# Missing Routine Work: Automation and the Life Cycle

Natalie Duncombe *

Government Accountability Office

December 5, 2024

**Abstract**

I study the effect of automation in manufacturing on workers at different stages of the life cycle. Using a difference-in-difference design and administrative data from Germany, I estimate earnings outcomes by age following an expansion in automatic machines. The analysis exploits variation in the use of routine manual tasks across occupations. Earnings losses for the youngest cohort of workers in manufacturing occupations following an expansion in automatic machines are four times as large as for the oldest cohort of workers.

* The findings and conclusions of this paper are those of the author and are not purported to represent the views of the U.S. Government Accountability Office. Data access for this project was provided by the Research Data Centre (FDZ) of the German Federal Employment Agency at the Institute for Employment Research (IAB) under project numbers fdz2140 and fdz2285.

# 1 Introduction

The expansion of automatic machines has the potential to reshape both labor markets and the nature of work. Automation has varying consequences across occupations but in general reduces the need for workers to perform routine tasks. While automatic machines offer the possibility of enhancing worker productivity, these machines can also be disruptive and lead to periods of labor market transition that may be difficult for certain workers. In this paper, I study how missing routine work in the labor market brought on by automation affects workers' earnings across the life cycle.

I focus on the expansion of automatic machines in 1980s Germany and these machines' role in replacing routine manual tasks previously performed by workers. It is well documented that older workers see larger earnings losses than young workers when forced to switch occupations. Based on this fact, we might conclude that older workers will face larger earnings declines following an expansion in automation that leads workers to reallocate across occupations. However, using administrative data from Germany and a difference-in-difference design, I show that it is young workers in highly routine occupations whose earnings decline the most after the expansion of automatic machines.

I begin by documenting evidence of how automatic machines affect earnings by age. Using a difference-in-difference analysis, I use variation in exposure at the occupation level to estimate the effect of automatic machine expansion in 1980s Germany on earnings.[1] For each occupation, I create a measure of treatment intensity that captures exposure to automatic machines using information on the shares of routine manual tasks from survey data. I then combine these occupation-level measures with worker employment histories in administrative data to analyze the labor market outcomes for workers at different stages of the life cycle. I

[1]Note that my main results are for the sample of men in Germany. In my full analysis, I will also present and discuss differences in the results for women.

find that the relative earnings decline between higher- and lower-treatment intensity workers was four times larger for the youngest cohort than for the oldest cohort. This result is surprising, given that we typically think of young workers as more resilient to labor market shocks.

There is a large empirical literature examining the labor market consequences of automation. The results from these studies vary with some finding earnings losses, particularly for lower-education workers (Akerman et al., 2015; Humlum, 2021), some finding earnings gains (Aghion et al., 2021; Hirvonen et al., 2022), and some finding no effect (Doms et al., 1997). Methodologically, by focusing on occupational exposure, my empirical strategy has elements similar to those of the strategy in Acemoglu and Restrepo (2022). They study the exposure of different demographic groups to routine tasks to measure the wage impacts of labor-replacing technologies in cross-sectional data starting in the 1980s from the United States. While they consider variation in exposure by gender and education, I document new evidence of the effect of automation on workers by age.

This paper falls within a long line of work focusing on the labor market consequences of automation. This literature has highlighted the role of automation technologies, and how they replace routine tasks, as a driver of earnings inequality and job polarization in the employment distribution (Autor et al., 2003; Acemoglu and Autor, 2011; Autor and Dorn, 2013; Goos et al., 2014).[2] Much of the literature on the labor market consequences of automation has been devoted to modeling the role of routine-biased technological change in accounting for shifts in the cross-sectional wage and employment structure (see Autor (2022) for a recent review). This paper builds on the prior literature by bringing new evidence and understanding of how automation affects workers across the life cycle.

My reduced-form exercise provides evidence of the differential consequences of automation

[2]Job polarization refers to the shift in employment away from occupations in the middle of the earnings distribution and into occupations in the upper and lower tails of the distribution.

across the life cycle. Other work analyzing the effects of automation by age is scarce. My finding that young workers saw larger relative earnings declines than older workers is similar to the findings of Dauth et al. (2021), who empirically study the labor market implications of exposure to industrial robots across local labor markets also using data from Germany. My empirical approach differs from theirs in terms of methodology and treatment measure. They use direct measures of industrial robot adoption at the industry-level and employ a shift-share approach. My empirical results may seem out of sync with the findings of another related paper Kogan et al. (2019), which emphasize that the loss of specific human capital leaves older workers facing larger consequences from technological change. However, Kogan et al. (2019) focus on highly displacing technologies and are thus more likely to select on workers forced to switch occupations.

