Automation and Disability: How Functional Limitations Shape Vulnerability to Technological Change*

Hoa Vu† Michelle Yin‡

May 8, 2026

Abstract

Broad demographic categories used to identify populations at risk from technological change can aggregate workers with opposing exposures and obscure who is actually affected. We document this for disability and industrial robot adoption. Instrumenting U.S. robot exposure across 722 commuting zones with industry-level penetration in seven European economies, we find that robot adoption over 1993 to 2014 reduces employment by 3.4 percentage points for workers with disabilities and 3.7 for workers without, effects that are statistically indistinguishable. This aggregate similarity masks substantial heterogeneity by functional limitation. Workers with sensory impairments experience employment declines of 5.5 percentage points, roughly 50 percent above the aggregate, while workers with cognitive difficulties show no significant effect. Within-disability heterogeneity exceeds the between-group difference, suggesting that vulnerability to automation is determined by the interaction between specific functional limitations and the task content that robots displace, rather than by disability status per se.

Keywords: Industrial Robots; Automation; Functional Limitations; Disability; Technological Change

JEL Codes: J14, J24, O33

* We thank participants at the 2025 Southern Economic Association Annual Meeting for their helpful comments and suggestions. All remaining errors are our own.

† School of Education and Social Policy, Northwestern University. Email: hoa.vu@northwestern.edu

‡ School of Education and Social Policy, Northwestern University. Email: michelle.yin@northwestern.edu

1. Introduction

A large literature documents that automation displaces workers in routine-intensive occupations (Acemoglu & Restrepo, 2020; Dauth et al., 2021; Graetz & Michaels, 2018). Policy discussions routinely translate this task-based finding into demographic terms, treating group membership (race, gender, disability status) as a reliable proxy for automation exposure. We show that this translation can be misleading. Using the case of disability and industrial robot adoption, we document that disability status alone is a poor predictor of automation vulnerability. The relevant predictor is the interaction between specific functional limitations and the task content that robots displace.

Workers with disabilities comprise 12.8% of the U.S. population and face persistent structural barriers in the labor market, with employment rates roughly 30-40 percentage points below those of workers without disabilities (Houtenville & Boege, 2020; Yin et al., 2026). But "disability" is not a functional category, it is an administrative and survey category that pools populations with fundamentally different capacities: sensory impairments, physical impairments, cognitive difficulties, psychological conditions, and combinations of these. These functional profiles map onto occupational task demands in very different ways. If industrial robots primarily displace sensory-motor and routine manual tasks such as machine operation, visual inspection, and manual assembly, then workers whose impairments overlap with these task demands face concentrated exposure, while workers with cognitive or psychological impairments, sorted into different occupational niches, may face little.

We provide some of the first evidence on this question, estimating the causal effect of industrial robot adoption on employment outcomes for workers with disabilities and systematically decomposing effects by functional limitation type. Our empirical strategy follows the shift-share instrumental variables design of Acemoglu and Restrepo (2020), exploiting variation in robot exposure across 722 U.S. commuting zones and instrumenting U.S. adoption with industry-level robot penetration across seven European economies. Under the Borusyak et al. (2022) identification framework, industry-level robot adoption reflects global technological advances orthogonal to local U.S. labor demand shocks, providing credible identification of causal effects.

In the aggregate, robot adoption reduces employment for workers with disabilities by 3.4 percentage points over 1993–2014, statistically indistinguishable from the 3.7 percentage point effect for workers without disabilities. Taken at face value, this similarity might suggest that disability is irrelevant to how automation reshapes labor markets. Decomposing by functional limitation, we find that workers with sensory impairments experience employment declines of 5.48 percentage points, roughly 50 percent larger than the aggregate effect and significantly above the non-disabled baseline. Workers with cognitive difficulties show no statistically significant impact. Within-disability heterogeneity is larger in magnitude than the between-group difference.

Functional limitations, interacted with task content, are therefore the relevant unit of analysis. Disability status averages over distinct experiences, understating vulnerability for workers whose impairments align with automated tasks and overstating it for workers whose impairments do not. More broadly, our results suggest caution in using broad demographic categories to identify at-risk populations. Demographic categories matter for automation exposure only insofar as they predict where workers are located in task space.

Our findings complement Di Giacomo and Lerch (2024), who document that robot exposure increases transitions into disability insurance programs among the broader workforce. They examine how automation produces disability classification; we examine how automation affects workers already living with functional limitations. Together, these results point to fiscal pressure on disability programs from both the inflow margin (newly displaced workers entering SSDI) and the stock margin (reduced employment among existing beneficiaries). Both papers exploit the same underlying source of variation (European robot adoption rates interacted with U.S. industry composition), so the two sets of estimates are not mechanically independent, but the conceptual distinction between the margins is meaningful for policy.

The heterogeneous effects we document also have implications for program design. Retraining programs and workplace accommodation initiatives that treat workers with disabilities as a uniform beneficiary category will misallocate resources, underserving the subgroup with sensory impairments most exposed to robot displacement. Back-of-the-envelope calculations suggest robot adoption over 1993–2014 displaced approximately 213,000 workers with disabilities, generating roughly $620 million in annual fiscal costs through SSDI enrollment and associated Medicare spending, alongside $6.7 billion in lost annual earnings, components of a broader displacement toll estimated at $274 billion in annual earnings losses across all affected workers.

Several additional findings are worth noting. Effects concentrate among workers without college degrees, consistent with education providing partial protection against automation. The age profile is U-shaped: young workers (16 to 24) and older workers (55 to 64) experience the largest effects, while prime-age workers show smaller impacts, suggesting automation disrupts labor market entry and accelerates exit but has more muted effects on established employment relationships. Adjustment operates primarily through employment rather than wages; workers with disabilities who remain employed show minimal wage effects, implying that the relevant margin is employment versus non-employment rather than wage compression.

The remainder of the paper proceeds as follows. Section 2 reviews relevant literature. Section 3 describes the data, including our construction of consistent disability measures across the 1990–2019 period. Section 4 presents the identification strategy and Section 5 reports results. Section 6 concludes with policy implications.

2. Literature Review and Contributions

2.1. Robots and Labor Markets

Automation displaces workers and depresses wages. In the United States, Acemoglu and Restrepo (2020) estimate that one additional robot per thousand workers reduces employment-to-population ratios by 0.2 percentage points and wages by 0.42%. These effects are concentrated in manufacturing-intensive commuting zones and routine occupations. Using data from 17 countries, Graetz and Michaels (2018) document substantial productivity gains from robot adoption, approximately 0.36% per year, which partly offset employment losses through lower prices and expanded demand.

The existing literature identifies education and age as the primary axes of heterogeneity. Acemoglu and Restrepo (2022) show that automation-driven task displacement falls disproportionately on workers with lower educational attainment, though college-educated workers are not fully insulated. This pattern is consistent with the canonical task-based framework (Acemoglu and Autor, 2011) where routine-intensive occupations are concentrated among workers without college degrees. Education shifts the incidence of displacement but does not eliminate exposure.

Age introduces a second dimension. Dauth et al. (2021) find that younger workers in more robot-exposed regions face sharply higher separation risks, consistent with newer workers lacking the firm-specific capital and institutional protections that buffer incumbents against displacement. In Germany, incumbent manufacturing workers largely absorb robot exposure through within-firm task reallocation rather than separation, reflecting strong vocational training systems and internal labor markets. Nordic countries achieve similarly smooth adjustment through active labor market policies and robust public employment services (OECD, 2025). These cross-country differences indicate that the distributional consequences of automation depend on institutional context, not technology alone.

Labor market institutions determine how these vulnerabilities translate into realized outcomes. In Germany, incumbent manufacturing workers largely absorb robot exposure through within-firm task reallocation rather than separation, reflecting strong vocational training systems and internal labor markets (Dauth et al., 2021). Nordic countries achieve similarly smooth adjustment through active labor market policies and robust public employment services (OECD, 2025). These cross-country differences indicate that the distributional consequences of automation are not technologically determined, they depend on whether institutions successfully match displaced workers to new roles.

The existing literature has established that automation vulnerability depends on the match between individual characteristics and occupational task content. It provides a cleaner identification of which functional capacities drive automation vulnerability.

2.2. Disability Employment: Structural Barriers and Program Dynamics

Employment disparities between workers with and without disabilities have persisted at 30-40 percentage points for three decades despite legislative intervention (Houtenville & Boege, 2020). This persistence reflects structural factors that interact directly with labor demand shocks. Bound & Waidmann (2002) demonstrate that disability program growth from 1984-2000 stemmed primarily from deteriorating labor market conditions for low-skilled workers rather than changes in population health. Their decomposition attributes 40% of SSDI enrollment growth. Autor & Duggan (2003) formalize this relationship, showing that state employment contractions were associated with increased disability insurance applications. These findings establish that labor demand shocks, including those from automation, translate into disability program enrollment through a well-documented channel.

Legislative efforts to close the disability employment gap have produced mixed results. Acemoglu & Angrist (2001) and DeLeire (2000) exploit the Americans with Disabilities Act as a natural experiment and find post-implementation employment declines, suggesting accommodation mandates can generate unintended hiring costs that offset their intended benefits. Schur et al. (2014) document that workplace accommodations are important for workers with disabilities, and in principal accommodation could function as accommodation, reducing physical barriers in ways that expand employment opportunities. Whether it does so in practice is an empirical question.

2.3. Disability Insurance Programs and Labor Supply

U.S. disability insurance can transform temporary shocks into permanent labor force exit. Maestas et al., 2013 exploit examiner variation in SSDI approval rates to identify causal effects, finding that the employment rate of marginal beneficiaries would have been on average 28 percentage points higher two years later absent SSDI benefits.

The dynamics of program participation prove particularly relevant for understanding technological displacement. Some applicants face long delays, for example, average wait times reach approximately two years and three months for appeals taken to an administrative law judge hearing, during which they cannot engage in substantial gainful activity (Autor, 2011). This enforced non-employment generates skill depreciation and employer stigma that persist even if benefits are ultimately denied. Among SSDI beneficiaries, return to work is rare, SSA longitudinal data show only around 28% of new beneficiaries return to work over a multi-year period (Social

Security Administration Bulletin, 2011). Labor demand shocks can therefore have persistent effects on labor force participation through the disability insurance channel. A displacement event that leads to SSDI entry may result in permanent withdrawal from the labor force.

2.4. Robots and Disability: Complementary Perspectives

Direct evidence on how robot adoption affects workers with disabilities remains scarce. Di Giacomo and Lerch (2024) provide the closest prior work, documenting that industrial robot exposure contributes to disability program take-up among the general workforce. Exploiting regional variation in robot adoption, they find that each additional robot displaces approximately four workers (often younger displaced workers going back to college or training), and 39.1% retire early. They report that middle-aged displaced workers increasingly enroll in disability insurance, while older workers opt for early retirement and younger workers return to school.

Our paper complements Di Giacomo and Lerch (2024) along two dimensions. First, whereas they examine transitions into disability status following displacement, we examine employment outcomes for individuals already living with functional limitations, a population for whom the interaction between impairment type and task content is directly observable. Second, our disaggregation by functional limitation type allows us to identify which specific capacities drive automation vulnerability, a question their framework does not address. Together, these perspectives reveal a compounding dynamic that automation both pushes non-disabled workers toward disability programs and reduces employment among those already with disabilities. The policy implications differ depending on which margin dominates, and our heterogeneity results suggest the relevant question is not whether workers with disabilities are more or less vulnerable on average, but which functional limitations map onto automation-susceptible tasks.

3. Data and Measurement

3.1. Industrial Robot Data

We measure robot adoption using data from the IFR, which provides annual information on robot shipments and operational stock by country, industry, and year for approximately 50 countries from 1993 onward. The IFR defines industrial robots as automatically controlled, reprogrammable manipulators programmable in three or more axes—excluding single-purpose machinery such as conveyor belts, cranes, and elevators. This definition ensures consistency across countries and time periods. Robot adoption in the United States has increased dramatically over our sample period. The operational stock rose from approximately 38,000 units in 1993 to 381,964 units by 2023, with robot density in manufacturing reaching 295 robots per 10,000 employees (IFR, 2024). This growth varies substantially across industries.

Following Lerch (2025), we classify industries into three categories based on adoption intensity: (i) high robot-intensive manufacturing (automotive, basic metals, electronics, food and beverages, metal products, and plastics and chemicals); (ii) low robot-intensive manufacturing (industrial machinery, minerals, paper and printing, shipbuilding and aerospace, textiles, wood and furniture); and (iii) non-manufacturing (agriculture, construction, education/research, mining, services, and utilities), given their limited robot adoption.

