Routine-Biased Technological Change and Endogenous Skill Investments*

Danyelle Branco Bladimir Carrillo Wilman Iglesias

May 9, 2023

Abstract

We investigate how individuals alter their educational investments in response to routine-biased technology. We find that individuals growing up in robot-impacted areas are more likely to complete a bachelor's degree and experience a relative increase in earnings. Changes in the skill premium and opportunity cost appear to drive these effects. To interpret these findings, we estimate a model of endogenous skill acquisition where changes in the demand and supply of skills shape the path of earnings. Counterfactual simulations suggest that endogenous human capital accumulation cannot undo most of the earnings effects of automation unless there are sufficiently generous educational subsidies.

JEL codes: I21, J23, J24

Keywords: automation; industrial robots; skill acquisition; college attainment; labor market

*Contact information: Branco: Sao Paulo School of Economics-FGV, Rua Itapeva 474 Sao Paulo-SP, 01332-000 Brazil (e-mail: danyellebranco@gmail.com). Carrillo: Sao Paulo School of Economics-FGV, Rua Itapeva 474 Sao Paulo-SP, 01332-000 Brazil (e-mail: blacaber8@hotmail.com). Iglesias: University of Nebraska-Lincoln, Lincoln, NE 68588 (e-mail: wilman.iglesias@huskers.unl.edu). We are especially appreciative to Pascual Restrepo for detailed comments that greatly improved the paper, particularly the structural analysis. We also thank Paulo Alvarate, Daniel Araujo, David Autor, Martin Beraja, Diogo Britto, Carlos Charris, Bruno Ferman, Kirabo Jackson, Renata Narita, Mathew Notowidigdo, Breno Sampaio, Bryan Stuart, Daniel Tannenbaum, and participants at various conferences and seminars for helpful comments and suggestions. We are grateful to Bryan Stuart for kindly sharing his data on the exposure to the 1980-82 recession across counties. We are solely responsible for this paper's contents. This study was financed in part by the Universidade Federal de Pernambuco and *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* —Brazil (CAPES) —Finance Code 001. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

How do agents adjust to automation technologies? Despite an extensive literature examining the labor market consequences of automation, little is known about how individuals alter their skill investments in response to these technological innovations. Answers to this question have important policy implications. If individuals endogenously accumulate more human capital in response to new technologies that lower the opportunity cost of and raise the returns to skill acquisition, then concerns over educational policies aimed at mitigating the disruptive effects of automation may be overstated. But if there are factors that limit this skill response, such as financial frictions and lack of academic readiness, then policies that remove such frictions or improve the learning process could be desirable. In this paper, we exploit the unprecedented penetration of robotics technology in the United States to provide detailed empirical evidence on this important question.

While previous research has undoubtedly advanced our understanding of the relationship between labor market conditions and human capital, existing studies have often focused on aggregate shocks that affect educational incentives primarily at the bottom of the skill distribution, including shocks stemming from the construction industry (Charles et al., 2018), agricultural sector (Shah and Steinberg, 2017; Carrillo, 2020), trade (Atkin, 2016; Greenland and Lopresti, 2016) as well as from immigration and offshoring (Hickman and Olney, 2011). What is different about the recent advances in automation technologies is that they tend to disproportionately affect routine-intensive occupations that are toward the middle range of the skill distribution (Goos and Manning, 2007; Autor and Dorn, 2013). It is not obvious that the implications of previous studies are immediately applicable to these settings, where additional skill investments in the form of a bachelor's-level education are more expensive, require a more complex set of skills, and take a significant period of training. Consistent with this higher bar, a quick look the data reveals that while 88 percent of individuals over 25 years of age had a high school diploma in 2015 in the United States, only 32 percent had completed a bachelor's degree.[1] Youths forgoing college education may be just those who are credit-constrained (Lovenheim, 2011), too impatient (Cadena and Keys, 2015; Lavecchia et al., 2016), or lack the foundational skills to succeed in college (Goldin and Katz, 2009). Recent work by Athreya and Eberly (2021) concisely highlights the importance of the latter and argue that:

> *"In the absence of improved college readiness .... the continuing long-standing trends in skill-biased technological change can be expected primarily to increase earnings inequality rather than college attainment."*

Whether or not routine-biased technology induces skill acquisition therefore remains an open question that we address in this paper. Our analysis focuses on industrial robots, which are reprogrammable machines that can perform a variety of routine tasks, ranging from painting and assembly to packaging, without requiring any human operator. With the incorporation of sophisticated

[1] See https://nces.ed.gov/programs/digest/d15/tables/dt15_104.10.asp, last accessed on May 2, 2023.

sensor and machine vision systems, robotics technology advanced dramatically in the 1990s. Following the substantial decline in the price of an industrial robot, there was a sharp and discontinuous rise in robot adoption rates since 1993 in the United States, with an increase of 120 percent from 1992 to 1995 alone and 200 percent to the end of the 1990s (Figure 1). This was in contrast to a relatively flat trend in adoption rates in previous years. Acemoglu and Restrepo (2020) document that this unexpected, sudden, and salient technological shock had negative effects on the earnings of routine-intensive workers. We investigate the consequences of this technological shock for the college decisions of individuals growing up in impacted labor markets.

The paper proceeds in three steps. We first characterize the impacts of robots on college attainment. Our research design exploits variation in the intensity of robot penetration across locations and the timing of birth cohort exposure in a difference-in-differences empirical strategy. We construct a Bartik-like measure of exposure to robots based on pre-existing industry mix across locations and industry-specific robot penetration, following the neat approach developed by Acemoglu and Restrepo (2020) and recent methodological advances in shift-share designs (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022; Adao et al., 2019). We assign individuals to robot exposure intensities based on their state of birth, assuming that the state where an individual was born is the same as the one where she or he grew up in an intent-to-treat strategy. We show that this assignment is reasonable and that there is a great deal of variation in exposure intensities across states.[2] We then compare the outcomes of cohorts exposed before, during, and after their critical college-going ages in states with varying robot penetration intensities. Under the common trends assumption that more- and less-exposed areas would have followed similar trends over time across birth cohorts in the absence of the robot shock, our estimates can be given a causal interpretation.

We find a visually clear and statistically significant increase in the likelihood of having a bachelor's degree in areas housing the industries with greater robot penetration. Higher-versus-lower exposed areas exhibit statistically identical trends for approximately 50 years and begin to diverge only when new birth cohorts exposed to the robot shock before or during the typical college-going ages enter the economy. Our estimates imply that the change in the bachelor's degree completion rate from the older- to younger-exposed cohorts in a highly affected state like Ohio was 1.5 percentage points (or 5 percent) more positive compared to the change between the same cohorts in a more mildly affected state like Montana. This effect comes entirely from individuals who otherwise would have completed exactly high school or attended a two-year college, or those on the relevant margin in the middle of the skill distribution.

We document extensively that our estimates are very unlikely to be capturing mean-reverting dynamics, differential trends in manufacturing employment, or differences in trends related to a diverse set of initial socioeconomic and demographic characteristics. We also show that our estimated effects are not confounded by other major shocks to the labor market, such as offshoring, import competition, other technological shocks, and the recession of the early 1980s, or by major

[2] As we discuss in Section 2 in more detail, approximately 80 percent of individuals reside in the same state as the one where they were born during their college-going ages.

social programs, such as school finance reforms, war-on-poverty programs, and Medicaid. Finally, we show that the results are robust to the specification tests recommended by recent contributions in the literature on shift-share designs, such as excluding industries with the largest Rotemberg weights (Goldsmith-Pinkham et al., 2020) and considering inference procedures that account for spatial correlation across areas with similar sectoral shares (Borusyak et al., 2022; Adao et al., 2019).

