WHITE PAPER

AI, Labor, and the Social Contract

Rethinking the Relationship Between Technology, Enterprises, and Society

May 2025
Strategic Analysis Series

CONTENTS

Table of Contents

A NOTE ON FRAMING

The Mechanism That Matters: Task Elimination

Much of the existing discourse on AI and labor frames the challenge as automation — the replacement of humans by machines performing the same tasks. That framing, while historically useful, obscures the more precise and more consequential mechanism at work in the current AI era. [4,20]

AI does not primarily replace workers. It eliminates tasks. And in some cases — as a secondary and more visible effect — it displaces jobs. The distinction is critical, because task elimination is the primary driver of seat compression, and it operates invisibly. [2,18]

How the mechanism works

When AI removes a task from a human workflow — a report that no longer needs writing, an analysis that no longer needs a analyst, a customer interaction that no longer requires a live agent — no individual loses their job in that moment. The role still exists. But it now requires fewer people to perform it. [5,10]

Multiply that across every function within an organization. Then multiply it across every organization in every industry. The aggregate effect is a steady, quiet contraction in the total number of economically viable human roles — not because of dramatic displacement events, but because each role simply requires fewer people over time. [1,6]

Job displacement — where AI eliminates a role entirely — does occur, and it is the more politically visible phenomenon. But it is task elimination that drives the deeper structural compression. No announcement is made. No crisis moment occurs. The seat count contracts through attrition, restructuring, and roles that are never backfilled. [2,18]

	Task Elimination	Job Displacement
Mechanism	Task Elimination	Job Displacement
What happens	Tasks removed from human workflows	Entire roles eliminated
Visibility	Invisible — accumulates silently	Visible — people lose jobs
Primary driver of seat compression?	Yes — highest impact	Secondary impact
How it manifests	Fewer people needed per role over time	Specific functions closed
Example	Team of 12 analysts becomes a team of 7	Entire data entry department eliminated
Policy response	Difficult to see or legislate against	Triggers immediate political reaction

The simultaneity of this process across industries removes the historical safety valve. In prior technological transitions, workers displaced in one sector could migrate to expanding sectors. When task elimination is occurring across every sector concurrently — healthcare, legal, finance, logistics, marketing, software, education — there is no adjacent industry absorbing the overflow. The total seat count contracts across the entire economy at once. [8,19]

> ***"AI does not primarily replace workers. It eliminates tasks. Fewer tasks mean fewer humans needed — quietly, across every industry, simultaneously."***

This distinction — task elimination as the primary mechanism, job displacement as the more visible but secondary one — is the analytical foundation on which the remainder of this paper rests.

Executive Summary

Artificial intelligence represents a profound departure from prior technological transitions. Unlike earlier waves that primarily shifted which tasks humans performed, AI increasingly eliminates tasks from human workflows entirely — quietly reducing how many people are needed per role, per function, and per organization. When this occurs simultaneously across every major industry, the aggregate effect is a compression of the total number of economically viable human seats. [5,10,6]

The central economic question of the AI era is not whether new jobs will emerge. History strongly suggests they will. The more pressing — and largely unanswered — question is whether enough economically meaningful roles will emerge quickly enough, at sufficient wages, and in sufficient volume to absorb large-scale simultaneous seat compression across multiple sectors. [8,19,12]

If AI materially compresses the total number of economically secure human participation opportunities relative to population size, the resulting strain on traditional labor markets, tax systems, and social safety nets may force governments, enterprises, and society to fundamentally rethink the economic and social contract. [8,12,15]

> ***"The long-term risk of AI may not be mass unemployment — it may be the hollowing out of economically secure seats relative to population and expectations."***

This paper examines the structural forces driving this shift, the mechanisms through which AI-driven task elimination and seat compression could reshape existing institutions, and the emerging policy and enterprise response frameworks that may define the next era of economic governance.

SECTION ONE

What Makes This AI Wave Different

Throughout modern economic history, the prevailing response to automation anxiety has been consistent: technology destroys some jobs, but ultimately creates new ones. For most of the industrial and digital revolutions, that argument proved largely correct. Agricultural mechanization fed factories. Factory automation fed service economies. Digital software created vast new industries in technology, finance, consulting, logistics, and communications. [13,4]

But that historical pattern rested on critical assumptions that may no longer hold.

Four Characteristics That Distinguish AI

Prior automation waves were characterized by specificity — displacing particular categories of labor in defined sectors over extended time horizons. AI simultaneously exhibits four properties that no prior technology wave has combined: [4,20]

- Cross-industry applicability — AI systems are not sector-specific tools. The same underlying models and platforms can operate across healthcare, legal services, financial analysis, software development, design, education, marketing, customer support, and logistics.
- Concurrent deployment — Unlike the decades-long rollout of prior industrial revolutions, AI is being adopted simultaneously across virtually every major sector of the global economy.
- Cognitive task elimination rather than purely physical displacement — For the first time, AI removes knowledge and cognitive tasks from human workflows — areas previously considered relatively insulated from displacement risk.
- Near-zero marginal cost of replication — Once an AI capability is developed and deployed, reproducing that capability at scale costs almost nothing. This fundamentally changes the economics of knowledge work.

