

Advanced Technology Adoption: Selection or Causal Effects?

By DARON ACEMOGLU, GARY ANDERSON, DAVID BEEDE, CATHERINE BUFFINGTON, ERIC CHILDRESS, EMIN DINLERSOZ, LUCIA FOSTER, NATHAN GOLDSCHLAG, JOHN HALTIWANGER, ZACHARY KROFF, PASCUAL RESTREPO, NIKOLAS ZOLAS*

Advanced technologies, including robotics, artificial intelligence, and software systems, are thought to be spreading rapidly in industrialized economies. In Acemoglu et al. (2022b), we used the 2019 Annual Business Survey (ABS) to provide a comprehensive overview of the adoption of AI, robotics, dedicated equipment, specialized software, and cloud computing for US firms in all sectors during 2016–2018.

Our work documented these facts:

- 1) The share of adopting firms remains low for AI and robotics (3.2% and 2% of firms) and rises to 19.6% to 40.2% for equipment and software.
- 2) Adoption concentrates in large firms.
- 3) As a result, a high share of workers is exposed to these technologies, especially in manufacturing. For example, 12-64% over US workers and 22-72 % of US manufacturing workers are exposed to these technologies.
- 4) A significant share of adopters, ranging from 30% for software to 65% for robotics by employment weight, report using these advanced technologies

for automation. In total, 30.4% of US workers and 52% of manufacturing workers are employed at firms using these technologies for automation.

- 5) Consistent for the use of these advanced technologies for automation, adopters have higher labor productivity but significantly lower labor shares.
- 6) Firms associate these technologies with an increase in firm demand for skills and not necessarily with an expansion in employment levels.

This paper revisits the second fact—the reasons why firms adopting advanced technologies are larger. In principle, this could be for two different reasons. Either adoption of advanced technologies **causally** expands employment. Or this fact could be due to **selection**. For example, already large firms may have a greater likelihood of adopting advanced technologies because of fixed costs, or firms that are growing fast for other reasons may also be better at adopting and using these technologies.

These two explanations have different implications. The former would suggest that advanced technologies contribute to employment growth, at least at the firm level (the industry-level implications could be very different from the firm-level ones as pointed out in Acemoglu, Lelarge and Restrepo, 2020; Koch, Manuylov and Smolka, 2021). The latter would weigh in favor of limited employment gains even in adopting firms and would caution against firm-level explorations using ordinary least squares or event-study strategies to uncover the effects of advanced technology adoption.

Our results favor the selection interpretation. Using data from the Longitudinal Business Survey (LBS), we document that adopters were already large and

* Acemoglu: MIT. Anderson: National Center for Science and Engineering Statistics. Beede, Buffington, Dinlersoz, Foster, Goldschlag, Kroff, Zolas: US Census Bureau. Childress: George Mason University. Haltiwanger: University of Maryland. Restrepo: Boston University. Any opinions and conclusions expressed herein are those of the authors and do not reflect the views of the US Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this data product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. DRB Approval Numbers: CBDRB-FY21-058, CBDRB-FY21-316, CBDRB-FY22-057, CBDRB-FY22-ESMD006-011, CBDRB-FY22-411, CBDRB-FY23-034, CBDRB-FY23-112. DMS # 7508509.




FIGURE 1. ADOPTION OF AI (LEFT) AND ROBOTICS (RIGHT) FOR FIRMS IN DIFFERENT SIZE AND AGE CATEGORIES.

Source: *Annual Business Survey, 2019.* The figure plots robot adoption rates for AI and robotics by firm age and size percentiles within detail six-digit industries. See Acemoglu et al. (2022b) for similar figures for the remaining technologies.

growing faster before AI, robotics, cloud computing, and specialized software systems became broadly available.¹ We also find that employment trends at adopting firms remained largely unchanged after the widespread use of these technologies. Persistent size and growth differences between adopters and non-adopters imply that firm-level estimates of the effects of advanced technologies must be interpreted with caution.

I. Adoption and Firm Size

We first provide graphical evidence on the relationship between firm size and adoption of AI and robotics. We focus on these technologies because they have received considerable attention in recent empirical work. Figure 1 plots adoption rates for firms in 36 size and age categories, defined in terms of employment and age percentiles within de-

tailed six-digit industries.² The figure also reports the average adoption rate for firms in each size class.

Adoption rises with size for all technologies in the ABS. 5.5% of firms in the top percentile of their industries' employment distribution use AI, 5.1% use robots, 31.4% use dedicated equipment, 67.4% use specialized software, and 63.5% use cloud computing. In contrast, the adoption rate among firms in the 50th to 75th percentile of industries' employment distribution is much lower: 3.1% for AI, 1.7% for robots, 18.6% for dedicated equipment, 39.6% for specialized software, and 33.4% for cloud.