We then look at the childhood-exposure effects on labor market earnings. The data indicate that cohorts exposed to robots at the beginning of the life cycle experienced an increase (or a smaller decline) in their incomes *relative* to late-exposed cohorts. This income effect disappears entirely once we account for the relationship between robots and college attainment, suggesting that education is the driving force behind this income effect. It is important to note that these results *do not* imply that automation is good on net for younger cohorts. The introduction of robots had negative labor market impacts on all individuals, but this effect has been smaller for younger cohorts who could alter their educational decisions.[3]

In the second part of the paper, we provide evidence of the mechanisms underpinning our findings. We estimate that labor markets with greater exposure to robots saw a rise in the premium from having a bachelor's degree relative to a two-year college and high-school degree. This premium effect has been paralleled fairly closely by a meaningful decline in the opportunity cost of college-going, as proxied by the average labor market earnings a young adult without college training receives. The magnitude of this effect is particularly large: the state in the 75th percentile of the exposure to robot distribution experienced a decline of approximately 17 percent in the labor market income a young adult worker receives. We rule out alternative explanations related to the supply-side of education and local government responses, including changes in the net cost of colleges, college revenues from public appropriation, or government expenditures in education.[4]

We interpret these findings as robust evidence that individuals endogenously redirect their human capital investments toward areas that are less susceptible to automation. While this evidence is elucidating, the magnitude of this reduced-form result is not straightforward to interpret. As such, our reduced-form analysis leaves many questions unanswered about the adjustment of the economy to changes in technology. For example, it says little about whether the endogenous educational response to the robot-induced shock is of the right order of magnitude to significantly counterbalance the disruptive effects of technology, or about how policy could influence this process. To investigate these questions, we structurally estimate a simple model of human capital investments where changes in the demand and supply of skills shape the long-run evolution of earnings. We

[3]Late cohorts were the ones feeling the bulk of the displacement effects created by robots. This raises the concern that our results may be driven by biases due to older cohorts experiencing the scarring effects from job losses they incurred during their earlier working life. As we shall see, our results hold (and become stronger) even when we compare adults in the labor market but that grew up in places with varying degrees of exposure to robots, suggesting that biases due to scarring effects are unlikely to be a major issue.

[4]These results, however, do not rule out the possibility that policymakers did respond to the adoption of robotics technologies but at the national level, an effect that would not be identified by our cross-location empirical strategy. But this does not necessarily affect the interpretation of our results and parameter of interest.

view this structural analysis as exploratory in nature but useful because it provides an intuitive way to interpret our findings while highlighting avenues for future areas of investigation.

In the final part of the paper, we introduce the structural model, discuss the identification conditions and perform counterfactual exercises. In the model, individuals are heterogeneous in terms of initial income, the childhood market of residence, and taste for college attendance. Individuals decide whether or not to attend college based on college costs and expected lifetime income, with the latter partially influenced by automation. We estimate the model using the simulated method of moments and document that it is successful in replicating our basic reduced-form results as well as the baseline share of individuals with a bachelor's degree. Moreover, the model predicts effects of changes in the skill returns on bachelor attainment that are quantitatively and qualitatively in line with recent quasi-experimental evidence (Abramitzky et al., Forthcoming).

Our simulations suggest that the endogenous educational response is not of the right order of magnitude to significantly offset the long-run decline in earnings induced by changes in technology. Indeed, this mechanism can reduce the adverse earnings effect of robots by only 33 percent. We then conduct a series of policy counterfactuals to explore the role of subsidies in enhancing the college response to technology. We find that a reform that increases the coverage and value of college grants can offset the earnings effects of technology by approximately 92 percent over the long term. Taken together, these counterfactual exercises suggest that endogenous human capital accumulation is unlikely to undo the adverse effects of automation on labor markets unless there are sufficiently generous educational subsidies.

These results naturally raise the question of why there has been little progress in the aggregate trends of college attainment for cohorts entering the labor market after the 1980s despite the rapid adoption of different routine-biased technologies. We believe that the most plausible explanation is that there have been changes in other important factors offsetting the skill-acquisition incentives brought by these new technologies. This point has already been highlighted by Goldin and Katz (2009) and quantitatively analyzed by Castro and Coen-Pirani (2016), who demonstrate that the sharp rise in tuition costs faced by recent cohorts can explain a substantial portion of the slowdown in aggregate college attainment.[5]Note that this does not imply that the human capital response to technology has not been important. Our analysis suggests that college attainment would have likely increased at a slower rate or even declined in the absence of endogenous skill investments.

Our findings contribute to a vast literature on the impacts of technology. This literature has documented extensively that routine-biased technologies adversely affect the demand for unskilled labor (see Jaimovich and Siu (2019) for an overview of the literature).[6]We contribute by providing

[5]The important role of rising tuition costs is consistent with recent experimental evidence documenting that financial aid, which reduces the costs of college attendance, has a fairly large causal effect on bachelor's degree attainment (Angrist et al., Forthcoming). Castro and Coen-Pirani (2016) also explore the role of declining learning ability and find that it also accounts for an important fraction of the slowdown in college attainment. Other studies employing an analogous approach reach similar conclusions (Jones and Yang, 2016; Donovan and Herrington, 2019).

[6]Pioneering studies in this literature include Katz and Murphy (1992), Krueger (1993), and Autor et al. (1998). Subsequent work provides more detailed evidence on the role of computers (Burstein et al., 2019), information and communication technologies (Michaels et al., 2014; Akerman et al., 2015; Hjort and Poulsen, 2019), industrial robots (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Dauth et al., 2021) and artificial intelligence and machine

evidence on whether, and how these technologies induce skill acquisition. While education has often been emphasized as an important factor mitigating the consequences of labor-displacing technologies (Katz and Murphy, 1992; Goldin and Katz, 2009), there has been little effort to estimate the key parameters governing this endogenous response. The parameters we estimate can serve as inputs to discipline models of the economy that consider changes in skill-biased technologies and endogenize education.[7]

This paper also adds to a large body of work linking aggregate economic conditions and human capital accumulation. As discussed above, the bulk of this literature have focused on labor market shocks affecting educational incentives among individuals in the lower end of the skills spectrum. These studies unsurprisingly tend to find no significant effects on bachelor attainment (e.g., Hickman and Olney, 2011; Charles et al., 2018). A recent strand of this literature explores how students alter their college major choices when exposed to major-related economic shocks (Han and Winters, 2020; Weinstein, 2022) and during the business cycle and the Great Recession (Ersoy, 2020; Blom et al., 2021). These studies therefore focus on the intensive margin of college decisions, among individuals who have already overcome non-trivial barriers associated with attending college. These estimates are not therefore necessarily generalizable to college enrollment decisions. We contribute to the literature by providing one of the first pieces of evidence on how the extensive margin of bachelor's-level college education reacts to a major labor market shock altering its net returns.

Finally, our study is most closely related to Di Giacomo and Lerch (2022), who estimate how college enrollment changes in commuting zones that were more and less exposed to robots. A major difference between our analysis and theirs is that we study the impacts of robots on *individuals* rather than on *places*. This is important because the latter may reflect compositional effects due to endogenous migration responses, and the evidence suggests that such migration effects are salient in this context (Acemoglu and Restrepo, 2020; Faber et al., 2022). Because the exposure to robots is based on an individual's place of birth in our intent-to-treat analysis, determined before the robot-induced shock, selective migration is not a concern. This may explain why they find results that are stronger and more in line with ours when they repeat their analysis at the state level, where migration is less likely to be an issue.

Our analysis also differs from Di Giacomo and Lerch (2022) in other substantive aspects. First, because we look at cohorts several years after the typical college-going ages, our analysis captures effects on completed rather than ongoing human capital decisions. While college enrollment outcomes are informative, they alone do not tell us the entirety of the story. If dropout is pervasive, the long-run effects on completed human capital could be limited or even null. Second, we examine how exposure to robots during the critical ages of college decisions affects subsequent earnings in the labor market, which may shed light on whether and by how much acquiring more human capital helps to mitigate the effects of automation. Finally, we move beyond the reduced-form analysis and structurally estimate a model of human capital investments to investigate deeper questions about

learning technologies (Babina et al., 2020; Acemoglu et al., 2020).

[7]Examples include Galor and Moav (2000), Adão et al. (2020), Caselli and Manning (2019), Hémous and Olsen (2022), and Guerreiro et al. (2022).

the adjustment of the economy to changes in technology.