This combination is historically unprecedented. And it materially changes the analytical framework through which societies must assess the long-term labor market implications of technological advancement. [20,4]

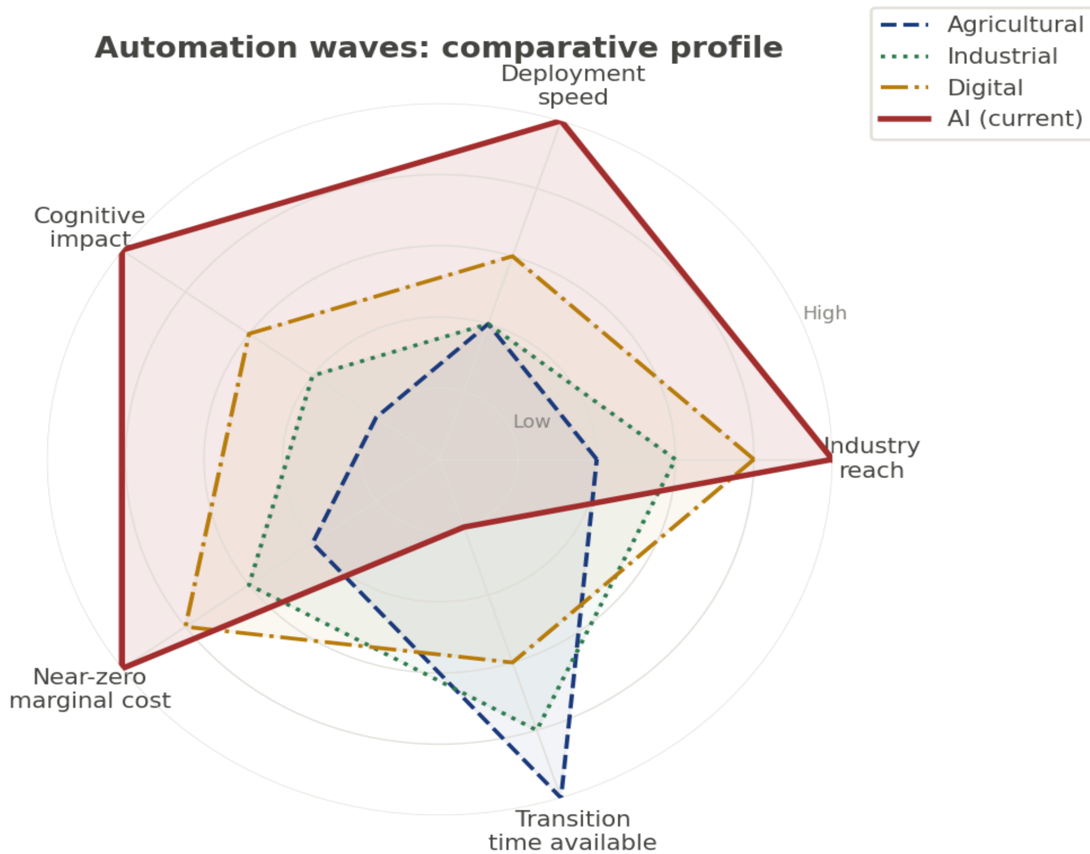

Figure 1: *Comparative profile of four automation waves across five structural dimensions*

The radar chart above illustrates how AI (red) scores at the maximum across all five dimensions simultaneously — a combination no prior wave achieved. Agricultural and industrial automation scored high on transition time available but low on cognitive impact. Digital software expanded industry reach but still left significant time for labor markets to adapt. AI collapses all five variables at once. [4,13]

The Replacement Surface Area

If AI systems can increasingly perform analysis, coding, customer support, legal drafting, research, design, logistics optimization, medical interpretation, and eventually portions of managerial coordination, the breadth of potential labor displacement is qualitatively different from anything prior technology waves produced. [6,10,5]

Previous technology transitions primarily augmented cognitive work. They gave knowledge workers better tools. AI increasingly has the capacity to replicate elements of the cognitive work itself. [14,2]

That distinction creates what might be termed a dramatically expanded task elimination surface — the proportion of total human economic activity from which AI can quietly remove tasks within a given time frame. [1,18,6]

A ten-year reassessment by Frey and Osborne from the University of Oxford's Oxford Internet Institute clarifies where the limits of this surface lie. In their 2024 update, they find that generative AI substantially expands automation potential in virtual and remote settings, while in-person interaction and genuinely novel creative judgment remain areas where human comparative advantage persists. Critically, their analysis suggests that the more transactional a task becomes, the more susceptible it is to automation — a finding directly relevant to the breadth of cognitive and administrative work at risk in enterprise settings. [26]

SECTION TWO

The Core Economic Tension

Modern economic systems are largely constructed around broad human labor participation. Income taxes, payroll taxes, healthcare funding, retirement systems, consumer demand, and social mobility all implicitly assume that large proportions of the population will participate meaningfully in the wage economy. [8,12,17]

AI introduces a structural tension into this architecture. [8,17]

The Productivity-Participation Divergence

As AI capabilities mature, enterprise productivity and profitability may rise dramatically — while labor demand compresses. This creates a potentially destabilizing divergence: [7,1]

- Enterprise output and value generation increase
- Capital returns accelerate
- But labor demand grows more slowly — or contracts in specific categories
- Wage growth weakens across broad middle-income segments
- Traditional payroll and income tax bases erode
- Government fiscal capacity to fund safety nets weakens