To quantify local labor market exposure to automation, we construct a Bartik-style robot exposure measure, following Acemoglu & Restrepo (2020). We begin by computing the growth in robots per worker for each industry as:

$$\Delta R_{i,t} = \frac{R_{i,t} - R_{i,90}}{L_{i,90}}$$

where $R_{i,t}$ is the stock of robots in industry i at time t, and $L_{i,1990}$ is baseline industry employment in 1990. We then build measures of robot exposure in the U.S. at the commuting zone level using a shift-share approach following Acemoglu & Restrepo (2020). The shift component is the adjusted change in the stock of robots at the industry level. The share component refers to the baseline employment share of industry i in commuting zone c in 1990:

$$\text{U.S. Robot Exposure}_{c,(t_0,t_1)} = \sum_{i \in I} \frac{L_{c,i,90}}{L_{c,90}} \times \left[\frac{R_{i,t_1} - R_{i,t_0}}{L_{i,90}} - \Delta ln(Y_{i,t}) \frac{R_{i,t_0}}{L_{i,90}}\right]$$

where $\frac{L_{c,i,90}}{L_{c,90}}$ denotes the share of commuting zone c's employment in industry i in 1990, and $\Delta ln(Y_{i,t})$ is overall industry output growth, used to adjust for robot adoption that is driven by output growth.

We make several adjustments to this measure following Acemoglu and Restrepo (2020). First, about 20–30% of robots in the IFR data are not classified into a specific industry. We redistribute unclassified robots across industries in proportion to each industry's share of the classified robots in a given year, ensuring total counts align. Second, the IFR reports North American totals for some industries rather than U.S.-specific counts; since the U.S. accounts for over 90% of North American robot demand during our period, we use the North America data as a proxy for U.S. adoption. Any measurement noise would likely produce attenuation bias: specifically, we use an "automation intensity" measure that nets out robots adopted due to expanding industry size (so that, for example, a region specialized in an industry that grew output rapidly isn't unfairly coded as high robot exposure just because of scale rather than automation intensity). This adjustment nets out robots adopted due to expanding industry scale, so that a region specialized in a fast-growing industry is not assigned high robot exposure solely because of output growth.

Our analysis covers three sub-periods: 1990–2000, 2000–2010, and 2010–2019. We treat these as three stacked observations per CZ, yielding 722 CZs × 3 periods = 2,166 CZ-period observations. Differencing over decades removes time-invariant regional factors. Figure 1 illustrates substantial variation in exposure to industrial automation across commuting zones. The highest exposure is concentrated in states such as Michigan, Indiana, and Ohio, sectors that have adopted industrial robots at accelerated rates. On average, U.S. commuting zones experienced an increase of 1.58 robots per 1,000 workers between 1993 and 2014; exposure at the 99th percentile reaches as high as 7 robots per 1,000 workers.

3.2. Labor Market Data and Disability Definition

Our individual-level data on employment and disability status come from the U.S. Decennial Census and American Community Survey (ACS). We use the 1990 Decennial Census as the baseline and the ACS annual microdata for years 2000 onward (through 2019). The ACS is an annual survey (1% sample of the population) that replaced the long-form Census questionnaire, and it provides detailed information on demographics, employment status, occupation, earnings, and disability status.

Measuring disability consistently over time is a challenge because the survey questions changed in 2000. In 1990, the Census asked about work-limiting disabilities in a narrow format. Starting in 2000, the ACS introduced six questions covering distinct domains of functional difficulty: *cognitive difficulty* refers to serious difficulty concentrating, remembering, or making decisions; *ambulatory difficulty* reflects serious difficulty walking or climbing stairs; *independent living difficulty* captures challenges in performing errands alone, such as visiting a doctor's office; *self-care difficulty* denotes difficulty with activities like dressing or bathing; and *vision or hearing* *difficulty* indicates serious difficulty seeing even with glasses or hearing even with a hearing aid. An individual is coded as having a disability if they report *any* of these difficulties.

However, not all of these questions were asked in earlier years. The 1990 Census captured work-limiting disability (a binary indicator for health conditions limiting or preventing work) along with measures of self-care and mobility. The 2000-onward questions are broader. To create a consistent disability measure across our entire sample period, we use the two domains that are available *both before and after 2000*: Independent Living and Self-Care difficulties. These were asked in 1990 (though phrased slightly differently) and continue to be asked post-2000. We define a person as having a disability if they report either an independent-living difficulty or a self-care difficulty. For years 2000 and beyond, we also include anyone reporting the other four difficulties (ambulatory, cognitive, vision, hearing) as having a disability, since those meet standard disability definitions. For 1990 (and our 1990–2000 change), we rely on Independent Living and Self-Care to identify disability, whereas for 2000–2019 we use all six. This approach may under-count people with disabilities in 1990 relative to 2000. A person with only a hearing difficulty in 1990, for example, would not be identified.

The definitional change introduces a specific threat: if the 1990 measure undercounts workers with sensory or cognitive limitations relative to the 2000+ measure, then the 1990-2000 first difference for these subgroups mechanically reflects measurement expansion rather than employment change. Two features of our design mitigate this concern. First, our stacked first-differences specification includes period fixed effects that absorb any level shift in disability prevalence common across commuting zones between the 1990 and 2000 measurement regimes. The identifying variation is cross-CZ within-period, not cross-period within-CZ. Second, we restrict the sample to 2000–2019 and apply a strictly consistent definition using all six disability questions throughout (Column 7 of Table 6), yielding estimates that are statistically indistinguishable from the full-sample results (-0.925 versus -1.050 for workers with disabilities, p-value for difference > 0.7). The disability-type decomposition in Table 2 uses only the 2000-2019 subsample for cognitive, ambulatory, and sensory difficulties (n = 1,444 CZ-periods rather than 2,166), so the heterogeneity results are estimated within the consistent-definition window and is not affected by the pre-2000 measurement change.

Using the above definition, our sample includes anyone of working age (16–64) who reports at least one of the relevant difficulties. In 1990, this constituted about 8% of the working-age population. By 2019, the share reporting any disability is around 10.3%, reflecting both demographic shifts and the expanded scope of survey questions. We control for baseline disability prevalence across commuting zones in our regressions, so regional variation in reporting pattern should not confound our results.

3.3. Sample Construction and Descriptive Statistics

We merge robot exposure measures to individual Census/ACS data at the commuting zone-year level, excluding those in institutional settings (e.g., nursing homes, prisons) and agricultural workers and self-employed individuals, and focus on working adults ages 16–64.

After applying sample restrictions and sample weights, we aggregate the microdata to the commuting zone × year × disability status level to obtain cell means of employment rates and other outcomes. We then form first differences between the baseline and end of each period for each CZ and group. For example, for the 1990–2000 period, we compute the change in the employment rate of people with disabilities in CZ *c* from 1990 to 2000, and similarly for people without disabilities. These changes are the dependent variables in our regression analysis. First-differencing removes persistent regional differences that may be correlated with robot adoption.

Table B-1 presents summary statistics for our analysis sample. Panel A presents employment rates for workers with disabilities, and Panel B presents rates for those without disabilities. In 1990, the employment rate was 35.67 percent for workers with disabilities and 69.97 percent for those without disabilities. By 2019, these rates increased to 37.36 and 71.68 percent, respectively. The 1990 and 2019 change is 1.69 percentage points increase for workers with disabilities and 1.71 percentage points for workers without disabilities. Relative to their baseline means, these changes correspond to an approximate 5% increase in employment for workers with disabilities and 2.4% for workers without.

3.4. Additional Measurement Considerations

Two measurement issues warrant discussion. First, we assign robot exposure based on the CZ of residence and assume workers work in the same local labor market. Workers, however, may live and work in different CZs. If a substantial number of people with disabilities commute out of their home CZ for work (or vice versa) and robotization affects jobs in the work CZ, our measure of exposure might not perfectly align with where employment effects show up. CZs are designed to capture local labor markets, and long-distance commuting is less common among workers with disabilities. Any resulting measurement error likely biases our effects toward zero.

Second, our disability measure captures a broad range of severities, from individuals with mild difficulties who are still in the labor force to those with severe impairments that may preclude work entirely. When we observe a drop in the employment-to-population ratio for people with disabilities, it could reflect increased unemployment or reduced labor force participation or a combination. Sample sizes at the CZ level for unemployment among people with disabilities are small, so we focus on the employment-to-population ratio.

4. Empirical Strategy

4.1. Baseline Specification

We estimate the impact of industrial robot adoption on employment using a stacked first-differences approach across 722 commuting zones over three periods (1990-2000, 2000-2010, 2010-2019). Our baseline specification is:

$$y_{c,(t_0,t_1)} = \alpha + \beta\ \text{U.S. Robot Exposure}_{c,(t_0,t_1)} + \gamma X_{c,(t_0,t_1)} + \delta_s + \theta_t \times \mu_d + \varepsilon_{c,(t_0,t_1)} \quad (1)$$

where $y_{c,(t_0,t_1)}$ is the change in employment rate for people with or without disabilities in commuting zone c between year t_0 and t_1. The key explanatory variable, $\text{U.S. Robot Exposure}_{c,(t_0,t_1)}$, measures the change in robots per thousand workers in commuting zone c, constructed as described in Section 3. We standardize this measure to have mean zero and standard deviation one for ease of interpretation. The specification includes state fixed effects δ_s to account for time-invariant regional factors, census division-by-period fixed effects to control for differential regional trends ($\theta_t \times \mu_d$), and a vector of controls for the industrial and occupational composition of employment, sociodemographic characteristics, the demographic composition of industries and occupations within commuting zones in 1990, and the China trade shock, measured following Autor et al. (2013)[4]. Following Lerch (2025), we hold these controls constant at 1990 values to avoid contamination from endogenous adjustments to robot adoption. Standard errors are clustered at the state level to account for spatial correlation in labor market shocks. We also report standard errors following Adão et al. (2019) as a robustness check, which account for correlated shocks across regions and industries.

[4] Including the China trade shock as a control is important because Autor et al. (2013) show that U.S. regions more exposed to Chinese import competition experienced substantial increases in disability benefit uptake. Controlling on this measure therefore demonstrates that the estimated robot effect on people with disabilities is robust to China trade exposure.

4.2. Identification Challenges

Two endogeneity concerns arise: reverse causality and omitted variables. Local labor market conditions might drive robot adoption. Tight labor markets or rising accommodation costs could accelerate automation, generating positive bias if robots are adopted where employment is growing, or negative bias if adopted in declining regions. Unobserved factors correlated with both robot adoption and employment could also generate bias. For example, regions experiencing broader technological transformations such as increased software adoption may simultaneously invest more in robotics and exhibit different employment trends, making it difficult to isolate the effect of robots.

4.3. Instrumental Variables Strategy

To further mitigate the concerns of confounding factors that may be correlated with both the industry-level spread of robots in the U.S and labor market outcomes among PWD and isolate robot adoption driven by exogenous advances in robotics (i.e., supply shocks), we construct an instrument by exploiting the industry-level spread of robots in other economies, which are meant to proxy improvements in the world technology frontier of robots (Acemoglu and Restrepo, 2020). Adopting the Borusyak et al. (2022) framework, we interpret technology shocks as originating from the global engineering frontier (the “shock”) rather than from the location of production (the “share”).

We use the average industry-level growth of robots in seven European countries (Denmark, Finland, France, Italy, Spain, Sweden, and the United Kingdom) available in the IFR data over the same period. These cross-industry growth rates proxy for technological breakthroughs at the global frontier that differentially affect industries’ feasibility of automation. Under the Borusyak et al. (2022) framework, identification relies on the assumption that industry-level robot growth rates are quasi-random, conditional on industry-level controls and weighted by average exposure. This approach allows the exposure shares to be endogenous (e.g., Detroit differs from San Francisco in industrial composition), while requiring that differential robot adoption across industries, such as faster growth in automotive relative to textiles, is driven by technological progress rather than U.S.-specific labor market conditions. Our instrument is formally defined as follows:

$$\text{EU Robot Exposure}_{\mathrm{c,(t_0,t_1)}} = \sum_{\mathrm{j \in EU}} \frac{1}{7} l_{c,j}^{70} \sum_{\mathrm{i \in EU}} \left[\frac{R_{j,t_1}^{i} - R_{j,t_0}^{i}}{L_{j,90}^{i}} - \Delta ln(Y_{j,(t_0,t_1)}^{i}) \frac{R_{j,t_0}^{i}}{L_{j,90}^{i}} \right]$$

where R_{j,t_1}^{i} is the stock of robots in country $i \in EU7$ at time t in industry j. The share component uses employment shares from 1970 which predate the introduction of industrial robots (Acemoglu and Restrepo, 2020).