The rest of the paper is organized as follows. Section 2 describes the data and variable definitions, while Section 3 presents the empirical strategy and basic findings. Section 4 investigates potential threats to identification. Section 5 examines the effects of exposure to robots during the critical college-going ages on subsequent labor market earnings. Section 6 provides evidence of the likely mechanism behind our results. Section 7 structurally estimates a model of human capital investments to evaluate the quantitative importance of the results. Section 8 concludes.

2 Data and Variable Definitions

In this section, we provide an overview of the data sources, present the robot exposure variable, and other variable definitions. Our basic analysis uses data from the American Community Survey (ACS) and data on robots from Acemoglu and Restrepo (2020). We also use other data sources that are described throughout the paper.

ACS microdata. We use data on the ACS for the years 2001 to 2019, a nationally representative sample of the population conducted annually by the US Census Bureau. The ACS provides rich demographic characteristics (including gender, age, race, and state of birth) as well as basic socioeconomic information (such as education and earnings). A compelling feature of these data relative to other population surveys is their enormous sample sizes, covering on average between 1.5 and 3 million individuals per year.[8] Our analysis compares cohorts exposed and unexposed to the dramatic advance in robotics during the 1990s and 2000s, which depends on when and where they were born. Our main outcome of interest is an indicator for bachelor's degree completion, the level of education that is less prone to experience the displacement consequences created by robots.[9]

Robot data. We use the measure of robot exposure built by Acemoglu and Restrepo (2020) at the state level, a level of aggregation discussed in detail below. For each state, we compute the robot exposure as the adjusted change in the stock of robots in that state's industries, weighted by each industry share in the state's baseline employment:

$$\text{Robot penetration}_s = \sum_{j \in \chi} \overbrace{\ell_{js}}^{\text{Industry share}} \underbrace{\left(\frac{\Delta M_j}{L_{jb}} - \lambda_j \frac{M_{jb}}{L_{jb}} \right)}_{\text{Robot Penetration}} \tag{1}$$

where ℓ_{js} is the initial employment share of industry j in state s, which we calculate using the census conducted in 1970 to capture the long-term industrial composition that was prevailing before the

[8]The number of people sampled changed sharply in 2005 and onward, going from 1.1 to more than 2.8 million individuals. This discontinuity is visible in our estimation sample (see Appendix Figure A.1). We have no reason to believe it has important implications for our identification strategy. The results are essentially the same if we exclude the ACS conducted before 2005.

[9]We use "bachelor attainment" and "bachelor completion" interchangeably throughout the paper.

major advance in automation. The variable $\Delta M_j = M_{j\tau} - M_{jb}$ is the change in the number of robots in each industry between the base year b and final year τ, normalized by the number of workers L_{jb}. In the model of automation developed by Acemoglu and Restrepo (2020), the labor market effects are related to the change in the number of robots per thousand workers after adjusting for the growth rate of output λ_j of each industry (captured by the expression $\lambda_j M_{jb}/L_{jb}$). We keep this adjustment term for consistency with their conceptual framework.[10] Data on robots come originally from the International Federation of Robotics (IFR), which is consistently available since 1993 for all industries aggregated into 19 consistent categories across 50 countries. We use the 1993 to 2007 period to measure the adjusted penetration of robots, a period that corresponds to the intense adoption of robots in the United States and the timing of college decisions of younger birth cohorts in our sample.

A concern with using realized penetration of robots in the United States is unobservable shocks correlated with that robot adoption. For example, reductions in the profitability of manufacturing certain products using traditional manufacturing techniques would like both lower demand in locations intensive in those industries and increase the adoption of robots in those industries. In this case, our results could be driven by the market effects induced by other factors that reduce profitability (e.g., foreign competition) rather than by the adoption of robots itself. To mitigate this concern, we follow Acemoglu and Restrepo (2020) and construct the robot penetration variable in equation (1) using data of average robot adoption in the top 5 non-US countries with greater advances in robotics (Denmark, Finland, France, Italy, and Sweden), which strongly predict US robot penetration (Figure 2).[11] This approach exploits variation coming from global advances in robotics technology rather than specific idiosyncratic US shocks. Much of the robotics advances occurred first in these countries and are thus unlikely to be driven by future factors hitting particular industries in the United States. Therefore, focusing on this measure of robot penetration allows us to isolate a source of variation plausibly independent of individuals' schooling decisions. Our baseline exposure variable focuses on this measure of European-based robot penetration, but we also present results using the observed US robot penetration as the key independent variable.

Main analysis sample. The enormous sample sizes in the ACS allow us to focus the analysis on the specific cohorts of interest while retaining a sufficient sample size. We focus on the 1953-83 birth cohorts, which include individuals who made their college decisions before and during the advance in robotics and are not too young or old to observe their outcomes consistently in the 2000s and 2010s. Our analytical sample restricts to adults born in the mainland of the United States and above age 30 at the survey time.[12]We exclude individuals residing in institutional group

[10]Data on the growth rate of output of each industry and baseline employment level in each industry are originally obtained from the Euro KLEMS database (Jäger, 2016). See Section A.2 of the Appendix for further details.

[11]This group excludes Germany, which is well ahead of the United States and thus is less relevant for robot adoption trends in the latter. We will examine the robustness of our results to alternative constructions of the exposure to robots, which consider expanding the top 5 to include Germany and other countries.

[12]This restriction excludes individuals from Hawaii and Alaska, so the resulting sample includes all individuals born in one of the remaining 48 states or the District of Columbia. In the ACS, the District of Columbia is considered a separate state. This sample restriction also excludes immigrants (about 10 percent of the observations), as it is not

quarters to increase consistency between the different rounds of the ACS, a restriction that results in dropping about 3 percent of the sample.[13]We pool all of the ACS rounds into a single file to increase the precision of our estimated results.[14] Our basic sample consists of approximately 15.3 million observations.

Geographic unit. We assign robot exposure intensity to individuals assuming that the place where they were born is the same as the one where they grew up, so our analysis is an intent-to-treat design. This requires that we choose the geographical level of the robot exposure measure. In principle, one would measure robot exposure at the county level. However, information on birthplace is only available at the state level in the ACS microdata, and thus, we are unable to match individuals with a measure of robot exposure at smaller geographies than a state. Therefore, we construct our measure of robot exposure at the state-of-birth level.[15,16]

While state divisions are relatively large geographic units, they have important strengths when studying the effects of robots on educational choices. Because mobility between states is much less frequent than between other smaller geographies,[17] the state of birth provides a more reasonable approximation of the place where individuals were residing during their childhood and college-going years. This reduces noise in our assignment of childhood exposure due to migration. Consistent with this notion, we find that approximately 80 percent of the birth cohorts in our sample were still residing in their state of birth when they were between ages 15 to 18, the critical ages when college decisions are formed.[18] In addition, many individuals attend college outside their county of birth, and therefore, their outlook on future job prospects is likely to be shaped by a wider geographic area beyond their granular place of birth. As a consequence, estimates based on measures of exposure to robots at highly disaggregated levels, such as a county, may underestimate the importance of

possible to infer whether or not they were exposed to automation technologies in the United States.

[13]The first rounds of the ACS conducted between 2001 and 2005 did not cover persons in group quarters. Hence, by excluding individuals in institutional group quarters in subsequent ACS rounds, we increase the consistency between ACS years.

[14]Since we have a fixed number of birth cohorts in our sample, the composition of these birth cohorts in the sample varies with the survey year. Younger cohorts are mechanically more likely to be observed in more recent survey years (see Appendix Figure A.3). In Section 3, we show that the results are essentially the same if we restrict the estimation sample to the 2015-19 survey years where all birth cohorts of interest are observed.

[15]An alternative possibility would be to assign robot exposure intensities based on an individual's place of residence at survey time rather than that of birth. This would allow us to explore variation in robot exposure at a fine geographic scale. We do not pursue this approach because, unlike the birthplace which is determined prior to future technological advances, the actual location of residence may reflect endogenous responses to contemporaneous trends in robot adoption.

[16]Because the original analysis of pre-trends in labor market outcomes by Acemoglu and Restrepo (2020) is at the commuting zone level, one may be worried about the possibility of differential trends at the state level. In Appendix C, we check for pre-trends in the main labor market outcomes used by Acemoglu and Restrepo (2020) at the state level. Reassuringly, we find no evidence that labor market trends between 1970 and 1990 across states are significantly correlated with robot penetration.

