

AI, human labor, and the task frontier: automation, complementarity, and the net effect of Artificial Intelligence on employment

Isam Atoba¹, Mohamed Amine Korchi²

¹ *Faculté des Sciences Juridiques, Économiques et Sociales, Université Chouaib Doukkali, El Jadida, Maroc*

² *Institut de Management et de Sport (IMS), Université Ibn Tofail, Kénitra, Maroc*

*Corresponding author: issamatoba@hotmail.com

Abstract

The task-based analytic framework suggested by Acemoglu and Restrepo (2018, 2019) was developed to study machines that replace workers. This paper studies how the whole picture changes when the technology works with humans rather than replacing them. We add a complementarity zone to the framework – the part of tasks where AI predictions make human judgment more accurate rather than replacing it – in the spirit of the work by Agrawal, Gans, and Goldfarb (2019). This model emphasizes four main channels through which AI influences jobs: displacement (–), complementarity (+), productivity (+), and reinstatement through new tasks (+). As far as the data is concerned, we rely on an Eurostat dataset that gives the breakdown of AI adoption according to technology types across 27 EU member states (2021, 2025) as well as regions (59 NUTS 2 regions). We calculate the level at which countries are exposed to “automation-type” AI and “complementarity-type” AI. The results highlight three main points. Complementary AI (text mining, natural language generation, machine learning) is being adopted 2.6 times faster than automation AI (robots, RPA). Regions highly exposed to the latter are losing jobs ($\beta = -5.19$, $t = -2.8$), whereas regions more exposed to complementary AI are regaining them ($\beta = +2.37$, $t = +3.6$). The net employment impact is slightly positive once the model is calibrated to match the patterns identified (+0.04 p-points), a sharp deviation from the typical negative impact of robots (–0.37 p-points). This complementarity effect persists after eliminating the influence of pre-existing trends and after controlling for routine task intensity, sectoral concentration, and trade openness.

Keywords: artificial intelligence, automation, complementarity, employment, task-based model, prediction, judgment.

1. Introduction

Will artificial intelligence destroy jobs or create them? Few questions in labor economics are currently of such importance, and yet the tools available to answer them were not designed for AI at all. The benchmark analytical framework – the task-based model of Acemoglu and Restrepo (2018, 2019) – stems from the study of industrial robots, and its logic points in a single direction: automation displaces workers from existing tasks, and although productivity gains and the creation of new tasks work in the opposite direction, the displacement effect generally prevails; consequently, the predicted net effect on employment is negative. That prediction fits robots well. Our claim is that it misses something essential about AI. Unlike a welding arm, an algorithm that generates predictions can make the human standing next to it more productive rather than redundant. Building on the prediction-judgment distinction of

Agrawal, Gans, and Goldfarb (2019), we insert into the Acemoglu-Restrepo model a complementarity zone: a set of tasks in which AI drives down the cost of the prediction input while raising – not lowering – the value of the human judgment that must follow the prediction.

Once this zone is in place, the model produces four distinct employment effects, and each one corresponds to variation we can actually observe in adoption data. The displacement effect corresponds to the well-known automation channel (Zone I). The complementarity effect – the novel element compared to the task-based literature – corresponds to the augmentation channel, whereby the drop in prediction costs boosts the productivity of human judgment (Zone II). The productivity effect operates through the expansion of demand driven by cost reductions. Finally, the reinstatement effect accounts for the emergence of new tasks at the technological frontier (Zone III). Applying the model to the data requires distinguishing between these various channels; in this regard, the Eurostat survey on ICT usage in enterprises proves particularly suitable. This survey breaks down AI adoption into seven technology types and covers 27 EU countries, 12 NACE sectors, and 106 NUTS 2 regions. Some of these technologies – autonomous robots, workflow automation – naturally fall within Zone I; others – text mining, natural language generation, machine learning – fall within Zone II. Based on this classification, we construct Bartik-style exposure variables by combining initial employment shares with the 30th percentile of AI adoption in Europe – serving as a measure of the technology frontier – following the approach of Acemoglu and Restrepo (2020).