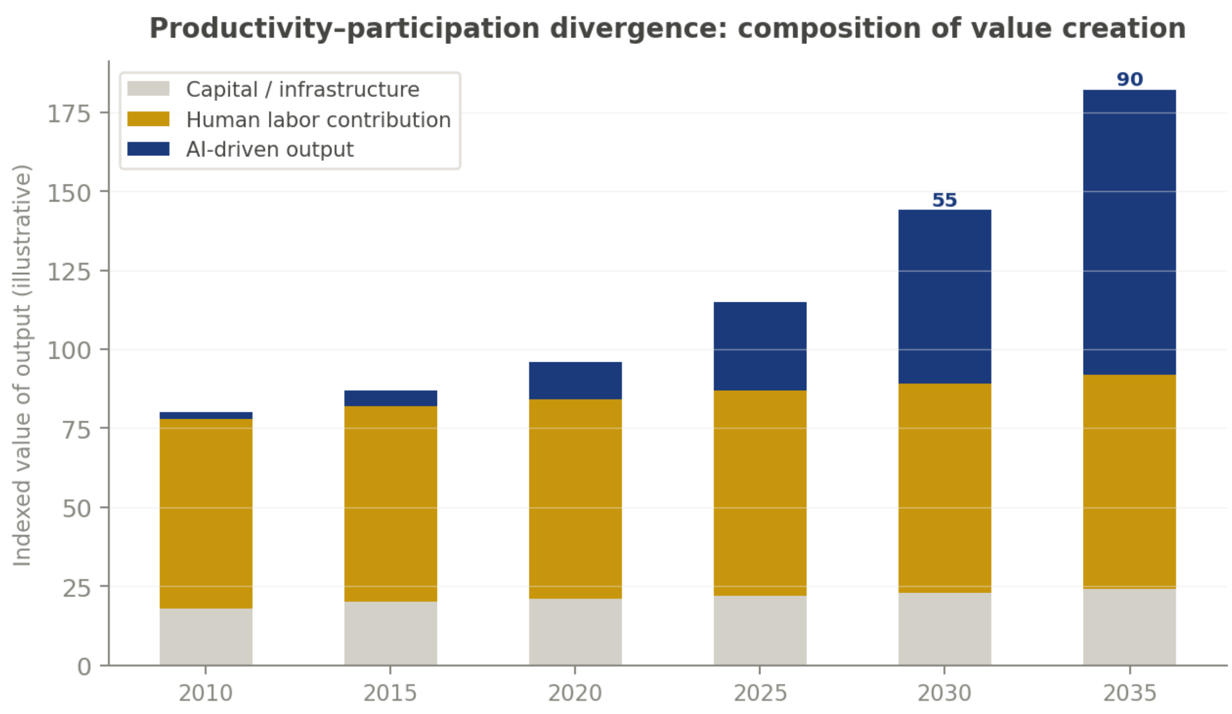

Figure 2: *Composition of value creation 2010–2035: AI-driven output versus human labor contribution*

The chart above illustrates how AI-driven output (navy) grows from a negligible share of total value creation in 2010 to the dominant component by 2035 — while human labor contribution (gold) remains relatively flat. This is the core structural tension: enterprise output expands dramatically, but the human share of that value diminishes proportionally. [7,10]

The critical insight is that an economy can theoretically become extraordinarily productive while simultaneously requiring fewer people to participate meaningfully in value creation. High GDP growth and deteriorating labor market conditions are not mutually exclusive — and that combination would be historically novel. [3,17,7]

> ***"An economy can become extraordinarily productive while simultaneously requiring fewer people to participate meaningfully in value creation."***

The Seat Count Problem

The most important macroeconomic variable in an AI-driven economy may not be GDP growth or aggregate productivity. It may be the total number of economically viable human participation opportunities — what might be termed available seats. [8,19,12]

The societal challenge is not simply whether AI creates new jobs. The relevant questions are far more precise: [8,19]

- Will enough new jobs emerge to absorb displaced workers across multiple sectors simultaneously?
- Will they emerge fast enough to prevent prolonged structural unemployment?
- Will they provide sufficient wages to maintain broad middle-class economic stability?
- Will they be accessible to workers across geographic regions and educational backgrounds?

Prior transitions unfolded over decades, giving labor markets, educational institutions, and governments time to adapt. AI-driven displacement may occur at a pace that outstrips those traditional adjustment mechanisms. [13,4,18]