[17]For example, according to the 2000 Census, only 7 percent of individuals declared they moved between states during the last previous five years. By contrast, about 30 percent moved between administrative divisions that are smaller than a state. Between-state mobility accounts for less than 20 percent of the overall internal migration.

[18]We use the censuses conducted in 1980, 1990, and 2000 to track birth cohorts' place of residence at different moments in time. By its decennial nature, the population census does not allow us to observe all cohorts when they were ages 15 to 18.

the college premium channel.

Importantly, there is a great deal of variation in the intensity of robot exposure across states, as shown in Figures A.4 through A.6.[19] The difference between the 25 and 75th percentiles in the robot exposure intensity distribution is approximately 55 percent, and the difference between the 10 and 90 percentiles is more than 230 percent.

3 Research Design and Main Findings

Our analysis exploits geographic and time variation in robot adoption in a difference-in-differences research design. The first difference is over time across birth cohorts, as some individuals were exposed to the global advance in robotics before, during, or after their college-going ages depending on when they were born. The second difference is across locations, as robot adoption differs substantially across regions depending on their industrial composition. Thus, our analysis compares individuals who were younger and older during the advance in robotics in more and less exposed areas. The key difference between this approach and the standard two-group/two-period difference-in-differences is that we use a continuous measure of "treatment" intensities given by the Bartik-like variable of robot exposure described above.

It is important to emphasize that, in the presence of important general equilibrium effects, our approach does not identify the "pure" effects of robots. The adoption of robots could imply significant changes in the organization of work and induce the adoption of other routine-biased technologies. For example, manufacturing firms adopting industrial robots may also adopt other technologies that work in tandem with robots to improve production efficiency, such as automated conveyor belts, autonomous guided vehicles, or software robots to automate data entry tasks. Hence, we interpret the measure of robot penetration as a proxy for the routine-biased technology shocks that affect different industries differently during this time period.

3.1 Basic Specification

To estimate the effects of robots on human capital, we use a baseline specification that takes the form:

$$\begin{aligned} S_{ist} = \alpha \ + \ \beta \, \text{Robot penetration}_s \times \text{Post}_t \\ + \ \mathbf{X}'_{ist}\Omega \ + \ \sum_{z \in \mathbf{Z}} \Phi_z(z \times \mathbf{FE}_t) + \ \mathbf{FE}_s + \ \mathbf{FE}_t + \ \xi_{ist} \end{aligned} \tag{2}$$

where S_{ist} is the outcome of interest for individual i born in state s and birth cohort t. All models include fixed effects for state-of-birth ($\mathbf{FE}_s$) and birth-cohort ($\mathbf{FE}_t$). Since we are using all of the ACS rounds pooled into a single file, we include a detailed set of survey-year$\times$age fixed effects. The

[19] While our state-level analysis comes at a cost in terms of loss of variation, much of the variation in the commuting-zone level data in fact stems from differences between (rather than within) states. Figure A.7 illustrates this visually. Remarkably, state fixed effects account for about 75 percent of the overall cross-commuting zone variation in robot exposure intensity. This suggests that our state-level analysis captures a substantial portion of the relevant identifying variation.

vector $\mathbf{X}'_{ist}$ includes a set of basic demographic characteristics such as gender and race. The term $\sum_{z\in\mathbf{Z}} \Phi_z(z \times \mathbf{FE}_t)$ controls for interactions between birth-cohort fixed effects and a full set of 1990 state characteristics $\mathbf{Z}$. Basic state-level demographic characteristics include the log of population, the share of the population over 65 years of age, the share of the population under 5 years of age, the share of blacks, and the share of the population that is urban. To account for any possible convergence (or divergence) effects in human capital across states, we also control for interactions between 1990 state college attainment and birth-cohort fixed effects. The 1990s also witnessed other major shocks affecting US labor markets. Following Acemoglu and Restrepo (2020), we control for the share of manufacturing employment in 1990, the share of light manufacturing employment in 1990 (textile industry and the paper, publishing, and printing industry), the share of employment in routine jobs in 1990, and a measure of exposure to imports from China, all of which interacted with birth-cohort fixed effects.[20] The residual term, ξ_{ist}, is clustered at the state-of-birth level to allow for serial correlation across birth cohorts.

The key independent variable of interest is given by the interaction between our time-invariant measure of robot exposure intensity (Robot penetration$_s$) and an indicator for the cohorts exposed to the dramatic advance in robotics technology during their college-going years (Post$_t$). Figure 1 provides descriptive evidence that the dramatic advance in robot adoption started in the early 1990s and became particularly salient around 1995.[21] If individuals have altered their human capital decisions in response to automation, one would expect these effects to emerge between the 1972 and 1977 cohorts, which were of college-going ages during the mid-1990s. Theory suggests that this college response should naturally be stronger for younger cohorts, but there is no precise prediction about when exactly these effects could begin to manifest. To guide our definition of Post$_t$ and summarize our findings in table format parsimoniously, we adopt a hands-off approach that is similar in spirit to Goodman-Bacon (2021). In particular, we estimate model (2) for all possible definitions of Post$_t$ and choose the one that maximizes the R^2, following the idea of structural break tests (Hansen, 2001). As documented in Appendix D, the breakpoint that best captures the pattern of college responses in the data is the 1974 birth cohort. Thus, we use this definition throughout the paper.

Identification. A causal interpretation of our results rests crucially on the assumption that the outcomes of individuals from areas that experienced different robot penetration intensities would have followed similar trends over time across birth cohorts in the absence of the global advance in robotics. Note that the identifying assumption does not require that low- and high-exposed areas are similar in observable or unobservable factors, but requires that such factors evolve similarly over time. By conditioning on state and birth-year fixed effects, the parameter of interest is identified

[20]We use the measure of exposure to imports from China developed by Autor et al. (2013). They construct this measure at the commuting zone level. For our analysis, we re-constructed this measure at the state level. The share of employment in routine jobs is defined as in Autor and Dorn (2013): routine occupations that are in the top employment-weighted third of routine task-intensity.

[21]Consistent with this descriptive evidence, in Appendix B, we document that the effects of robots on labor markets became particularly pronounced by the mid-1990s.

from within-state differences between cohorts that were exposed at younger and older ages to robots after partialling out shocks common to all states. The interaction of a wide range of pretreatment state characteristics with birth-cohort trends reduces the risk of differential trends driven by other factors. More importantly, we will show that states disproportionately exposed to robots were on similar trends for approximately 50 years before the birth cohorts exposed to robots during their college-going ages.

The inclusion of manufacturing shares (interacted with birth-cohort fixed effects) means that we are exploiting variation in robot penetration within manufacturing industries across states. The comparison is thus between younger and older cohorts from states with similar exposure to manufacturing employment but that differ in their overall robot penetration. This is important because robot penetration was much higher in manufacturing industries, and these industries have experienced a secular decline that started before the recent progress in robotics technology. Controlling for the baseline manufacturing employment ensures that our estimates are not confounded by some long-run common causal factor behind the general manufacturing decline correlated with trends in human capital.

3.2 Visual Evidence

The evidence suggests that the recent advances in robotics became salient by the mid-1990s. Therefore, the first birth cohorts who could have altered their human capital investments in response to the robot-induced shock range between the 1972 and 1977 cohorts, since they were between 18 and 23 years old during this period. To evaluate this premise and the plausibility of the common trends assumption, we present results from estimating a fully flexible version of equation (2) that replaces the $Post_t$ dummy with birth-cohort indicators:

$$\begin{aligned} S_{ist} = \alpha \ + & \sum_{t \in T} \beta_t \text{ Robot penetration}_s \times \mathbb{1}\{\tau = t\} \\ & + \ \mathbf{X}'_{ist}\Omega \ + \sum_{z \in \mathbf{Z}} \Phi_z(z \times \mathbf{FE}_t) + \ \mathbf{FE}_s + \ \mathbf{FE}_t + \ \xi_{ist} \end{aligned} \tag{3}$$

where $\mathbb{1}\{\cdot\}$'s are indicators for birth cohorts, which are equal to one if the observation falls in birth year t. The cohort of comparison is 1953, the oldest group in our sample. By estimating the effects separately for each birth cohort, this approach allows us to assess the plausibility of the identification condition visually and transparently without imposing any parametric structure. There is no reason to expect higher- versus lower-exposure states to have differential trends before the 1972 cohort, as these cohorts had largely completed their schooling investments by the mid-1990s. Thus, if the identifying assumption is valid, the path of the cohort-specific coefficients $\{\beta_t\}$ should be flat prior to 1972 and begin to diverge only around this date. Any clear tendency toward improving schooling before 1972 would suggest that our results may simply reflect pre-existing differential trends across more-and less-exposed states.