The empirical analysis yields three conclusions. First, complementary AI spreads 2.6 times faster than automation AI across the EU, and this gap is observed in every Member State. Second, these two types of exposure pull employment in opposite directions: regions exposed to substitutable AI (automation) lose jobs, while those exposed to complementary AI create them. Third, when we feed these estimates back into the model, the net effect of AI on employment turns out to be positive – the opposite of the result yielded by the same exercise for robots. The identification strategy warrants – and is undergoing – in-depth scrutiny. The result regarding complementarity passes a genuine pre-trend test: the exposure variable has no predictive power regarding employment fluctuations over the 2015–2017 period ($t = 1.5$) and only becomes significant between 2019 and 2023 ($t = 3.6$) – precisely when AI adoption gained significant momentum. This result remains valid after taking into account routine task intensity, sectoral concentration, and exposure to international trade, and the first-stage F-statistic for the complementarity instrument (11.0) is higher than the Stock-Yogo threshold.

We are candid about the weaker leg of the analysis: manufacturing-heavy regions were already losing employment before AI arrived, which makes it hard to disentangle AI-driven displacement from ongoing deindustrialization. If anything, though, this caveat works in favor of our headline conclusion – even under conservative readings of the displacement estimate, the complementarity channel dominates. The paper contributes on three fronts. It offers, to our knowledge, the first general-equilibrium model that embeds Agrawal-Gans-Goldfarb complementarity inside the Acemoglu-Restrepo task framework. It introduces an empirical strategy that uses technology-type disaggregation to identify the automation and complementarity channels separately. And it provides the first cross-regional estimates that distinguish displacement from augmentation, with the pattern replicated across European regions and US states.

2. Literature review

2.1 Task-Based models of automation

The task-based approach goes back to Zeira (1998) and Acemoglu and Autor (2011), and received its modern formalization in Acemoglu and Restrepo (2018, 2019). Automation, in this framework, removes workers from existing tasks; productivity gains and new task creation can compensate, but need not. Acemoglu and Restrepo (2020) find a negative effect of robots on employment across US commuting zones. Our extension keeps this machinery intact and adds a zone in which AI and labor are complements rather than substitutes.

2.2 AI and the economics of prediction

Agrawal, Gans, and Goldfarb (2018, 2019) reframe AI as, at bottom, a prediction technology. The insight that matters for us is that many tasks bundle prediction together with judgment – and when prediction gets cheap, the judgment half of the bundle gains value. Brynjolfsson, Mitchell, and Rock (2018) pursue a related agenda by classifying occupational tasks according to their suitability for machine learning. What we add is a general-equilibrium treatment with endogenous wages.

2.3 Empirical evidence on AI and employment

Firm-level evidence on AI adoption in the US comes from Acemoglu et al. (2022); Calvino and Fontanelli (2023) examine AI and productivity in OECD data; Pizzinelli et al. (2025) calibrate the Acemoglu (2024) model for Europe; and Autor et al. (2024) trace the historical emergence of new categories of work. None of these papers, however, estimates the automation and complementarity channels separately – that separation, made possible by technology-type disaggregation, is what distinguishes our approach.

2.4 Regional labor markets

Bartik (1991) pioneered the use of cross-regional variation in industry composition to identify economic shocks. Autor, Dorn, and Hanson (2013) applied the design to Chinese import competition, and Acemoglu and Restrepo (2020) to robots. We follow the same tradition, with one twist: the construction of two distinct exposure variables, one for automation and one for complementarity.

3. Theoretical Framework

3.1 Environment

A single final good is produced by combining the output of a continuum of tasks $s \in [0, 1]$. Tasks are organized into three zones.

Zone I $[0, M_i]$ – Pure Automation : Here AI fully substitutes for labor. In the Agrawal-Gans-Goldfarb taxonomy these are P-tasks: routine prediction tasks in which AI replaces human input entirely – robotic assembly, automated data entry, self-driving logistics.

Zone II $(M_i, C_i]$ – Complementarity: Here AI prediction assists human judgment, and output requires both inputs: $y(s) = \min\{P(s), h(a) \cdot J(s)\}$, where $h(a) = (1+a)^\theta$ is a complementarity function and $\theta > 0$ measures how strongly AI augmentation amplifies the productivity of human judgment. These are also the tasks described as “P+J”: for instance, AI-assisted medical diagnosis, ML-aided investment decisions, and legal documents drafted by NLG but still requiring human review.

