economic hardship, suggesting that the loss of work itself—*independent of its financial implications*—constitutes a material health risk [24].

Communities that have experienced concentrated and rapid industrial displacement—coal-dependent regions in Appalachia, manufacturing cities in the post-industrial American Midwest, former steel communities in the United Kingdom—have demonstrated predictable patterns of deteriorating community health: elevated rates of depression and anxiety, increased substance use disorders, reduced civic participation, declining educational attainment in subsequent generations, and documented reductions in life expectancy [25]. These outcomes are not attributed solely to income loss; they reflect the disintegration of the social and institutional fabric that employment sustains.

This evidence base implies that any policy or design framework that treats AI-driven labor displacement as primarily a redistribution problem—solvable by income transfer alone—is addressing a necessary but insufficient component of the challenge. Effective transition support must address the social, purposive, and relational dimensions of work alongside the financial ones.

### **VI. Policy Frameworks for Human-Centered Transition**

#### ***A. Workforce Retraining***

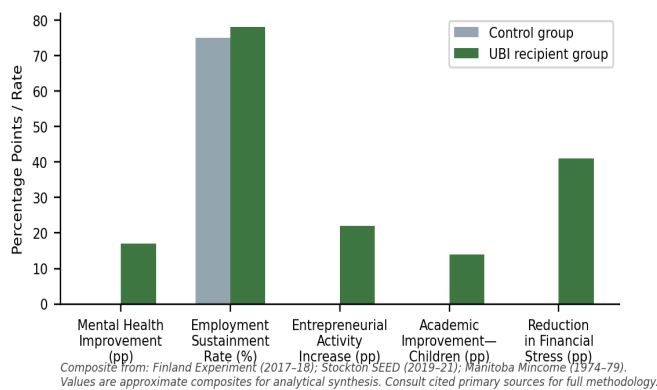
Workforce retraining programs are the most widely implemented policy response to technological displacement. Their effectiveness varies substantially with design. Programs that provide skills demonstrably in market demand, adequate income replacement during training, and geographic flexibility in delivery show

measurable positive employment outcomes. Generic or supply-led programs disconnected from employer demand show substantially weaker results [26]. Germany's Kurzarbeit short-time work scheme, which subsidizes reduced working hours to retain workers during economic transitions, and Denmark's "flexicurity" model—which combines flexible hiring and firing with generous income replacement and active labor market policy—are frequently cited as models demonstrating that labor market flexibility and worker security are not in conflict [26].

### B. Universal Basic Income: Evidence from Pilots

Universal Basic Income (UBI) has been the subject of several controlled or quasi-experimental pilots since 2016. The Finnish Basic Income Experiment (2017–2018) provided 2,000 unemployed individuals with unconditional monthly payments. Recipients reported significant improvements in mental wellbeing, trust in institutions, and life satisfaction, with no significant reduction in employment relative to controls [27]. The Stockton SEED program (2019–2021) in California found that recipients experienced higher rates of full-time employment than controls, alongside improvements in physical and mental health measures [28]. The Manitoba Mincome experiment (1974–1979) found reduced hospitalization rates and improved secondary school completion among recipients [29]. GiveDirectly's ongoing Kenya program has generated findings of sustained positive economic and wellbeing effects at scale [30].

Aggregate economic critiques of UBI—primarily inflationary risk and sustainability at full-population scale—remain subjects of active economic debate. The pilot evidence does not resolve these macro-level questions; it does, however, consistently refute the prediction that unconditional income support reduces recipients' motivation to work or creates adverse behavioral incentive effects [27][28].



**Fig. 4.** Composite summary of outcomes from major UBI pilot programs. Values are approximate composites for analytical synthesis. Sources: Finnish Basic Income Experiment [27]; Stockton SEED [28]; Manitoba Mincome [29].

### C. Portable Benefits and Work-Week Reduction

Portable benefits—employment-adjacent protections (healthcare, pension, parental leave) that attach to the worker rather than the employment relationship—have been proposed and partially implemented in several jurisdictions as a mechanism for maintaining social protection through increasingly non-standard labor market participation patterns [26]. Separately, the four-day work week has been evaluated in large-scale trials in Iceland (2015–2019), the United Kingdom (2022), Japan, and Australia. The UK trial,

involving 61 companies and approximately 2,900 workers, found sustained or improved productivity in 92% of participating companies, alongside significant improvements in worker health, burnout reduction, and retention [31]. This approach distributes the productivity dividend as time rather than solely as income, and addresses labor distribution challenges by spreading available work across a broader workforce.