4.4. Instrument Validity

Figure 2 demonstrates strong first-stage relevance, with a positive correlation of 0.48 between European and U.S. robot adoptions. The first-stage F-statistic of 101.13 far exceeds conventional thresholds for weak instruments. Global technological diffusion produces correlated adoption patterns across advanced economies, making a strong first stage expected.

For the exclusion restriction to hold, European robot adoption must affect U.S. disability employment only through its effect on U.S. robot adoption. The primary threat would be foreign robot adoption proxying for a global industry trend that also directly affects U.S. labor markets. One concern is trade competition: if European industries adopt robots and become more efficient, they might increase exports to the U.S., affecting U.S. employment. We address this concern by controlling for trade exposure. In particular, in robustness checks (Table 6), we include controls for import competition from Europe using a shift-share measure of U.S. imports from Europe in the spirit of Autor et al. (2013). Our estimates remain stable, suggesting trade competition does not drive our results.

A second concern is that foreign robot adoption could be correlated with other global technology trends (such as computerization) which also diffuse to the U.S. We include division-by-period fixed effects to control for common shocks that affect broad regions of the U.S. equally. Furthermore, by focusing on differences across CZs within the U.S., we are implicitly controlling for any nationwide time effects (such as the overall productivity boom of the late 1990s) which would affect treated and untreated regions alike. A remaining concern is whether commuting zones with higher subsequent robot exposure were already on differential employment trends prior to robot adoption. While we cannot implement a formal pre-trends test with our stacked first-differences design, we note that our specification controls for 1990 baseline levels of industrial composition, occupational structure, and demographic characteristics, which absorb pre-existing differences correlated with future robot exposure. The stability of our estimates across alternative instrument constructions and sample restrictions (Table 6) provides indirect evidence against confounding pre-trends.

4.5. Shift-Share Diagnostics

Our instrument combines two sources of variation: pre-determined 1970 industry employment shares and European robot adoption shifts. Two frameworks offer complementary perspectives on shift-share instruments. Goldsmith-Pinkham et al. (2020, hereafter GPSS) treat the shares as instruments, requiring exogeneity of each industry's employment share. Borusyak et al. (2022, hereafter BHJ) treat the shifts as the source of exogenous variation, requiring that industry-level robot adoption is quasi-random conditional on industry controls. We adopt the BHJ framework: identification relies on the exogeneity of the shifts, not the shares.

We implement four sets of diagnostics, reported in Section 5.6. First, we compute Rotemberg weights (GPSS) to characterize which industries contribute most to the identifying variation. Second, we report Adão, Kolesár, and Morales (2019) standard errors that account for correlated shocks across regions sharing similar industry composition. Third, we conduct a leave-one-industry-out exercise to assess sensitivity to any single sector. Fourth, we examine the correlation between industry-level shocks and observables. Under our BHJ framework, concentrated Rotemberg weights are expected when one industry dominates robot adoption and do not threaten identification; we report the GPSS decomposition for transparency rather than as a validity test of the identifying assumption.

5. Results

5.1. Main Effects on Employment: Disability Status as a Noisy Proxy

Table 1 presents our estimates of robot exposure's impact on employment. We begin with the aggregate comparison between workers with and without disabilities to establish the baseline against which the disaggregated results should be read, and because it illustrates precisely why disability status is an inadequate proxy for automation exposure. Panel A reports OLS estimates, which suggest minimal effects. A one-standard-deviation increase in robot exposure reduces employment by 0.15 percentage points for workers with disabilities (not statistically significant) and 0.40 percentage points for those without disabilities ($p < 0.01$). The implied change in the disability employment gap is essentially zero (-0.003 pp).

Panel B presents our preferred IV estimates. Robot exposure reduces employment by 1.05 percentage points for workers with disabilities ($p < 0.05$) and 1.16 percentage points for those without ($p < 0.01$). Scaled over the actual increase in robot exposure between 1993 and 2014 (1.58 robots per thousand workers), these imply employment declines of 3.39 and 3.74 percentage points respectively[5]—economically significant magnitudes representing roughly 10 percent of baseline employment for workers with disabilities. The IV estimates are 3-7 times larger than OLS, suggesting substantial downward bias from endogenous robot adoption or measurement error.

Robot adoption produces employment losses for workers with and without disabilities that are statistically indistinguishable (p-value for equality = 0.82), leaving the disability employment gap unchanged. At the aggregate level, disability employment tracks overall labor market conditions rather than diverging from them.

This aggregate similarity is not evidence that disability is irrelevant to automation. The disability category pools workers with sensory impairments, concentrated in sensory-motor and routine manual tasks that industrial robots directly displace, alongside workers with cognitive impairments, sorted into occupational niches that are substantially more robot-resistant over this period. These opposing exposures average out in the aggregate, producing a mean effect that is comparable to the non-disabled population but masks heterogeneity that is both theoretically meaningful and policy-relevant. Unpacking this composition is the central analytical task of what follows.

5.2. Heterogeneity by Disability Type

Table 2 disaggregates effects by specific functional limitation type.[6] To interpret magnitudes, we scale the estimated coefficients by dividing by the standard deviation of the robot exposure measure (0.49) and multiplying by the average increase in the stock of robots over 1993

[5] To obtain the effect of the adoption of robots between 1993-2014, we divide the estimated coefficients by the standard deviation of the robot exposure measurement (0.49) and multiple with the average increase in the stock of robots between 1993-2014 (1.58).

[6] The sample sizes in Table 2 vary across columns because data on cognitive, ambulatory, and vision or hearing difficulties are unavailable for 1990. As a result, these outcomes include only two periods: 2000–2010 and 2010–2019.

to 2014 (1.58). Effects vary substantially across disability types, from -1.70 percentage points for vision or hearing difficulties to no significant effect for cognitive difficulties.

Workers with sensory disabilities (vision or hearing difficulties) experience the largest employment losses: the IV coefficient of -1.70 percentage points scales to -5.48 percentage points over 1993–2014, roughly 50 percent above the aggregate effect. This potentially reflects challenges adapting to new technologies or workplace reorganization following automation. Communication barriers may impede retraining or transition to service occupations.

Workers with ambulatory difficulties experience moderate effects (-0.71 pp, scaling to -2.29 percentage points over 1993 to 2014). While robots reduce physical demands in some occupations, displaced workers with mobility limitations may face barriers entering new industries or occupations requiring different physical capabilities. Workers with cognitive difficulties show no statistically significant effect (-0.27 pp, scaling to -0.87 pp), suggesting these workers are either employed in less-exposed occupations or that cognitive limitations do not amplify automation's employment effects. Independent living difficulties (-0.56 pp, scaling to -1.81 pp) and self-care difficulties (0.17 pp, scaling to 0.55 pp, not significant) show mixed results, consistent with the heterogeneous nature of these categories.

These patterns challenge uniform narratives about automation and disability. Rather than helping or harming workers with disabilities equally, effects depend on the specific interaction between functional limitations and automation's task-specific impacts. For example, if robots replace visual inspection tasks, not only do some sighted workers lose jobs, but workers who are blind (who may hold roles in other parts of the production line) do not necessarily benefit. The only positive coefficient (+0.517) for self-care is not statistically significant. The net effect reinforces existing patterns of exclusion for workers with sensory and certain physical impairments.

It also indicates that future technologies (such as AI and service robots) could have similarly uneven effects if the interaction between functional limitations and task displacement is not considered in program and technology design.

5.3. Heterogeneous Effects by Worker Characteristics

Table 3 presents IV estimates disaggregated by education, age, race, gender, and disability status.

Education. Among workers with disabilities, those without college degrees experience employment declines of 1.13 percentage points compared to 0.89 for graduates, a 27 percent difference. Education provides partial insulation, likely by broadening access to roles that robots have not displaced. The education gradient is steeper among workers without disabilities (1.46 versus 0.49 pp), indicating that a college degree offers considerably less protection for workers with disabilities than for their non-disabled counterparts. This pattern suggests that the occupational transitions through which education typically buffers displacement, moving into robot-resistant roles, are partially blocked for workers with disabilities.

*Ag*e. Employment effects follow a U-shaped age profile where the youngest workers (16–24: -1.33 pp) and oldest workers (55–64: -1.06 pp) experience the largest declines, while prime-age workers show smaller and often insignificant effects. For young workers with disabilities, automation forecloses entry into routine manufacturing before stable employment relationships are established. For older workers, displacement more readily translates into permanent labor market exit. Prime-age workers, already embedded in employment relationships, are more insulated at both margins. Panel B shows a similar U-shape among workers without disabilities, with effects ranging from -1.60 pp for ages 16 to 24 to -0.81 pp for ages 35 to 44, consistent with the entry and exit margins being sharper for workers with disabilities.

Race. White workers with disabilities experience larger employment declines than non-White workers (-1.40 versus -0.57 pp), while the pattern reverses for workers without disabilities (-0.89 versus -2.06 pp). This cross-group reversal is consistent with differential occupational sorting by race and disability status. We interpret these estimates cautiously because sample sizes for non-White workers with disabilities are small. The results nonetheless indicate that technology shocks can have disparate impacts across racial groups.

Gender. Among workers with disabilities, men experience larger employment losses than women (-1.27 versus -1.00 pp). The gender gap is narrower within the disability population than among non-disabled workers (-1.73 versus -0.76 pp in Panel B). This compression reflects disability-driven occupational sorting constraining both men and women with disabilities into a more similar range of roles, reducing the gender-based occupational differentiation that drives the larger gap in Panel B.

5.4. Industry Heterogeneity

Table 4 examines how employment effects vary across industry categories. Effects among workers with disabilities (Panel A) concentrate in high robot-intensive manufacturing, automotive, metals, electronics, with employment declining by 0.70 percentage points. Low robot-intensive manufacturing shows negligible effects. The comparison with Panel B is informative here in a specific way. Workers without disabilities show the same directional pattern, losses concentrated in high robot-intensive manufacturing (-1.16 pp), small positive effects in low-intensive manufacturing, but at larger magnitudes. The proportional similarity across panels suggests that the industry distribution of where workers with disabilities were employed, not disability status per se, is what determined exposure.

5.5. Wage Effects

Table 5 presents wage effects, measured as changes in log average wages at the commuting zone-demographic group level. The effects are modest for workers with disabilities and concentrated among specific subgroups. Workers with disabilities without college degrees experience wage declines of 0.128%, while college graduates show no significant effect. Women with disabilities face larger wage declines than men (-0.09% versus -0.038%), consistent with gender segregation into different occupations within exposed industries.

For workers without disabilities, wage effects are both larger and more pervasive. Those without college degrees experience wage declines of 1%, nearly eight times the effect for similarly educated workers with disabilities. This difference partly reflects lower baseline wages for workers with disabilities, providing less room for downward adjustment, and greater likelihood of exiting employment entirely rather than accepting wage cuts. The gender wage gap widens among workers with disabilities (women's wages fall more) while narrowing among those without disabilities. This differential pattern suggests interactions between automation, gender, and disability in wage determination.

Overall, the wage analysis suggests that adjustment to robot adoption among workers with disabilities operates primarily along the employment margin rather than through wage reductions. Employment effects are substantial, whereas wage effects are small and imprecisely estimated. This pattern is consistent with evidence that shorter unemployment durations among workers with disabilities often reflect higher rates of labor force exit rather than rapid reemployment (Maroto and Pettinicchio, 2025). Adverse labor market shocks are more likely to induce transitions out of employment for workers with disabilities than downward wage adjustments, underscoring that the relevant margin of adjustment is employment versus non-employment rather than wage compression.

5.6. Robustness Checks

We perform a variety of additional checks to ensure the stability and validity of our results as presented in Table 6. Across all specifications, the estimated coefficients remain negative and close in magnitude to the baseline results, reinforcing the stability of our findings. Column (1) reproduces the baseline estimate from Table 3 as a reference point.

Trade Competition. A potential threat to our identification strategy is that the European robot adoption also affects U.S. labor market outcomes through product market competition. Following Lerch (2025), we control for international product market competition using a shift-share measure of U.S. imports from Europe (Autor et al., 2013). Column (2) shows effects of –1.14 for workers with disabilities, and –1.23 for workers without, suggesting that trade competition does not account for the observed employment declines. We also include the China trade shock as a control in Equation (1), following Autor et al. (2013), who show that U.S. regions more exposed to Chinese import competition experienced substantial increases in disability benefit uptake. Our results are robust to controlling for China trade exposure.