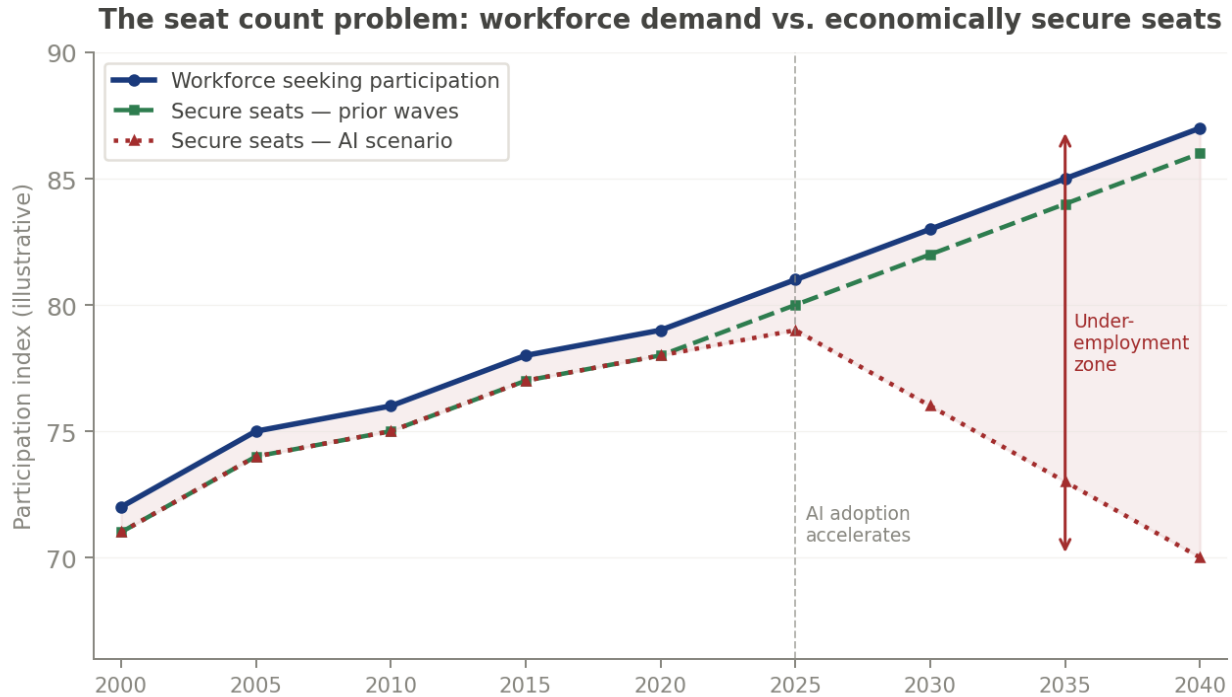

Figure 3: *The seat count problem: workforce demand versus economically secure seats under AI adoption*

The chart above contrasts two scenarios. In prior technological transitions (green dashed), secure seats kept pace with workforce growth — the gap remained negligible. In the AI scenario (red dotted), seats begin to diverge downward after 2025 as task elimination accelerates across industries, creating a growing underemployment zone: people still working, but in less secure, lower-wage roles that do not fully replace what was lost. [1,18]

The Risk of Chronic Underemployment

The greatest near-to-medium-term risk may not be absolute mass unemployment in the traditional sense. It may instead be a more insidious pattern: [9,12,8]

- Persistent underemployment across large segments of the workforce
- Downward wage pressure in previously stable middle-income roles
- Fragmented gig-style cognitive work replacing structured employment
- Fewer high-paying professional pathways for mid-career workers
- Increasing divergence between capital owners and wage earners

In that scenario, people may still work — but the number of economically rewarding seats will have contracted relative to population expectations. That distinction carries enormous implications for social cohesion, political stability, and the long-term functioning of consumer economies. [9,3,17]

SECTION THREE

The Historical Pattern of Institutional Obligation

When institutions become systemically influential — when their operations materially reshape the conditions under which large populations live and work — societies have historically responded by expanding the obligations those institutions are expected to carry.

This pattern has repeated across multiple transformative technologies and industries.

Precedents Across Industry

- Railroads — Initially celebrated as pure engines of economic progress, railroads were eventually regulated as infrastructure with obligations to public access, pricing fairness, and territorial service.
- Banking and Financial Services — Financial institutions, after the repeated shocks of the 20th century culminating in the 2008 crisis, were subjected to systemic regulation reflecting their role as critical economic infrastructure.
- Utilities — Energy, water, and telecommunications providers inherited public service obligations once their infrastructure became essential to societal functioning.
- Industrial firms — As environmental consequences became undeniable, industrial enterprises absorbed regulatory obligations reflecting their externalized costs on broader society.
- Social media platforms — As dominant platforms came to influence democratic discourse, privacy, and social dynamics at scale, regulatory pressure expanded accordingly.

In each case, the underlying logic followed a consistent pattern: once an institution materially reshapes society at scale, it accumulates obligations beyond pure profit optimization.

The table below maps each historical precedent to its triggering event, the regulatory obligation imposed, the time it took to emerge, and the direct parallel to AI enterprises.

Industry	Triggering event	Obligation imposed	Time	AI parallel
Railroads (1800s–1900s)	Monopoly pricing; rural service gaps	Public access mandates, fare controls, service obligations	Decades	AI access equity; algorithmic pricing scrutiny
Banking (Post-2008)	Systemic collapse; public bailouts required	Capital requirements, stress tests, systemic institution rules	Years	AI systemic risk frameworks; model audits
Utilities (Early 1900s)	Essential services; natural monopoly dynamics	Universal service, rate regulation, reliability standards	Decades	AI as essential infrastructure; access mandates
Industrial firms (Mid-1900s)	Environmental damage at scale	Emissions limits, cleanup obligations, environmental taxes	Decades	AI 'labor externality' taxes; displacement levies
Social media (2010s–now)	Election interference; privacy breaches	Content moderation, data rules, competition oversight	Years	AI transparency, accountability, anti-concentration

Industry	Triggering event	Obligation imposed	Time	AI parallel
AI enterprises (Next?)	Labor market compression at scale	Productivity taxes, retraining funds, social contributions	TBD	— (this paper's thesis)

Table 1: Historical precedents for enterprise obligation — from industry to AI

Why AI May Accelerate This Dynamic

AI firms — and potentially large enterprises leveraging AI to systematically compress labor markets — may increasingly face similar institutional pressure. The logic is straightforward: [20,3,15]

If your systems reshape society at scale — if they materially alter labor markets, compress cognitive employment, concentrate productivity gains, and affect the economic security of large populations — then you may eventually inherit obligations that reflect that systemic influence. [20,16]

That does not necessarily mean nationalization, extreme socialist policy, or punitive regulatory regimes. It more likely manifests incrementally, through a broadening of enterprise responsibility frameworks.