Panel A of Figure 3 plots the estimated coefficients of β_t and corresponding 95 percent confidence intervals for each birth cohort. The coefficients are shown in percentage points for ease of reading.

The pattern is fairly clear. We do not observe any clear tendency toward improving or deteriorating schooling before the 1972 cohort. After the 1972 cohort, we observe a meaningful, visually clear, and gradual increase in bachelor's degree completion in higher-versus-lower-exposed areas. This is the pattern that one would expect to observe if the identification condition holds. The gradual increase in the estimated effect is natural, given that younger cohorts spent more childhood years exposed to robots and given the increasing adoption of robots over time. The effects begin to emerge with the 1974 cohort but become highly significant only after the 1977 birth cohort.

The 1974 birth cohort was 18 in 1992, at which point they would be preparing to graduate from high school and make decisions about starting college in 1993. While many students, especially at less selective universities, take more than four years to graduate from college, the fraction who delay their entry into college is relatively small 12-16 percent.[22] This suggests that these effects do not come solely from increased enrollment, since the adoption of robots experienced the largest increase in the mid-1990s and had barely begun in 1992. It could be that the 1974 cohort is equally likely to enroll in college, but they are more likely to persist in college as the shock begins during their time in college. For example, youths initially enrolled in associate's degree programs may later pursue a bachelor's degree, a practice occurring in 31 percent of cases (Shapiro et al., 2017).[23] In the next subsection, we will provide evidence consistent with this possibility.

Longer Pretrends. To provide further evidence on the identification assumption, we repeat the flexible model (3) using the census conducted in 1990. With these data, we are able to look at longer birth cohorts whose college attainment was measured before the robot-induced shock. If the identification condition holds, these past trends in college attainment should be unrelated to the robot penetration across states. We construct the estimation sample following the same logic as in the baseline sample by including individuals between ages 30 and 66 at the time of the census —or cohorts born between 1924 and 1960. The results are displayed in Panel B of Figure 3. Consistent with the identification condition, the placebo coefficients are statistically insignificant and small in magnitude without any clear tendency toward improving or deteriorating college attainment.

In sum, the evidence indicates that the pre-robot cohorts were on similar trends for at least since 1924, approximately 50 years. This is striking given that this period was marked by a series of pivotal social, cultural, and economic factors plausibly affecting educational investments, including the Great Depression, World War II, baby boom, civil rights movements, and significant expansion of social safety net programs. Notably, these birth cohorts also experienced a secular increase in schooling attainment, with bachelor's degree completion rates nearly tripling between the 1924 and 1970 birth cohorts. Yet, these remarkable trends appear to be largely unrelated to the recent trends

[22] Based on the Beginning Postsecondary Students Longitudinal Study, for people beginning their postsecondary education in 1995-1996, 16 percent delayed their entry from high school to college for those enrolling in public four-year universities, and this figure is 12 percent at private not-for-profit four-year institutions (National Center for Education Statistics, 2005).

[23] See the first row in columns 3 and 4 of Table 2 in (Shapiro et al., 2017). Among a full cohort of first-time students who started their postsecondary studies at community colleges 852,439 —approximately 268,749 transferred to four-year institutions (31 percent).

in robot-induced changes in routine-biased technologies. In the absence of the recent advances in robotics technology, higher-versus-lower-exposed states most likely would have experienced similar trends before, during, and after the robot-exposed birth cohorts.

3.3 Baseline Estimates

The flexible estimates provide visually clear evidence that the advance in robotics taking place since the mid-1990s has had significant impacts on college attainment. We now focus on the parametric, parsimonious model (2) to summarize magnitudes in tables and perform specification checks.

These results are reported in Table 1. Column (1) presents results from a specification that incorporates the basic set of fixed effects, individual demographic characteristics as well as 1990 state demographic characteristics interacted with birth-cohort fixed effects. The coefficient of interest is estimated at 0.32 with a standard error of 0.07 and significant at the 1 percent level. Column (2) adds baseline shares of employment in manufacturing and light manufacturing interacted with birth cohort fixed effects. The inclusion of these controls has little impact on the coefficient of interest, going to 0.36. Columns (3) and (4) control for the exposure to imports from China and the baseline share of employment in routine jobs, again interacted with birth cohort fixed effects. We now observe effects that are slightly larger in magnitude, with the coefficient of interest standing at 0.40 (standard error=0.18). Column (5) presents results from our preferred specification, which includes interactions between baseline state college attainment and birth cohort fixed effects. The coefficient of interest becomes somewhat larger in magnitude and much more precisely estimated. This suggests that the inclusion of these additional controls substantially reduces sampling variation and any convergence or divergence effect cause us to underestimate the effects of robots.

Quantitatively, the results from the preferred specification imply that the change in the bachelor's degree completion rate from the older- to younger-exposed cohorts in a highly affected state like Ohio was 1.5 percentage points more positive compared to the change between the same cohorts in a more mildly affected state like Montana. Relative to the sample mean, this effect represents an increase of approximately 5 percent.

In our main analysis, we use all rounds of the ACS pooled into a single file to increase power. An advantage of this approach beyond the gains in precision is that it reduces biases due to mortality attrition in the older cohorts, as individuals are included in the sample only if they are alive at the time of the survey.[24] However, by construction, younger cohorts are disproportionately underrepresented in the estimation sample.[25] As a robustness check, we restrict the estimation sample to the ACS conducted between 2015 and 2019, where all the birth cohorts are observed with similar likelihood. This restriction reduces sample size by approximately 66 percent, yet both

[24]For example, if more educated individuals in the older birth cohorts are more likely to survive at the time they are observed in the survey (as previous studies suggest (Lleras-Muney, 2005)), then it may change the composition of the sample. This issue is largely absent when including all rounds of the ACS, beginning since 2001, because we observe the outcomes of the older cohorts at younger ages when mortality risk is relatively low and because education changes very little with age after formal schooling is completed.

[25]This is illustrated in Figure A.1, which plots the share of each birth cohort in the estimation sample.

the point estimate and standard error are not appreciably affected (column 6, Table 1). This is likely due to the fact that, although the nominal sample size shrinks substantially, the number of clusters is unchanged.

In Appendix Table E.1, we explore heterogeneity with respect to gender. We do not observe statistically meaningful differences in the estimated coefficients between males and females. At first glance, this lack of heterogeneous effects could seem surprising, given that Acemoglu and Restrepo (2020) document effects of robots on labor market outcomes that somewhat larger for males (though not always statistically significant). However, these forces may be counterbalanced whether female education responds more strongly to a same change in labor market conditions. This possibility is plausible in view of the recent evidence by Charles et al. (2018) documenting that a labor market shock, which similarly affected male and female employment opportunities, had larger effects on female education. This possibility is also consistent with the rapid growth in female schooling during our study period.

Because the underlying variation that we exploit is across states and birth cohorts, the use of individual-level data in the analysis effectively weights state-cohort clusters by population size. As discussed by Solon et al. (2015), such weighting is appropriate in the presence of heteroskedasticity, but it may be counterproductive in its absence. We follow their practical recommendations and report both weighted and unweighted estimates. To estimate the unweighted version of model (2), we first estimate a regression of bachelor's degree attainment on individual-level characteristics (gender, race) and then collapse the residuals into state-of-birth and year-of-birth cells. We then estimate model (2) using these state-cohort level data. Appendix Table E.2 shows that the unweighted estimate is extremely similar to the weighted one (0.48 versus 0.44) but notably less precise. This confirms the strengths of our main approach.