Alternative Instrument Construction. Columns (3)–(5) examine sensitivity to alternative constructions of the instrument. Column (3) expands the baseline EU7 measure to include Germany. We exclude Germany in the baseline because its robot adoption is advanced relative to other countries. Including Germany yields slightly smaller effects (–0.85 for workers with disabilities, –1.00 for workers without) but both remain statistically significant. Column (4) follows Acemoglu and Restrepo (2020) and uses an EU5-based instrument comprising Denmark, Finland, France, Italy, and Sweden. The results (–1.11 and –1.17) closely match the baseline. Column (5) employs an instrument that omits the adjustment for industry growth, $\Delta ln(Y^{i}_{j,(t_0,t_1})$. The estimates (–0.94 and –0.96) are consistent with the baseline.

Excluding Outliers. Column (6) excludes Detroit (MI), the commuting zone with the highest exposure to robots, to assess whether the results are driven by this single outlier. For both groups, the coefficients increase in magnitude (–1.76 and –1.41) but lose precision, indicating that Detroit contributes to the estimates but does not drive them.

Rotemberg Weight Decomposition. We report the GPSS decomposition for transparency, not as a validity test of our instrument. Under the Borusyak et al. (2022) framework we adopt, identification relies on the exogeneity of industry-level shifts (EU7 robot adoption rates), not on the exogeneity of employment shares. Rotemberg weight concentration is expected under Borusyak et al. (2022) when one industry has both large employment shares and large robot adoption growth, and does not threaten identification. We treat the decomposition as a way to characterize where identifying variation comes from and to assess whether results are driven by coherent or idiosyncratic patterns.

We first implement the Goldsmith-Pinkham et al. (2020) decomposition. The Bartik 2SLS estimator can be written as a weighted average of 19 industry-specific just-identified IV estimators, $\hat{\beta} = \sum_k \hat{\alpha}_k \hat{\beta}_k$, where the Rotemberg weights $\hat{\alpha}_k$ measure each industry's contribution to identification. Table B-2 reports these weights alongside the industry-specific IV estimates. Identification is highly concentrated in the automotive sector, which receives a weight of $\hat{\alpha}_{\text{auto}} = 0.964$, so that 96.4 percent of the 2SLS estimate is driven by variation in 1970 automotive employment shares interacted with EU7 robot adoption rates. This concentration reflects the empirical reality that industrial robot adoption in the United States over this period was dominated by the automotive sector. Our estimates therefore primarily capture the employment effects of automotive-sector automation, which is the most policy-relevant margin given the concentration of workers with sensory and physical impairments in manufacturing production. Petrochemicals is a distant second at $\hat{\alpha}_{\text{ptrc}} = 0.032$. Negative Rotemberg weights, often a source of concern in shift-share designs, are economically negligible, summing to only -0.0006 across all 10 negatively-weighted industries. No industry is pulling the estimate in the opposite direction from the others.

The industry-specific IV estimates $\hat{\beta}_k$ further support this interpretation. For the automotive sector, the primary source of identifying variation, the just-identified estimate is $\hat{\beta}_{\text{auto}} = -0.735$ for workers with disabilities and $\hat{\beta}_{\text{auto}} = -1.109$ for workers without disabilities. Both are negative and consistent with the aggregate 2SLS estimates. Estimates for the remaining industries are imprecise, reflecting their minimal contribution to overall identification.

We also implement a leave-one-industry-out exercise in which we reconstruct the instrument 19 times, each time excluding one industry from the national exposure measure and re-estimating the baseline specification. This provides a transparent check on whether the main findings are sensitive to any single sector. The results are plotted in Figure B-1. For workers with disabilities, the distribution of leave-one-out coefficients is clustered around the baseline estimate and remains negative across all 19 samples. For workers without disabilities the coefficients are somewhat more dispersed but also remain negative throughout. As anticipated given the large Rotemberg weight on automotive, dropping that industry attenuates the point estimate (consistent with automotive being the primary source of identifying variation), but does not change its sign, and all 19 leave-one-out coefficients remain negative. This attenuation is expected and uninformative about instrument validity under BHJ: removing the industry that contributes the most variation mechanically reduces precision and magnitude. What matters is whether the industry-level shocks are plausibly exogenous, not whether identification is concentrated. Taken together, the GPSS decomposition and the leave-one-out exercise confirm that identification is concentrated in the automotive sector, that the direction of the effect is not an artifact of a single sector, and that no individual industry is pulling the estimate in the opposite direction.

Adão-Kolesár-Morales (AKM) Standard Error. A potential concern in shift-share settings is that conventional region-level (in our context state-level) clustering standard errors in shift-share designs are often too small, leading researchers to find "statistically significant" results that are actually just noise. This is because in a shift-share design, regions that have similar industry structures (e.g., two commuting zones that both rely heavily on the "Auto" industry) will have highly correlated residuals, even if those commuting zones are geographically distant. Standard geographic clustering (by state) fails to account for this industry-based correlation. Adão, Kolesár, and Morales (2019) develop an inference procedure that accounts for correlation in residuals driven by common exposure to industry shocks. Their approach therefore produces standard errors that are robust to correlated shocks across regions and industries. We report these AKM standard errors in Table 1 (square brackets). The AKM-robust SEs do not change our conclusions. In the IV specification (Panel B) the coefficient for both workers with and without disabilities remains statistically significant at conventional levels. We also note that under the BHJ framework, the effective number of independent shocks contributing to identification is small (reflecting automotive dominance), which reinforces the importance of the AKM standard errors that account for correlated industry-level shocks. The BHJ-specific identification assumption is that industry-level robot adoption in Europe is quasi-random conditional on industry controls; the automotive sector's technological trajectory, driven by global engineering advances in welding, painting, and assembly automation, is plausibly exogenous to U.S. local labor demand conditions.

Consistent Disability Definition. Column 7 restricts the sample to 2000-2019, applying all six ACS disability questions consistently throughout. The disability employment effect is -0.925

(SE 0.288, $p < 0.01$), compared to -1.050 in the full sample. The 12 percent attenuation is consistent with the 1990-2000 period contributing modestly to the full-sample estimate, not with the pre-2000 measurement difference driving the result. The disability-type decomposition in Table 2 uses only this 2000-2019 subsample for all disability types that require the expanded ACS questions (cognitive, ambulatory, sensory). The within-disability heterogeneity (sensory -1.70 versus cognitive -0.27) is therefore estimated on a definitionally consistent sample throughout and does not depend on the pre-2000 measurement bridge.

5.7. Welfare Calculations

We provide back-of-the-envelope calculations to illustrate the employment and fiscal consequences of robot adoption. These calculations are illustrative and focus on distributional and fiscal costs, abstracting from potential productivity gains or consumer surplus generated by automation.

In 2019, there were approximately 6.3 million workers with disabilities and 156 million without disabilities. Applying the cumulative 1993 to 2014 displacement rates to the 2019 workforce under the assumption that the displacement effect is persistent and continues to operate on the current population, our estimates suggest that robot adoption displaced approximately 213,570 workers with disabilities (3.39 pp × 6.3 million) and 5.83 million workers without disabilities (3.74 pp × 156 million).

The displacement of these workers represents not only lost earnings but also fiscal costs through disability insurance programs. Using average annual earnings from the U.S. Bureau of Labor Statistics, displaced workers with disabilities lost approximately \$6.73 billion in annual income (213,570 × \$31,500; Office of Disability Employment Policy DOL, 2025), while displaced workers without disabilities lost \$267.22 billion (5.83 million × \$45,800; Office of Disability

Employment Policy DOL, 2025). These direct earnings losses likely understate total welfare effects, which also include fiscal externalities through increased reliance on transfer programs.

We assume that 10% of displaced workers with disabilities transition to Social Security Disability Insurance (SSDI), generating 21,357 new beneficiaries, following Di Giacomo and Lerch (2024) who estimate it for the general displaced population. For workers already living with disabilities, the SSDI transition rate following displacement is almost certainly higher, implying that our fiscal cost estimate is likely a lower bound. With average annual benefits of $15,000 and associated Medicare spending, these new SSDI entrants generate approximately $620 million in annual fiscal costs.

Lost labor income also reduces tax revenue. Applying a conservative combined income and payroll tax rate of 20 percent to the estimated $274 billion in foregone earnings implies approximately $55 billion in lost annual tax revenue. Focusing on workers with disabilities alone, foregone tax revenue associated with earnings losses is approximately $1 billion per year.

At the individual level, the present discounted value of displacement can be substantial. For a worker with a disability displaced at age 50 who subsequently enters SSDI until full retirement age, the discounted value of disability and Medicare benefits alone exceeds $350,000 under standard discounting assumptions. Accounting for foregone tax contributions raises the total social cost per displaced SSDI entrant to over $400,000. While not all displaced workers enter SSDI, even temporary nonemployment or early retirement implies meaningful fiscal and welfare costs.

Workers with disabilities face particularly severe displacement consequences (Maroto & Pettinicchio, 2025) and higher probability of permanent labor force exit (Collischon et al., 2025). These disparities suggest that standard unemployment insurance and retraining programs may be inadequate for displaced workers with disabilities. Intergenerational effects compound these costs, as children of displaced workers experience 9% lower annual earnings (Oreopoulos et al., 2008).These calculations illustrate the magnitude of displacement costs and motivate policy responses that address displacement costs without impeding productivity-enhancing adoption. Our calculations do not account for offsetting benefits such as productivity gains or consumer surplus from lower prices, which would be required for a complete welfare analysis.

6. Conclusion

This paper began with a simple observation: disability status, as recorded in surveys and used in policy discussions, is not a functional category. It pools workers with sensory impairments, physical limitations, cognitive difficulties, and psychological conditions whose occupational positions, task exposures, and adjustment capacities differ substantially. Treating this administrative aggregate as a proxy for automation vulnerability produces misleading inferences in both directions.

The aggregate employment effect of robot adoption on workers with disabilities, 3.4 percentage points over 1993 to 2014, is statistically indistinguishable from the 3.7 percentage point effect on workers without disabilities. This similarity reflects composition, not uniform impact. Workers with sensory impairments lose 5.48 percentage points of employment, roughly 50 percent above the aggregate, while workers with cognitive difficulties show no significant effect. The within-disability variation exceeds the between-group difference. Disability status accounts for relatively little of the meaningful heterogeneity; functional limitations alongside task content are the primary source of variation. The age profile reinforces this conclusion as effects follow a U-shaped pattern, concentrating among labor market entrants and older workers while prime-age workers show smaller impacts. Combined with education gradients showing effects concentrated among workers without college degrees, the evidence indicates that vulnerability is determined by the intersection of functional limitations, career stage, and occupational task content.

Our findings complement Di Giacomo and Lerch (2024), who show that robot exposure pushes non-disabled workers toward disability insurance programs. We show that the same exposure reduces employment among those already living with a disability. Automation thus generates fiscal pressure on disability programs from both margins simultaneously, increasing inflows through displacement of the general workforce while reducing employment of existing beneficiaries. Back-of-the-envelope calculations suggest these effects displaced over 200,000 workers with disabilities, generating $6.7 billion in lost annual earnings and $620 million in annual fiscal costs through SSDI enrollment and associated Medicare spending.

Several limitations warrant acknowledgment. The disability measurement change in 2000 requires care in interpretation, though results are robust to consistent definitions throughout. We estimate reduced-form effects without identifying specific mechanisms; whether displacement occurs through direct task substitution, reduced hiring, or accelerated separation remains an open question. Our welfare calculations are illustrative; a complete accounting would include productivity gains from automation, general equilibrium effects on wages and prices, and non-pecuniary costs of displacement.

Artificial intelligence and service robots will extend automation beyond manufacturing into retail, healthcare, and hospitality, sectors employing many workers with disabilities. Our evidence from industrial robots provides an early indication of what may follow. The comparable displacement rates we document indicate that workers with disabilities compete in the same labor markets and face similar technological pressures as all workers. Disability employment is a labor economics question, not solely a social policy concern. Whether technological progress reduces or reinforces barriers for workers with disabilities depends on how functional limitations interact with the specific tasks that new technologies displace.