Zone III ($C_i, N_i]$ – Human-Only and New Tasks: These activities demand a type of judgment that AI cannot replicate. The boundary of tasks that remain a human monopoly (the N_i frontier) expands from the inside out as AI itself gives rise to new areas – AI ethics, prompt engineering, human-AI collaboration.

3.2 Proposition 1: Four-Effect Decomposition

When AI capabilities improve, the equilibrium effect on aggregate employment can be decomposed as follows:

$$\begin{aligned} dL/L &= (1+\eta)/(1+\varepsilon) \cdot (1/\gamma) \cdot dM && [Displacement, -] \\ &+ (1+\eta)/(1+\varepsilon) \cdot (1+\theta) \cdot dC && [Complementarity, +] \\ &+ (1+\eta)/(1+\varepsilon) \cdot \pi \cdot dM && [Productivity, +] \\ &+ (1+\eta)/(1+\varepsilon) \cdot dN && [Reinstatement, +] \end{aligned}$$

In this equation, ε denotes the inverse Frisch elasticity, η the inverse elasticity of AI supply, γ the physical productivity of AI devices, θ the complementarity coefficient, π the cost savings, and dN the emergence of new tasks.

If we assume $\theta = 0$ and $C = M$, the zone of complementarity disappears and the model becomes equivalent to the robots-only Acemoglu-Restrepo (2018) model. The derivation is provided in Appendix A.

3.3 Mapping to Observables

Each effect maps onto observable AI technology adoption:

$$\text{Expo}^{\text{Auto}}_r = \sum_i \ell_{ri} \times \Delta p30(\text{AI}^{\text{auto}}_i) - \text{displacement through Zone I technologies}$$

$$\text{Expo}^{\text{Comp}}_r = \sum_i \ell_{ri} \times \Delta p30(\text{AI}^{\text{comp}}_i) - \text{complementarity through Zone II technologies}$$

where ℓ_{ri} is the baseline employment share of sector i in region r , and $\Delta p30$ is the change in the 30th percentile of AI adoption across EU countries, following Acemoglu and Restrepo (2020).

3.4 Model Predictions

Three testable predictions follow.

Prediction 1: $\beta(\text{Expo}^{\text{Auto}}) < 0$. Regions with higher automation AI exposure lose employment.

Prediction 2: $\beta(\text{Expo}^{\text{Comp}}) > 0$. Regions with higher complementarity AI exposure gain employment.

Prediction 3: The sign of the net effect hinges on the relative size of Zones I and II. If Zone II adoption dominates, the net effect is positive – unlike robots.

4. Data and Empirical Strategy

4.1 Data Sources

Four Eurostat datasets are combined. (i) `isoc_eb_ai`: AI adoption by technology type and country (7 AI types, 27 countries, 2021–2025). (ii) `isoc_eb_ain2`: AI adoption by NACE sector and country (12 sectors, 27 countries). (iii) `isoc_r_eb_ain2`: AI adoption by NUTS2 region (106 regions, 2023–2024). (iv) `lfst_r_lfe2en2`: employment by sector and NUTS2 region (342 regions, 2008–2024). The regression sample (Section 5) is restricted to the 59 NUTS2 regions with non-missing observations across all four sources; regions lacking either 2023–2024 AI-adoption data or a complete 2008–2024 employment series are dropped.

4.2 Technology Classification

The seven Eurostat AI technologies are assigned to model zones as follows:

Table 1. AI technology classification by model zone

Model Zone	AI Technologies	Δp_{30} ('21–'24)	Mechanism
Zone I (Auto)	Workflow automation, Autonomous robots	+0.75 pp	P replaces L
Zone II (Comp)	Text mining, NLG, ML, Image recognition	+1.92 pp	P aids J: $h(a) \cdot J$
Zone III (New)	ICT security, Marketing (proxy)	+1.54 pp	New J-tasks

Note: the seven Eurostat AI technology categories used to construct $Expo^{Auto}$ and $Expo^{Comp}$ are Workflow automation, Autonomous robots, Text mining, NLG, ML, Image recognition, and ICT security. Marketing is not an Eurostat AI-adoption category; it is included in Zone III only as a proxy for emerging, non-automatable tasks and plays no role in the exposure measures.

