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

Reviewer 2

February 11, 2026

v3

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 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. Version 3 has 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 investigate the effect of industrial robots on US labor markets, offering a significant empirical contribution to the debate on automation and employment. The authors construct a measure of exposure to robots for US commuting zones, instrumenting US adoption with trends from European countries to isolate technological supply shocks. Their headline conclusion is that industrial robots, unlike other forms of capital, exert a distinct negative influence on employment and wages. Specifically, they estimate that locally, the addition of one robot per thousand workers reduces the employment-to-population ratio by 0.39 percentage points and wages by 0.77% (p. 2191). Furthermore, using a calibrated structural model to account for trade and spillovers, they simulate aggregate effects, claiming that "one more robot per thousand workers reduces the aggregate employment-to-population ratio by about 0.2 percentage points and wages by 0.42%" (p. 2192).

The credibility of the local reduced-form estimates is supported by the authors' careful distinction between automation and other forms of capital deepening. They show that while exposure to robots is negatively correlated with employment, exposure to IT capital and general capital intensity is not, lending empirical weight to their theoretical argument that automation involves a unique "displacement effect" (p. 2228). However, the identification strategy faces a notable challenge regarding the source of variation. The authors acknowledge that the automotive industry accounts for 67% of the cross-commuting zone variation in robot exposure (p. 2226). This heavy reliance on a single sector raises the concern that the instrument might capture unobserved shocks specific to the global auto industry—such as demand shifts or non-robot technological changes—rather than a pure exogenous automation shock. While the authors demonstrate that results hold when the automotive sector is excluded (p. 2227), the dominance of this one industry in the main specification suggests the estimates may be less generalizable to the broader economy than the language of "automation technologies" implies.

The transition from local estimates to the headline aggregate claims introduces further uncertainty. The aggregate figures (0.2 percentage points employment loss) are not direct econometric estimates but outputs of a structural model calibrated using the local IV results. This calibration implies an inverse Frisch elasticity of labor supply of 0.17, corresponding to a wage elasticity of labor supply of approximately 5.9 (p. A-18). As the authors note, this is consistent with some macro estimates, but it is substantially higher than standard microeconomic estimates, which typically find labor supply to be much less elastic. If the local employment response is indeed this large relative to the wage response, it suggests that the mechanism of displacement may involve frictions or discouragement effects that standard equilibrium models do not fully capture. Consequently, the aggregate simulations should be viewed as model-dependent projections rather than observed national outcomes.

Furthermore, the robustness of the magnitude of the effects appears sensitive to the time period and control variables. The negative employment effect diminishes by approximately 40% when the analysis is extended from 1990–2007 to 1990–2014 (p. 2237). This attenuation suggests that the labor market may adjust over the longer term, or that the Great Recession introduced noise that dilutes the signal. Similarly, the magnitude of the local employment estimate drops by roughly 30% when controls for "light manufacturing" are excluded (p. 2216). While including these controls is methodologically defensible to account for secular declines in specific industries, the sensitivity of the point estimates indicates that the precise size of the robot effect is somewhat malleable depending on the specification.

Ultimately, this article provides compelling evidence that industrial robots caused displacement in US manufacturing hubs during the 1990s and 2000s. The claim that this specific technology reduced local labor demand is credible. However, the extrapolation to aggregate national job losses relies on structural assumptions that may overstate the elasticity of labor supply, and the results appear specific to a period of intense manufacturing adjustment. The findings confirm that automation is distinct from capital accumulation, but the long-run magnitude of the "technological unemployment" predicted by the model remains subject to considerable uncertainty.

The Bottom Line

Acemoglu and Restrepo provide strong evidence that industrial robots displaced workers and lowered wages in heavily exposed US local labor markets between 1990 and 2007. However, the article's headline claims regarding aggregate national job losses are simulations derived from a model, not direct measurements, and rely on assumptions about labor supply elasticity that are high compared to standard micro estimates. Additionally, the estimated negative employment effects weaken significantly when the analysis is extended to 2014, suggesting that the labor market impacts of automation may be more transient or complex than the initial displacement effects suggest.