Finally, my work brings new insights into our understanding of age and reallocation outcomes following automation or other shocks. Prior work highlights that young workers account for an important share of workers who reallocate during periods of structural change (Cociuba and MacGee, 2018). In models, it is not uncommon to consider young worker's mobility as a means by which they can more easily avoid the earnings consequences of automation (Guerreiro et al., 2022; Kogan et al., 2019). My results highlight that higher mobility is not always a guarantee that young workers will escape the *earnings* consequences of automation.

This paper proceeds as follows. Section 2 discusses the background and data and section 3 contains the reduced-form analysis. Section 4 discusses possible mechanisms. Finally, section 5 concludes.

# 2 Background and Data

## 2.1 Background

In my reduced-form analysis, I use the expansion in the use of automatic machines in the 1980s as exogenous variation that allows me to study the effect of automation. At this time, there was an acceleration in the adoption of machines that incorporated computer-related technologies. The computerization of machines improved the range of tasks that machines could accomplish and made such machines more affordable. Automatic machines in the 1980s had advances such as sensors, and personal computer-based technologies improved information storage, retrieval, and analysis. For instance, the number of operational industrial robots per 10,000 workers in Germany expanded starting from less than on at the beginning of the 1980s to around 9 per worker by the start of the 1990s IFR and UNECE (2004). While industrial robots receive a lot of attention in the discussion of automatic machines, computer numerical control machines were also expanding at this time. Automatic machines in production led to the replacement of routine tasks requiring manual dexterity (Groover et al., 1986). I exploit variation in the degree to which occupations involved these types of routine manual tasks in my analysis.

## 2.2 Data

For the reduced-form exercise, I combine occupation-level measures of exposure to automatic machines with administrative data from Germany.[3] For the worker panel, I use the Sample of Integrated Labor Market Biographies (SIAB) provided by the Research Institute of the Federal Employment Service (IAB) (Antoni et al., 2019). The SIAB is a 2-percent random

[3] I refer to the country of study here as Germany, but more specifically, during the time period that I study, it was the Federal Republic of Germany (or West Germany, as it was known colloquially).

sample of all individuals in Germany who were employed and subject to social security over my sample period. It provides information on a worker's employment history along with key demographic variables. For my occupation-level measure of exposure, I use the Qualification and Career Survey from 1979 provided by the German Federal Institute for Vocational Training (BIBB) and the IAB (BIBB and IAB, 1983). Further description of how the I use the SIAB to construct an annual panel and how variables are constructed including sample restrictions can be found in the online appendix. A comparison of the demographic composition of each data set can also be found in the online appendix.

The administrative data provides me with key information on workers' labor market outcomes. I focus on workers over the period from 1975 to 1989. Focusing on this time period allows me to analyze the consequences of the rise in automation in the 1980s.[4] I focus on workers aged 25–50 in 1979. This means that the youngest workers will have spent some time in the labor force prior to the 1980s and older workers will be less likely to retire over the sample period, making for better comparisons across age groups. My restriction to this time period also helps me avoid potential issues in the data related to the reunification of Germany in the early 1990s. The main outcome variable that I am interested in is annual earnings. I follow Dauth and Eppelsheimer (2020) and construct an annual data set from the original employment spell data. Earnings are censored above the social security contribution threshold, so I impute wages above this threshold. My main outcomes of interest is log annual earnings. Annual earnings is calculated as the average daily wage times the number of days employed in a given year. This measure includes spells of full-time and part-time employment. I do not observe hours worked.

To analyze the consequences of automation, I need a measure of exposure to the expansion

[4] Acemoglu and Restrepo (2022) use a similar cut-off date in their analysis. Although their data are from the United States, they have data prior to the 1970s, which I do not, and show that there was no significant differences between wage patterns for routine versus other workers prior to the 1980s.

in automatic machines. I create an occupational measure of exposure to automatic machines–which I refer to as *treatment intensity*–using information on occupation tasks included in the BIBB survey. To determine potential exposure in the 1980s, I use the survey from 1979. Since automation alters the composition of job tasks over time, using the 1979 survey gives me a measure of treatment intensity based on information only on tasks performed before the arrival of new technologies in the 1980s. The survey covers approximately 30,000 individuals. In the survey, respondents are asked to report from a list which activities they do in their job. There were changes to the task list in the later surveys that make cross-year analysis of occupational task content difficult. The data set does not include information about how much time workers devote to the different activities in their jobs. The survey also records what tools and machines the worker uses in her job from a set list, which I use to validate my measure.