## VII. Open-Source AI as Democratic Infrastructure

The question of who controls AI systems that shape consequential decisions is a dimension of the employment debate that receives insufficient attention. When AI systems used in hiring, performance management, credit assessment, criminal justice, or public benefit administration are proprietary, independent verification of their behavior is structurally unavailable. Research has documented systematic bias and error in proprietary AI systems used in consequential contexts, identifiable only when researchers were able to audit the systems through indirect means [32].

Open-source AI development—in which code, training data documentation, and model architecture are publicly available—enables verification by independent researchers, civil society, and affected communities. It reduces the access differential between large organizations and smaller enterprises or public institutions. And it provides a practical accountability mechanism: if a system makes decisions affecting employment, housing, or health, those affected can examine how the system works rather than relying solely on the representations of its developer [32][33].

The productivity dividend from AI should, in part, accrue to the public—both because AI development relies on public data, publicly funded research, and publicly educated labor, and because the social infrastructure that makes AI deployment possible is itself a public good. Open-source AI development is one mechanism by which this public benefit can be realized without requiring legislative action in every jurisdiction.

## VIII. Discussion

The evidence reviewed in this paper supports several conclusions that, taken together, constitute a coherent framework for human-centered AI development and transition policy.

First, the aggregate long-run employment effect of AI is unlikely to be catastrophically negative based on both historical analogy and current projections. However, the distributional effects during the transition period are likely to be significantly inequitable absent deliberate intervention. Workers in routine cognitive and physical roles—disproportionately lower-wage workers with less access to retraining resources—bear the primary displacement risk. The workers most likely to capture the productivity dividend from AI adoption are those with existing educational credentials and organizational access to AI tools [8][9].

Second, the augmentation/replacement distinction is a design choice with real consequences. Systems designed to maintain human decision authority—where AI informs and humans decide—distribute the benefits of AI capability while preserving human agency, accountability, and professional development. Systems designed to minimize human involvement in the interest of cost reduction sacrifice these dimensions.

Third, the non-income dimensions of employment—identity, structure, belonging, purpose—are not adequately addressed by

income replacement policies alone. Effective transition frameworks must include mechanisms for meaningful occupational transition, not merely financial bridging.

Fourth, transparency and accountability in AI systems are preconditions for democratic legitimacy in AI-affected labor markets. When algorithmic systems shape hiring, performance evaluation, benefit allocation, and resource distribution, the absence of public accountability is inconsistent with democratic principles of governance.

Limitations of this analysis should be noted. This paper presents a synthetic review of published evidence rather than primary empirical research. Projection data are inherently uncertain; labor market forecasting has a documented record of significant error. Sectoral case analyses are illustrative rather than exhaustive. The policy evidence on UBI, in particular, draws on pilot programs that differ substantially in scale, population, and design from any potential full-implementation program. Readers are encouraged to consult cited primary sources for methodological detail.

## IX. Conclusion

Artificial intelligence will continue to transform labor markets across virtually every sector of the global economy. The central finding of this paper is that this transformation is not technologically predetermined in its human consequences. The outcomes for workers, communities, and societies depend on governance choices, design decisions, and policy frameworks that remain, in significant part, within human control.

The evidence supports a framework centered on five principles: (1) AI systems should be designed to augment human capability rather than eliminate human roles wherever feasible; (2) the productivity dividend from AI deployment should be broadly shared; (3) human decision authority should be preserved at points of consequential judgment; (4) AI systems affecting employment and welfare should be subject to independent audit; and (5) transition support must address the psychological and social dimensions of labor displacement alongside the financial.

Collaboration between human capability and machine intelligence is not merely a preferable alternative to replacement—it is demonstrably more productive across a broad range of domains, more consistent with democratic accountability, and more likely to generate the broad-based human flourishing that justifies technological development in the first instance. Building this future requires deliberate choices from AI developers, employers, policymakers, and individuals. The evidence reviewed here suggests those choices are available to us—and that the cost of not making them deliberately will be borne, as has historically been the case, disproportionately by those least able to absorb it.

## Glossary of Terms

### Artificial Intelligence (AI)

Computer systems designed to perform tasks typically requiring human intelligence, including pattern recognition, natural language processing, decision-making, and content generation. AI is not a single technology but a family of computational approaches.