2SLS estimates. We next present two-stage least squares (2SLS) estimates where our European-based robot penetration measured is used as an instrumental variable for the observed US robot penetration. These results are presented in Appendix Table E.3. Column (1) documents a powerful first-stage relationship, as had already been noted in Figure 2, with the F-statistics well above the conventional weak instrument threshold of 10. Notably, the F-statistics is also well above 100, which is important in view of the recent advances in instrumental variable estimation, suggesting that it may be the relevant threshold for weak instruments (Lee et al., 2022). Turning to the second stage, we find that the 2SLS estimate is comparable to the reduced-form coefficient and corresponding OLS estimates, both in magnitude and statistical significance.

Additional robustness checks. We perform several additional sensitivity tests, all of which are presented in the Online Appendix to save space. We examine the robustness of the basic results to: *i*) alternative forms of constructing the robot penetration measure (Appendix Table E.4); *ii*) excluding outlier observations based on regression residuals as well as on Cook's distance measures (Appendix Table E.5); *iii*) excluding industries with the largest Rotemberg weights (Appendix Table E.6), as recommended by Goldsmith-Pinkham et al. (2020); *iv* including baseline employment rate and 1990-2008 change in state demographics interacted with birth-cohort fixed effects (Appendix Table E.7); and v) alternative inference procedures (Appendix Table E.8), including standard errors that account for spatial correlation across areas with similar sectoral shares (Borusyak et al., 2022; Adao et al., 2019).

The distribution of education. We next investigate the sources of the gains in bachelor's degree completion. As shown in Table 2, the robot-induced increase in bachelor attainment is driven by a combination of the extensive and intensive margins: lower likelihood of having a high school degree (extensive margin) and an associate's degree (intensive margin). Workers with these education categories fall in the middle of the skill distribution and are more likely to engage in occupations highly prone to be automated by robots.[26] As a result, youths who otherwise would have completed high school or an associate's degree had incentives to pursue a bachelor's degree-level education. From a causal perspective, we take this pattern in the data as an indication that our findings are unlikely to be the product of unobservable factors affecting all individuals in the bottom and middle of the skill distribution similarly. The fact that the effects are driven by a combination of the intensive and extensive margins suggests that both the opportunity cost and premium channels are at play. The former is more likely to account for the extensive margin effect, whereas the latter for the intensive margin effect.

We also learn other important insights from the table. Among individuals who have at most a high school degree, we can distinguish between those who have less than one year of college credit and those who have more than one year of college credit but no degree. As one can infer from Table 2, there is an increase in the likelihood of having more than one year of college credit but not a degree. This is consistent with the effects coming in part from individuals who otherwise would not have started a bachelor's degree at all. The negative effect on associate's degree completion is also consistent with persistence in college, as some youths may be enrolling in or transferring to a bachelor's degree program during the course of their associate's degree training.[27]

Overall, the results of this section show that exposure to robots leads to a significant improvement in bachelor's degree completion, an effect driven by youths in the middle of the skill distribution who otherwise would have attained an education level prone to automation by robots. We next turn to major identification concerns and provide further insights into the effects of robots. To save space, we focus on the indicator of bachelor's degree completion in the remaining analyses.

[26] Even workers with associate's degrees are disproportionately engaged in occupations replaceable by robots, such as assemblers of electrical equipment, automobile mechanics, cementing and gluing machine operators, machinists, and machine operators. According to the 1990 Census, approximately 22 percent of jobs performed by workers with an associate's college degree are "replaceable" by industrial robots, as defined in Acemoglu and Restrepo (2020). The same figure for workers with a high-school degree is 38 percent, and about 7 percent for those with a bachelor's degree.

[27] The ACS data provide information on the highest degree completed, so an individual with both an associate's and bachelor's degree is recorded as having a bachelor's degree. Therefore, these results may reflect individuals starting a bachelor's rather than an associate's degree, individuals transferring from an associate's to a bachelor's degree before completing the former, individuals enrolling in a bachelor's degree after completing an associate's degree, or a combination of these three possibilities. We are unable to identify the separate importance of each mechanism.

4 Robustness Checks and Identification Concerns

In this section, we investigate potential threats to the validity of our findings, including possible mean reversion and other shocks coinciding with the advances in robotics.

4.1 Preexisting Trends and Mean Reversion

While the magnitude of our results is virtually unchanged when we flexibly control for differences in trends correlated with baseline college attainment levels, and while we find no evidence of pre-trends, one might still be worried about the possibility that our estimates are capturing some pre-existing convergence (or divergence) effect in human capital across states. We perform several additional exercises to address this issue in Table 3.

Mean reversion. Our baseline specification includes interactions between college attainment levels in 1990 and birth cohort fixed effects. Thus, this model accounts to a great extent for any possible mean-reverting (or diverging) dynamics in college attainment taking place around the onset of recent advances in robotics technology. As an additional check, we consider interactions between birth cohort fixed effects and 1970 college attainment levels, approximately 30 years before the surge in the adoption of robots. Column (2) of Table 3 shows that, if anything, the estimated effect becomes slightly larger, with the coefficient of interest going from 0.48 to 0.50. Column (3) goes a step further and controls rather for the 1970-1990 change in college attainment interacted with birth cohort fixed effects. Once again, the point estimate becomes slightly larger in magnitude and remains highly significant.

State-specific pretrends. Another way to investigate whether pre-existing mean-reverting dynamics could explain our findings is to directly control for pre-robot state-specific linear trends. To do so, we first estimate state-specific linear trends using data covering the pre-robot cohorts, which leads us to estimate a slope coefficient $\hat{\kappa}_s$ for each state. We then extrapolate the pre-robot trends in our baseline specification using the following augmented specification:

$$\begin{aligned} S_{ist} = \alpha \ + \ \beta \text{ Robot penetration}_s \times \text{Post}_t \ + \overbrace{\sum_{s \in \Theta} \hat{\kappa}_s \mathbf{1}\left[\nu = s\right] \cdot t}^{\text{state-specific pre-trends}} \\ + \ \mathbf{X}'_{ist}\Omega \ + \sum_{z \in \mathbf{Z}} \Phi_z(z \times \mathbf{FE}_t) + \ \mathbf{FE}_s + \ \mathbf{FE}_t + \ \xi_{ist} \end{aligned} \tag{4}$$

By including state-specific pre-trends, we account for underlying linear time trends in college attainment potentially correlated with the exposure to robots across states.[28] As shown in column (4) of Table 3, the inclusion of these pretrends has virtually no impact on our results, with both the coefficients and standard errors nearly identical to the baseline. This is perhaps unsurprising given

[28] An alternative approach is to control for the interaction between a cohort trend and state-of-birth dummies. We do not consider this approach because these trends may mechanically bias our estimates in the presence of varying treatment effects across birth cohorts (Lee and Solon, 2011; Wolfers, 2006).

the essentially flat pre-robot trends documented in Figure 3.

4.2 Examining Within Region Variation

While the results above are very reassuring, one could still be concerned that they are simply capturing that on average northern states are more exposed to robots and that the north diverged from the rest of the United States for other reasons. As a robustness check, we incorporate a rich set of region-of-birth $\times$ birth-cohort fixed effects (column 5, Table 3). With this more demanding specification, the impact of robots is identified not from comparisons between northern states and other regions but rather from differences between states within the same region. Therefore, we can rule out any form of mean-reverting (or diverging) trends across regions. While the inclusion of this detailed set of fixed effects substantially reduces the variation in the data, which is natural as there are fewer states within each region, the results are strikingly similar to the baseline. Although the precision of the estimation is substantially reduced as expected, the coefficient of interest remains almost unchanged and statistically significant at the 5 percent level.

4.3 Other Coincident Shocks

The results from the previous subsections are striking and support a causal interpretation of our estimates. However, the identification condition could still be violated if there were other important changes coinciding with the recent advances in robotic technology. We now consider several important contemporary shocks and provide direct evidence that they are unlikely to generate the specific pattern of exposure effects we document.