# 3 Analysis

## 3.1 Measuring Treatment Intensity

For my difference-in-difference analysis, I exploit differences in treatment intensity across occupations. I observe 115 occupations in the SIAB. Since I do not observe measures of automatic machine adoption directly, I use information on tasks from the BIBB survey data to measure each occupation's treatment intensity.

I start by categorizing the tasks in the BIBB data. I focus on the 1979 survey, which records the use of 84 distinct tasks ranging from conducting research to operating machines to performing care activities. In the data, workers report from a list which tasks they perform. I first categorize which of these tasks can be considered routine manual, nonroutine manual, routine cognitive, nonroutine cognitive, and nonroutine interpersonal. I then construct the

Figure 1: Measures of Treatment Intensity by Occupation and Gender

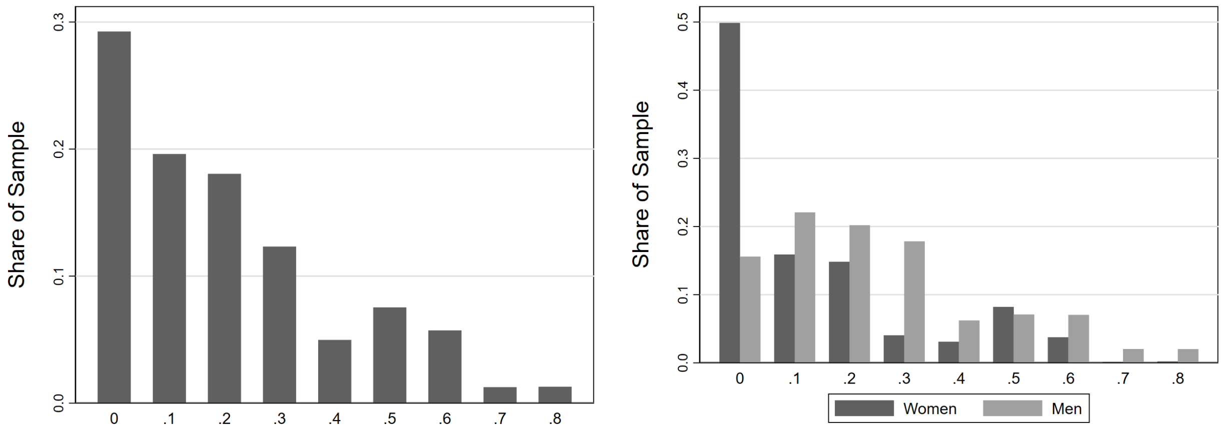


(a) Distribution of Treatment Intensity (b) Distribution of Treatment Intensity by Gender

Figure (a) plots the distribution of employment over by each occupation's treatment intensity measure ($D_j$) Figure (b) reports the same by gender. Source: SIAB and BIBB.

measure of treatment intensity. I follow the work of Spitz-Oener (2006) and Mihaylov and Tijdens (2019) to categorize the tasks. The online appendix contains the full mapping of tasks to categories. Respondents can choose multiple tasks, and the median number of reported tasks is 3. Since many workers report only a small numbers of tasks, I aggregate the tasks to the occupational level to construct the measures of treatment intensity.

I define the treatment intensity in each occupation as the share of routine manual (RM) tasks performed by workers in each occupation:

$$D_j = \frac{\text{\# of RM tasks performed by workers in occupation j}}{\text{total \# of tasks performed by workers in occupation j}}$$

Figure 1a plots the employment shares across the measures of treatment intensity. It is interesting to note that we do observe that men and women differ in their average exposure because of the difference in their occupational choices. Figure 1b plots the employment shares by gender across routine manual tasks shares. We see that women are less represented than men in occupations with high shares of routine manual tasks. These differences in exposure lead me to separately consider outcomes by gender as well as age.

The occupations with the highest treatment intensity tend to be production occupations, such as drillers, metal producers, welders, and sheet metal pressers. However, all occupations in my analysis use routine manual tasks to some degree. The online appendix reports the treatment intensity value by occupation. That all occupations use routine manual tasks motivates my decision to use a difference-in-difference design with a continuous treatment. To validate the use of these measures to account for the expansion of automatic machines, I analyze how well they correspond with the future use of automatic machines. The online appendix shows how the shares of routine manual tasks correspond with higher eventual adoption of industrial robots by workers in these occupations. I also perform a validation exercise where I regress my measure of treatment intensity on the level of adoption of industrial robots by the 1990s, following Autor et al. (2003).