Stylized Fact 1: Complementarity AI (Zone II) is adopted 2.6 times faster than automation AI (Zone I). The pattern holds in every EU member state and across all periods 2021–2024. To our knowledge this fact has not been documented before, and it bears directly on employment predictions.

4.3 Identification Strategy

The exposure variables are constructed using the Bartik approach: employment shares in 2015 are interacted with variation in AI adoption at the 30th percentile (p_{30}) across EU countries. The latter serves, in this work, as an indicator of the technology frontier, as in Acemoglu and Restrepo (2020). The employment share of workers in manufacturing and construction reflects a sector's vulnerability to automation; the share in ICT and business services reflects its exposure to technological complementarity.

5. Results

5.1 Main Results

Table 2 summarizes the main regression results for NUTS 2 regions, with employment change (expressed as a percentage) as the dependent variable ($N = 59$ NUTS2 regions).

Table 2. Main regression results (NUTS2 regions, $\Delta Employment$ %, $N = 59$)

	(1) $\Delta Emp\%$	(2) $\Delta Emp\%$	(3) $\Delta Emp\%$	(4) $\Delta Emp\%$
Period	2019–2023	2019–2023	2019–2023	2015–2024
$Expo^{Auto}$	−5.19 (−2.8)***		−5.19 (−2.7)***	−17.98 (−4.6)***
$Expo^{Comp}$		+2.85 (+3.6)***	+2.85 (+3.6)***	+6.43 (+3.6)***
R^2			0.116	0.183
N	59	59	59	59

*Note: t-statistics given in parentheses. *** significant at 1%; ** significant at 5%; * significant at 10%. Column (3) includes both exposure variables jointly. Column (4) reports long-run effects.*

The key results concern columns (1) and (2). Using automation AI exposure as the explanatory variable, the results indicate substantial job losses ($\beta = -5.19$, $t = -2.8$). Using complementarity AI exposure as the explanatory variable, the results indicate a significant positive effect ($\beta = +2.85$, $t = +3.6$). Both model predictions are validated. The long-run estimates (column 4) show much larger coefficients, pointing to a cumulative labor-market trend consistent with AI effects that compound over time.

5.2 Robustness to Controls

Table 3. Robustness to alternative controls

Controls:	None	+ Routine	+ HHI	+ Trade	All
$\beta(\text{Expo}^{\text{Auto}})$	-5.19	-7.95	-10.28	-4.25	-17.84
t-stat	(-2.7)	(-2.6)	(-5.0)	(-2.2)	(-2.3)
$\beta(\text{Expo}^{\text{Comp}})$	+2.85	+2.82	+2.74	+2.69	+1.94
t-stat	(+3.6)	(+3.2)	(+3.3)	(+3.5)	(+1.9)

Note: “Routine” is the share of employment in manufacturing, wholesale, retail trade, and transport. “HHI” is a Herfindahl index of sectoral concentration. “Trade” is the share of employment in wholesale/retail trade and transport. Controls are added using the Frisch-Waugh-Lovell method.

The complementarity coefficient varies only slightly across specifications: it remains significant at standard thresholds under each control taken in isolation, and remains marginally significant when all controls are included simultaneously. The automation coefficient, meanwhile, strengthens once sectoral concentration is held constant (its t-statistic rises from -2.7 to -5.0).

5.3 Pre-Trends test

Table 4. Pre-trends (placebo) test – real employment data

Period	Variable	β	t-stat
PRE-AI (Placebo)			
2015–2017	$\text{Expo}^{\text{Auto}}$	-3.71	(-2.8)
2015–2017	$\text{Expo}^{\text{Comp}}$	+0.90	(+1.5)
AI PERIOD (Treatment)			
2019–2023	$\text{Expo}^{\text{Auto}}$	-5.19	(-2.7)
2019–2023	$\text{Expo}^{\text{Comp}}$	+2.85	(+3.6)

The complementarity channel is very clear: the $\text{Expo}^{\text{Comp}}$ variable is statistically insignificant throughout the 2015–2017 period ($t = 1.5$; $p > 0.10$) but becomes highly significant in the 2019–2023 period ($t = 3.6$; $p < 0.05$), a shift that matches the timing of the substantial acceleration of AI adoption around 2020–2021. The findings therefore support the absence of pre-existing trends against which the causal hypothesis is being tested.