Import competition and offshoring. While our baseline specification controls for exposure to imports from China, the United States also experienced an unprecedented increase in imports from Mexico during the 1990s with the signing of the North American Free Trade Agreement (NAFTA) in 1994. Previous studies have shown that the NAFTA had important consequences for US labor markets (Hakobyan and McLaren, 2016). As a robustness check, we control for interactions between the exposure to imports from Mexico and birth cohort fixed effects (Table 4, column 2). The coefficient of interest remains virtually unchanged. Another major factor affecting US labor markets is offshoring —the practice of reallocating production processes to other countries—, which expanded significantly during the 1990s and 2000s with the development of communication technologies (Hummels et al., 2018). To investigate the sensitivity of our results to this phenomenon, we use the degree of task "offshorability" in an industry from Autor and Dorn (2013) and baseline industrial composition of employment across states to construct a Bartik-like measure of exposure to offshoring. We then repeat the baseline specification controlling for the exposure to offshoring, interacted with birth cohort fixed effects (Table 4, column 3). Both the point estimate and its standard error remain strikingly unaltered.

Other technologies. The remarkable progress in robotics technology we study coincided with other important routine-biased technological advancements. These include information and computer technologies. To control for these trends, we generate Bartik measures of exposure to information technology capital and computer intensity across states, as in Acemoglu and Restrepo (2020). Controlling for the exposure to these technologies (interacted with birth cohort fixed effects) has no material impact on our results. If anything, the magnitude of the results becomes slightly larger. It seems unlikely that our results are simply reflecting the effects of these major technological shocks.

1980-82 recession. Many of the post-robot cohorts were in their early childhood years during the recession between 1980 and 1982, whose severity varied substantially across regions. The illuminating work of Stuart (2022) shows that exposure to this recession in the first years of life led to poorer adult outcomes later in life, including reduced educational attainment. In light of this evidence, for the recession to be a threat to our identification strategy, it would need to have differentially affected states that were *less* exposed to robots. In practice, the correlation between robot exposure and recession severity as a measure in Stuart (2022) is fairly weak, and if anything, the recession was slightly more severe in states with *greater* exposure to robots. Not surprisingly, considering this pattern in the data, the inclusion of the recession severity measure interacted with birth-cohort effects yields coefficients of β very close to the baseline (column 3, Table 4).

Social reforms. A final consideration is the adoption of major reforms and safety net programs during the second half of the 20th century, many of which have been shown to have important implications for educational attainment. A major change in educational policy was the school finance reforms across states that began in the early 1970s and accelerated in the 1980s, which led to a substantial increase in K–12 education spending and improvements in educational attainment (Jackson et al., 2016). Other important social reforms include the war on poverty programs implemented during the late-1960s and 1970s, including Head Start, Food Stamp, and Community Health Centers.[29]During this period, Medicaid was also introduced for the first time in some states and Goodman-Bacon (2021) documents that it had important long-run consequences for human capital.

While the adoption of these programs differed across states and affected many of the cohorts in our estimation sample,[30] Appendix Table E.9 documents that the post-robot cohorts from states with greater robot penetration are not significantly more likely to have been exposed to these programs in childhood. Thus, these programs cannot explain much of the gains in bachelor attainment

[29]Johnson and Jackson (2019) and Hoynes et al. (2016) document that Head Start and Food Stamp respectively lead to improvements in long-run adult outcomes. Bailey and Goodman-Bacon (2015) provide evidence that Community Health Centers of improvements in health outcomes, particularly of the elderly, but they do not examine other socioeconomic outcomes such as education or labor market outcomes.

[30]The shares of the population exposed in the relevant childhood years of individuals in our sample range from 17 to 58 percent depending on the program. Data on Community Health Centers come from Bailey and Goodman-Bacon (2015), Head Start from Bailey et al. (2021), Food Stamp from Hoynes et al. (2016), and School Finance Reforms from Jackson et al. (2016).

we report in Table 1. Consistent with this notion, controlling for the exposure to these programs has very little impact on our estimates (columns 4-6, Table 4). In Appendix Table E.10, we control for the influence of these programs in a more flexible fashion by including program-year×birth-cohort fixed effects. Once again, these controls do not materially affect our results.

4.4 Evidence from IPEDS

In our final robustness exercise, we use data from the Integrated Postsecondary Education Data System (IPEDS). The IPEDS is a comprehensive source of information on enrollments and other aspects of postsecondary education, managed by the National Center for Education Statistics (NCES). An important benefit of the IPEDS relative to ACS is that, since they come from administrative registers, measurement error is less likely to be an issue. However, a major drawback of IPEDS is that it is not possible to precisely measure childhood exposure to robots using this database because it does not record information on an individual's place-of-birth, but only the number of enrollments in an institution. Therefore, we assume that the childhood place of residence of enrollees corresponds to the location of the institution where they are currently enrolled. This naturally introduces measurement error, but as discussed in Section 2 this is likely to be less of an issue at the state level.

Using these data, we estimate the following first-difference model of changes in bachelor's degree enrollment rates:

$$\Delta y_{s,90-08} = \alpha \ + \ \gamma \text{Robot penetration}_s \ + \mathbf{Z}_s'\Omega \ + \ \xi_s \tag{5}$$

where s indexes state and $\mathbf{Z}$ represents controls for baseline state characteristics, manufacturing shares, the share of employment in routine jobs, and the exposure to imports from China. We use standard errors robust to heterokedasticity and the observations are weighted by population size.

The results from estimating the specification from equation (5) are presented in column 1 of Table 5. Consistent with our baseline results, we find that states with greater exposure to robots experienced a differential increase in bachelor's degree enrollments. While this result is less precise when compared to those obtained from the ACS, it is statistically significant at conventional levels. As a falsification exercise, we regress *past* changes in bachelor enrollment rates on *future* exposure to robots. The estimated coefficient is about 8 times smaller in magnitude and far from significant (column 2), showing that relationship between robots and bachelor enrollment was absent immediately before the advance in robotics technology.

Overall, these results support the basic picture presented so far. While the estimated coefficient is not directly comparable to our baseline approach, it implies that the change in the bachelor's degree enrollment rate in a highly affected state like Ohio was 5.3 percentage points more positive compared to the change in a more mildly affected state like Montana.

5 Effects on Labor Market Earnings

An important question is whether robots affected the path of income of cohorts exposed to them in childhood. Answers to this question may shed light on whether and by how much college education mitigates the displacement effects of robots. To examine this question systematically, we estimate our model (2) for several measures of income as dependent variables. A complication with this exercise is that the introduction of robots contemporaneously affected the labor market outcomes of both younger and older adults. As such, late cohorts were the ones feeling the bulk of the displacement effects created by robots and they might still be experiencing the scarring effects from job losses they incurred during their earlier working life. In this case, our estimates would be biased toward finding a relative improvement in the labor market incomes of younger cohorts, even in the absence of a causal relationship. To mitigate this concern, we control for a detailed set of state-of-residence $\times$ birth-cohort fixed effects. With these additional controls, the parameter of interest is identified from the comparison between individuals within the same labor market but that grew up in different places during their education years.

These results are presented in Table 6. Columns (1)-(2) look at the log total personal income from all sources in the previous year. Columns (3) and (4) focus on log earned income, which includes the income earned from wages or a person's own business in the previous year. Columns (5) and (6) present results with the log income wages as the dependent variable, which is each respondent's total pre-tax wage and salary income received as an employee in the previous year. As one can infer from the table, cohorts exposed to robots in childhood see an increase (or a smaller decline) in their labor market income relative to older cohorts exposed later in the life cycle. The change in the income from older-exposed to younger-exposed cohorts in a highly affected state like Ohio was 2.3-2.6 percent more positive (or less negative) compared to the change between the same cohorts in a more mildly affected state like Montana.

Role of education. It is inherently interesting to understand to what extent education shapes the income effects we find. In principle, education is the most plausible explanation behind these results but they could also have arisen in the absence of an educational response if, for example, reallocation to less robot-exposed sectors is easier for individuals in the early stage of their labor market careers. To explore the role of education, we perform a mediation-style analysis by controlling for bachelor attainment in the income regressions and establishing the extent to which the estimated coefficient of interest is reduced. The results from this exercise are presented in the even-numbered columns of Table 6. Once the association between robots and education is accounted for, the magnitude of the income effects drops massively and loses all of its statistical significance. While this exercise must be interpreted with caution since education is a "bad control" affected by the exposure to robots, the picture is striking and suggests that education is likely the most important driver of the income effects.