The historical record also cautions against assuming that long-run benefits will arrive smoothly. Frey's examination of the Industrial Revolution at the Oxford Martin School demonstrates that transformative technology can generate extraordinary long-run prosperity while simultaneously producing decades of wage stagnation, rising inequality, and political instability during the transition — the period he terms "Engels' Pause." The gap between aggregate output growth and real worker welfare in that era is directly analogous to the productivity-participation divergence this paper identifies for the AI era. The lesson is not that disruption is permanent, but that the transition period can span an entire working generation. [27]

SECTION FOUR

Emerging Policy and Enterprise Response Frameworks

Several distinct mechanisms may emerge as governments, enterprises, and international institutions grapple with AI-driven labor compression. These are not mutually exclusive — most plausible futures involve combinations of multiple approaches operating in parallel. [8,12,15]

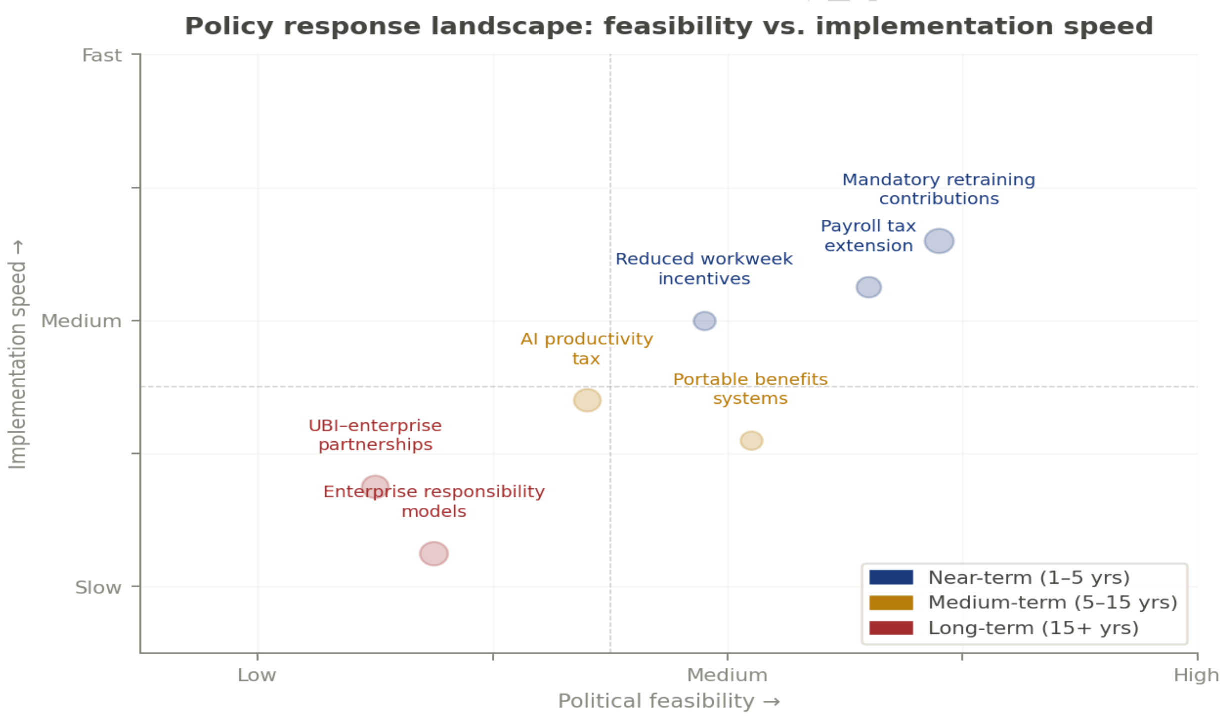

Figure 4: *Policy response landscape: political feasibility versus implementation speed*

The matrix above positions seven policy mechanisms by two critical variables: political feasibility and implementation speed. Near-term mechanisms (navy) such as mandatory retraining contributions and payroll tax extensions sit in the high-feasibility, high-speed quadrant — they build on existing frameworks and are politically actionable. Medium-term mechanisms (amber) require new legislative architecture. Long-term mechanisms (red) such as Universal Basic Income, UBI-enterprise partnerships and enterprise responsibility models face significant political headwinds but may become necessary if structural underemployment becomes persistent.