II. Firm Employment Histories

The previous sections documented sizeable differences in employment level between adopting and non-adopting firms (for robotics and AI). We now explore whether employment histories, both in terms of levels and trends, differ between adopters and

¹These statements refer to employment. We document in Acemoglu et al. (2022b) that firms adopting advanced technologies increase their sales, while reducing their labor share, and this accounts for a nontrivial portion of the increasing concentration of sales in firms with low labor shares. The same pattern for French manufacturing is documented in Acemoglu, Lelarge and Restrepo (2020).

²We assign firms to their main six-digit NAICS industry in terms of payroll across all its establishments. Employment percentiles are defined based on the employment distribution in each industry. By construction, Figure 1 isolates differences in adoption rates across firms of different size operating in the same narrowly defined industry and controls for size differences between manufacturing and non-manufacturing firms.




FIGURE 2. EMPLOYMENT TRENDS FOR ESTABLISHMENTS IN ROBOT-USING FIRMS AND OTHERS FOR 1978–2018.

Source: *Annual Business Survey, 2019 and Longitudinal Business Survey (1978–2018)*.: The figure plots the inverse hyperbolic sine of employment in establishments associated with firms using robots in the 2018 ABS (lines with circles) and associated with non-robot users in the 2018 ABS (dashed lines). For each cohort, we report employment numbers for the years following its entry into the LBD.

non-adopters.

Because LBD does not contain consistent information on firm-establishment histories, we create a pseudo-firm establishment panel that tracks employment in all establishments associated with each firm in the ABS technology module in 2018. We then conduct our empirical analysis at the level of these establishments between 1978 and 2018.³

Figure 2 focuses on the differential employment histories of adopters and non-adopters for robotics for illustration purposes. It plots the evolution of average employment by cohort for establishments for adopting and non-adopting firms.⁴ The figure reveals three key patterns. First, establishments in adopting firms are initially larger (have higher employment) than es-

tablishments in non-adopting firms. These size differences are present at an early age and grow over time, especially for early cohorts. Second, differences in employment levels and growth rates precede the period of rapid robot adoption in the US, which took place in the late 90s and early 2000s. Third, employment dynamics of adopters' establishments seem unaffected by rising adoption of robots in recent decades.

To explore these patterns for all technologies, we turn to the following regression model:

$$(1) \quad y_{j,i,c,t} = \alpha_c + \beta_{i,t} + \gamma_c \times \text{Adopter}_j + \delta_t \times \text{Adopter}_j + \epsilon_{j,i,c,t},$$

for an establishment j in industry i , cohort c , in year t . The left-hand side variable is the inverse hyperbolic sine (IHS) of establishment employment, which allows us to include zeros in our analysis. The right-hand side variables are cohort dummies α_c , industry by year dummies $\beta_{i,t}$, which account for differences in employment trends by four-digit industries, and cohort and

³In particular, this pseudo-panel follows the same establishments over time, even though some of these establishments may not have belonged to the firm in question in the past. See Foster et al. (2016) for more details on this strategy to track activity of firms back in time.

⁴The first year in the LBD is 1976. We do not observe the exact age of establishments that existed at this point and assign them to a “pre-77” cohort.

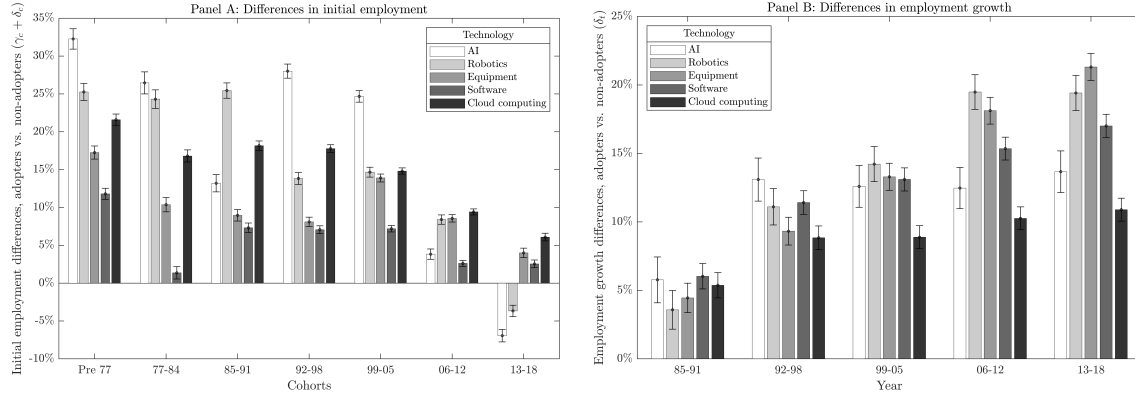


FIGURE 3. DIFFERENTIAL EMPLOYMENT DYNAMICS FOR ESTABLISHMENTS IN ADOPTING FIRMS RELATIVE TO OTHERS.

