A Report on “Robots and Jobs: Evidence from US Labor Markets” by Acemoglu and Restrepo (2020)

Reviewer 2

February 07, 2026

v2

isitcredible.com

Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made here. Any errors in the report 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, D., and Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, Vol. 128, No. 6, pp. 2188–2244.

URL: https://doi.org/10.1086/705716

Abstract Summary: This study examines the effects of industrial robots on US labor markets using a theoretical model and empirical analysis across commuting zones. The findings show robust negative effects of robot exposure on both employment and wages, distinct from other forms of capital deepening.

Key Methodology: Theoretical model of automation and tasks; empirical analysis using a Bartik-style shift-share measure of exposure to robots, instrumented by European robot adoption trends (IV/2SLS) across US commuting zones (1990–2007).

Research Question: What are the equilibrium effects of industrial robots on employment and wages in local US labor markets?

Editor's Note

Version 2 of this report has benefited from Pascual Restrepo's feedback. In Version 1, Reviewer 2 claimed that there was a calibration error in Acemoglu and Restrepo's general equilibrium model—this was incorrect and has been removed from this version of the report. In addition, a human researcher working for The Catalogue of Errors Ltd has added a potential issue relating to the high implied Frisch elasticity of labor supply. The report has also been written by an improved model of Reviewer 2. The Catalogue of Errors Ltd is very grateful to Professor Restrepo for his feedback on Version 1.

Summary

Is It Credible?

Acemoglu and Restrepo present a landmark investigation into the equilibrium impact of industrial robots on US labor markets. The article advances the headline claim that, unlike general capital deepening or factor-augmenting technologies, automation via robots creates a "displacement effect" that reduces employment and wages. Specifically, the authors state that "one more robot per thousand workers reduces the aggregate employment-to-population ratio by about 0.2 percentage points and wages by about 0.42%" (p. 2188). This assertion challenges the conventional economic wisdom that technological improvements generally raise labor demand, positing instead that automation creates distinct winners and losers by directly replacing human labor in specific tasks.

The credibility of the local labor market effects is supported by a rigorous research design. The authors employ a Bartik-style instrumental variable strategy, isolating robot adoption driven by technological advances in Europe to predict exposure in US commuting zones (pp. 2200–2203). This approach effectively addresses concerns that US robot adoption might simply reflect domestic labor shortages or other local shocks. The analysis convincingly demonstrates that commuting zones with higher exposure to robots experienced relative declines in employment and wages between 1990 and 2007 (Table 2, p. 2214). Crucially, the article empirically distinguishes robots from other forms of capital; controlling for IT capital and overall capital deepening does not diminish the negative impact of robots, supporting the theoretical distinction between displacement and augmentation (Table 6, p. 2230).

However, the transition from these local estimates to the headline aggregate claims requires significant caution. The aggregate figures cited in the abstract—the 0.2 percentage point decline in employment and 0.42% decline in wages—are not direct econometric estimates but the output of a calibrated structural model (pp. 2238, A-17). The validity of these aggregate numbers depends entirely on the model's parameters. Notably, calibrating the model to match the empirical results requires an inverse of the wage elasticity of labor supply (ε) of 0.17, implying a labor supply elasticity of approximately 5.9 (p. 2239). This is substantially higher than standard microeconomic estimates, which typically hover below 1. This discrepancy suggests the model may be missing important adjustment channels, or that the reduced-form estimates capture effects beyond simple labor supply responses, such as severe friction in labor reallocation.

Furthermore, the robustness of the primary employment estimate shows some sensitivity to specification choices. The magnitude of the negative employment effect relies partly on controlling for the decline of "light manufacturing" industries; excluding this control reduces the estimated impact by approximately 30% (p. 2216, footnote 20). Additionally, the identification strategy is heavily dependent on the automotive industry, which accounts for 67% of the variation in robot exposure (p. 2226). While the authors show that results hold when excluding the auto sector (Table 5, p. 2226), the heavy reliance on a single industry for identifying variation raises questions about whether the results capture a general technological phenomenon or specific dynamics within US auto manufacturing. Finally, the negative employment effects appear to attenuate over time; when the analysis is extended to 2014, the impact shrinks considerably, suggesting that the labor market may eventually adjust to the shock (Table 7, p. 2232).

The Bottom Line

Acemoglu and Restrepo provide compelling evidence that the introduction of industrial robots caused significant, localized job displacement and wage stagnation in US manufacturing hubs between 1990 and 2007. However, the headline claim regarding national aggregate job losses is a model-based simulation rather than a direct measurement, and it relies on assumptions about labor supply that differ sharply from standard economic estimates. While the distinction between automation and other forms of capital is a vital contribution, the effects are heavily concentrated in the automotive sector and appear to weaken over longer time horizons, suggesting that the permanent displacement of labor may be less severe than the initial shock implies.