1. AI Productivity Taxation

Governments may introduce taxation models tied to AI-enabled productivity gains, autonomous labor systems, or labor displacement at scale. While the terminology may vary — and direct "robot taxes" have faced significant political resistance — the underlying economic logic represents a likely direction for fiscal adaptation. [15,3,12]

Plausible mechanisms include taxation on AI-generated output volumes, taxation on the labor cost savings attributable to AI adoption, and surcharges applied to companies demonstrating AI-driven reductions in workforce size relative to revenue growth. [15,12]

2. Mandatory Enterprise Social Contributions

Large enterprises could be required to fund or co-fund social stabilization programs as a condition of operating AI-driven systems that contribute to labor displacement. This could manifest as an extension of existing payroll-tax frameworks adapted for AI-era conditions: [11,12,19]

- Workforce retraining and reskilling contributions
- Lifelong education fund participation
- Transition assistance for displaced workers
- Healthcare coverage portability obligations
- Regional economic stabilization program funding

3. Public-Private Income Stabilization

As structural underemployment becomes more persistent, governments may increasingly pressure enterprises to participate directly in income stabilization mechanisms: [15,8,12]

- Wage stabilization programs in high-displacement sectors
- Incentivized reduced-workweek adoption
- Portable benefits and labor-sharing frameworks
- Universal basic income-adjacent pilot programs co-funded with large AI-leveraging enterprises

4. Enterprise Responsibility Models

Over time, dominant AI enterprises and heavy AI-adopting organizations may be formally treated as systemically important institutions — analogous to how financial regulators identify systemically important financial institutions — with accompanying obligations and oversight frameworks. [16,20,3]

This could manifest through enhanced reporting requirements on labor impact, regulatory approval processes for large-scale AI-driven workforce reductions, mandatory participation in national workforce transition frameworks, and governance obligations related to AI decision-making systems. [16,11]

SECTION FIVE

The Counterarguments and Optimistic Scenarios

Any serious analysis of AI-driven labor market transformation must account for the significant forces that may counterbalance displacement dynamics, as well as genuine uncertainties in long-run outcomes. [11,19]

The Innovation and Abundance Case

Optimistic scenarios envision AI driving an era of broadly distributed abundance — dramatically lower costs for goods and services, near-free access to high-quality education and healthcare diagnostics, accelerated scientific discovery, and the emergence of entirely new industries around AI that dwarf current employment categories. [4,10,7]

Proponents of this view note that prior predictions of technology-driven unemployment have consistently been too pessimistic, and that human creativity and the emergence of entirely new industries reliably absorb displaced workers over time. [13,18]

Competitiveness and Innovation Risks

Governments and enterprises alike may resist aggressive regulatory intervention on the grounds that: [11,19,20]

- Overly restrictive frameworks slow AI development and erode competitive positioning
- Enterprises subject to heavy AI-linked taxation may relocate to more favorable jurisdictions
- Early-stage regulation may constrain beneficial AI applications that could ultimately expand employment

The Honest Uncertainty

Intellectual honesty requires acknowledging genuine uncertainty about long-run outcomes. The history of automation has repeatedly confounded pessimistic predictions. New industries, new forms of human-AI collaboration, and unforeseen economic structures may yet produce favorable outcomes that today's analytical frameworks cannot fully anticipate. [13,18,4]

The critical concern, however, is not the long-run equilibrium — it is the transition period. Societies often react to disruption after pressure becomes visible rather than anticipating it. And transitions of this magnitude are rarely smooth.

The New Jobs Evidence: What the Strongest Counterargument Actually Shows

The most credible and empirically grounded challenge to this paper’s central thesis comes from Autor, Chin, Salomons, and Seegmiller’s landmark 2024 study of occupational change from 1940 to 2018. Their analysis of the US labor market finds that approximately 60 percent of employed workers in 2018 held jobs that did not exist in 1940 — occupations created by technological change itself. This evidence provides the strongest empirical foundation for the optimistic view that AI will generate new categories of work that cannot currently be anticipated. [21]

However, the same authors note that the pace of new task creation has slowed materially since the 1980s, even as technology adoption has accelerated — a finding that complicates the simple historical optimism narrative. The mechanism that generated new jobs in prior eras — the expansion of labor-intensive new industries — may operate more slowly when the new industries themselves are built on labor-saving architectures from the outset. This is precisely the analytical gap between historical precedent and AI-era dynamics that this paper's framing attempts to address.

Brynjolfsson, Rock, and Syverson's work on what they term the "productivity J-curve" offers a further important nuance. They argue that AI-driven productivity gains are real but temporarily invisible in aggregate statistics because the reorganization of work, skills, and processes required to capture them takes years to complete — creating a lag between technological adoption and measurable output growth. Under this framing, the productivity benefits of AI are not absent; they are deferred. This is a substantively different claim from arguing that disruption will be limited, and it does not resolve the transition-period risk that this paper identifies as central. [22]

The Policy Adjudication Problem: Why Frameworks Must Be Assessed, Not Only Listed

The policy framework section of this paper identifies four plausible mechanisms — productivity taxation, mandatory social contributions, income stabilization, and enterprise responsibility models. Each is legitimate; but presenting them as a catalog without adjudicating between them understates the real obstacles each faces.

