The Economics of Artificial Intelligence: An Integrated Analysis of Daron Acemoglu's Research Program

Synthesized from 30 publications (2018–2026)
by Daron Acemoglu and coauthors

March 2026

Abstract

This report synthesizes thirty publications by Daron Acemoglu and coauthors that collectively constitute the most influential research program on the economics of artificial intelligence. Drawing on the full text of these papers, the analysis traces how Acemoglu's task-based framework—distinguishing displacement from reinstatement, easy from hard tasks, and automation from augmentation—generates a set of precise, quantified predictions about AI's economic impact. The central findings are striking: AI will likely increase total factor productivity by no more than 0.53–0.66% over the next decade; 50–70% of the rise in US wage inequality since 1980 is attributable to task displacement; even technologies that improve low-skill worker productivity can paradoxically increase inequality through ripple effects; and AI-generated "new bad tasks" (social media manipulation, deepfakes) could appear to add 2% to GDP while actually reducing welfare by 0.72%. The Nobel Lecture's utility-technology possibilities frontier reveals a deeper truth: the tech sector's *ideology*—not merely its economic power—can push society to a lower frontier where everyone is worse off. Policy intervention through tax reform, gradual regulation, and deliberate redirection of AI toward worker-complementary applications can reverse this trajectory.

Contents

1 Introduction **2**

2 The Task-Based Framework **2**

2.1 Why Factor-Augmenting Models Fail 2

2.2 The Core Model: Displacement and Reinstatement 2

2.3 Low-Skill and High-Skill Automation 3

2.4 Automation and Rent Dissipation 3

2.5 Automation and Polarization 3

3 Empirical Evidence **4**

3.1 Robots and Jobs in US Labor Markets 4

3.2 Competing with Robots: Firm-Level Evidence from France 4

3.3 Selection vs. Causal Effects in Technology Adoption 4

3.4 The 2019 Annual Business Survey . . . 4
3.5 Demographics as a Driver of Automation . . . 5
3.6 50–70% of US Wage Inequality from Task Displacement . . . 5

4 The Macroeconomic Case for Skepticism **5**
4.1 Hulten's Theorem and the 0.66% Estimate . . . 5
4.2 Easy vs. Hard Tasks . . . 5
4.3 The Inequality Paradox . . . 6
4.4 New Bad Tasks: GDP Up, Welfare Down . . . 6
4.5 Wage and Inequality Effects . . . 6

5 The Direction of Innovation **6**
5.1 Why the Market Gets It Wrong . . . 6
5.2 Quantifying the Distortion . . . 7
5.3 The Tax Code Bias . . . 7
5.4 Historical Precedent: Ricardo's Pivot . . . 7

6 The Digital Platform Economy **7**
6.1 Too Much Data . . . 7
6.2 Behavioral Manipulation . . . 7
6.3 Online Misinformation . . . 8
6.4 Digital Advertising and Naïve Users . . . 8

7 The Nobel Lecture: Institutions, Technology, and Prosperity **8**
7.1 The Utility-Technology Possibilities Frontier . . . 8
7.2 Moves Along vs. Shifts of the Frontier . . . 8
7.3 Three Mechanisms of Persistence . . . 8
7.4 AI as a Critical Juncture . . . 9

8 Regulating AI **9**
8.1 Optimal Adoption Is Slow and Convex . . . 9
8.2 Building Pro-Worker AI . . . 9

9 Integrating the Research Program **9**
9.1 Theme 1: Technology's Direction Is a Choice, Not a Destiny . . . 9
9.2 Theme 2: Aggregate Effects Are Modest; Distributional Effects Are Enormous . . . 10
9.3 Theme 3: Firm-Level Evidence Is Misleading . . . 10
9.4 Theme 4: The Digital Economy Amplifies Harms . . . 10
9.5 Theme 5: Demographics Complicate the Simple Story . . . 10
9.6 Theme 6: We Are at a Critical Juncture . . . 10

10 Critical Assessment **10**

11 Conclusion **11**

1 Introduction

Daron Acemoglu, the 2024 Nobel laureate in Economics, has produced the most comprehensive and theoretically grounded body of work on the economics of artificial intelligence. His research program spans foundational theory, empirical evidence from multiple countries, macroeconomic forecasting, digital platform analysis, and institutional political economy—all unified by a single intellectual thread: the economic consequences of AI are determined not by the technology's capabilities but by the *direction* in which it is developed, which depends on incentives, institutions, and ideology.