Potential Issues

High implied Frisch elasticity of labor supply: Acemoglu and Restrepo calibrate their structural model using the IV estimates from Table 7 to recover two key parameters: ε, the inverse of the wage elasticity of labor supply, and η, the inverse of the elasticity of robot supply. Their preferred estimates for 1990–2007 yield $\varepsilon = 0.17$, implying a Frisch elasticity of $\phi = 1/\varepsilon \approx 5.88$. They write that this estimate "is in line with the 'macro' Frisch elasticities that are consistent with the observed short-run movements in wages and employment (see table 1 in Chetty et al. 2011)" (pp. 2239–2240, A-18). Nonetheless, Chetty et al. report that micro-level estimates of the Frisch elasticity average 0.82, while macro-level Real Business Cycle models typically require an elasticity of 2.84 to match business cycle fluctuations. Chetty et al. conclude that "models that require a Frisch elasticity of aggregate hours above 1 are inconsistent with micro evidence" and recommend a preferred calibration of $\phi = 0.75$ (2011, pp. 4–5). The Frisch elasticity of 5.88 implied by Acemoglu and Restrepo's estimates is more than double the highest macro benchmark reported by Chetty et al. and nearly eight times their recommended value. One interpretation is that the IV coefficients may be biased in ways that inflate the implied elasticity. The ratio of the employment coefficient to the wage coefficient implies a reduced-form labor supply elasticity that is already above standard estimates before any structural model enters the picture. It could be that the wage estimates suffer from composition bias. If robots disproportionately displace lower-paid workers within demographic cells, the average wage of the remaining employed workers understates the true decline in wage offers. Exclusion restriction violations may also play a role in biasing the IV coefficients. In addition, the structural model may be too parsimonious. Even if the IV coefficients are unbiased, the model used to translate them into structural parameters contains only a handful of adjustment channels: labor supply preferences, local robot supply, trade in goods, a nontradable sector, and mobile capital. It omits migration, housing market adjustment, search frictions, firm entry and exit, local fiscal multipliers, and endogenous technology responses. Consistent with this concern, Acemoglu and Restrepo report that when the estimation window is extended to 1990–2014, the calibration yields $\varepsilon \approx \$\ 0.39$ and $\phi \approx 2.56$—a substantial shift that suggests the exercise may not be recovering a stable structural parameter (p. 2240, footnote 30). Alternatively, the discrepancy may reflect a genuine feature of local labor markets. It is possible that the employment response to a persistent local labor demand shock is substantially more elastic than the response to the individual-level variation typically studied in the micro literature on labor supply. Displacement from automation may differ from a marginal wage change: affected workers lose jobs entirely, and the resulting adjustment involves search frictions, discouragement, and local demand multipliers that compound the initial displacement. If so, the high implied Frisch elasticity is not a sign of misspecification but an empirical finding that challenges existing estimates of labor supply responsiveness.

Instrumental variable validity: The article's identification strategy may be compromised by its heavy reliance on a single industry. The instrument, which uses robot adoption trends in other advanced economies to predict adoption in US commuting zones, is intended to isolate an exogenous technology shock. However, its variation is overwhelmingly driven by the automotive industry. The authors acknowledge that "the share of employment in the automotive industry explains 67% of the cross-commuting zone variation in exposure to robots" and that this industry contributes the vast majority of the identifying variation in the main estimates (p. 2226, footnote 23). They concede that this means the "estimates may be sensitive to other shocks affecting local labor markets specializing in the automotive industry during this period" (p. 2226, footnote 23). Because the global automotive sector is subject to common shocks—such as shifts in consumer demand, trade policy, or non-robot technological advances—that could simultaneously affect robot adoption in Europe and employment in US auto-centric regions, the instrument may violate the exclusion restriction by capturing these unobserved factors. While the authors conduct robustness checks by separating the automotive industry from others (Table 5, p. 2226), the core identification remains tied to a single sector, which may limit the generalizability of the findings.

Nature of the headline aggregate findings: The article's most prominent quantitative claims are not direct econometric findings but are the output of a calibrated structural model. The abstract's claim that "One more robot per thousand workers reduces the employment-to-population ratio by 0.2 percentage points and wages by 0.42%" is a simulation result (p. 2188). The direct instrumental variable (IV) estimates are substantially larger. The authors are transparent about this, stating, "To explore these aggregate implications, we need to make further assumptions on cross-commuting zone spillovers (and this suggests greater caution in interpreting these aggregate estimates than the local effects...)" (p. 2238). However, the validity of this simulation rests on several parameters calibrated from sources that are not peer-reviewed academic literature, including popular media reports, a TV show's website, and a Boston Consulting Group report (p. 2239, footnote 29). Presenting these model-dependent calculations in the abstract, rather than the direct econometric estimates, risks overstating the certainty of the evidence for the article's aggregate claims.

Temporal stability of the employment effect: The negative employment effect weakens substantially over a longer time period, a finding that is not prominently discussed in the main text. The primary analysis focuses on the 1990–2007 period, for which the IV coefficient for the effect on the employment-to-population ratio is -0.388. When the analysis is extended to 1990–2014, this effect shrinks by over a third to -0.250 (Table 7, p. 2232). This result is discussed only in a footnote, where the authors suggest it "might reflect the fact that as wages have continued to adjust in the affected commuting zones, some of the initial employment response may have been reversed" (p. 2237, footnote 25). The limited prominence given to this result in the abstract, introduction, and conclusion is a notable omission. The finding suggests a more complex dynamic where an initial employment shock may partially reverse over the longer run, which alters the article's primary message about the persistent displacement effects of robots.