Rodrik and Sabel's work on building a "good jobs economy" makes a persuasive case that the more productive policy direction is not redistribution after the fact but restructuring labor markets to ensure that technological gains flow into employment quality rather than purely into returns to capital. Their framework suggests that enterprise-level investment in workforce skill development, supported by active industrial policy, is more likely to produce durable outcomes than taxation-and-transfer models that leave the underlying labor market structure intact. [23]

On the specific question of AI taxation, Furman and Seamans' analysis for the National Bureau of Economic Research concludes that direct "robot taxes" are likely to be economically inefficient and practically unenforceable in their naive form — but that targeted payroll tax reforms, combined with wage insurance for displaced workers, represent a more tractable near-term policy tool. Their work suggests the feasibility problem with AI taxation is not the concept but the measurement: determining which productivity gains are attributable to AI adoption, rather than to capital investment or management quality, is analytically difficult and legally contested. [24]

The EU AI Act (2024), the world's first comprehensive AI regulatory framework, offers a real-world test case for the enterprise responsibility model this paper describes. Its risk-based architecture — imposing tiered obligations on AI systems based on potential societal harm — demonstrates that governments are already moving toward the institutional obligation frameworks this paper projects, even if labor market impacts are not yet the primary regulatory trigger. Critically, the EU framework's implementation has revealed the core political economy challenge: multinational enterprises can relocate development activities to less regulated jurisdictions, creating regulatory arbitrage that weakens the effectiveness of any single-jurisdiction approach and underscoring the need for international coordination that Section 6 of this paper identifies. [25]

A further complication for both policy adjudication and the optimistic counterargument is that standard economic statistics may systematically undercount the benefits AI is already generating — as well as some of its costs. Coyle and Poquiz at the University of Cambridge argue in an NBER working paper that generative AI creates significant productivity and quality improvements in workplace tasks that do not appear in GDP or productivity accounts, because they manifest as faster completion, higher output quality, or process changes that evade conventional measurement. This measurement gap cuts in two directions: it suggests that pessimistic readings of AI's productivity impact may be too early, but it also means that policymakers are attempting to design regulatory responses to labor market shifts they cannot yet fully quantify. [28]

Finally, a dimension of AI-driven labor market disruption that this paper's analysis does not fully address is the intergenerational dimension. Research from IESE Business School identifies a structural risk that goes beyond current-cohort displacement: if AI automates entry-level tasks, younger workers are deprived of the early career interactions through which tacit workplace knowledge is transmitted. Ide's model shows that this can raise short-run productivity while simultaneously slowing long-run growth and welfare — because future generations inherit a workforce with weaker foundational skills. This intergenerational externality is not captured by headline employment or wage statistics and presents a further argument for proactive enterprise responsibility frameworks rather than waiting for the aggregate damage to become visible. [29]

SECTION SIX

Implications and Strategic Considerations

For Governments

Governments face a dual mandate: enabling the productivity gains of AI adoption while protecting the structural foundations of wage-based societies. This may require proactive investment in adaptive safety net frameworks, fiscal innovation to replace eroding payroll tax bases, and international coordination to prevent regulatory arbitrage. [8,12,15]

For Enterprises

Enterprises capturing disproportionate AI-driven productivity gains should anticipate an evolving regulatory and social environment. Organizations that proactively engage in workforce transition programs, demonstrate measurable commitment to equitable AI adoption, and participate in policy design will likely navigate the transition period more effectively than those that do not. [11,19,20]

There is also an important long-term economic self-interest dimension: consumer economies depend on consumers having purchasing power. Enterprises that benefit from AI while simultaneously eroding the wage base of their customer populations are ultimately operating against their own long-run commercial interests. [17,7]

For Society

The most important societal variable may not be the aggregate economic outcomes AI produces. It may be whether the benefits of AI-driven productivity are distributed broadly enough to maintain the political legitimacy and social cohesion upon which stable, functioning economies depend. [16,20,15]

> ***"AI may represent not merely another technological transition, but a restructuring of the economic and social contract itself."***

Conclusion

The question of whether AI will create or destroy jobs misses the deeper structural challenge. AI is not simply another productivity-enhancing technology operating within existing economic frameworks. It may be a force capable of compressing the total volume of economically secure human participation opportunities at a pace and scale that existing institutions were not designed to absorb. [6,8,10]

The historical record provides a clear pattern: when technologies become systemically influential, the institutions that deploy them inherit obligations that reflect that influence. Banks, utilities, telecommunications providers, industrial firms, and social media platforms all followed this trajectory. [20,16]

AI-driven enterprises may be next. Not necessarily through ideological redistribution, but through the practical necessity of maintaining social stability, fiscal function, and the consumer demand that advanced economies require to operate. [15,3,12]

The defining policy and governance questions of the AI era are therefore not primarily technological. They are fundamentally about the social contract — about what obligations institutions owe to the societies whose labor markets, economic structures, and social cohesion they materially reshape. [20,16,15]

Those questions do not yet have settled answers. But they will increasingly define the political economy of the coming decades.

References

The following works represent the core academic, institutional, and policy literature informing the analysis presented in this paper. Readers seeking to explore specific arguments in greater depth are directed to the sources below.

[1] Acemoglu, D. & Restrepo, P. (2017). Robots and Jobs: Evidence from US Labor Markets. NBER Working Paper No. 23285. National Bureau of Economic Research. nber.org/papers/w23285

[2] Acemoglu, D. & Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. Journal of Economic Perspectives, 33(2), 3–30. aeaweb.org — DOI: 10.1257/jep.33.2.3

[3] Acemoglu, D. (2021). Harms of AI. NBER Working Paper No. 29247. National Bureau of Economic Research. nber.org/papers/w29247

[4] Brynjolfsson, E. & McAfee, A. (2014). The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. W. W. Norton & Company. Available via booksellers and university libraries

[5] Brynjolfsson, E., Li, D. & Raymond, L. (2023). Generative AI at Work. NBER Working Paper No. 31161. National Bureau of Economic Research. nber.org/papers/w31161