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 Frisch 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). 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 Frisch elasticity is driven by the ratio of the employment coefficient ($\beta_L = -0.39$) to the wage coefficient ($\beta_W = -0.77$). The ratio of these two coefficients alone implies a reduced-form labor supply elasticity of approximately 0.5, but the structural interpretation requires converting the employment-to-population ratio change into a log change in hours, which results in the much higher implied elasticity. 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 $\varepsilon = 0.39$ ($\phi \approx 2.56$)—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. Possibly displacement from automation differs from a marginal wage change: affected workers may 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.

The instrumental variable's validity may be compromised by its heavy reliance on the automotive industry: The article's identification strategy hinges on the assumption that European robot adoption trends affect US local labor markets only through their influence on US robot adoption. This assumption is potentially threatened because the instrument's variation is overwhelmingly driven by a single, globally integrated industry: automotive manufacturing. 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 the industry has very high Rotemberg weights, indicating that it 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 numerous common shocks—such as shifts in consumer demand, oil price volatility, or non-robot technological advances—that could simultaneously drive robot adoption in both Europe and the US while also directly affecting employment in auto-centric commuting zones, the instrument may be capturing these unobserved common shocks rather than an exogenous technology shock. The authors attempt to address this concern by separating the automotive industry from others and find a statistically significant and similarly-sized negative effect for robot exposure in non-automotive industries (Table 5, p. 2227). While this provides some evidence of generalizability, the fact remains that the primary source of identification for the main specification comes from a single sector, which poses a potential threat to the exclusion restriction.

The article's headline aggregate effects are model-derived simulations, not direct econometric findings: The article's most prominent quantitative claims, including those in the abstract ("One more robot per thousand workers reduces the employment-to-population ratio by 0.2 percentage points and wages by 0.42%"), are not direct empirical estimates (p. 2192). Instead, they are the output of a calibrated structural model that uses the local instrumental variable (IV) estimates as just one of several inputs. 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. For instance, a key parameter for the productivity ratio of robots to human labor is justified using popular media reports and a TV show's website, while the parameter for cost savings from robots is sourced from a Boston Consulting Group report (p. 2239, footnote 29). By presenting these model-dependent calculations in the abstract alongside direct econometric findings without clearly distinguishing their origins, the article may give a misleading impression about the nature and certainty of the evidence for its aggregate claims.

The negative employment effect weakens substantially over a longer time period, a finding that is not prominently discussed: The article's main analysis focuses on the 1990–2007 period, for which it finds a strong negative effect of robots on employment. However, when the analysis is extended to 1990–2014, this effect shrinks by approximately 40%, while the negative wage effect becomes larger. The IV coefficient for employment changes from -0.388 in the 1990–2007 period to -0.250 in the 1990–2014 period (Table 7, p. 2237). This finding is reported in the tables but 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). This finding is reported in the tables but is discussed only in a footnote, representing a notable omission from the main narrative in the abstract, introduction, and conclusion. This finding suggests a more complex dynamic where robots may cause a significant initial employment shock that partially reverses over the longer run, accompanied by a persistent decline in wages, which alters the article's primary message.

The magnitude of the main employment finding appears sensitive to a specific control variable: The main estimate of the effect of robots on employment is sensitive to the inclusion of a control variable for the baseline share of employment in "light manufacturing" (defined as the textile, paper, publishing, and printing industries). The authors acknowledge this sensitivity in a footnote, stating 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 long-differences employment coefficient of -0.448 (Table 2, column (4), p. 2214) falls to -0.295 when this control and the female share of manufacturing employment are excluded (Table A11, Panel A, column (1), p. A-33), a reduction of over 30%. The authors justify including the control because these industries were on a steep downward trend for reasons unrelated to robots. While this is a valid reason to include the control to avoid omitted variable bias, the sensitivity of the point estimate to this choice suggests that the precise magnitude of the article's main finding is dependent on specific modeling decisions.

