A Review of "Robots and Jobs: Evidence from US Labor Markets" by Acemoglu and Restrepo (2020)

Reviewer 2

January 14, 2026

v1

Disclaimer

This report was generated by Reviewer 2, an automated system that uses large language models to assess academic texts. It has been read and approved by a human editor on behalf of The Catalogue of Errors Ltd. The report's goal is to facilitate the discovery of knowledge by identifying errors in the existing literature. Comments can be made here. Any errors will be corrected in future revisions.

> I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.
>
> Plato, *The Apology of Socrates*, 21d

> To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.
>
> Karl Popper, 'Knowledge and the Shaping of Reality'

Overview

Citation: Acemoglu, Daron, and Pascual Restrepo. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*. Vol. 128, No. 6, pp. 2188–2244.

Abstract Summary: This study investigates the effects of industrial robots on US labor markets, showing theoretically that robots may reduce employment and wages and that their local impacts can be estimated using variation in exposure to robotics. The authors estimate robust negative effects of robots on employment and wages across commuting zones.

Key Methodology: Empirical analysis using a Bartik-style measure of exposure to robots, instrumented by analogous measures constructed from robot penetration trends in European countries (EURO5), applied to US commuting zones.

Research Question: What are the effects of industrial robots on US labor markets?

Summary

Is It Credible?

Acemoglu and Restrepo's analysis of the impact of industrial robots on US labor markets is a highly credible, methodologically rigorous contribution that fundamentally alters the economic understanding of automation. The article succeeds in its primary goal: distinguishing the "displacement effect" of robots from the general productivity benefits of capital deepening. By exploiting plausibly exogenous advances in robotics technology across European industries to instrument for US adoption, the authors provide compelling causal evidence that industrial robots reduced employment and wages in exposed local labor markets between 1990–2007. While the local reduced-form estimates are robust and persuasive, the article's subsequent attempt to quantify aggregate national effects relies on a structural calibration that contains mathematical inconsistencies and modeling choices that do not fully align with the empirical findings.

The strength of the article lies in its careful empirical identification. The authors construct a Bartik-style instrument using robot adoption in European countries (the "EURO5") to isolate technological supply shocks from US-specific demand shocks. This strategy effectively addresses endogeneity concerns, a claim bolstered by the demonstration that exposure to robots is uncorrelated with pre-1990 labor market trends. Crucially, the article empirically validates its theoretical core: unlike IT capital or general capital deepening—which are shown to have neutral or positive effects—robots specifically depress labor demand. This finding empirically grounds the theoretical distinction between "automation" (replacing tasks) and "augmentation" (making tasks more productive). The authors also transparently address the concentration of robot adoption in the automotive industry, demonstrating through robustness checks that the negative effects persist, albeit with less precision, even when this sector is isolated.

However, the credibility of the article strains somewhat when bridging the gap between these robust local estimates and the aggregate national quantification. The transition relies on a structural model calibrated to match the local estimates, but there appears to be a calculation error or inconsistency in this calibration; the reported parameters for labor supply elasticity do not mathematically align with the structural equations and inputs provided. Furthermore, the model relies on a "productivity effect" to offset displacement, yet the empirical analysis fails to detect any significant employment gains in non-exposed sectors locally. The article acknowledges this absence as "surprising," but it proceeds to model aggregate effects using assumptions of market clearing and labor supply elasticity that may not fully capture the frictions evident in the reduced-form results. Specifically, the model assumes labor supply adjusts along the intensive margin (hours), whereas the empirical results show the adjustment is primarily driven by labor force withdrawal (the extensive margin).

Ultimately, the reality revealed by this article is that the arrival of industrial robots represented a distinct negative shock to labor demand for blue-collar workers in the US, a shock that was not compensated by immediate growth in other local sectors. The estimated loss of approximately 400,000 jobs is a model-derived figure subject to the calibration issues noted, but the direction and mechanism of the effect are established with high confidence. The article proves that automation is not merely another form of capital accumulation; it possesses a unique capacity to displace labor that, at least in the medium term, outweighs its productivity spillovers.

The Bottom Line

This article provides convincing evidence that industrial robots negatively affected US employment and wages between 1990–2007, distinct from the effects of general capital or IT. The identification strategy, using European adoption trends, is robust and successfully isolates the displacement caused by automation. While the local estimates are highly credible, the aggregate national quantification should be interpreted with caution due to inconsistencies in the model calibration and a reliance on theoretical assumptions about labor supply that do not perfectly match the observed data. Nevertheless, the core finding—that robots displace rather than augment human labor in specific tasks—is solidly supported.