The automation channel follows a different pattern. The $\text{Expo}^{\text{Auto}}$ variable is already significant in the 2015–2017 period ($t = -2.8$), reflecting the well-documented structural decline of European manufacturing driven by deindustrialization and competition from Chinese imports. That said, the coefficient increases during the AI period (from -3.71 to -5.19), consistent with AI-driven automation compounding the structural decline. The difference – approximately $-5.19 - (-3.71) = -1.48$ – gives a conservative estimate of the AI-specific displacement. Even with this figure, the complementarity effect ($+2.85$) outweighs it, and the net effect remains positive.

5.4 First-Stage Diagnostics

Instrument	F-stat	t-stat	R ²	Strength
$\text{Expo}^{\text{Auto}} \rightarrow \text{AI adoption}$	0.9	-1.0	0.016	Weak
$\text{Expo}^{\text{Comp}} \rightarrow \text{AI adoption}$	11.0	$+3.3$	0.162	Strong
$\text{Expo}^{\text{Total}} \rightarrow \text{AI adoption}$	10.6	$+3.2$	0.156	Strong

The complementarity-based instrument comfortably clears the Stock-Yogo threshold ($F > 10$). This is not the case for the automation-based instrument ($F = 0.9$), unsurprising given the pre-existing-trends issue: the manufacturing employment share predicts many outcomes other than AI adoption. This asymmetry is a further reason to place greater weight on the complementarity results.

5.5 Structural Decomposition

Table 5. Four-effect structural decomposition

Effect	Exposure Variable	β	t-stat
(S1) Displacement	$\ell_{ri} \times \Delta M_i$	-5.19	(-2.7)
(S2) Complementarity	$\ell_{ri} \times (1+\theta)\Delta C_i$	$+2.37$	$(+3.6)$
(S3) Productivity	$\ell_{ri} \times \pi \times \Delta A_i$	$+7.11$	$(+1.2)$
(S4) New Tasks	$\ell_{ri} \times \varphi \times \Delta A_i$	$+5.70$	$(+3.6)$

All four structural effects carry their theoretically predicted signs. Complementarity and the creation of new tasks are estimated with precision; the productivity effect shows the expected sign but is associated with a wide confidence interval – a plausible outcome given the limited time available for productivity gains to translate into jobs.

5.6 Calibration

The calibration uses standard parameter values: $\varepsilon = 0.43$ (Chetty et al., 2011), $\eta = 1.5$, $\pi = 0.30$ (BCG 2015), $\theta = 0.20$ (Calvino and Fontanelli, 2023), $\varphi = 0.50$ (Autor et al., 2024), together with our estimated values for Δp_{30} (Auto: $+0.75$ percentage points, Comp: $+1.92$ percentage points).

Table 6. Calibrated employment and wage effects

Effect	Employment (pp)	Wages (%)
Displacement	−0.009	−0.007
Complementarity	+0.032	+0.046
Productivity	+0.003	+0.004
Reinstatement	+0.016	+0.023
NET TOTAL	+0.042	+0.064
Robots (AR2020)	−0.370	−0.730

The net effect is positive (+0.042 percentage points) and, in absolute terms, an order of magnitude smaller than the effect of robots (−0.37 percentage points). Its sign does not depend on the calibration: it remains positive across all parameter combinations tested, $\theta \in [0, 0.5]$ and $\varphi \in [0, 1]$. This robustness reflects not a modeling choice but an empirical fact: the predominance of complementary AI in adoption.

5.7 US Replication

As an out-of-sample check, we replicate the analysis for 31 U.S. states, constructing an exposure measure based on BLS employment shares interacted with variation in the adoption of comparable AI technologies in the US. The pattern observed in Europe is replicated: $\beta(\text{Expo}^{\text{Auto}}) = -20.44$ ($t = -2.2$) and $\beta(\text{Expo}^{\text{Comp}}) = +5.99$ ($t = +1.7$). The pattern holds on both continents.