### 3.2 Difference-in-Difference Specification

I exploit the variation in treatment intensity as measured by routine manual task shares to analyze the effect of increased automation on earnings for workers across occupations. To do this, I adopt a DID-with-continuous-treatment strategy considering the 1980s an unanticipated event from the perspective of workers in the 1970s. I designate a worker's treatment group by her 1979 occupation. The DID design compares outcomes of workers across treatment occupation groups before and after 1980. To analyze consequences by age, I divide workers into 5-year bins based on their year of birth. I restrict the sample so that the oldest workers were 50 and the youngest 25 in 1980. This is to avoid exit due to early retirement, which in Germany could effectively start at age 60 due to its generous unemployment system during this time period. After implementing these restrictions on my sample, I am left with 3,149,817 person-year observations.

In my main analysis, I estimate a difference-in-difference model with continuous treatment using a two-way fixed effect (TWFE) specification. The main estimated coefficient from this specification is typically interpreted as the average causal response (ACR). That is, the estimated coefficient is taken to represent the causal effect of a unit change in treatment intensity. However, recent work by Callaway et al. (2021), shows that the estimated coefficient is equivalent to the ACR only when a strong parallel trends assumption holds; otherwise, the estimate is susceptible to what they call "selection bias." In my case, for the strong parallel trends assumption to hold, workers in both less routine and more routine occupations would have to have seen the same evolution of outcomes as they would have had they been in the less routine occupation. As Callaway et al. (2021) discuss, there is no test for the strong parallel trend assumption, so I am left to test for "normal" parallel trends. In this case, the parallel trends test examines whether the lower-dose units have the same evolution of untreated potential outcomes as that of higher-dose units. Before running my main specification, I focus on this test.

To compare outcomes across age groups, I split the sample into five-year cohort bins. I run the specification separately for each cohort and for men and women. This is equivalent to a specification where I interact cohort and gender with each variable. I estimate a DID with a two-way fixed effects regression:

$$Y_{it} = \beta D_{j(i)} \cdot Post_t + x'_{it}\gamma + \alpha_i + \theta_t + \varepsilon_{it} \tag{1}$$

In the specification, $Y_{it}$ is the log annual earnings of worker $i$ in year $t$. $D_j$ is the measure of the worker's treatment intensity determined by her occupation $j(i)$ in 1979. $Post_t$ is a dummy for the post-treatment period, which is 1980–1989. The variable $x_{it}$ is a vector of individual-level controls including experience, experience squared, and education. Finally, $\alpha_i$ is individual fixed-effects, which will absorb the treatment group fixed-effects, and $\theta_t$ is

Figure 2: Estimated Average Causal Response by Age and Gender

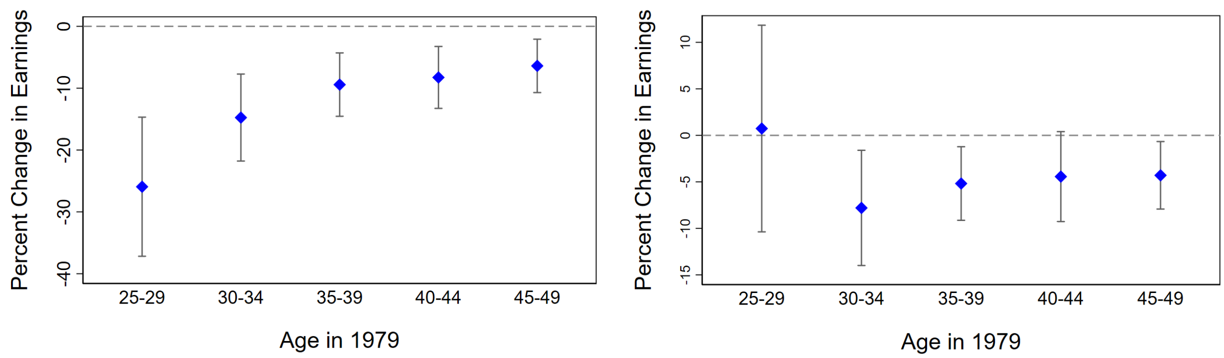


(a) Men

(b) Women

Panel (a) plots the estimated ACR as a percent change in earnings ($\beta \times 100$) for men by age from the TWFE specification in equation (1) with log annual earnings as the outcome variable and the routine manual task share as the measure of treatment intensity. The ACR is the effect of a one unit change in treatment intensity. The lines represent the 95-percent confidence intervals. Panel (b) plots the same for women.

year fixed effects. I cluster standard errors at the treatment occupation level.