The theoretical framework omits the creation of new tasks, potentially biasing the model toward negative employment effects: The article's theoretical model is built around a framework where technology either automates and displaces labor from existing tasks or augments labor in remaining tasks. The model does not explicitly incorporate a key countervailing force: the creation of entirely new tasks, job roles, and 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 by citing their other work on the "reinstatement effect" of new tasks and noting in the conclusion 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.

The research design does not account for inter-regional spillovers in its local estimates: The article's empirical strategy 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 neighboring low-exposure zones, thereby contaminating the control group and potentially biasing the local estimates. While the article attempts to account for these spillovers in its separate, model-based calculation of aggregate effects, the primary local IV estimates should be interpreted as partial equilibrium effects that do not capture the full range of general equilibrium adjustments.

The article's conclusions may over-generalize from industrial robots to broader automation and AI: The article motivates its research with broad concerns about "automation technologies" and "artificial intelligence" (pp. 2189, 2240). However, its empirical analysis is based on a specific technology—"industrial robots" as defined by the International Federation of Robotics—with identifying variation coming predominantly from manufacturing sectors during the 1990s and 2000s. The article's conclusion then generalizes these findings back to the broader initial framing, discussing the implications for "artificial intelligence, and other automation technologies" (p. 2240). While it is standard to connect specific findings to a wider debate, the direct evidence presented is narrow, and extrapolating from the effects of industrial arms in factories to the potential societal impacts of modern AI and software automation may not be warranted.

Minor inconsistencies in presentation: There are several minor inconsistencies in the text. First, the main results are from regressions weighted by commuting zone population, but the article also reports unweighted results, which show slightly larger negative point estimates for both employment and wages (Table 2, p. 2214). For example, the unweighted employment coefficient is -0.516 compared to the weighted estimate of -0.448. The authors describe these as "similar" (p. 2218), which is a reasonable characterization as the difference is small relative to the standard errors and likely statistically insignificant, but the potential for heterogeneity across more and less populous commuting zones is not discussed. Second, a calculation in footnote 26 (p. 2238) for the total number of jobs lost appears to be inconsistent with the inputs provided in the same footnote. The text states that an increase of 120,000 robots is predicted to have reduced employment by 756,000 jobs, implying a per-robot job loss of 6.3. However, the same footnote calculates the per-robot job loss as "about six" and provides a formula that yields 6.5. This minor discrepancy is likely due to rounding of the underlying coefficients but creates some confusion regarding the precise magnitude of the estimated local effect.

Future Research

Investigation of long-run adjustment mechanisms: Future work should investigate the discrepancy between the strong negative employment effects observed in the 1990–2007 period and the attenuated effects found when extending the sample to 2014. Research could focus on identifying whether this reversion is due to the "reinstatement" of labor in new tasks, as posited by theory, or simply the absorption of displaced manufacturing workers into lower-productivity service sectors. This would require longitudinal worker-level data to track employment trajectories over decades rather than years.

Diversification of identification strategies: Given the heavy reliance of the current identification strategy on the automotive industry, future research should seek to exploit variation in automation technologies that are less concentrated in a single sector. This could involve analyzing the diffusion of service robots, warehousing automation, or AI-driven software, where adoption patterns may differ significantly from industrial manufacturing. Establishing whether the "displacement effect" holds for technologies that permeate the service sector is crucial for generalizing these findings.

Micro-foundations of labor supply elasticity: To address the implausibly high implied Frisch elasticity in the structural calibration, researchers should estimate the employment response to automation shocks at the individual level, distinguishing between voluntary labor force exit, involuntary unemployment, and migration. Understanding whether the large employment response is driven by wage elasticity or by other frictions (such as skills mismatch or geographic immobility) would allow for more accurate calibration of models used to simulate aggregate effects.

isitcredible.com