### Augmentation

The use of AI to extend human capability rather than replace human roles. In augmentation, the human retains decision authority and professional function; AI handles computational or data-intensive components.

### Automation

The use of technology to perform tasks with reduced or eliminated human involvement. Automation is distinguished from augmentation by the degree to which the human role is preserved or eliminated.

### Labor Displacement

The elimination of jobs or tasks due to technological change. Displacement may be temporary (requiring occupational transition) or permanent within a given sector.

### Large Language Model (LLM)

An AI system trained on large text corpora to generate, summarize, translate, and respond to natural language. LLMs are the architecture underlying widely deployed conversational AI tools.

### Lump of Labor Fallacy

The erroneous assumption that there is a fixed quantity of work available in an economy. Economic evidence consistently shows that technological productivity gains expand the frontier of possible production, generating new categories of demand and work.

### Open-Source AI

AI systems whose code, architecture, and training data documentation are publicly available for inspection, modification, and redistribution. Open-source AI enables independent verification of system behavior.

### Post-Quantum Cryptography

Cryptographic algorithms designed to resist attacks from quantum computers, which can break widely used asymmetric cryptographic systems (RSA, ECC). NIST has standardized post-quantum algorithms under FIPS 203, 204, and 205.

### Productivity Dividend

The economic gains resulting from technology-driven increases in output per unit of labor or capital input. The distribution of this dividend among workers, shareholders, and the public is a central question of AI governance.

### Self-Determination Theory (SDT)

A motivational framework (Deci & Ryan, 1985) identifying three fundamental human psychological needs: autonomy, competence, and relatedness. Meaningful employment substantially satisfies all three; job displacement threatens all three.

### Standard of Living

A composite measure of material wellbeing including income, housing, healthcare access, education, safety, and leisure. Maintenance of the standard of living is proposed in this paper as a non-negotiable floor in AI transition policy.

### Universal Basic Income (UBI)

A policy in which all citizens receive a regular, unconditional cash transfer regardless of employment status. Pilot evidence suggests positive effects on wellbeing, employment, and entrepreneurial activity at tested scales.