Sensitivity of the main employment estimate: The magnitude of the main employment finding is sensitive to the inclusion of a specific control variable. The baseline estimate of the effect of robots on employment depends on controlling for the baseline share of employment in "light manufacturing" (textiles, paper, publishing, and printing). The authors acknowledge in a footnote that the employment estimates "are about 30% smaller in specifications that do not control for light manufacturing" (p. 2216, footnote 20). A comparison of the tables confirms this: the baseline employment coefficient of -0.448 (Table 2, p. 2214) falls to -0.295 when this control is excluded (Table A11, p. A-33). While the authors provide a valid justification for including the control—that these industries were declining for reasons unrelated to robots—the sensitivity of the point estimate to this choice suggests that the precise magnitude of the article's main finding is dependent on this specific modeling decision.

Scope of the theoretical framework: The article's theoretical model omits the creation of new tasks, a potentially important countervailing force to job displacement. The model is built around a framework where technology either automates and displaces labor from existing tasks or augments labor in remaining tasks (pp. 2193–2198). It does not explicitly incorporate a mechanism for the creation of entirely new tasks, job roles, or industries where labor has a comparative advantage. The "productivity effect" in the model is limited to increasing labor demand in non-automated tasks within existing industries. The authors acknowledge this limitation in the conclusion by citing their other work on the "reinstatement effect" of new tasks and noting that some general equilibrium effects "might emerge only slowly" (p. 2241). Nonetheless, the empirical analysis and its interpretation are guided by this displacement-centric model, which may not fully capture the long-run dynamics of technological change.

Interpretation of the local estimates: The article's primary empirical strategy does not account for inter-regional spillovers, meaning the local estimates should be interpreted as partial equilibrium effects. The research design estimates the effect of robot adoption within commuting zones, treating each as an independent unit. The authors acknowledge that this approach does not capture all equilibrium responses, particularly spillovers across regions through trade, supply chains, or migration (p. 2191, footnote 2). For example, a negative shock in a high-exposure commuting zone could reduce its demand for goods and services from other zones, potentially biasing the local estimates. While the article attempts to account for these spillovers in its separate, model-based calculation of aggregate effects, the core local IV estimates do not capture the full range of general equilibrium adjustments.

Generalization of the findings: The article's conclusions may over-generalize from the specific case of industrial robots to the broader categories of automation and artificial intelligence. The empirical analysis is based on data for "industrial robots" as defined by the International Federation of Robotics, with identifying variation coming predominantly from manufacturing sectors during the 1990s and 2000s (p. 2200). The article's conclusion, however, generalizes these findings back to the broader initial framing, discussing the implications for "artificial intelligence, and other automation technologies" (p. 2240). Extrapolating from the effects of industrial arms in factories to the potential societal impacts of modern AI and software automation may not be warranted without further evidence.

Minor methodological and presentation issues: Several smaller issues are present in the analysis and its presentation. First, the article does not discuss the heterogeneity implied by the difference between its weighted and unweighted results. The unweighted regressions show systematically larger negative effects for both employment and wages, suggesting that the impacts of robot exposure may be more severe in less populous commuting zones, a point that is not explored (Table 2, p. 2214). Second, a calculation in footnote 26 for the total number of jobs lost appears to be internally inconsistent. The text states that an increase of 120,000 robots reduced employment by 756,000 jobs, implying a per-robot job loss of 6.3. However, the same footnote provides a formula that yields a per-robot job loss of 6.5, a minor discrepancy likely due to rounding (p. 2238).

Future Research

Diversification of identification sources: Future work could improve upon the identification strategy by moving beyond the heavy reliance on the automotive industry. Research could exploit firm-level data or granular supply chain linkages to construct instruments that are not dominated by a single sector's trends. This would help determine if the displacement effects observed are truly a general property of automation technologies or an artifact of the specific restructuring of the global auto industry during the 1990s and 2000s.

Long-run adjustment mechanisms: To address the temporal instability of the estimates, researchers should investigate the long-run adjustment mechanisms that allow labor markets to recover. Since the negative employment effects weaken when the sample extends to 2014, future studies could focus on the "reinstatement effect"—the creation of new tasks and industries—which the current model acknowledges but does not fully empirically characterize. Quantifying the speed and magnitude of this reinstatement mechanism is essential for understanding the net long-term impact of automation.

Structural model refinement: Future research should aim to reconcile the high implied labor supply elasticity with micro-level evidence. The current structural model could be expanded to include frictions such as geographic immobility, retraining costs, or housing market constraints. By explicitly modeling these barriers to adjustment, researchers could produce aggregate counterfactuals that do not require implausibly high labor supply elasticities to fit the data, thereby providing more credible estimates of the national impact of automation.

isitcredible.com