This analysis synthesizes thirty publications spanning 2018–2026, organized into six pillars: (1) the task-based theoretical framework, (2) empirical evidence on robots and AI, (3) macroeconomic assessment, (4) the direction of innovation, (5) the digital platform economy, and (6) the institutional political economy of technology.

2 The Task-Based Framework

2.1 Why Factor-Augmenting Models Fail

The canonical skill-biased technical change (SBTC) model of Katz and Murphy (1992) treats technology as factor-augmenting: it multiplies the effective labor supply of skilled workers, with the college premium determined by relative supply and the elasticity of substitution σ. In "Unpacking Skill Bias" (2020), Acemoglu and Restrepo show this model is fatally limited. Without technological *regress*, the SBTC model predicts that unskilled wages should be rising—yet they have declined for four decades. Moreover, generating the observed college premium increase between 1963 and 1987 would require TFP growth of at least 1.9% per annum, whereas US TFP grew only 1.2%. The SBTC model's core mechanism—substitution of more productive skilled workers for less productive unskilled ones, mediated by σ—simply cannot match the data.

The task-based alternative resolves both problems. Changes in the skill premium are driven by the *task content of production*—which tasks are reallocated between factors—rather than by factor-augmenting productivity. Critically, the effect on inequality is "completely driven by the set of tasks (weighted by their effective productivity) that unskilled labor loses relative to the entire set of tasks previously performed by these workers" and "is not mediated by the elasticity of substitution." This decoupling of inequality from productivity is the framework's key conceptual insight: automation can generate "large swings in the skill premium from very small changes in TFP." Acemoglu calculates that if automation reduces task costs by 30% (as with industrial robots), the increase in the college premium between 1963 and 1987 "can be explained with as little as 0.54 percent per annum growth in TFP"—less than half of what the SBTC model requires.

2.2 The Core Model: Displacement and Reinstatement

In "Automation and New Tasks" (2019), production of a unique final good requires N tasks combined via a CES aggregator with elasticity $\sigma \leq 1$ (tasks are gross complements). Tasks below threshold I are automated (capital); tasks above I use labor. The wage equation decomposes into:

- **Displacement effect** (negative): automation expands I, taking tasks from labor.
- **Productivity effect** (positive): cost savings from substituting cheaper capital for labor increase output.
- **Reinstatement effect** (positive): new tasks expanding N create new labor demand and always increase wages and the labor share.

When $\sigma < 1$ (empirically ≈ 0.5), task complementarities that raise A_L may not increase wages at all—because more abundant labor-intensive tasks see their prices fall faster than their physical output rises. Acemoglu shows that when $\sigma \approx s_K$ (the capital share, ≈ 0.4), "the overall impact [on wages] will be essentially zero." Only new task creation unambiguously helps workers.

2.3 Low-Skill and High-Skill Automation

"Low-Skill and High-Skill Automation" (2018) introduces a richer comparative advantage structure. Capital performs both simple tasks $i \in [0, I_L]$ (competing with low-skill labor) and complex tasks $i \in [J, I_H]$ (competing with high-skill labor), with a disjoint set in between performed by labor. This abandonment of the "supermodular" single-crossing assumption is necessary to study AI, which automates tasks across the entire complexity distribution.

The model's most important novel result is the **ripple effect**: when the threshold M separating low-skill from high-skill tasks falls below J (so the two types of labor directly compete for some tasks), high-skill automation displaces high-skill workers who then "compete with low-skill labor in other tasks and displace this latter group." Both types of automation can therefore depress wages for *all* workers simultaneously when capital is scarce. The ripple effects vanish only when comparative advantage around M is very strong ($\varepsilon \to \infty$) or when the two labor types are buffered by capital-produced tasks ($M > J$).

Nevertheless, the **inequality** effects always go in the intuitive direction: low-skill automation increases inequality; high-skill automation decreases it. In the long run, with capital accumulation restoring the rental rate, automation always increases the total wage bill—but can still depress the wage of the directly displaced factor.

2.4 Automation and Rent Dissipation

"Automation and Rent Dissipation" (2025) introduces a mechanism that makes the market failure even worse: firms automate not just to save costs but to *dissipate worker rents.* When workers earn quasi-rents (wages above their outside option), firms adopt automation technologies that are *less productive than human workers*—because eliminating worker bargaining power is profitable even when the technology itself is not. This generates *excessive automation*: technologies that reduce both output and wages get adopted because they shift surplus from labor to capital.