I start by reporting the results for men. Figure 2 reports estimates of $\beta$ for each cohort of men, which–if the proper assumptions hold–we can consider the ACR. The units of the coefficient when multiplied by 100 give us the ACR as a percentage of annual earnings, which is what I plot in Figure 2. We see differences in outcomes by age when comparing the results for the oldest and youngest cohorts. A one-standard-deviation increase in treatment intensity equals a 19 percentage point increase in the routine manual task share. This is approximately similar to a worker's moving from a being a locksmith (routine manual share of 52 percent) to a welder (routine manual share of 71 percent). The point estimate implies that a one-standard-deviation increase in treatment intensity leads to a 4-percent relative decline in annual earnings for the youngest cohort of workers (25–30) and closer to a 1-percent relative decline for the oldest cohort (45–50).

I next consider the patterns by age for women. Figure 2b reports the results. For the oldest four cohorts, women see losses with a similar pattern to those of men, but the magnitude of the losses is more muted. However, we do not see the same declines in earnings for the youngest cohort of women that we do for the youngest cohort of men. This result could reflect that, for a given dose, men and women have different responses, or it may reflect other differences in how men and women adjust to following exposure to automation.

These results could also reflect limitations of the data or selection into the sample. Women make up a smaller portion of the sample and work more in part-time or marginal positions. For approximately 20 percent of women, the primary employment spell reported corresponds to a part-time position, while the same is true for only 1 percent of men. The reporting requirements for these marginal positions varied over my treatment period, making the observations for women more subject to selection bias. Young women, around the age of having children, may have even less labor force attachment. Further analysis with better data would be required to completely tease out the heterogeneity in the consequences of automation across gender.

## 3.3 Validity

I test for parallel trends by means of an event study. This allows me to examine how well 1980 captures the beginning of the treatment period. I run the following specification:

$$Y_{it} = \sum_{t=1975, t\neq 1979}^{1989} \beta^t \cdot D^t_{j(i)} + x'_{it}\gamma + \alpha_i + \theta_t + \varepsilon_{it} \tag{2}$$

In the specification, $Y_{it}$ is log annual earnings for worker $i$ in year $t$. $D^t_{j(i)}$ is the measure of the worker's treatment intensity determined by her occupation $j(i)$ in 1979.[5] This measure is equal to zero prior to 1980, and all results are in reference to 1979. The variable $x_{it}$ is a vector

[5]For robustness, I also run a similar specification designating treatment shares based on each worker's occuaption in 1975 and get similar results.

Figure 3: Estimated Average Causal Responses from the Event Study by Gender

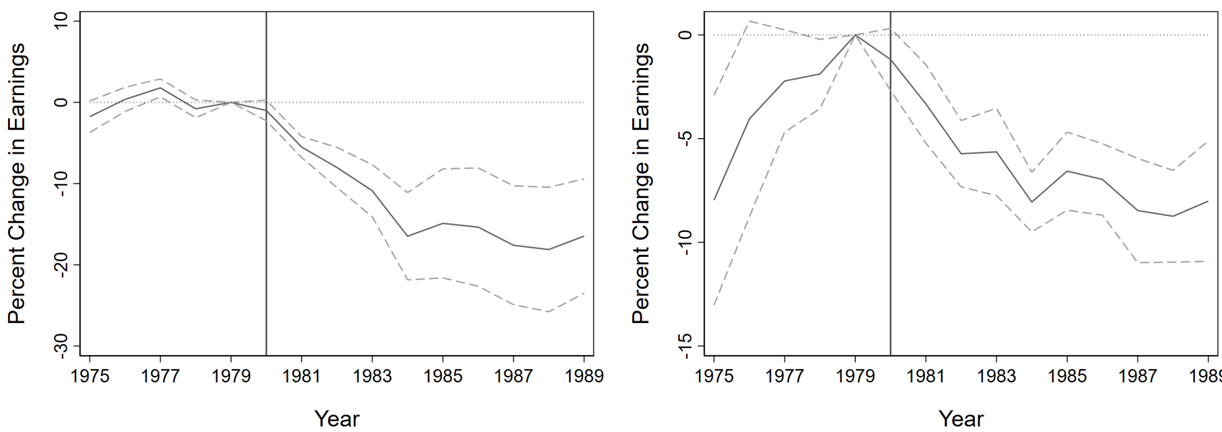


(a) Men

(b) Women

These figures plot the estimated ACR as a percent change in earnings ($\beta \times 100$) from the event study specification in equation (2) with log earnings as the outcome variable and the routine manual task share as the measure of treatment intensity. The ACR is the effect of a one unit change in treatment intensity. The dotted lines represent the 95-percent confidence intervals. Panel (a) reports the results for men, and panel (b) reports the results for women.

of individual-level controls including experience, experience squared, and education. Finally, $\alpha_i$ is individual fixed-effects, which will absorb the treatment group fixed-effects based on the worker's occupation in 1979, and $\theta_t$ is year fixed effects. I also run the specification separately for men and women. This is equivalent to a specification where I interact gender with each variable. I cluster standard errors at the treatment occupation level.