Source: *Annual Business Survey, 2019 and Longitudinal Business Survey (1978–2018)*.: Panel A plots estimates of $\gamma_c + \delta_c$ (from equation 1), which measures the differential establishment employment size for adopter firms relative to non-adopters. Panel B plots δ_t , which measures the differential establishment employment growth for adopter firms relative to non-adopters.

growth effects depending on adopter status (as measured by the adopter dummy Adopter_i). These terms allow adopters to have different initial levels (differences by cohort) and different growth dynamics (different time effects).

Figure 3 depicts estimates from equation (1) separately for the five technologies in the ABS. Since the interactions captured by the γ_c terms compare establishments of firms that start at different points in time, Panel A presents estimates of $\gamma_c + \delta_c$, which compare the initial establishment size of adopting firms of cohort c to the size of non-adopting firms at the time of entry. The results in this panel show that, consistent with our discussion for robotics adoption in Figure 2, initial size (establishment employment) of adopting firms is significantly greater than the size of non-adopters at the same point in time. For example, establishments at robot-adopting firms from the 1977–1984 cohort were initially 24.3% larger than establishments of firms not adopting robotics technology. The same difference is 14.7% for robot-adopting firms from the 1999–2005 cohort.

Panel B depicts the estimates of δ_t , which measures the differential (establishment) employment growth of adopting firms. It confirms that establishment employment for adopters grew more rapidly than for non-adopters. For example, from 1978–

1984 to 1992–1998, establishments of robot-adopting firms expanded their employment by 11.1% more than non-adopters. Notably, for most technologies, these differential growth experiences long predated the periods of high adoption in the US as a whole. Indeed, we know that robotics, AI, specialized software systems, and cloud computing were not spreading rapidly before the late 1990s.⁵ For example, the adoption of AI concentrates in the 2016–2018 period (see Acemoglu et al., 2022a), while robot adoption gained prominence in the late 90s and 2000s (see Acemoglu and Restrepo, 2020). Yet, establishments of AI and robot-adopting firms were larger and grew more rapidly than those of non-adopters decades before these periods.

Panel B also shows that the differential employment growth of adopters relative to non-adopters is unaffected by the increased adoption of these technologies in recent years. If anything, establishments in adopting firms grew at more comparable rates to establishments in non-adopting firms in recent years. For example, the yearly growth differential for establishments in robot-adopting firms relative

⁵The exception is dedicated equipment, such as computer-numerically-controlled (CNC) machines, whose widespread adoption dates back to the early 70s and is studied in detail in Boustan, Choi and Clingsmith (2022).

to non-adopters went from 0.8% per year in 1978-1998 to 0.4% in 1999-2018.

III. Discussion

Figures 2 and 3 show that establishments in adopting firms were initially larger and grew more rapidly than non-adopters, even before the adoption of advanced technologies intensified in recent years. These patterns support the view that adopters of advanced technologies are differentially selected and were already large and on differential growth trajectories.

The figures also document that the difference in employment dynamics between adopting firms' establishments and others remain largely unchanged or become less pronounced in recent years, as adoption intensifies. This is the opposite of what one would expect if advanced technologies caused adopting firms to expand their employment. Instead, it points to small or negative effects of automation technologies on firm employment trajectories.

The possibility that technology does not lead to large employment expansions at adopting firms aligns with the fact that a significant share of adopters report using advanced technologies for automation. In contrast to other applications of advanced technologies, automation reduces production cost by displacing workers from their roles, creating an ambiguous effect on firm-level employment. This possibility also aligns with firms' self-assessments on the effects of these technologies, which point to ambiguous effects of advanced technologies on employment levels (Acemoglu et al., 2022b).

One challenge when interpreting our findings is that we do not know the exact adoption date of these technologies. Currently, the ABS data only tell us whether a firm used a technology in 2016–2018. Future waves of the ABS technology module will measure year of adoption, providing a more accurate picture of how technology changes firm employment dynamics.

REFERENCES

Acemoglu, Daron, and Pascual Re-

strepo. 2020. "Robots and jobs: Evidence from US labor markets." *Journal of Political Economy*, 128(6): 2188–2244.

Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo. 2020. "Competing with Robots: Firm-Level Evidence from France." *AEA Papers and Proceedings*, 110: 383–88.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022a. "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics*, 40(S1): S293–S340.

Acemoglu, Daron, Gary W Anderson, David N Beede, Cathy Buffington, Eric E Childress, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas. 2022b. "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey." National Bureau of Economic Research Working Paper 30659.

Boustan, Leah Platt, Jiwon Choi, and David Clingingsmith. 2022. "Automation After the Assembly Line: Computerized Machine Tools, Employment and Productivity in the United States." National Bureau of Economic Research Working Paper 30400.

Foster, Lucia, John Haltiwanger, Shawn Klimek, C.J. Krizan, and Scott Ohlmacher. 2016. "The Evolution of National Retail Chains: How We Got Here." In *Handbook on the Economics of Retailing and Distribution*, ed. Emek Basker. Elgar.

Koch, Michael, Ilya Manuylov, and Marcel Smolka. 2021. "Robots and firms." *The Economic Journal*, 131(638): 2553–2584.