2.5 Automation and Polarization

"Automation and Polarization" (2026) shows that automation of routine middle-skill tasks can explain the "hollowing out" of the wage distribution—the simultaneous growth of high-wage and low-wage employment at the expense of middle-wage jobs—as a single technological force rather than multiple independent trends.

3 Empirical Evidence

3.1 Robots and Jobs in US Labor Markets

The landmark "Robots and Jobs" (2020) estimates that one additional robot per thousand workers reduces the employment-to-population ratio by 0.2 percentage points and wages by 0.42%. Between 1990 and 2007, robot adoption accounts for 400,000 lost manufacturing jobs. The effects concentrate among workers without college degrees, and there is **no evidence of offsetting employment gains** in other industries within the same commuting zone.

3.2 Competing with Robots: Firm-Level Evidence from France

"Competing with Robots" (2020) reveals a crucial distinction between firm-level and market-level effects. Robot-adopting French firms see 20% higher value added, 4.3 percentage point lower labor shares, and—surprisingly—*higher* employment (10.9%). But this is entirely due to adopters expanding at competitors' expense. A 10 percentage point increase in robot adoption by competitors reduces their employment by 2.5% and value added by 2.1%. The **overall market-level effect** is negative: a 20 percentage point increase in robot adoption (the sample average) is associated with a 3.2% decline in industry employment. About 80% of the decline in the manufacturing labor share covariance term is accounted for by robot adopters being large, expanding further, and experiencing sizable labor share declines.

3.3 Selection vs. Causal Effects in Technology Adoption

"Advanced Technology Adoption" (2023) delivers a sobering finding: the positive association between firm size and AI/robot adoption is driven by **selection, not causation**. Using the Longitudinal Business Database from 1978–2018, Acemoglu et al. show that establishments at adopting firms were already larger and growing faster *decades before* these technologies existed. Robot adopters from the 1977–1984 cohort were 24.3% larger at entry. Employment growth differentials "remained largely unchanged or became less pronounced" after adoption intensified. This means firm-level studies overstating positive employment effects of AI should be interpreted with great caution.

3.4 The 2019 Annual Business Survey

The ABS reveals that only **3.2% of US firms used AI** during 2016–2018, and 2% used robotics. But because adopters are disproportionately large, 12.6% of US workers are employed at AI-using firms and 15.7% at robotics-using firms. Adopters have 11.4% higher labor productivity, lower labor shares, and 55–65% report using these technologies for automation. Integration costs are substantial—for industrial robots, they can be four times the equipment cost.

3.5 Demographics as a Driver of Automation

"Demographics and Automation" (2021) shows that population aging drives robot adoption through directed technology: as middle-aged workers become scarce, their wages rise, making automation more profitable. Aging explains **35% of cross-country variation** in robot adoption. A 20 percentage point increase in aging (roughly the Germany–US difference) produces 0.15 more robots per thousand workers per year, accounting for 50% of the robot adoption gap between the countries. The relationship is causal (instrumenting with historical birth rates) and concentrates in industries relying on middle-aged workers.

3.6 50–70% of US Wage Inequality from Task Displacement

"Tasks, Automation, and the Rise in US Wage Inequality" (2022) provides the most comprehensive decomposition of the US wage structure from 1947–2016, finding that **between 50 and 70 percent** of changes since 1980 are attributable to task displacement. At the industry level, a 10% increase in displacement during 1987–2016 is associated with an 8% increase in relative demand for college workers.

4 The Macroeconomic Case for Skepticism

4.1 Hulten's Theorem and the 0.66% Estimate

"The Simple Macroeconomics of AI" (2024) applies Hulten's theorem to discipline AI productivity forecasts. For a competitive economy with constant returns to scale, the TFP change from AI equals:

$$d \ln \text{TFP} = \bar{\pi} \times \text{GDP share of tasks impacted by AI}$$

where $\bar{\pi}$ is average cost savings. Using Eloundou et al.'s (2024) exposure estimates: 20% of US labor tasks are exposed to AI. Applying Svanberg et al.'s (2024) feasibility filter (23% can be profitably automated within 10 years), the GDP share of tasks actually impacted is $0.23 \times 0.200 = 4.6\%$.