References

Acemoglu, D., & Angrist, J. D. (2001). Consequences of Employment Protection? The Case of the Americans with Disabilities Act. *Journal of Political Economy*, *109*(5), 915–957. https://doi.org/10.1086/322836

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

Adão, R., Kolesár, M., & Morales, E. (2019). Shift-Share Designs: Theory and Inference. *The Quarterly Journal of Economics*, *134*(4), 1949–2010. https://doi.org/10.1093/QJE/QJZ025

Autor, D. (2011). *The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options* (No. W17697; p. w17697). National Bureau of Economic Research. https://doi.org/10.3386/w17697

Autor, D. H., Dorn, D., & Hanson, G. H. (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. H., & Duggan, M. G. (2003). The Rise in the Disability Rolls and the Decline in Unemployment. *The Quarterly Journal of Economics*, *118*(1), 157–206. https://doi.org/10.1162/00335530360535171

Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, *89*(1), 181–213. https://doi.org/10.1093/RESTUD/RDAB030

Borusyak, K., Hull, P., & Jaravel, X. (2025). A Practical Guide to Shift-Share Instruments. *Journal of Economic Perspectives*, *39*(1), 181–204. https://doi.org/10.1257/jep.20231370

Bound, J., & Waidmann, T. (2002). Accounting for Recent Declines in Employment Rates among Working-Aged Men and Women with Disabilities. *The Journal of Human Resources*, *37*(2), 231. https://doi.org/10.2307/3069646

Bureau of Labor Statistics. (2020). *PERSONS WITH A DISABILITY: LABOR FORCE CHARACTERISTICS—2019*.

Bureau of Labor Statistics. (2025). *People with a Disability: Labor Force Characteristics—2025*.

Centers for Medicare and Medicaid Services. (2019). *2019 Medicare Trustees Report*.

Collischon, M., Hiesinger, K., & Pohlan, L. (2025). Disability onset and labor market outcomes. *Socio-Economic Review*, mwaf021. https://doi.org/10.1093/ser/mwaf021

DeLeire, T. (2000). The Wage and Employment Effects of the Americans with Disabilities Act. *The Journal of Human Resources*, *35*(4), 693. https://doi.org/10.2307/146368

Di Giacomo, G., & Lerch, B. (2024). *Robots and Non-participation in the US: Where Have All the Workers Gone?* SSRN. https://doi.org/10.2139/ssrn.4874380

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, *110*(8), 2586–2624. https://doi.org/10.1257/aer.20181047

Houtenville, A., & Boege, S. (2020). 2019 Annual Report on People with Disabilities in America. *Institute on Disability, University of New Hampshire*.

Lerch, B. (2025). From Blue- to Steel-Collar Jobs: The Decline in Employment Gaps? *American Economic Journal: Macroeconomics*, *17*(1), 126–160. https://doi.org/10.1257/mac.20220051

Maestas, N., Mullen, K. J., & Strand, A. (2013). Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt. *American Economic Review*, *103*(5), 1797–1829. https://doi.org/10.1257/aer.103.5.1797

Maroto, M., & Pettinicchio, D. (2025). First Fired, Last Hired, and Lower Paid: Re-employment Outcomes Among Displaced Workers With Disabilities, 2007–2021. In F. Hyseni, L. Schur, D. Kruse, & P. Blanck (Eds.), *Disability and the Future of Work* (pp. 25–42). Emerald Publishing Limited. https://doi.org/10.1108/S1479-354720250000016003

OECD. (2025). *The Role of Public Employment Services in Promoting an Inclusive Nordic Common Labour Market*. OECD Publishing. https://doi.org/10.1787/54cec4d6-en

Office of Disability Employment Policy DOL. (2025). *Median Annual Earnings Map*. https://www.dol.gov/agencies/odep/research-evaluation/MAEmap

Oreopoulos, P., Page, M., & Stevens, A. H. (2008). The Intergenerational Effects of Worker Displacement. *Journal of Labor Economics*, *26*(3), 455–000. https://doi.org/10.1086/588493

Schur, L. (2002). The Difference a Job Makes: The Effects of Employment among People with Disabilities. *Journal of Economic Issues*, *36*(2), 339–347. https://doi.org/10.1080/00213624.2002.11506476

Social Security Administration. (2019). *Annual Statistical Report on the Social Security Disability Insurance Program, 2019*.

Social Security Administration Bulletin. (2011). Longitudinal Statistics on Work Activity and Use of Employment Supports for New Social Security Disability Insurance Beneficiaries. *Social Security Administration Research, Statistics, and Policy Analysis*. https://www.ssa.gov/policy/docs/ssb/v71n3/v71n3p35.html?utm_source=chatgpt.com

Yin, M., Seo, R., & Vu, H. (2026). The labor market effects of subminimum wage elimination: Evidence from a national analysis. *Labour Economics*, *100*, 102884. https://doi.org/10.1016/j.labeco.2026.102884

Figures

Figure 1 - Increase in Robot Exposure Across Commuting Zones between 1993-2014

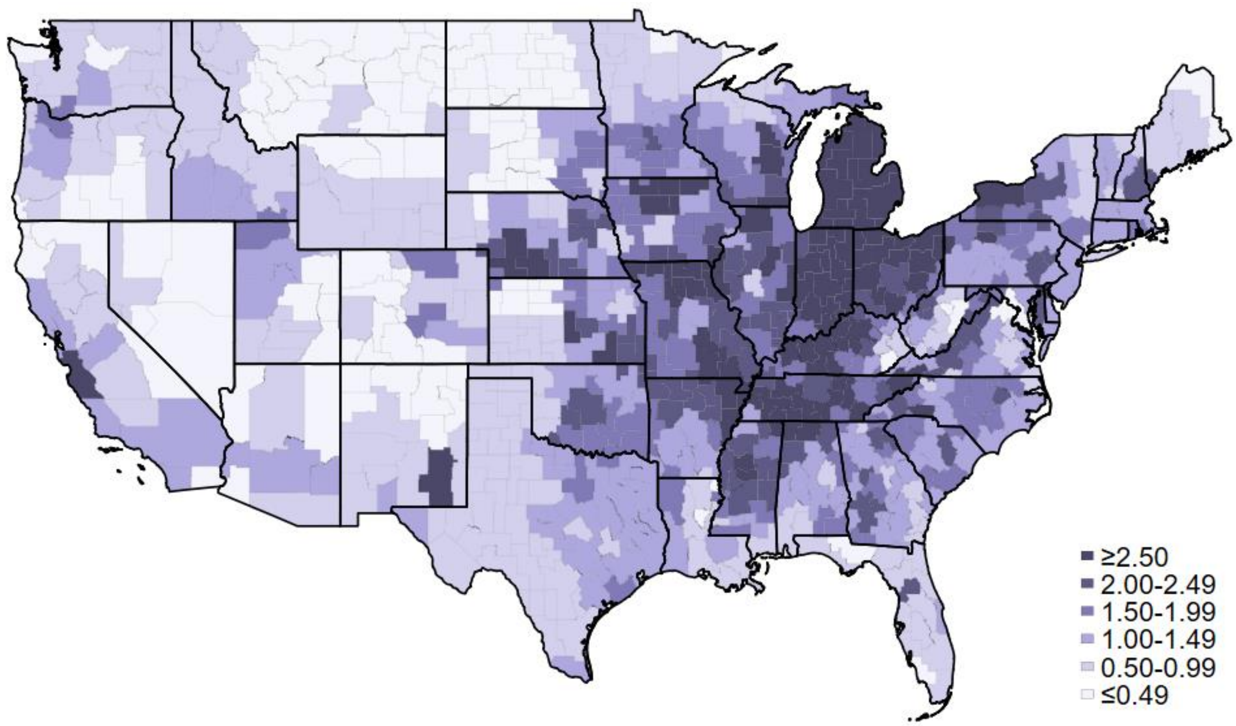

Notes. This figure shows variation in exposure to industrial automation across commuting zones in the U.S. per 1,000 workers. The map shades commuting zones by exposure intensity, with darker areas indicating greater exposure.

Figure 2 – First Stage Relationship

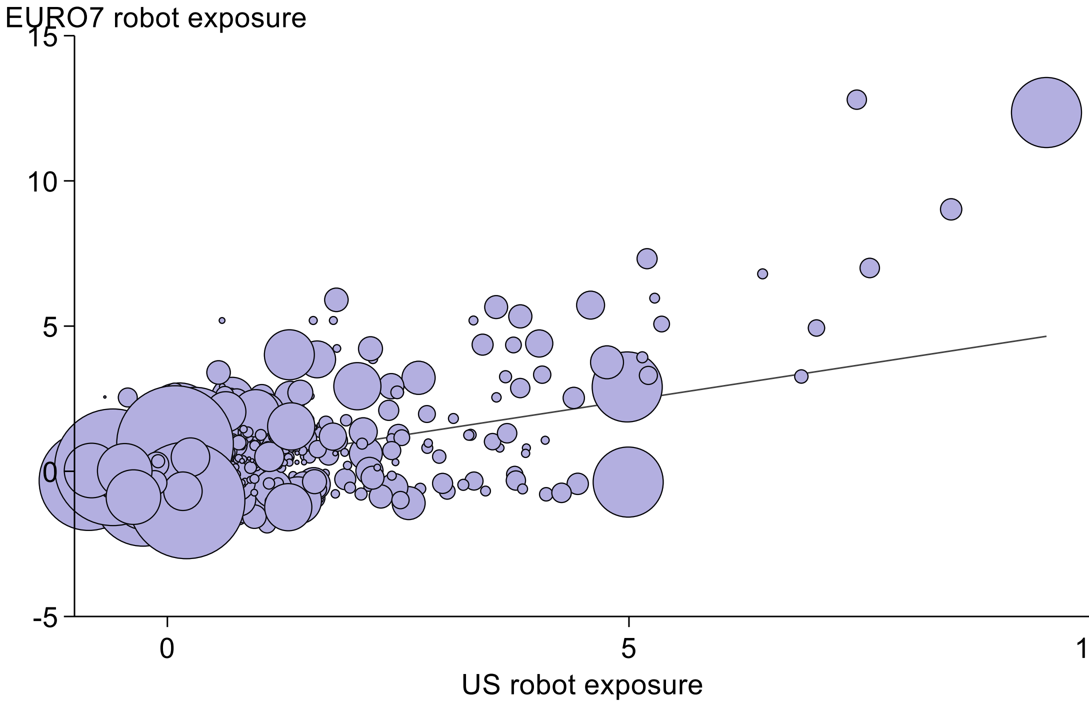

Notes. The figure illustrates the relationship between robot exposure in EURO7 countries and in the United States for 1993–2014. The estimated coefficient is 0.48 (standard error 0.045), with a first-stage F-statistic of 101.13. Each circle represents a commuting zone in one of the three sample periods. The solid line shows a regression weighted by 1990 commuting zone population, and marker sizes reflect the 1990 commuting zone population.

Tables

Table 1 – Effects of Robots on Employment

	Outcome: Employment		
	People with disabilities	People without disabilities	Employment gap
	(1)	(2)	(3)
Panel A. OLS estimates			
Robot exposure	-.1548	-.3984***	-.0028*
	(.1228)	(.09)	(.0015)
Observations	2166	2166	2166
R-squared	.8329	.7699	.6617
Panel B. IV estimates			
Robot exposure	-1.0501**	-1.1617***	-.0003
	(.4386)	(.186)	(.0039)
	[.4239]	[.1522]	[.0056]
Observations	2166	2166	2166
R-squared	.0266	-.0153	.2031

Notes. The table presents the estimates of the impact of exposure to robots on employment rates by disability status. Data are drawn from the American Community Survey. Exposure to robots is standardized. Changes are expressed in percentage points of the working-age population of the respective demographic group and are multiplied by 100. There are three time periods (1990–2000, 2000–2010, and 2010–2019) and 722 commuting zones. All regressions include state fixed effects, division-by-year fixed effects, and controls for industrial and occupational composition of employment, sociodemographic characteristics, and the demographic composition of industries and occupations within commuting zones in 1990. Standard errors, reported in parentheses, are clustered at the state level. Standard errors following Adão, Kolesár, and Morales (2019) are reported in square brackets. Negative R-squared values in IV specifications are expected when IV-predicted values explain less within-sample variation than the controls alone and do not indicate model misspecification.