Figure 3 plots the results of the event study separately for men and women with routine manual tasks as the measure of treatment intensity. The units of the coefficient when multiplied by 100 give us the ACR as a percent of annual earnings, which I plot in Figure 3a. We see that, for men, 1980 is a good proxy for the start of the treatment period for exposure to automation of routine manual tasks. In Figure 3b, the estimated coefficients prior to 1980 suggest that women in highly exposed occupations were already receiving lower earnings than workers in less exposed occupations. This informs my decisions to consider men and women separately in my analysis. We also see that the magnitude of losses for men

Figure 4: Analysis of Mobility Outcomes

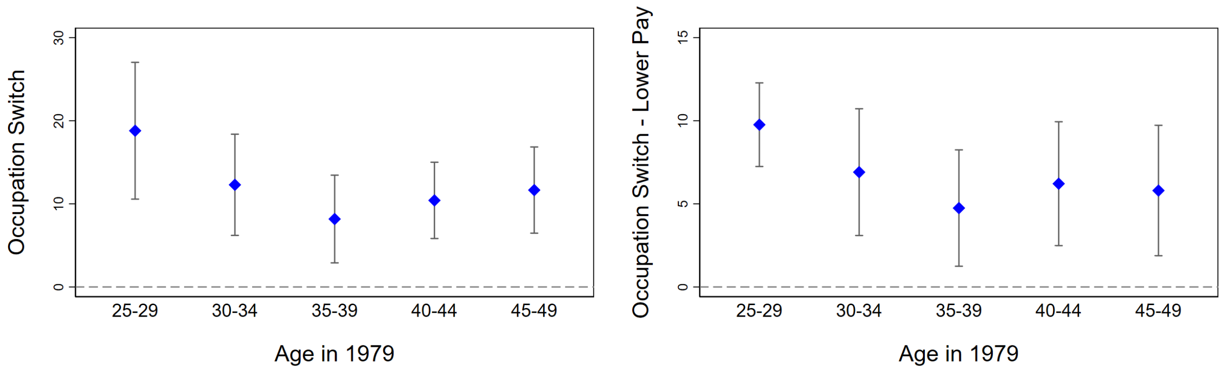


(a) Leave 1979 occupation (b) Move to lower paying occupation

Panel (a) plots the estimated ACR ($\beta \times 100$) for men from the TWFE specification in equation (1) with an indicator equal to one if the worker is in an occupation different from the one he was in in 1979 and the routine manual task share as the measure of treatment intensity. The lines represent the 95-percent confidence intervals. Panel (b) plots the estimated ACR for men from the TWFE specification in equation (1) with an indicator equal to one if the worker is in an occupation with a lower average wage than the one he was in in 1979 and the routine manual task share as the measure of treatment intensity. The lines represent the 95-percent confidence intervals. Source: SIAB and BIBB.

is approximately twice as large as for women at its peak.

# 4 Possible Mechanisms

## 4.1 Mobility

Macroeconomic models of automation often highlight the importance of worker mobility for "escaping" the consequences of automation. Thus, it is of additional interest to understand whether and where workers move following technological change. For this analysis, I focus on the sample of men.

First, I consider the effect of technological change on occupational moves. I create an indicator equal to one when the worker is in an occupation that differs from his occupation in 1979. I then run the same TWFE regression to estimate the DID. Figure 4a reports the resulting coefficients across cohorts. Workers in routine manual occupations do have higher rates of occupational mobility across all cohorts, with slightly higher occupational exits from the youngest cohort.

I can also observe some characteristics of the occupations that workers move to. I first rank occupations based on the average earnings of workers in the occupation in 1979. I then create an indicator for when a worker's current occupation was ranked lower in terms of average earnings from the occupation that she was in when I designate the treatment in 1979. Running the main specification with this indicator as the outcome, as reported in Figure 4b, I see that highly routine manual workers tended to move into lower-paying occupations than those that they started in. While these moves to lower-paying occupations could be a reason for the higher earnings losses of young workers, they do not explain why young workers move to these lower-paying occupations to begin with.