For cost savings, Acemoglu uses the average of Noy and Zhang (2023) (40% faster task completion) and Brynjolfsson et al. (2023) (14% faster), yielding 27% labor cost savings, or 14.4% overall cost savings after adjusting for industry labor shares. Thus: TFP gains $= 0.046 \times 0.144 \approx 0.66\%$ over 10 years, or 0.064% annual TFP growth.

4.2 Easy vs. Hard Tasks

Acemoglu introduces a crucial distinction. **Easy tasks** have reliable, observable outcome metrics and simple action-outcome mappings (boiling an egg, writing standard subroutines). **Hard tasks** involve complex, context-dependent factors without clear success metrics (diagnosing a persistent cough, creative writing). Existing experimental evidence comes from easy tasks. Using a classification procedure (verb analysis, IWA coding, LDA topic modeling, gradient-boosted trees), 72.6% of exposed tasks are easy.

Assuming hard-task cost savings are one-quarter of easy-task savings (7% vs. 27% labor cost savings), the adjusted estimate becomes: $0.033 \times 0.144 + 0.013 \times 0.037 \approx 0.53\%$ TFP over 10 years.

4.3 The Inequality Paradox

Acemoglu proves formally that even when AI improves low-skill worker productivity (reducing λ_H/λ_U), inequality can *increase*. The mechanism: when $\sigma < 1$, higher productivity in low-skill tasks reduces task prices more than it raises physical output. If the effect is large enough, high-skill workers are reallocated away from these tasks, and both types of workers compete in higher-indexed tasks where the relative wage is determined by $\omega > \lambda_H/\lambda_U$. The skill premium thus rises despite the improvement in low-skill productivity. This means "the overall inequality implications of AI cannot be directly deduced from its effects on the performance of workers of different skills in a given set of tasks."

4.4 New Bad Tasks: GDP Up, Welfare Down

AI generates new tasks with negative social value—deepfakes, manipulative social media, addictive algorithms, malicious cyber attacks. Using Bursztyn et al.'s (2023) experiment with college students: users would pay \$53/month to keep using social media, but simultaneously pay \$19/month to get everyone off it. Thus for every \$53 of revenue, there is a net negative welfare effect of \$19, a proportionate damage of $-19/53 \approx -0.36$.

Combining this with revenues from social media/digital ads (\$460 billion, 1.64% of GDP) plus IT security spending (\$78 billion, 0.28%), Acemoglu estimates that AI-powered bad tasks could add approximately 2% to measured GDP while **reducing actual welfare by 0.72%** in consumption-equivalent units.

4.5 Wage and Inequality Effects

Using 500 demographic groups and the propagation matrix from Acemoglu and Restrepo (2022), the paper finds AI exposure is more equally distributed across demographic groups than pre-AI automation. Nevertheless: (1) between-group inequality increases slightly (standard deviation of log wages rises from 0.35 to 0.36); (2) low-education women, especially white native-born, experience real wage declines; and (3) the capital share increases by 0.31 percentage points, widening the gap between capital and labor income. GDP increases 1.4% (with easy/hard distinction) while average wages rise only 0.77%—confirming that capital captures disproportionate gains.

5 The Direction of Innovation

5.1 Why the Market Gets It Wrong

"The Wrong Kind of AI?" (2019) argues that AI is being developed as an automation platform when it should be a platform for creating new tasks. Acemoglu identifies education, healthcare, and augmented reality in manufacturing as areas where AI could create "new, high-productivity tasks for labour" rather than replacing it. The market failures are multiple: externalities in innovation, competing paradigms where the leading one locks in even if inferior, withdrawal of government from steering R&D, the "ethos" of Silicon Valley companies that "ceaselessly tried to automate everything," tax subsidies for capital over labor, and the fact that worker-complementary AI requires missing complementary inputs (teacher training, hospital reorganization).

5.2 Quantifying the Distortion

"Distorted Innovation" (2023) estimates the socially optimal automation ratio is 18% lower than equilibrium ($n^{SP}/n^{EQ} \approx 0.82$), with welfare losses of about 1% in consumption-equivalent terms from the pecuniary externalities of automation ($\tilde{\tau}_2 = 0.07$).

5.3 The Tax Code Bias

"Does the US Tax Code Favor Automation?" (2020) shows effective tax rates on labor exceed those on capital, creating artificial incentives to substitute machines for workers even when efficiency gains are minimal.