*** p<0.01, ** p<0.05, * p<0.1

Table 2 - Heterogeneity in the Effects on Robots by Disability Type – IV estimates

	Outcome: Employment					
	All people with disabilities	Cognitive difficulty	Ambulatory difficulty	Independent living difficulty	Self-care difficulty	Vision or hearing difficulty
	(1)	(2)	(3)	(4)	(5)	(6)
Robot Exposure	-1.0501**	-.2741	-.7074*	-.5593*	.1732	-1.6996***
	(.4385)	(.3179)	(.3641)	(.2831)	(.1872)	(.5615)
Observations	2166	1444	1444	2166	2166	1444
R-squared	.0266	.0339	.0639	.0208	.0369	.0308

Notes. The table presents the estimates of the impact of exposure to robots on employment rates by disability status, education, age, race, and gender. Data are drawn from the American Community Survey. Exposure to robots is standardized. Changes are expressed in percentage points of the working-age population of the respective demographic group and are multiplied by 100. There are three time periods (1990–2000, 2000–2010, and 2010–2019) and 722 commuting zones. For cognitive, ambulatory, and vision or hearing difficulties, data are unavailable for 1990; hence, only two periods (2000–2010 and 2010–2019) are included for these outcomes. All regressions include state fixed effects, division-by-year fixed effects, and controls for industrial and occupational composition of employment, sociodemographic characteristics, and the demographic composition of industries and occupations within commuting zones in 1990. Standard errors, shown in parentheses, are clustered at the state level.
*** p<0.01, ** p<0.05, * p<0.1

Table 3 - Heterogeneity in the Effects on Robots by Education, Age, Race, and Gender – IV estimates

	Outcome: Employment										
	Education		Age					Race		Gender	
	College degree	Less than college	16-24	25-34	35-44	45-54	55-64	Whites	Non-Whites	Women	Men
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Panel A. People with disabilities											
Robot Exposure	-.885**	-1.134**	-1.3333***	-.8421	-1.2797*	-1.0546**	-1.0565***	-1.4026***	-.5718	-1.0046**	-1.2699***
	(.3813)	(.4762)	(.4607)	(.9782)	(.7047)	(.4659)	(.2805)	(.3887)	(.391)	(.4734)	(.4435)
Observations	2083	2166	2138	2166	2164	2165	2166	2166	2106	2166	2166
R-squared	.0107	.0448	.0226	.0193	.0075	.0176	.0114	.0008	.0442	.0101	.0455
Panel B. People without disabilities											
Robot Exposure	-.4903***	-1.4645***	-1.5964***	-1.1285***	-.9427***	-.8095***	-1.4967***	-.8936***	-2.0578***	-.7608***	-1.7326***
	(.0879)	(.2281)	(.296)	(.2499)	(.1874)	(.1626)	(.1686)	(.1231)	(.3387)	(.1295)	(.2538)
Observations	2166	2166	2166	2166	2166	2166	2166	2166	2166	2166	2166
R-squared	.0135	-.0179	.0244	.0103	.0212	.0135	-.0105	.0175	-.0214	.0464	-.0297

Notes. The table presents the estimates of the impact of exposure to robots on employment rates by disability status, education, age, race, and gender. Data are drawn from the American Community Survey. Exposure to robots is standardized. Changes are expressed in percentage points of the working-age population of the respective demographic group and are multiplied by 100. There are three time periods (1990–2000, 2000–2010, and 2010–2019) and 722 commuting zones. All regressions include state fixed effects, division-by-year fixed effects, and controls for industrial and occupational composition of employment, sociodemographic characteristics, and the demographic composition of industries and occupations within commuting zones in 1990. Standard errors, shown in parentheses, are clustered at the state level.
*** p<0.01, ** p<0.05, * p<0.1

Table 4 - Heterogeneity by Industry – IV estimates

	Outcome: Employment		
	Industry		
	High robot-intensive manufacturing	Low robot-intensive manufacturing	Non-manufacturing
	(1)	(2)	(3)
Panel A. People with disabilities			
Robot Exposure	-.6995***	.2329*	-.0187
	(.2435)	(.1218)	(.2391)
Observations	2166	2166	2166
R-squared	.0936	.1006	.0968
Panel B. People without disabilities			
Robot Exposure	-1.1603***	.2729*	.262
	(.3087)	(.1501)	(.2181)
Observations	2166	2166	2166
R-squared	.2355	.2445	.2307

Notes. The table presents the estimates of the impact of exposure to robots on employment rates by workers' disability status and industry. *High robot-intensive manufacturing* refers to manufacturing industries that have adopted robots at relatively high rates. These include Automotive, Basic Metals, Electronics, Food and Beverages, Metal Products, and Plastics and Chemicals. *Low robot-intensive manufacturing* comprises manufacturing industries with relatively low levels of robot adoption. They include Industrial Machinery, Minerals, Paper and Printing, Shipbuilding and Aerospace, Textiles, Wood and Furniture, and Miscellaneous Manufacturing. *Non-manufacturing* encompasses six residual sectors where robot adoption has been minimal compared to manufacturing. These sectors include Agriculture, Construction, Education and Research, Mining, and Services and Utilities. Data are drawn from the American Community Survey. Exposure to robots is standardized. Changes are expressed in percentage points of the working-age population of the respective demographic group and are multiplied by 100. There are three time periods (1990–2000, 2000–2010, and 2010–2019) and 722 commuting zones. All regressions include state fixed effects, division-by-year fixed effects, and controls for industrial and occupational composition of employment, sociodemographic characteristics, and the demographic composition of industries and occupations within commuting zones in 1990. Standard errors, shown in parentheses, are clustered at the state level.
*** p<0.01, ** p<0.05, * p<0.1

Table 5 – Effects of Robots on Wages – IV estimates

	Education		Age					Race		Gender	
	College degree	Less than college	16-24	25-34	35-44	45-54	55-64	Whites	Non-Whites	Women	Men
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Panel A. People with disabilities											
Robot Exposure	.0004	-.128**	-.0037	-.0131	-.0475***	-.0287	-.0345*	-.1193	-.0082	-.09**	-.0376
	(.0268)	(.0598)	(.0148)	(.026)	(.0151)	(.0303)	(.0199)	(.0863)	(.1453)	(.043)	(.0435)
Observations	2166	2166	2166	2166	2166	2166	2166	2166	2166	2166	2166
R-squared	.0224	.0478	.0249	.0195	.0125	.0187	.0244	.0596	.0081	.022	.0381
Panel B. People without disabilities											
Robot Exposure	-.1019	-1.0009***	-.1007	-.7653***	.0378	-.0785	-.1961***	-.9783**	-.1245	-.396***	-.7068***
	(.0999)	(.2715)	(.1106)	(.0686)	(.108)	(.0576)	(.0632)	(.4472)	(.2927)	(.0924)	(.1171)
Observations	2166	2166	2166	2166	2166	2166	2166	2166	2166	2166	2166
R-squared	.4068	.0851	.0409	.0476	.0631	.0493	.0463	.1414	.3263	.0594	.0589

Notes. The table presents the estimates of the impact of exposure to robots on wages by disability status, education, age, race, and gender. Data are drawn from the American Community Survey. Exposure to robots is standardized. Changes are expressed in percentage points of the working-age population of the respective demographic group and are multiplied by 100. There are three time periods (1990–2000, 2000–2010, and 2010–2019) and 722 commuting zones. All regressions include state fixed effects, division-by-year fixed effects, and controls for industrial and occupational composition of employment, sociodemographic characteristics, and the demographic composition of industries and occupations within commuting zones in 1990. Standard errors, shown in parentheses, are clustered at the state level.

*** p<0.01, ** p<0.05, * p<0.1

Table 6 – Robustness Checks

	Baseline	Adding control for U.S. imports from seven European countries	Alternative IV: EU7 + Germany	Alternative IV: EU5 (Acemoglu and Restrepo, 2020)	Alternative IV: Without adjustment for industry growth	Excluding Detroit	Sample restricted to 2000–2019
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. People with disabilities							
Robot Exposure	-1.0501**	-1.1387**	-.848**	-1.1092*	-.9433**	-1.7645*	-.925***
	(.4386)	(.4774)	(.3622)	(.552)	(.3904)	(.8968)	(.288)
Observations	2166	2166	2166	2166	2166	2163	1444
R-squared	.0266	.0375	.036	.0236	.0319	.0142	.024
Panel B. People without disabilities							
Robot Exposure	-1.1617***	-1.2321***	-1.0039***	-1.1749***	-.9629***	-1.4109***	-1.2144***
	(.186)	(.1929)	(.1564)	(.2331)	(.141)	(.3439)	(.1525)
Observations	2166	2166	2166	2166	2166	2163	1444
R-squared	-.0153	-.0013	.0157	-.0084	.0205	-.0263	.1053

Notes. The table presents the estimates of the impact of exposure to robots on employment by disability status. Data are drawn from the American Community Survey. Exposure to robots is standardized. Changes are expressed in percentage points of the working-age population of the respective demographic group and are multiplied by 100. There are three time periods (1990–2000, 2000–2010, and 2010–2019) and 722 commuting zones. All regressions include state fixed effects, division-by-year fixed effects, and controls for industrial and occupational composition of employment, sociodemographic characteristics, and the demographic composition of industries and occupations within commuting zones in 1990. Standard errors, shown in parentheses, are clustered at the state level.

*** p<0.01, ** p<0.05, * p<0.1

Appendix A: Welfare Effects of Industrial Robot Adoption

A.1 Overview
This appendix provides detailed calculations of the welfare effects of industrial robot adoption. We combine our employment estimates with parameters from the literature to calculate: (1) direct earnings losses and (2) fiscal costs through disability insurance programs.

A.2 Direct Earnings Losses
Workers with Disabilities
Employment loss calculation:
- Baseline PWD workforce (2019): 6.3 million (Bureau of Labor Statistics, 2020)
- Employment effect: -3.39 percentage points (Authors' estimates)
- Workers displaced: 6,300,000 × 0.0339 = 213,570 workers

Annual earnings loss:
- Average annual earnings for PWD 2019-2023: $31,500 (Office of Disability Employment Policy DOL, 2025)
- Total annual loss: 213,570 × $31,500 = $6.73 billion

Workers without Disabilities
Employment loss calculation:
- Baseline PWOD workforce (2019): 156 million (Bureau of Labor Statistics, 2020)
- Employment effect: -3.74 percentage points (Authors' estimates)
- Workers displaced: 156,000,000 × 0.0374 = 5,834,400 workers

Annual earnings loss:
- Average annual earnings for PWOD: $45,800 (Office of Disability Employment Policy DOL, 2025)
- Total annual loss: 5,834,400 × $45,800 = $267.22 billion

Total Direct Earnings Loss: $273.95 billion annually

A.3 Disability Insurance Program Costs
SSDI Enrollment
Following Di Giacomo and Lerch (2024), we assume 10% of displaced PWD transition to SSDI:
- New SSDI beneficiaries: 213,570 × 0.10 = 21,357 workers

Annual SSDI benefit costs:
- Average annual SSDI benefit: $15,000 (Social Security Administration, 2019)
- Annual cost: 21,357 × $15,000 = $320 million

Medicare Costs
SSDI beneficiaries become eligible for Medicare after 24 months:
Annual Medicare costs (starting year 3):
- Average annual Medicare cost per beneficiary: $14,000 (Centers for Medicare and Medicaid Services, 2019)
- Annual cost: 21,357 × $14,000 = $300 million

Total Fiscal Cost: $620 million annually

A.4 Lost Tax Revenue
Applying a 20% combined income and payroll tax rate:

- Total lost earnings: \$273.95 billion
  Estimated tax revenue loss: \$55 billion annually
- PWD lost earnings: \$6.73 billion
  PWD tax revenue loss: \$1.35 billion annually

A.5 Present Discounted Value of Fiscal Burden

Assuming a 3% real discount rate and average SSDI/Medicare duration of 15 years:

- Present value of fiscal cost per SSDI entrant:
  SSDI (\$15,000) + Medicare (\$14,000) = \$29,000/year
- 15-year annuity, 3% discount rate ≈ 11.94 annuity factor $\left(=\frac{1-1.03^{-15}}{0.03}\right)$
- NPV of fiscal cost per SSDI entrant: \$29,000 × 11.94 = \$346,260
- Aggregate NPV of fiscal cost (21,357 entrants): \$7.4 billion

A.6 Summary

Total Annual Cost of Robot-Induced Labor Displacement (1993–2014)

Component	All Displaced Workers	Workers with Disabilities
Earnings losses	\$273.95B	\$6.73B
Fiscal cost (SSDI + Medicare)	\$0.62B	\$0.62B
Foregone tax revenue	\$55.00B	\$1.35B

It is worth noting that these calculations yield three distinct cost estimates that should be interpreted as complementary rather than additive. First, robot adoption over 1993–2014 generated approximately \$273.95 billion in annual earnings losses borne directly by displaced workers, of which \$6.73 billion falls on workers with disabilities. Second, the fiscal cost to government through SSDI enrollment and associated Medicare spending amounts to \$620 million annually, with a present discounted value of \$7.4 billion over a 15-year horizon. Third, foregone tax revenue of approximately \$55 billion annually (\$1.35 billion for displaced workers with disabilities). This forgone tax revenue and earnings losses are mechanically related: foregone taxes are a share of lost earning. Therefore, summing them involves some double-counting.