## 4.2 Exposed Tasks by Age

I first return to the BIBB survey data and consider variation in the types of tasks done by workers across age groups within occupations. If young workers perform more of the tasks performed by automatic machines this could be one reason for young workers seeing higher earnings losses. It could suggest that they are closer substitutes with automatic machines in production. One question in the BIBB asks "How often does your daily work require good manual dexterity?" and another asks workers to report the type of industrial machine that they report using at work. In the online appedix, I show the responses to the question on manual dexterity for workers using hand-operated machines (e.g., lathes), semi-automatic machines, and fully automated machines (e.g., industrial robots). In the data, the share of workers reporting that they "frequently" or "almost always" require manual dexterity in

Table 1: Manual Dexterity Requirements and Age

| | (1) | (2) | (3) |
|---|---|---|---|
| Age | -0.0022*** | -0.0014*** | -0.00087** |
| | (0.00026) | (0.00024) | (0.00026) |
| Constant | 0.55*** | 0.53*** | 0.48*** |
| | (0.010) | (0.0093) | (0.012) |
| Occ. FE | No | No | Yes |
| Controls | No | Yes | Yes |
| Adj. R2 | 0.0024 | 0.22 | 0.24 |
| N | 27,359 | 27,359 | 23,666 |

Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table reports the coefficients of regressions of manual dexterity across age. The outcome variable is an indicator for whether the worker reports "almost always" or "frequently" requiring manual dexterity in her job. Column (1) reports the results without occupation fixed effects and worker controls, column (2) adds worker controls, and column (3) adds both worker controls and occupation fixed effects.

their job declines as the level of automation of the industrial machine increases.

Next, I analyze how these affected tasks vary across age groups. To do this, I create an indicator variable for workers who report "almost always" or "frequently" requiring manual dexterity in their job. I then run a simple ordinary least squares (OLS) regression of these indicators on age and control for occupation, gender, and education.

The results are reported in Table 1. The results show that the share of workers reporting that they require high manual dexterity in their jobs declines with age. That older workers are in jobs with lower manual dexterity requirements is not entirely surprising given that manual strength can decline with age. This could suggest that one reason for young workers seeing higher earnings losses is because they are closer substitutes with automatic machines in production. However, the magnitudes of the estimated coefficients are relatively small. Even without controls or fixed effects, the estimated coefficient suggests that, with 10 additional years of age, the share of workers reporting being in high-dexterity jobs declines by only 2 percentage points.

## 4.3 Institutions

Germany has certain labor market institutions that may play a role in determining the differential earnings responses to automation by age. I focus on the role of collective bargaining and early retirement incentives.

Collective bargaining can prioritize incumbent workers, who tend to be older, over outsiders. I draw from Abraham and Houseman (1993) for background on the German labor market features described in this section. Wages for many workers in Germany fall under collective bargaining agreements. In the 1980s, approximately 40 percent of wage and salary earners in Germany belonged to a union. Industry-level agreements determine minimum wages in the respective industries. In the 1980s, when Germany's unemployment rate was elevated, some of these wages were considered "too high" and were cited as a potential reason for longer unemployment duration. In the 1980s, hours were also an issue addressed by collective bargaining. Collective bargaining agreements apply to all workers in a sector.

To test whether collective bargaining affected the response to technological change, I use data on German union coverage provided by the Institute of Economic and Social Research (WSI).[6] I test the importance of unions by comparing earnings and wage outcomes when I interact occupational exposure with union density. For union density, I create an indicator for industries that have above-median union density and interact that with the continuous treatment variable. Tables in the online appendix report the results. When interacted with the routine manual treatment, workers in sectors with higher union density do not have statistically different outcomes from those in low-density sectors. The data on union coverage are quite coarse, but they do not seem to indicate that unions are strong enough to drive the results across age groups.

During the 1980s, Germany implemented several plans to reduce excess unemployment

[6] I thank Tommaso Porzio for pointing me to these data.

via early retirement schemes that started at age 60. The duration of unemployment benefits for older workers was also expanded during this period, effectively allowing a worker starting at age 57 to use unemployment benefits as a bridge to early retirement.

Early retirement would affect my analysis if it led to selection into employment for the older workers whom I observe. Once workers retire, I do not observe them in my sample, although I do continue to observe individuals collecting unemployment benefits. I choose the dates and cohorts in my analysis such that workers over 60 are excluded from the sample. In the online appendix, I report the levels of attrition by cohort in the sample. I measure attrition based on the share of workers whom I observe in 1979 and also observe at the end of my sample period in 1989. I do see slightly higher attrition for the oldest cohort, which may reflect these early retirements. Attrition for the other cohorts is similar those, so I conclude that early retirement is not a main contributor to my differential results by age, which are driven by large losses for the youngest cohort.