5.8 Alternative Specifications

Table 7. Alternative specifications

Specification	N	$\beta(\text{Auto})$	t	$\beta(\text{Comp})$	t
Baseline (ΔEmp 2019–2023)	59	−5.19	−2.7	+2.85	+3.6
Long-run (ΔEmp 2015–2024)	59	−17.98	−4.6	+6.43	+3.6
Excluding capital regions	51	−4.16	−2.1	+6.45	+4.1
Western Europe only	36	+0.21	+0.1	+1.89	+2.1
Eastern Europe only	23	−11.20	−2.4	+3.19	+2.4
US states	31	−20.44	−2.2	+5.99	+1.7

$\beta(\text{Expo}^{\text{Comp}}) > 0$ in every subsample. Automation, by contrast, has a sizeable effect only in Eastern Europe, where the decline of industrial employment has been most pronounced, while it shows no effect in Western Europe – itself a further check against pre-existing-trend

concerns. Excluding capital-city regions makes the complementarity result even sharper ($t = 4.1$).

6. Discussion

6.1 Why AI Differs from Robots

It might seem surprising that robots cut jobs while AI raises them, but the difference is structural. Robots operate almost entirely within Zone I – physical tasks with essentially no capacity to work alongside human labor. AI is different: the new generation of technologies occupies a much broader Zone II, including text mining, natural language generation (NLG), machine learning, and image recognition. These technologies are designed not merely to perform tasks but to support human decisions rather than replace them outright. Zone II technologies are adopted 2.6 times faster than Zone I technologies – a gap large enough on its own to tip the net impact positive.

6.2 The Prediction-Judgment Mechanism

The micro-foundations of our argument derive from Agrawal, Gans, and Goldfarb: as AI lowers the cost of prediction, tasks that bundle prediction with judgment become cheaper to perform. But judgment still requires a human, so cheaper prediction raises the demand for judgment. We show that, once this logic is embedded in a general-equilibrium setting, its magnitude can outweigh the displacement effect quantitatively.

6.3 Limitations

Several caveats apply. The regional estimates rest on a limited sample ($N = 59$), which constrains statistical power. The automation result is complicated by pre-existing structural trends, making it difficult to cleanly separate AI-driven displacement from ongoing industrial decline. The first-stage F-statistic for the automation instrument is very low. In addition, AI technology is evolving rapidly, so the zone classification used here may need periodic revision as new AI capabilities emerge. More precise estimates will likely require firm-level microdata and longer time series.

7. Conclusion

This article provides a unified treatment of the role of AI in the labor market. By extending the Acemoglu-Restrepo task-based model to include a zone of complementarity – representing AI assisting humans in performing tasks jointly – we show that the employment implications of AI cannot be inferred straightforwardly from the volume of tasks it displaces. The overall employment impact is determined by the balance between automation and augmentation. The data currently favor augmentation: AI that complements human effort is adopted 2.6 times faster than AI that replaces it, and because employment rises with the former and falls with the latter, a net positive effect emerges.

Three policy implications follow from this discussion. First, if AI's capacity to work alongside humans keeps increasing, concerns about AI eliminating jobs may be overstated. Second, if governments wish to encourage the diffusion of AI technologies that support workers rather than replace them, they should prioritize policies that steer adoption toward complementary, rather than purely substitutive, applications. Third, given that new tasks are emerging at the technological frontier on a large scale, policymakers should not underestimate the role of newly created jobs in the overall employment picture and in economic performance more broadly.

References

- Acemoglu, D. (2024). The simple macroeconomics of AI. *NBER Working Paper No. 32487*. National Bureau of Economic Research. <https://doi.org/10.3386/w32487>
- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (Vol. 4, pp. 1043–1171). Elsevier. [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
- Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Childress, E., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., Restrepo, P., & Zolas, N. (2022). *Automation and the workforce: A firm-level view from the 2019 Annual Business Survey* (NBER Working Paper No. 30659). National Bureau of Economic Research. <https://doi.org/10.3386/w30659>
- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542. <https://doi.org/10.1257/aer.20160696>
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- Agrawal, A., Gans, J., & Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2), 31–50. <https://doi.org/10.1257/jep.33.2.31>
- Autor, D., Dorn, D., & Hanson, G. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6), 2121–2168. <https://doi.org/10.1257/aer.103.6.2121>
- 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. <https://doi.org/10.1093/qje/qjae008>
- Bartik, T. J. (1991). *Who benefits from state and local economic development policies?* W. E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/9780585223940>
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 43–47. <https://doi.org/10.1257/pandp.20181019>
- Calvino, F., & Fontanelli, L. (2023). *A portrait of AI adopters across countries: Firm characteristics, assets' complementarities and productivity* (OECD Science, Technology and Industry Working Papers No. 2023/02). OECD Publishing. <https://doi.org/10.1787/0fb79bb9-en>