## References

- [1] R. Maslej et al., "The AI Index Report 2024: Measuring Trends in Artificial Intelligence," Stanford University Human-Centered AI Institute (HAI), Stanford, CA, Apr. 2024. [Online]. Available: <https://aiindex.stanford.edu/report/>
- [2] E. P. Thompson, *The Making of the English Working Class*. London: Victor Gollancz, 1963.
- [3] J. E. Bessen, "How computer automation affects occupations: technology, jobs, and skills," Boston University School of Law Working Paper No. 15-49, 2016. [Online]. Available: <https://ssrn.com/abstract=2690435>
- [4] D. Autor, "Work of the past, work of the future," *AEA Papers and Proceedings*, vol. 109, pp. 1–32, 2019. doi: 10.1257/pandp.20191110
- [5] D. H. Autor, "Why are there still so many jobs? The history and future of workplace automation," *Journal of Economic Perspectives*, vol. 29, no. 3, pp. 3–30, 2015. doi: 10.1257/jep.29.3.3
- [6] D. Acemoglu and P. Restrepo, "The race between man and machine: implications of technology for growth, factor shares, and employment," *American Economic Review*, vol. 108, no. 6, pp. 1488–1542, 2018. doi: 10.1257/aer.20160696
- [7] E. Brynjolfsson and A. McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton, 2014.
- [8] World Economic Forum, "Future of Jobs Report 2025," Geneva, Switzerland: WEF, 2025. [Online]. Available: <https://www.weforum.org/reports/the-future-of-jobs-report-2025>
- [9] McKinsey Global Institute, "The future of work after COVID-19," McKinsey & Company, Feb. 2021, updated 2023. [Online]. Available: <https://www.mckinsey.com/featured-insights/future-of-work>
- [10] C. B. Frey and M. A. Osborne, "The future of employment: how susceptible are jobs to computerisation?" *Technological Forecasting and Social Change*, vol. 114, pp. 254–280, 2017. doi: 10.1016/j.techfore.2016.08.019
- [11] OECD, *OECD Employment Outlook 2019: The Future of Work*. Paris: OECD Publishing, 2019. doi: 10.1787/9ee00155-en
- [12] S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023. doi: 10.1126/science.adh2586
- [13] MIT Work of the Future, "AI and the Future of Work," Massachusetts Institute of Technology Industrial Performance Center, Cambridge, MA, Research Brief, 2024.
- [14] IBM Institute for Business Value, "Augmented work for an automated, AI-driven world," IBM Corporation, Armonk, NY, 2024.
- [15] M. D. Abràmoff et al., "Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices," *npj Digital Medicine*, vol. 1, art. 39, 2018. doi: 10.1038/s41746-018-0040-6
- [16] J. Jumper et al., "Highly accurate protein structure prediction with AlphaFold," *Nature*, vol. 596, pp. 583–589, 2021. doi: 10.1038/s41586-021-03819-2
- [17] C. Sinsky et al., "Allocation of physician time in ambulatory practice: a time and motion study in 4 specialties," *Annals of Internal Medicine*, vol. 165, no. 11, pp. 753–760, 2016. doi: 10.7326/M16-0961
- [18] D. Silver et al., "Cooling data centres with deep reinforcement learning," DeepMind Blog, 2016. Subsequent operational deployment results reported by Google LLC, 2018. [Online]. Available: <https://deepmind.google>
- [19] International Labour Organization, *World Employment and Social Outlook 2023: The Value of Essential Work*. Geneva: ILO, 2023.
- [20] ISC2, "ISC2 Cybersecurity Workforce Study 2023," ISC2, Clearwater, FL, 2023. [Online]. Available: <https://www.isc2.org/research/workforce-study>
- [21] National Institute of Standards and Technology, "Post-Quantum Cryptography Standardization: FIPS 203, 204, 205," NIST, Gaithersburg, MD, 2024. doi: 10.6028/NIST.FIPS.203
- [22] B. Means et al., "Evaluation of evidence-based practices in online learning: a meta-analysis and review of online learning studies," U.S. Department of Education, Office of Planning, Evaluation, and Policy Development, 2010.
- [23] E. L. Deci and R. M. Ryan, *Intrinsic Motivation and Self-Determination in Human Behavior*. New York: Plenum, 1985.
- [24] D. J. Roelfs, E. Shor, K. W. Davidson, and J. E. Schwartz, "Losing life and livelihood: a systematic review and meta-analysis of unemployment and all-cause mortality," *Social Science & Medicine*, vol. 72, no. 6, pp. 840–854, 2011. doi: 10.1016/j.socscimed.2011.01.005
- [25] A. Case and A. Deaton, *Deaths of Despair and the Future of Capitalism*. Princeton, NJ: Princeton University Press, 2020.
- [26] OECD, *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. Paris: OECD Publishing, 2023. doi: 10.1787/08785bba-en
- [27] K. Kangas et al., "The basic income experiment 2017–2018 in Finland: preliminary results," Ministry of Social Affairs and Health, Reports and Memorandums 2019:9, Helsinki, Finland, 2019.
- [28] A. Tubbs and E. West, "SEED: Stockton economic empowerment demonstration—findings from the first year," Stockton Economic Empowerment Demonstration, 2021. [Online]. Available: <https://www.stocktondemonstration.org>
- [29] E. Forget, "The town with no poverty: using health administration data to revisit outcomes of a Canadian guaranteed annual income field experiment," *Canadian Public Policy*, vol. 37, no. 3, pp. 283–305, 2011. doi: 10.3138/cpp.37.3.283
- [30] J. Haushofer and J. Shapiro, "The short-term impact of unconditional cash transfers to the poor: experimental evidence from Kenya," *Quarterly Journal of Economics*, vol. 131, no. 4, pp. 1973–2042, 2016. doi: 10.1093/qje/qjw025
- [31] 4 Day Week Global, "The Results Are In: The UK's Four-Day Week Pilot," Cambridge, UK: Autonomy / Cambridge University, 2023. [Online]. Available: <https://www.4dayweek.com/research>
- [32] J. Angwin et al., "Machine bias: there's software used across the country to predict future criminals. And it's biased against blacks," *ProPublica*, May 23, 2016. [Online]. Available: <https://www.propublica.org>
- [33] M. Mitchell et al., "Model cards for model reporting," in *Proc. ACM Conf. Fairness, Accountability, and Transparency (FAccT)*, New York, 2019, pp. 220–229. doi: 10.1145/3287560.3287596