[6] Frey, C. B. & Osborne, M. A. (2017). The Future of Employment: How Susceptible Are Jobs to Computerisation? Technological Forecasting and Social Change, 114, 254–280. DOI: 10.1016/j.techfore.2016.08.019 sciencedirect.com — DOI: 10.1016/j.techfore.2016.08.019

[7] Goldman Sachs Economics Research. (2023). The Potentially Large Effects of Artificial Intelligence on Economic Growth (Briggs/Kodnani). Goldman Sachs Global Investment Research, March 2023. Note: The full report is available via Goldman Sachs's client portal; the link below leads to a public summary article. goldmansachs.com/insights

[8] International Monetary Fund. (2024). Gen-AI: Artificial Intelligence and the Future of Work. IMF Staff Discussion Note SDN/2024/001. imf.org — Full PDF

[9] Katz, L. F. & Krueger, A. B. (2019). The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015. ILR Review, 72(2), 382–416. DOI: 10.1177/0019793918820008 journals.sagepub.com — DOI: 10.1177/0019793918820008

[10] McKinsey Global Institute. (2023). The Economic Potential of Generative AI: The Next Productivity Frontier. McKinsey & Company. mckinsey.com — Full Report

[11] MIT Work of the Future. (2020). The Work of the Future: Building Better Jobs in an Age of Intelligent Machines. MIT Task Force on the Work of the Future. workofthefuture.mit.edu

[12] OECD. (2023). OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market. OECD Publishing, Paris. oecd.org — Full Report

[13] 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. aeaweb.org — DOI: 10.1257/jep.29.3.3

[14] Autor, D., Levy, F. & Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. Quarterly Journal of Economics, 118(4), 1279–1333. nber.org/papers/w8337 — Open Access PDF

[15] Korinek, A. & Stiglitz, J. E. (2021). Artificial Intelligence, Globalization, and Strategies for Economic Development. NBER Working Paper No. 28453. nber.org/papers/w28453

[16] Mazzucato, M. (2021). Mission Economy: A Moonshot Guide to Changing Capitalism. Allen Lane. Available via booksellers and university libraries

[17] Piketty, T. (2014). Capital in the Twenty-First Century. Harvard University Press. hup.harvard.edu

[18] Autor, D. & Salomons, A. (2018). Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share. Brookings Papers on Economic Activity, Spring 2018. brookings.edu — Full Paper

[19] World Economic Forum. (2023). Future of Jobs Report 2023. World Economic Forum, Geneva. weforum.org — Full Report & PDF

[20] Acemoglu, D., Autor, D. & Johnson, S. (2023). Can We Have Pro-Worker AI? Choosing a Path of Machines in Service of Minds. CEPR Policy Insight No. 123. MIT Shaping the Future of Work Initiative. shapingwork.mit.edu — Open Access PDF

[21] Autor, D., Chin, C., Salomons, A. & Seegmiller, B. (2024). New Frontiers: The Origins and Content of New Work, 1940–2018. The Quarterly Journal of Economics, 139(3), 1399–1465. academic.oup.com — Oxford Academic

[22] Brynjolfsson, E., Rock, D. & Syverson, C. (2021). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. American Economic Journal: Macroeconomics, 13(1), 333–372. doi.org/10.1257/mac.20180386

[23] Rodrik, D. & Sabel, C. (2022). Building a Good Jobs Economy. In D. Allen et al. (Eds.), A Political Economy of Justice (pp. 61–95). University of Chicago Press. Working paper version (HKS RWP20-001, 2019) freely available via Columbia Law Scholarship Archive and SSRN (ssrn.com/abstract=3533430). scholarship.law.columbia.edu — Open Access PDF

[24] Furman, J. & Seamans, R. (2019). AI and the Economy. Innovation Policy and the Economy, 19, 161–191. National Bureau of Economic Research. nber.org/chapters/c14031

[25] European Parliament. (2024). Regulation (EU) 2024/1689 — The EU Artificial Intelligence Act. Official Journal of the European Union. Publications Office of the EU. eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689

[26] Frey, C. B. & Osborne, M. (2024). Generative AI and the Future of Work: A Reappraisal. Brown Journal of World Affairs, 30(1), 1–17. University of Oxford (Oxford Internet Institute / Oxford Martin School). ora.ox.ac.uk/objects/uuid:f52030f5-23eb-4481-a7f1-8006685edbae

[27] Frey, C. B. (2019). The Technology Trap: Capital, Labor, and Power in the Age of Automation. Princeton University Press. (Oxford Martin School, University of Oxford.) press.princeton.edu/books/hardcover/9780691172798/the-technology-trap

[28] Coyle, D. & Poquiz, J. L. S. (2025). Making AI Count: The Next Measurement Frontier. NBER Working Paper No. 34330. University of Cambridge (Bennett School of Public Policy). nber.org/papers/w34330

[29] Ide, E. (2025). Automation, AI, and the Intergenerational Transmission of Knowledge. IESE Business School Working Paper. SSRN 5360803. Department of Economics, IESE Business School, University of Navarra. ssrn.com/abstract=5360803

About This AI Assisted Paper

This white paper synthesizes a structured analytical exchange exploring the long-term labor market, fiscal, and societal implications of large-scale AI adoption. It is intended for policymakers, enterprise leadership, economists, and informed general readers engaged with questions at the intersection of technology, economics, and governance.