Appendix B: Additional Figures and Tables

Figure B-1 – Robustness to dropping each of the 19 industries in the sample

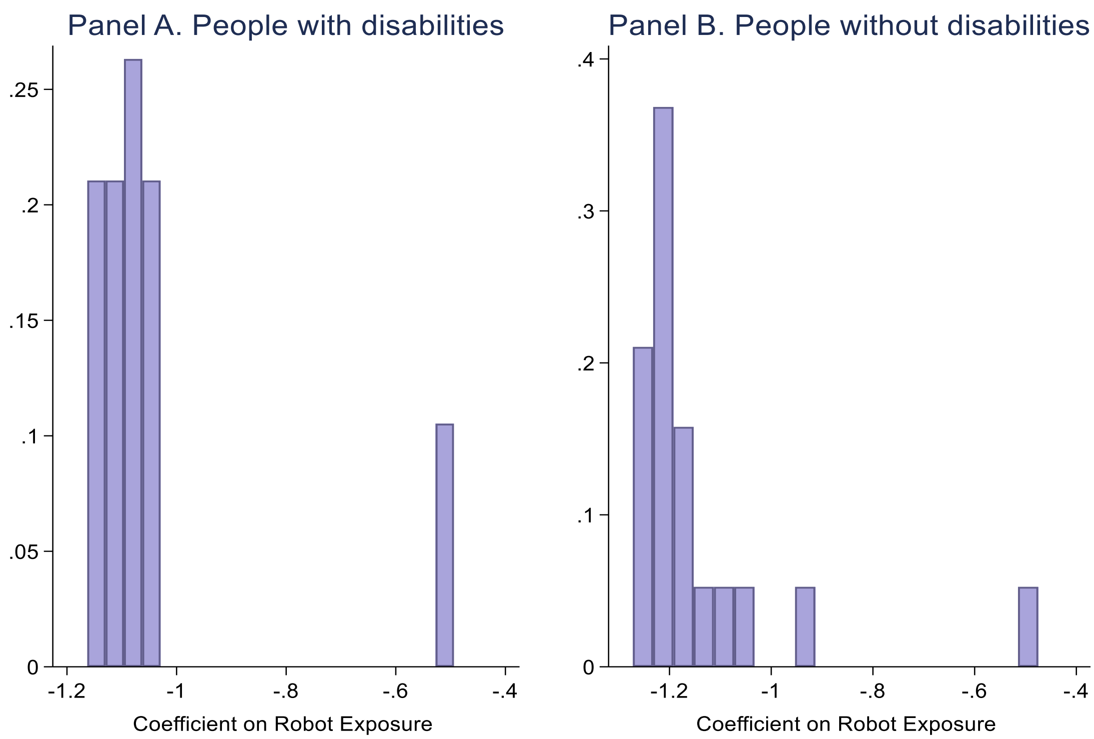

Notes. This figure explores the robustness of our findings to dropping each of the 19 industries from our sample. The vertical axis shows the fraction (proportion) of industries falling into each coefficient bin. Standard errors are clustered at the state-level.

Table B-1 – Summary Statistics

	1990	2019	1900-2019 change
	(1)	(2)	(3)
Panel A. People with disabilities			
Employment rate	35.67	37.36	1.69
Observations	722	722	722
Panel A. People without disabilities			
Employment rate	69.97	71.68	1.71
Observations	722	722	722

Notes. This table shows average employment for people with and without disabilities. Data are drawn from the American Community Survey.

Table B-2 – Rotemberg Weight Decomposition of the Shift-Share IV Estimate

Industry	Rotemberg Weight		$\hat{\beta}_k$ (PWD)		$\hat{\beta}_k$ (PWOD)		Negative weight
	$\hat{\alpha}_k$	%	Coef.	SE	Coef.	SE	
Panel A: Industries with non-negligible weight (|$\hat{\alpha}$k| > 0.001)							
Automotive	0.964	96.29	-0.735	0.249	-1.109	0.132	No
Petrochemicals	0.032	3.23	0.307	1.761	-1.447	1.241	No
Metal (products)	0.002	0.22	20.538	63.992	9.773	33.495	No
Metal (basic)	0.001	0.08	2.104	3.326	-1.216	1.992	No
Food	0.001	0.08	10.184	9.883	2.782	3.226	No
Furniture	-0.001	0.06	-0.773	1.124	1.519	1.157	Yes
Panel B: Industries with negligible weight (|$\hat{\alpha}$k| < 0.001)							
Metal (machinery)	≈0.000	0.006	-56.85	365.83	-25.33	163.89	No
Manufacturing (other)	≈0.000	0.001	1.07	3.00	3.75	2.57	No
Mineral	≈0.000	0.003	-0.10	2.26	-0.46	1.34	No
Electronics	≈0.000	0.004	18.96	47.19	6.55	20.55	No
Vehicles (other)	≈0.000	0.027	12.80	9.07	6.89	5.03	No
Paper	≈0.000	0.002	0.47	1.02	-0.45	0.42	Yes
Textiles	≈0.000	0.003	-0.69	1.44	-0.41	0.91	Yes
Mining	≈0.000	0.000	130.20	193.99	130.20	193.53	Yes
Services	≈0.000	0.000	-71.38	145.28	-4.09	20.37	No
Research	≈0.000	0.000	-14.51	22.42	-3.07	6.13	Yes
Utilities	≈0.000	0.000	2.41	1.70	5.28	2.65	Yes
Construction	≈0.000	0.000	-224.59	18206.53	-299.77	24351.64	Yes
Agriculture	≈0.000	0.000	9.36	9.21	-11.07	11.56	Yes

Notes: Rotemberg weights $\hat{\alpha}_k$ follow Goldsmith-Pinkham, Sorkin, and Swift (2020). The 2SLS estimate equals $\sum_k \hat{\alpha}_k \hat{\beta}_k$. $\hat{\beta}_k$ is the just-identified IV estimate using only industry k's 1970 employment share interacted with the EU7 robot adoption rate as the sole instrument. Standard errors clustered at the state level. PWD = people with disabilities; PWOD = people without disabilities. Negative weight = Yes if $\hat{\alpha}_k < 0$, indicating the industry contributes to identification in a perverse direction. Industries sorted by $|\hat{\alpha}_k|$ in descending order within each panel. Industries with $|\hat{\alpha}_k| < 0.001$ are grouped in Panel B; their individual $\hat{\beta}_k$ estimates are imprecise by construction as the instrument carries negligible identifying variation for those sectors.

Appendix C: Conceptual Framework

1. Production Technology and Task Structure

This framework builds on the task-based production model of Lerch (2025), which itself extends Autor et al. (2013). Lerch (2025) uses a two-skill (brawn/brain) Roy model to study the heterogeneous employment effects of industrial robots across gender and race/ethnicity groups, showing that occupational segregation drives differential displacement. The present framework departs from Lerch (2025) in two respects. First, rather than brawn and brain skills, we distinguish routine-automatable and complex task capacities, a framing better suited to the disability context, where the barrier to complex employment stems from historical exclusion and accommodation costs rather than biological comparative advantage. Second, we introduce three disability-specific distortion parameters (reduced skill support $\varepsilon_C^{D=1}$, mobility compression $\mu^{D=1}$, and accommodation costs $\phi^{D=1}$). With this in mind, consider a production technology that combines three factor inputs: routine-automatable labor L_R, complex labor L_C, and robot/automation capital K, to produce output Y:

$$Y_t = \left(K_t^{\rho} + L_{R,t}^{\rho}\right)^{\frac{\beta}{\rho}} L_{C,t}^{1-\beta} \tag{1}$$

with $\beta, \rho \in (0,1)$ and $\beta < \rho$. Automation capital K substitutes imperfectly for routine labor L_R with elasticity $1/(1-\rho) > 1$ and complement to complex labor L_C. The elasticity of substitution between the routine task bundle and complex labor is normalized to unity.

Automation capital is competitively supplied with a price falling exogenously over time due to technological progress:

$$p_t = \theta e^{-\delta t} \tag{2}$$

where $\theta = e^{\delta}$ is an efficiency parameter and $\delta > 0$ governs the rate of technological advance. This falling price is the causal force of the model.

Perfect competition implies labor is paid its marginal product. The first-order conditions yield endogenous labor demand:

$$\omega_{R,t} = \beta\left(K_t^{\rho} + L_{R,t}^{\rho}\right)^{\frac{\beta}{\rho}-1} L_{R,t}^{\rho-1} L_{C,t}^{1-\beta} \tag{3}$$

$$\omega_{C,t} = (1-\beta)\left(K_t^{\rho} + L_{R,t}^{\rho}\right)^{\frac{\beta}{\rho}} L_{C,t}^{-\beta} \tag{4}$$

where ω_R and ω_C are wages per efficiency unit in routine and complex tasks respectively.

2. Worker Heterogeneity and Occupational Sorting

The labor force consists of a unit continuum of workers $i \in [0,1]$, each endowed with a skill vector $\xi_i = [x_{R,i}, x_{C,i}]$ representing capacity in routine and complex tasks. Skills are distributed according to density $f(x_{R,i}, x_{C,i})$ with support $x_{j,i} \in [\varepsilon_j, 1+\varepsilon_j]$ for $j \in \{R, C\}$.

Each worker maximizes income by choosing among routine labor, complex labor, any convex combination of the two, or nonemployment:

$$U_i(\omega, x) = \max\{\omega_R x_{R,i}, \omega_C x_{C,i}, \omega_N\} \tag{5}$$

where ω_N is exogenous nonlabor income. This yields the sorting rule:

$$Worker\ i\ supplies \begin{cases} Rountine\ labor & if\ x_{R,i} > \bar{x}_R\ and\ x_{C,i} < x_{C,i}^* \\ Complex\ labor & if\ x_{R,i} > \bar{x}_R\ and\ x_{C,i} \geq x_{C,i}^*\ or\ x_{R,i} \leq \bar{x}_R\ and\ x_{C,i} > \bar{x}_C \\ No\ labor & if\ x_{R,i} \leq \bar{x}_R\ and\ x_{C,i} \leq \bar{x}_C \end{cases}$$

where $\bar{x}_R = \omega_N/\omega_R$, $\bar{x}_C = \omega_N/\omega_C$, and $x_{C,i}^* = (\omega_R/\omega_C)x_{R,i}$ is the indifference locus between routine and complex labor.

Aggregate labor supplies in efficiency units are:

$$L_R = \int_{\bar{x}_R}^{1+\varepsilon_R} \int_{\varepsilon_C}^{x_{C,i}^*} x_{R,i} f(x_{R,i}, x_{C,i}) dx_{C,i} dx_{R,i} \tag{6}$$

$$L_C = \int_{\varepsilon_R}^{\bar{x}_R} \int_{\bar{x}_C}^{1+\varepsilon_C} x_{C,i} f(x_{R,i}, x_{C,i}) dx_{C,i} dx_{R,i} + \int_{\bar{x}_R}^{1+\varepsilon_R} \int_{x_{C,i}^*}^{1+\epsilon_C} x_{C,i} f(x_{R,i}, x_{C,i}) dx_{C,i} dx_{R,i} \tag{7}$$

The share of workers in each employment state is:

$$N_R = \int_{\bar{x}_R}^{1+\varepsilon_R} \int_{\epsilon_C}^{x_{C,i}^*} f(x_{R,i}, x_{C,i})\, dx_{C,i} dx_{R,i},$$

$$N_C = \int_{\varepsilon_R}^{\bar{x}_R} \int_{\bar{x}_C}^{1+\varepsilon_C} f(x_{R,i}, x_{C,i})\, dx_{C,i} dx_{R,i} + \int_{\bar{x}_R}^{1+\varepsilon_R} \int_{x_{C,i}^*}^{1+\epsilon_C} f(x_{R,i}, x_{C,i})\, dx_{C,i} dx_{R,i}$$

$$N_N = \int_{\bar{x}_R}^{\bar{x}_R} \int_{\epsilon_C}^{\bar{x}_C} f(x_{R,i}, x_{C,i})\, dx_{C,i} dx_{R,i} = 1 - N_R - N_C$$

In equilibrium, wages adjust such that labor supply (equations 6 and 7) equals labor demand (equations 3 and 4).

3. Disability as a Skill Distribution Shifter

Now suppose the workforce contains two groups in equal proportions[7]: workers without disabilities ($D = 0$) and workers with disabilities ($D = 1$). Disability enters the model through three distinct channels that shift the skill distribution, reduce task mobility, and introduce accommodation costs.