5.4 Historical Precedent: Ricardo's Pivot

"Learning from Ricardo and Thompson" (2024) documents that David Ricardo initially told Parliament in 1819 that "machinery did not lessen the demand for labor," but reversed this view by 1821 after witnessing the devastation of handloom weavers. Wood's (1910) wage series show nominal weekly wages fell from 240 pence in 1806 to 99 pence in 1820. Real wages fell to roughly 25% of their golden-age peak. Despite this, 240,000–250,000 weavers remained in the occupation. Economy-wide real wages were largely flat from the 1770s to the 1830s. The Peterloo Massacre of 1819—where 60,000 people demonstrating for reform were violently dispersed—was directly connected to weaving worker discontent and likely influenced Ricardo's change of mind.

6 The Digital Platform Economy

6.1 Too Much Data

"Too Much Data" (2022) proves that data markets generate excessive sharing because of externalities: one user's shared data reveals information about non-sharing users through correlations. The platform acquires too much data at depressed prices. Under some conditions, social surplus is negative—shutting down data markets would improve welfare. A "decorrelation" scheme (a trusted intermediary that removes correlations between sharing and non-sharing users' data) can restore efficiency.

6.2 Behavioral Manipulation

"When Big Data Enables Behavioral Manipulation" (2025) introduces the concept of *glossiness*—products that temporarily appear better than their quality warrants. The platform's superior knowledge of glossiness states (from data on similar users) creates two opposing forces. The **helpfulness effect** (when glossiness is short-lived, high ρ): the platform steers users toward genuinely good products. The **manipulation effect** (when glossiness is long-lived, low ρ): the platform exploits users by offering low-quality glossy products. Acemoglu proves via the FKG–Harris inequality that for sufficiently small ρ, the platform strictly prefers to offer the glossy low-quality product—behavioral manipulation dominates. A "big data double whammy": more products give more opportunities to find glossy items, intensifying manipulation.

6.3 Online Misinformation

"A Model of Online Misinformation" (2024) shows that profit-maximizing platforms choose maximal connectivity when article reliability is high but fragmented, homophilic networks (echo chambers) when reliability is low. Greater polarization increases sharing and engagement, explaining platforms' incentives to promote divisive content.

6.4 Digital Advertising and Naïve Users

"Online Business Models" (2025) shows that advertising-based platforms exploit naïve users with heavy ad loads while excluding sophisticated users, and the inflated demand drives up monopoly prices, harming both types.

7 The Nobel Lecture: Institutions, Technology, and Prosperity

7.1 The Utility-Technology Possibilities Frontier

Acemoglu's 2025 Nobel Lecture introduces the **utility-technology possibilities frontier (UTPF)**—a curve in (u_p, u_r) space showing achievable utility combinations for poor and rich groups. The frontier can be **backward-bending**: excessive extraction by the elite can reduce even their own welfare (a political economy version of the Laffer curve). More importantly, the frontier can bend backward through *technology choices*: "if all technologies focus on automating work, this might ultimately lead to excessive automation, whereby even tasks that could be performed more productively by labor are now allocated to capital."

7.2 Moves Along vs. Shifts of the Frontier

Changes in fiscal tools (redistribution) move society *along* the frontier. But institutional and technology choices can *shift* the frontier inward or outward. Critically, "efficiency and distribution of the gains are often inseparable; many institutions simultaneously determine the shifts of the frontier and the location of society along the frontier."

7.3 Three Mechanisms of Persistence

Holdup: Without institutional constraints, governments expropriate investments, discouraging production and shifting the frontier inward.

Economic losers: Elites resist technologies that shift the frontier outward when the gains accrue to others—especially when fiscal tools cannot compensate them.

Political losers: Elites resist superior technologies because adoption threatens their political power through "political creative destruction." Even when the elite could benefit from point A on a higher frontier, the risk of losing power makes them prefer point C on a lower frontier.

7.4 AI as a Critical Juncture

The lecture's Figure 12 depicts four possible AI frontiers. Frontier 1 is the pre-AI status quo. Frontier 2 represents pro-worker AI that shifts the frontier outward. Frontier 3 represents *excessive automation*—the frontier shifts *inward*. Frontier 4 is the best case. The devastating insight: "an ideological bias in favor of one group can push us toward lower frontiers. In contrast, a bias in the distribution of power favoring one group can tilt the frontier in their favor but does not typically lead to a downward shift of the frontier." The tech sector's automation ideology is not just redistributive—it is destructive.