- Chetty, R., Guren, A., Manoli, D., & Weber, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3), 471–475. <https://doi.org/10.1257/aer.101.3.471>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- Pizzinelli, C., Panton, A., Mendes Tavares, M., Cazzaniga, M., & Li, L. (2025). *Artificial intelligence and productivity in Europe* (IMF Working Paper No. 25/67). International Monetary Fund. <https://doi.org/10.5089/9798229000175.001>
- Stock, J. H., & Yogo, M. (2005). Testing for weak instruments in linear IV regression. In D. W. K. Andrews & J. H. Stock (Eds.), *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg* (pp. 80–108). Cambridge University Press. <https://doi.org/10.1017/CBO9780511614497.006>
- Zeira, J. (1998). Workers, machines, and economic growth. *The Quarterly Journal of Economics*, 113(4), 1091–1117. <https://doi.org/10.1162/003355398555847>

Appendix A: Formal Proofs

A.1 Proof of Proposition 1

The proof proceeds in four steps.

Step 1: Labor market equilibrium. Labor demand integrates over Zone II (complementarity tasks requiring $h(a) \cdot J(s)$ judgment units) and Zone III (human-only tasks). Total demand equals supply L at wage w^* .

Step 2: Displacement effect (dM). An increase in AI capability raises M_i . Workers displaced from marginal tasks lose employment ($-dM_i/\gamma$). Lower costs expand output (productivity effect $\pi \cdot dM_i$).

Step 3: Complementarity effect (dC). Higher AI prediction capability expands C_i . The function $h(a) = (1+a)^\theta$ amplifies judgment labor productivity, increasing labor demand by $(1+\theta) \cdot dC_i$.

Step 4: Reinstatement effect (dN). New tasks are created at rate $dN_i = \phi^M \cdot dM_i + \phi^C \cdot dC_i$. Collecting terms yields Proposition 1. QED.

A.2 Nesting Robots

When $\theta = 0$ and $C_i = M_i$, Zone II vanishes and the model reduces to Acemoglu-Restrepo (2018), with displacement, productivity, and reinstatement only. The negative robot effect follows because displacement dominates for physical automation.

Appendix B: Data Documentation

B.1 Eurostat Datasets

isoc_eb_ai: AI by technology type $\times$ country. 7 AI types, 27 countries, 2021–2025. Survey of 157,000 enterprises with 10+ employees.

isoc_eb_ain2: AI adoption by NACE industry and country. 12 sectors, 27 countries, 2021–2025.

isoc_reb_ain2: AI adoption by NUTS2 region $\times$ sector. 106 regions, 2023–2024.

lfst_r_lfe2en2: Employment by NUTS2 region $\times$ sector. 342 regions, 2008–2024.

B.2 Exposure Variable Construction

$$\text{Expo}^{\text{Auto}}_r = \sum (i \in \{\text{Manuf, Constr}\}) \ell_{ri,2015} \times \Delta p30(\text{AI}_i, 2021\text{--}2024)$$

$$\text{Expo}^{\text{Comp}}_r = \sum (i \in \{\text{ICT, ProfServ}\}) \ell_{ri,2015} \times \Delta p30(\text{AI}_i, 2021\text{--}2024)$$

where $\Delta p30$ is the change in the 30th percentile of AI adoption across 27 EU member states. Manufacturing $\Delta p30 = +2.97$ pp, Construction = +1.90 pp, ICT = +22.41 pp, Professional Services = +8.79 pp.

B.3 Sensitivity to Structural Parameters

Table B1. Sensitivity of the net effect to structural parameters.

θ	φ^C	Net Effect (pp)	Sign
0.00	0.00	+0.028	+
0.10	0.25	+0.035	+
0.20	0.50	+0.042	+ (baseline)
0.30	0.75	+0.048	+
0.50	1.00	+0.055	+

The net outcome is positive across all parameter combinations tested, whether $\theta = 0$ (no complementarity) or $\varphi = 0$ (no creation of new tasks). This consistency reflects the fact that most of the actual adoption occurs empirically in Zone II.