Channel 1: Occupational Concentration via Skill Distribution Shifts. Historical exclusion from credential employment has concentrated workers with disabilities in routine-task-intensive

[7] For simplicity, we assume equal proportions across the two groups; in reality, approximately 23 percent of the U.S. working-age population reports a disability (Bureau of Labor Statistics, 2025).

roles (Schur, 2002). Formally, workers with disabilities have a comparative disadvantage in complex tasks, represented by a left shift in the support of the complex skill distribution:

$$x_{C,i}^{g} \in \left[\varepsilon_C^{g}, 1+\varepsilon_C^{g}\right] \quad \forall \mathrm{g} \in \{\mathrm{D}=0, \mathrm{D}=1\}$$

where $\varepsilon_C^{D=1} < \varepsilon_C^{D=0} = 0$. This mirrors the structure used by Lerch (2025) for racial/ethnic minorities and implies, by analogy with their Proposition 1, that workers with disabilities are overrepresented in routine labor and have lower baseline employment rates than non-disabled workers (the disability employment gap is positive):

$$EG^{(0,1)} = (1 - N_N^{D=0}) - (1 - N_N^{D=1}) = \int_0^{\bar{x}_C} \int_{\epsilon_C^{D=1}}^{0} f(x_{R,i}, x_{C,i})\, dx_{C,i} dx_{R,i} > 0 \tag{8}$$

Channel 2: Reduced Task Mobility. Beyond skill distribution differences, disability may constrain the ability to reallocate across the skill-task space. Define a mobility parameter $\mu^{g} \in [0,1]$ that scales the effective range of task reallocation for group g, where $\mu^{D=1} \leq \mu^{D=0} = 1$. Functional limitations — whether physical, sensory, or cognitive — compress the feasible task set, so the effective complex skill support available upon reallocation is:

$$\tilde{x}_{C,i}^{D=1} = \varepsilon_C^{D=1} + \mu^{D=1}\left(x_{C,i} - \varepsilon_C^{D=1}\right)$$

Workers with $\mu^{D=1} < 1$ cannot fully exploit the productivity effect of automation because their reachable task space is bounded below the non-disabled frontier.

Channel 3: Accommodation Costs. Retraining and job transitions involve fixed accommodation costs $\phi^{g} \geq 0$ for group g, where $\phi^{D=1} \geq \phi^{D=0} = 0$. These costs raise the effective nonlabor income threshold::

$$\tilde{\omega}_N^{D=1} = \omega_N + \phi^{D=1}$$

This shifts the nonemployment threshold upward to $\tilde{\bar{x}}_C^{D=1} = \tilde{\omega}_N^{D=1}/\omega_C > \bar{x}_C$, expanding the region of nonemployment for workers with disabilities relative to the non-disabled baseline.

4. The Three Automation Mechanisms and Their Disability Interactions

An exogenous decline in the automation price p_t activates three mechanisms simultaneously.

The Displacement Effect. As p_t falls, firms substitute capital for routine labor, reducing ω_R. Workers near the margin $\bar{x}_R$ are pushed into nonemployment. The marginal displaced worker satisfies $\omega_R x_{R,i} = \omega_N$, so the nonemployment threshold rises to $\bar{x}_R' > \bar{x}_R$. The displacement effect on employment is:

$$\left.\frac{\partial N_N}{\partial p}\right|_{disp} = -\int_{\bar{x}_R}^{\bar{x}_R'} \int_{\epsilon_C}^{x_{C,i}^{*}} f(x_{R,i}, x_{C,i})\, dx_{C,i} dx_{R,i} < 0 \tag{9}$$

For workers with disabilities, this effect is amplified through two forces. First, greater concentration in routine tasks ($\varepsilon_C^{D=1} < 0$) places more workers in the displacement zone. Second, reduced mobility $\mu^{D=1} < 1$ and elevated accommodation costs $\phi^{D=1} > 0$ prevent reallocation to complex tasks, so displaced workers flow into nonemployment rather than complex labor. The disability-specific displacement effect is therefore:

$$\left.\frac{\partial N_N^{D=1}}{\partial p}\right|_{disp} > \left.\frac{\partial N_N^{D=0}}{\partial p}\right|_{disp} \tag{10}$$

The Productivity Effect. The inflow of capital raises the productivity of complex labor, increasing ω_C. This draws workers from nonemployment and routine tasks into complex labor. The marginal entrant into complex labor satisfies $\omega_C x_{C,i} = \omega_N$, so $\bar{x}_C$ falls. The productivity effect on employment is:

$$\left.\frac{\partial N_N}{\partial p}\right|_{prod} = -\int_{\epsilon_R}^{\bar{x}_R} \int_{\bar{x}_C'}^{\bar{x}_C} f\left(x_{R,i}\, , x_{C,i}\right) dx_{C,i} dx_{R,i} > 0 \tag{11}$$

For workers with disabilities, this effect is attenuated. Reduced task mobility means the effective threshold reduction is $\tilde{\bar{x}}_C^{D=1} - \Delta\bar{x}_C \cdot \mu^{D=1}$, so fewer workers with disabilities can access newly viable complex tasks. Additionally, elevated accommodation costs $\phi^{D=1}$ maintain a higher effective reservation threshold $\widetilde{\omega}_N^{D=1}/\omega_C$, further limiting entry into complex labor:

$$\left.\frac{\partial N_N^{D=1}}{\partial p}\right|_{prod} < \left.\frac{\partial N_N^{D=0}}{\partial p}\right|_{prod} \tag{12}$$

The Reinstatement Effect. Technological progress also creates new task categories $\mathcal{T}^{\text{new}}$ with associated labor demand L^{new}, generating wage ω_T for workers able to perform them. This expands the feasible task set and may absorb displaced workers. The reinstatement effect on employment is:

$$\left.\frac{\partial N_{\mathcal{T}^{\text{new}}}}{\partial p}\right|_{reinst} = \int_{\mathcal{T}^{\text{new}}} \int_{\bar{x}_T}^{1+\varepsilon_T} f(x_{T,i}) dx_{T,i} d\,\mathcal{T} \geq 0 \tag{13}$$

Access to reinstated tasks depends on the match between $\mathcal{T}^{\text{new}}$ and the skill distribution of workers with disabilities. Denote this match quality by $\lambda^g \in [0,1]$, where $\lambda^{D=1} \leq \lambda^{D=0} = 1$ captures the degree to which newly created tasks are accessible given functional limitations. The disability-specific reinstatement effect is thus scaled by $\lambda^{D=1}$, and for workers whose limitations specifically impede new technology operation, sensory impairments interacting with digital interfaces, for instance, $\lambda^{D=1}$ may approach zero.

5. Heterogeneity by Disability Type

The parameters $(\varepsilon_C^{D=1}, \mu^{D=1}, \phi^{D=1}, \lambda^{D=1})$ vary systematically by disability type, generating a mapping from functional limitation to net employment effect.

Sensory Difficulties. Workers with visual or hearing difficulties interact with automation through all three channels. The displacement channel reduces employment as robots substitute for routine labor: as p falls, ω_R declines and workers near the margin $\bar{x}_R$ are pushed into nonemployment. The productivity channel partially restores employment as rising ω_C draws workers toward complex roles, but this offset is attenuated to the degree that sensory limitations restrict access to complex employment. The defining feature of this group is the reinstatement channel: automation-induced transitions toward digitally intensive human-machine interfaces, visual monitoring dashboards, and voice-activated systems may introduce new barriers rather than remove existing ones, so match quality is low, $\lambda^{\text{sens}} \approx 0$:

$$\left.\frac{\partial N^{sens}_{\mathcal{T}^{\text{new}}}}{\partial p}\right|_{reinst} = \lambda^{\text{sens}} \cdot \left.\frac{\partial N^{D=0}_{\mathcal{T}^{\text{new}}}}{\partial p}\right|_{reinst} \approx 0 \tag{14}$$

The net effect on employment within this group is:

$$\frac{\partial E^{\text{sens}}}{\partial \text{p}} = \underbrace{-\left.\frac{\partial N^{sens}_{\text{N}}}{\partial p}\right|_{disp}}_{\text{displacement}} + \underbrace{\left.\frac{\partial N^{\text{sens}}_{C}}{\partial p}\right|_{prod}}_{\text{attenuated productivity}} + \underbrace{0}_{\text{no reinstatement}} \tag{15}$$

With the reinstatement channel effectively closed, the productivity offset must fully counteract displacement for employment to be non-negative on net. The model therefore predicts that $\partial E^{\text{sens}} / \partial \text{p} > 0$, meaning employment falls as automation expands, and that the absence of any reinstatement offset makes this group a strong candidate for the largest adverse employment effect among the three groups considered here.

Physical and Mobility Difficulties. Workers with physical and mobility difficulties interact with automation through all three channels, with no single channel theoretically dominating. The displacement channel may be directly amplified for this group: industrial robots substitute specifically for physical and manual task content and physical limitations may have historically channeled workers into precisely these occupations, implying $\varepsilon^{\text{phys}}_{\text{C}} < 0$. The productivity channel operates as for other groups, rising ω_{C} draws workers toward complex roles, and unlike sensory-impaired workers, physical limitations do not inherently prevent access to digital or cognitive task interfaces, so this offset is less attenuated. The reinstatement channel introduces a countervailing force: when automation eliminates physically demanding tasks it may remove functional barriers that previously excluded these workers, raising match quality $\lambda^{\text{phys}} > 0$:

$$\left.\frac{\partial N^{phys}_{\mathcal{T}^{\text{new}}}}{\partial p}\right|_{reinst} = \lambda^{\text{phys}} \cdot \left.\frac{\partial N^{D=0}_{\mathcal{T}^{\text{new}}}}{\partial p}\right|_{reinst}, \ \lambda^{\text{phys}} > 0 \tag{16}$$

The net effect on employment within this group is:

$$\frac{\partial E^{\text{phys}}}{\partial \text{p}} = \underbrace{-\left.\frac{\partial N^{phys}_{\text{N}}}{\partial p}\right|_{\text{disp}}}_{\text{displacement}} + \underbrace{\left.\frac{\partial N^{phys}_{\text{C}}}{\partial p}\right|_{\text{prod}}}_{\text{productivity}} + \underbrace{\lambda^{\text{phys}} \cdot \left.\frac{\partial N^{D=0}_{\mathcal{T}^{\text{new}}}}{\partial p}\right|_{reinst}}_{\text{partial reinstatement}} \tag{17}$$

The sign of $\partial E^{\text{phys}} / \partial$pis ambiguous. If reinstatement and productivity effects are sufficiently large, plausible when automation removes physical barriers broadly, employment may rise on net. If displacement dominates, employment falls. Relative to the sensory group, the presence of a non-trivial reinstatement term ($\lambda^{\text{phys}} > \lambda^{\text{sens}} \approx 0$) and a less attenuated productivity channel suggests the adverse employment effect, if present, is likely smaller.

Cognitive Difficulties. Workers with cognitive disabilities interact with automation through all three channels, with the overall effect theoretically the most ambiguous of the three groups. The displacement channel is present but not obviously amplified: industrial robots substitute specifically for physical and manual task content, and cognitive or psychiatric limitations do not map directly onto this physical task dimension. The displacement term for this group is therefore closer to the non-disabled baseline than for physical difficulty workers, where the link between functional limitation and routine manual employment is more direct. The productivity channel operates as for other groups. The reinstatement channel introduces genuine ambiguity in both directions: match quality λ^{cog} is uncertain, as some reinstated roles requiring coordination and flexible response to novel interfaces may be demanding for this group, while others that are structured, predictable, or remote may be well-suited:

$$\left.\frac{\partial N_{\mathcal{T}^{\text{new}}}^{\text{cog}}}{\partial p}\right|_{reinst} = \lambda^{\text{cog}} \cdot \left.\frac{\partial N_{\mathcal{T}^{\text{new}}}^{D=0}}{\partial p}\right|_{reinst} \quad , \quad \lambda^{\text{cog}} \in [0,1] \tag{18}$$

The net effect on employment within this group is:

$$\frac{\partial E^{\text{cog}}}{\partial \text{p}} = \underbrace{-\left.\frac{\partial N_{\text{N}}^{\text{cog}}}{\partial p}\right|_{\text{disp}}}_{\text{near-baseline displacement}} + \underbrace{\left.\frac{\partial N_{\text{C}}^{\text{cog}}}{\partial p}\right|_{\text{prod}}}_{\text{productivity}} + \underbrace{\lambda^{\text{cog}} \cdot \left.\frac{\partial N_{\mathcal{T}^{\text{new}}}^{D=0}}{\partial p}\right|_{reinst}}_{\text{uncertain reinstatement}} \tag{19}$$

The sign of $\partial E^{\text{cog}} / \partial \text{p}$ is theoretically ambiguous. The displacement term is not strongly amplified relative to the non-disabled baseline, and the reinstatement offset while uncertain is not foreclosed as it is for sensory-impaired workers. The model therefore makes no directional prediction for this group.