# 5 Conclusion

In this paper, I examine how the automation of routine manual tasks affects workers across the life cycle. Using administrative data from Germany, I exploit variation in exposure to automatic machines across occupations to estimate the effect of automation on earnings by age. I find important new evidence that it is the earnings of young workers that are the most negatively affected by machines. These losses follow young workers even when they escape high-exposure occupations.

My work shows, that despite the increased mobility of younger workers, they may still be unable to escape the negative effects of automation. The most promising evidence suggests that young workers were have been more vulnerable to the expansion of automatic machines because they were called on more to perform manual tasks requiring strength. Institutional characteristics present in Germany such as strong unions did not appear to play a role. However, future work should explore in greater depth why young workers faced greater consequences following automation.

The age distribution of earnings losses following automation has important implications. As new automation technologies such as generative AI advance, there is early evidence that these jobs may also affects entry-level occupations.My work show, that these negative affects can have long lasting earnings consequences for early career workers with earnings not recovering even ten years after the introduction of new automation technologies. Policymakers may need to take more drastic actions when developing targeted education and training programs to help young workers adapt to the evolving job market.

# References

Abraham, K. and S. Houseman (1993). Job Security in America: Lessons from Germany. Books from Upjohn Press, W.E. Upjohn Institute for Employment Research.

Acemoglu, D. and D. Autor (2011). Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Elsevier.

Acemoglu, D. and P. Restrepo (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica 90*(5), 1973–2016.

Aghion, P., C. Antonin, S. Bunel, and X. Jaravel (2021). What Are the Labor and Product Market Effects of Automation? New Evidence from France.

Akerman, A., I. Gaarder, and M. Mogstad (2015). The Skill Complementarity of Broadband Internet. *The Quarterly Journal of Economics 130*(4), 1781–1824.

Antoni, M., P. v. Berge, and A. Ganzer (2019). Factually anonymous Version of the Sample of Integrated Labour Market Biographies (SIAB-Regionalfile) - Version 7517 v1.

Autor, D. (2022). The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty.

Autor, D. H. and D. Dorn (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review 103*(5), 1553–1597.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

BIBB and IAB (1983). Employment Survey 1979. Collected by the Federal Institute for Vocational Education and Training (BIBB) and the Institute for Employment Research (IAB). Published by GESIS Data Archive, Köln. ZA1243 Data Version 1.0.0, https://doi.org/10.4232/1.1243.

Callaway, B., A. Goodman-Bacon, and P. H. C. Sant'Anna (2021). Difference-in-Differences with a Continuous Treatment.

Cociuba, S. E. and J. MacGee (2018). Demographics and sectoral reallocations: A search theory with immobile workers.

Dauth, W. and J. Eppelsheimer (2020). Preparing the sample of integrated labour market biographies (SIAB) for scientific analysis: a guide. *Journal for Labour Market Research 54*(1), 10.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2021, December). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association 19*(6), 3104–3153.

Doms, M., T. Dunne, and K. R. Troske (1997). Workers, Wages, and Technology. *The Quarterly Journal of Economics 112*(1), 253–290.

Goos, M., A. Manning, and A. Salomons (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review 104*(8), 2509–2526.

Groover, M., M. Weiss, R. Nagel, and N. Odrey (1986). *Industrial Robotics: Technology, Programming, and Applications.* McGraw-Hill Book Company.

Guerreiro, J., S. Rebelo, and P. Teles (2022). Should Robots Be Taxed? *The Review of Economic Studies 89*(1), 279–311.

Hirvonen, J., A. Stenhammar, and J. Tuhkuri (2022). New Evidence on the Effect of Technology on Employment and Skill Demand.

Humlum, A. (2021). Robot Adoption and Labor Market Dynamics.

IFR and UNECE (2004). *World Robotics 2004: Statistics, Market Analysis, Forecasts, Case Studies, and Profitability of Robot Investment.* United Nations.

Kogan, L., D. Papanikolaou, L. Schmidt, and B. Seegmiller (2019). Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations.

Mihaylov, E. and K. G. Tijdens (2019). Measuring the Routine and Non-Routine Task Content of 427 Four-Digit ISCO-08 Occupations.

Spitz-Oener, A. (2006). Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics 24*(2), 235–270.