8 Regulating AI

8.1 Optimal Adoption Is Slow and Convex

"Regulating Transformative Technologies" (2024) shows that optimal adoption of a transformative technology is gradual, enabling learning about disaster risk. Paradoxically, **a higher growth rate leads to slower optimal adoption** when damages are large—because faster-growing technology creates larger potential damages, and the effective discount rate for future output declines (making delay less costly). A quantitative example with $g_N = 2.6\%$ and reasonable disaster parameters shows optimal adoption takes about **40 years**; at $g_N = 3.8\%$, it takes nearly **60 years**. Equilibrium adoption is inefficiently fast because firms internalize only private, not social, damages. Sector-specific Pigouvian taxes can restore optimality; when infeasible, a "regulatory sandbox" that restricts high-damage sectors until risk is low can improve welfare.

8.2 Building Pro-Worker AI

"Building Pro-Worker Artificial Intelligence" (2026) proposes: (1) redirect AI toward worker-complementary applications (better information for electricians, personalized education, healthcare decision support); (2) reform tax incentives to eliminate the capital-over-labor bias; (3) invest in complementary skills and organizational change; (4) strengthen labor institutions; (5) regulate AI deployment in high-stakes contexts.

9 Integrating the Research Program

9.1 Theme 1: Technology's Direction Is a Choice, Not a Destiny

Across all 30 papers, the central message: the same AI capabilities can automate or augment. The Nobel Lecture shows technology choices shift the UTPF itself. The wrong choices—driven by ideology, tax distortions, rent dissipation incentives, path dependence, and scientific culture—can make *everyone* worse off.

9.2 Theme 2: Aggregate Effects Are Modest; Distributional Effects Are Enormous

AI will add 0.5–1.5% to GDP over a decade (modest). But 50–70% of four decades of rising US inequality is from task displacement. The capital share rises by 0.31 percentage points from AI. Low-education women see real wage declines. New bad tasks could create negative welfare effects masked by GDP increases.

9.3 Theme 3: Firm-Level Evidence Is Misleading

Robot-adopting firms grow—but at competitors' expense. Overall industry employment declines. AI-adopting firms were already large and growing faster before AI existed. Firm-level evidence systematically overstates the positive employment effects of automation.

9.4 Theme 4: The Digital Economy Amplifies Harms

Data markets generate excessive sharing. Platforms manipulate through glossiness exploitation. Misinformation is profit-maximizing. Advertising models exploit behavioral biases. The combination concentrates power while degrading the information environment.

9.5 Theme 5: Demographics Complicate the Simple Story

Aging societies adopt robots partly to fill genuine labor shortages. Automation driven by demographics doesn't displace available workers—it substitutes for workers who are no longer there. Optimal automation levels depend on demographic context.

9.6 Theme 6: We Are at a Critical Juncture

Like colonialism and the Industrial Revolution, AI is a moment where small institutional differences generate enormous divergences. The Progressive Era eventually channeled industrialization toward shared prosperity—but it took 60–80 years. The question is whether we can do better this time.

10 Critical Assessment

1. **Are the macro estimates too conservative?** The 0.53–0.66% estimate assumes current AI capabilities and 10-year adoption horizons. If capabilities improve faster (as frontier models suggest), or if "hard" tasks become "easy" through architectural breakthroughs, the upper bound shifts. Acemoglu acknowledges this: "if AI helps create new tasks that increase productivity..., its consequences can be much more positive."
2. **Can the task framework capture general-purpose AI?** The framework excels for narrow, task-specific automation. General-purpose systems that simultaneously transform many tasks across the complexity distribution may not follow the "marginal task" logic.
3. **What about endogenous task creation?** The framework treats new tasks as largely

exogenous. If AI itself generates new labor-intensive tasks (as the internet created web development), displacement estimates may overstate net effects.

4. **The political feasibility gap:** The Nobel Lecture's "political losers" mechanism implies that the very firms benefiting from automation-focused AI will resist pro-worker alternatives. The policy agenda requires political power that displaced workers may lack.
5. **The speed problem:** The Industrial Revolution's harms lasted 60–80 years before reforms. If AI transforms economies in years rather than decades, the institutional response must be unprecedented in speed.
6. **Open-source disruption:** Acemoglu's framework focuses on corporate AI. Open-source models may change innovation direction dynamics in ways not yet modeled.