Specific Issues

Mathematical inconsistency in model calibration: There appears to be a calculation error or a discrepancy in the reported parameters used for the aggregate quantification. The article reports calibrated parameters for the inverse wage elasticity of labor supply (ϵ) and the preference parameter (ψ) that do not align with the solutions to the structural equations provided in the appendix when using the stated inputs. Solving the system with the reported values yields a significantly different elasticity. While this does not negate the qualitative finding of a negative aggregate effect, it introduces uncertainty regarding the precise magnitude of the claimed 400,000 job losses and the aggregate wage impacts (p. 2239, p. A-18).

Empirical absence of positive productivity spillovers: The theoretical framework posits that the negative displacement effect should be partially offset by a positive productivity effect, where automation lowers costs and increases labor demand in other sectors. However, the empirical analysis fails to find statistically significant employment gains in any non-automated industry or occupation locally. The authors acknowledge they do not estimate positive effects in other occupations, which contrasts with the theoretical reliance on these offsets for the general equilibrium calculation. This suggests that the friction preventing labor reallocation is stronger than the model assumes (p. 2190, p. 2233).

Subjective exclusion of key countries from IV instrument: The identification strategy excludes major robotics leaders like Germany and Japan from the instrumental variable construction. The authors justify excluding Germany because its adoption levels are "so far ahead" they might not be relevant for the US, and Japan due to data reclassification issues. While robustness checks including Germany show similar (though slightly smaller) negative coefficients, the exclusion of the most technologically advanced nation from a measure of "global technological frontier" is a subjective choice that maximizes the estimated negative impact (p. 2200, p. 2202, p. A-46).

Mismatch between modeled and observed labor supply adjustment: The structural model assumes labor supply adjusts based on a Frisch elasticity, implying changes in hours worked (intensive margin). However, the empirical results explicitly show that the employment decline is driven primarily by labor force withdrawal (extensive margin), with about three-quarters of the non-employed dropping out of the labor force entirely. The authors justify the high calibrated elasticity by citing macro literature, but the functional form of the model does not endogenously generate the participation drop observed in the data (p. 2219, pp. 2239–2240).

Concentration of variation in the automotive industry: The variation in the robot exposure instrument is heavily concentrated in the automotive sector, which explains 67 percent of the cross-commuting zone variation. While the authors provide robustness checks separating the automotive industry from others—finding negative effects in both—the precision of the estimates drops when the automotive sector is excluded. This heavy reliance on a single industry raises questions about whether the results are capturing a general phenomenon of automation or a specific shock to US auto manufacturing (p. 2226, p. 2227).

Discrepancy regarding high-skill complementarity: The task-based model predicts that automation should complement high-skill workers. However, the empirical results show no significant positive effect for workers with masters or doctoral degrees. The authors acknowledge this as surprising, suggesting that robots may not complement high-skill tasks in the same way other computer technologies do, or that spillovers were insufficient to boost demand. This finding challenges the universality of the "skill-biased technological change" narrative often applied to automation (p. 2233).

Sensitivity to baseline employment shares: The magnitude of the results is sensitive to the choice of baseline year for industry shares. Using 1990 shares instead of the baseline 1970 shares yields estimates that differ in magnitude, though they remain negative and significant. The authors prefer 1970 to avoid endogeneity, but the sensitivity indicates that the definition of "exposure" varies depending on how deep into history one anchors the industrial structure (p. A-46).

Minor data and presentation issues: There are several minor issues regarding data and presentation. The article notes that 30 percent of robots are unclassified in the IFR data and are allocated proportionally, which introduces measurement error (p. 2201). The calibration uses a tradable sector share (0.18) that differs slightly from the mean manufacturing share in the data (0.22) (p. 2239, p. A-31). The article claims "no systematic differences" between baseline and Borusyak et al. standard errors, yet the latter are notably larger, though significance is maintained (p. 2237, p. A-52). Additionally, Table 7 reports identical First-stage *F*-statistics for different instrument constructions, which is almost certainly a clerical error (p. 2232). Finally, the article relies on fixed 1990 commuting zone definitions and imputes early robot data for Denmark, which are standard but imperfect methodological choices (p. 2202, p. 2203).

Future Research

Endogenous extensive margin modeling: Future research should develop a structural model that explicitly incorporates the extensive margin of labor supply (participation decisions) rather than relying on a representative agent with a high Frisch elasticity. Given the empirical finding that robot exposure leads primarily to labor force withdrawal rather than unemployment or reduced hours, a model that endogenizes the decision to exit the workforce would provide a more accurate quantification of the aggregate welfare losses and the permanence of the displacement shock.

Investigating the missing productivity spillover: Research is needed to explain the empirical absence of the "productivity effect" at the local level. The current article assumes these spillovers exist but are perhaps diffused nationally or offset by local demand shocks. New work should utilize firm-level or more granular service sector data to trace the prices of non-tradable goods in robot-exposed zones. Determining whether the cost savings from automation are retained as rents, passed on as lower prices, or absorbed by transaction costs is crucial for understanding why the predicted employment expansion in non-automated sectors failed to materialize.

www.isitcredible.com