11 Conclusion

Acemoglu's research program provides the most rigorous framework available for thinking about AI's economic consequences. The message is neither optimistic nor pessimistic, but *agentive*: the trajectory is a choice. The current choice—excessive automation, so-so technologies, manipulative platforms, tax biases favoring capital—is not inevitable. It reflects specific market failures, institutional weaknesses, and ideological commitments that can be corrected.

The Nobel Lecture places AI within a millennia-long story of how institutions mediate the relationship between technology and prosperity. The same technology—the power loom, the steam engine, artificial intelligence—produces vastly different outcomes depending on institutional context. We are at a critical juncture. Whether AI becomes a tool for broad-based prosperity or a mechanism for concentrating power depends on choices being made now, by policymakers, technologists, and citizens. Acemoglu's work provides both the diagnostic framework and the policy roadmap. The question is whether we will use them.

Papers Analyzed

[1] Acemoglu, D. (2024). "The Simple Macroeconomics of AI." *Economic Policy*, 39(120), 3–46.
[2] Acemoglu, D. (2022). "Harms of AI." *Oxford Handbook of AI Governance.*
[3] Acemoglu, D., Autor, D., & Johnson, S. (2026). "Building Pro-Worker AI." Brookings.
[4] Acemoglu, D. & Restrepo, P. (2019). "The Wrong Kind of AI?" *Cambridge J. Regions, Econ. & Society.*
[5] Acemoglu, D. & Restrepo, P. (2019). "AI, Automation, and Work." *The Economics of AI.*
[6] Acemoglu, D., Ozdaglar, A., & Siderius, J. (2025). "AI and Social Media."
[7] Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). "AI and Jobs." *J. Labor Econ.*
[8] Acemoglu, D. & Restrepo, P. (2019). "Automation and New Tasks." *J. Econ. Perspectives.*

[9] Acemoglu, D. & Restrepo, P. (2022). "Tasks, Automation, and US Wage Inequality." *Econometrica*.

[10] Acemoglu, D. & Restrepo, P. (2025). "Automation and Rent Dissipation." *QJE*.

[11] Acemoglu, D. & Loebbing, J. (2026). "Automation and Polarization." *JPE*.

[12] Acemoglu, D. & Lensman, T. (2024). "Regulating Transformative Technologies." *AER: Insights*.

[13] Acemoglu, D. & Johnson, S. (2024). "Learning from Ricardo and Thompson." *Annual Rev. Econ*.

[14] Acemoglu, D., Kong, F., & Restrepo, P. (2025). "Tasks at Work." *Handbook of Labor Econ.*, v.6.

[15] Acemoglu, D. (2023). "Distorted Innovation." *AER P&P*.

[16] Acemoglu, D., Manera, A., & Restrepo, P. (2020). "Does the US Tax Code Favor Automation?" *BPEA*.

[17] Acemoglu, D. & Restrepo, P. (2020). "Robots and Jobs: US Labor Markets." *JPE*.

[18] Acemoglu, D. & Restrepo, P. (2021). "Demographics and Automation." *REStud*.

[19] Acemoglu, D. & Restrepo, P. (2024). "A Task-Based Approach to Inequality." *Oxford Open Econ*.

[20] Acemoglu, D. et al. (2025). "When Big Data Enables Behavioral Manipulation." *AER: Insights*.

[21] Acemoglu, D. & Restrepo, P. (2020). "Unpacking Skill Bias." *AEA P&P*.

[22] Acemoglu, D. & Restrepo, P. (2018). "Low-Skill and High-Skill Automation." *J. Human Capital*.

[23] Acemoglu, D. et al. (2025). "Online Business Models, Digital Ads, and User Welfare." *JPE*.

[24] Acemoglu, D., Ozdaglar, A., & Siderius, J. (2024). "Online Misinformation." *REStud*.

[25] Acemoglu, D. et al. (2022). "Too Much Data." *AEJ: Micro*.

[26] Acemoglu, D. (2025). "Institutions, Technology and Prosperity: Nobel Lecture." *AER*.

[27] Acemoglu, D. et al. (2024). "Automation and the Workforce: Firm-Level View." *NBER/U. Chicago*.

[28] Acemoglu, D., Lelarge, C., & Restrepo, P. (2020). "Competing with Robots: France." *AEA P&P*.

[29] Acemoglu, D. et al. (2023). "Advanced Technology Adoption: Selection or Causal?" *AEA P&P*.

[30] Acemoglu, D., Koster, H., & Ozgen, C. (2023). "Robots and Workers: Netherlands." *J. Labor Econ*.