In summary, cohorts exposed to robots at the beginning of the life-cycle experienced an increase in their incomes relative to late-exposed cohorts. Note that this *does not* imply that automation is good on net for younger cohorts. The introduction of robots had negative impacts on the labor market income of everyone, but this negative effect is smaller for younger cohorts who could alter their educational decisions.

6 Mechanisms

This section provides evidence on the likely mechanisms behind the college response to robots. The patterns in the data suggest that changes in the college premium and opportunity costs of attending college are at work.

6.1 Market Incentives

The adoption of industrial robots may alter the incentives to attend college by altering their opportunity costs and expected labor market premium. Robots may reduce the opportunity costs of attending college by reducing the average earnings a young unskilled person would receive in the market. At the same time, since the effects of routine-biased technological changes are persistent over time and felt heterogeneously across the skill distribution, it has the potential to affect the college premium and thus the attractiveness of college attendance.

To investigate these hypotheses, we measure log-changes in earnings and college premium using the census for 1990 and the ACS pooled across the years 2006-2008. We refer to the time window in the pooled ACS data simply as 2008. We assume that individuals attending college in a given year forgo immediate income gains equivalent to the average earnings of individuals aged 18-21 without any college training.[31] We measure the college premium as the earnings gap between older working adults (ages 22-65) with and without college training. We compute the average of these labor market measures within about 220,000 cells defined by demographic $\times$ state groups. The demographic groups are defined by gender ($\times 2$), age ($\times 48$), race ($\times 9$), and place-of-birth ($\times 52$).[32] For a given outcome y of demographic group g in state s, we estimate the following first-difference specification:

$$\Delta y_{gs,90-08} = \alpha \ + \ \gamma \text{Robot penetration}_s \ + \mathbf{Z}'_s \Omega \ + \ \alpha_g \ + \ \xi_{gs} \tag{6}$$

where $\Delta y_{gs,90-08}$ is the log-change in the labor market measure between 1990 and 2008. The α_g represents a detailed set of demographic group fixed effects, which help reduce concerns about possible compositional changes. Standard errors are clustered at the state level, and all regressions are weighted by the 1990 cell size.

The results from estimating (6) are presented in Table 7. Column (1) shows that states experiencing greater exposure to robots have seen a decline in the average earnings of young workers. The magnitude of this effect suggests a sizeable decline in the opportunity cost of attending college.

[31] We exclude currently enrolled students, as they have on average lower earnings than those already fully in the labor market.

[32] The place of birth corresponds to a state for US-born individuals and to a country for foreign-born people. We group the place of birth within a same category for individuals born outside of the United States.

The point estimate of -0.076 implies that the state in the 75th percentile of the exposure to robot distribution experienced a decline of approximately 17 percent in the labor market income a young adult worker receives.

Columns (2) to (5) also document a decline in the average earnings of older adult workers across all education groups (high school, associate's degree, and bachelor's degree).[33] This effect is smaller for workers with a bachelor's degree, so the exposure to robots increased the premium from having a bachelor's degree relative to high school and associate's degree. The increase in the skill premium from having a bachelor's degree relative associate's degree is significantly larger than that relative to high school. This is consistent with our findings in Table 2 showing that part of the increase in bachelor's degree attainment stems from individuals in the middle of the skill distribution.

We interpret the findings of this section as evidence supporting the hypothesis that the adoption of industrial robots altered the market incentives to invest in Bachelor's-level training by raising its relative returns and reducing its opportunity cost.

6.2 Parental Resources

Since the widespread adoption of robots led to a sizable decline in average income, it is natural to ask if this shock was large enough to translate into lower parental income. A decline in parental resources may limit the ability of credit-constrained parents to finance college, generating an effect that must work against the observed increase in college attainment we document. To explore the empirical importance of this effect, we estimate equation (6) using the log-change in total family income as the dependent variable. We limit the sample to families where parents co-reside with children aged 17-18, the timing of initial college decisions. We collapse the log-total income in 1990 and 2008 in each state by education categories in addition to the demographic cells defined in the previous subsection. We then control for the full set of demographic-cell fixed effects in our estimation.

The results are shown in column (7) of Table 7. As one can see, there is no evidence of an effect on parental income. The estimated coefficient is statistically insignificant and small in magnitude. A possible reason for this null result is that parents could have been able to offset the robot-induced income loss by increasing their overall labor supply or by enrolling in government assistance programs.[34] Identifying the importance of these responses is beyond the scope of this paper and a possible direction for future work. In any case, the evidence suggests that the null change in overall parental resources is not a mechanism working against the observed increase in college attainment we document.

[33] The fact that we observe a significant decline even in the earnings of individuals with a Bachelor's degree (though to a much lesser degree than other education groups) suggests that college education is not completely protective against the displacement consequences created by robots.

[34] Acemoglu and Restrepo (2020) provide evidence consistent with the latter. Specifically, they document that areas with greater exposure to robots saw a larger increase in the use of Social Security Administration retirement, disability benefits, and other government transfers.

6.3 Supply-Side Responses

One less apparent mechanism is changes on the side of supply. Institutions awarding bachelor's degrees may have responded to changing labor market conditions by altering tuition costs. Additionally, local and state governments may help facilitate access to college in response to a growing mass of young adults failing to find a job in affected labor markets. As a consequence, state administrations may have increased their investments in education and training programs or directly provided grants to students.

To explore these possibilities, we use state-level data on college tuition and fees as well as data on revenue from state and local appropriations available in the IPEDS database. We also use data on government expenditure on education and training assistance programs from the Regional Economic Information System. With these data, we estimate the effects of robots on tuition, revenue, and expenditure using a state-level first difference version of equation (6). As can be seen from Appendix Table E.11, there is no systematic evidence of statistically meaningful effects on these variables. We conclude that there is limited support for the interpretation that supply-side responses play an important role in explaining the improvements in college attainment.

7 Implications for Earnings Dynamics and Policy

The key message we take from the results presented thus far is that individuals tend to adjust to routine-biased technological shocks by redirecting their human capital investments toward skill areas that are less susceptible to automation. However, it is far from straightforward to interpret the magnitude of this finding based only on the reduced-form analysis. Ideally, one would like to understand whether these effects are of the right order of magnitude to significantly alter the adjustment of the economy to changes in technology: how big would the long-term effects of robots on earnings be in the absence of the endogenous educational response? Is this mechanism going to gain or lose relevance over time as more people become educated and enter the workforce? What is the role of policy in terms of subsidies in enhancing this endogenous response and determining the path of earnings? To interpret the magnitude of the findings, we develop and estimate a simple model of human capital investments where changes in the demand and supply of skills shape the evolution of earnings. We view this exercise as exploratory in nature since we are abstracting from other more complex forms of general equilibrium effects. Nevertheless, our analysis can be taken as a natural starting point for further research.

7.1 Model

This section introduces the model of human capital investments and key variable definitions. In the model, individuals are heterogeneous in terms of preferences and parental income. Additionally, individuals differ in their labor market of residence, defined as the state where they were residing at the time of their college decisions. We omit individual and state subscripts for simplicity and ease of readability.

7.1.A Human Capital

The model consists of two parts. In the first part, individuals are young adults (from 18 to 21 inclusive) and must decide whether to attend bachelor's-level college or enter the labor force. We can also think of this binary decision as individuals choosing between attending bachelor's-level college or pursuing a shorter post-secondary training and entering the labor force earlier. In the second part, all individuals enter the labor force and receive a labor market income that depends on their educational attainment. Therefore, the educational decisions made in the first period are irreversible.[35] To match the typical working life of 48 years and use periods of equal length, time is divided into 12 periods of 4 years each. Individuals are credit-constrained and can only borrow to finance the costs of college attendance. There are no savings, so consumption is equal to income in each period.[36] The choice of college attendance is represented by $s = \{0, 1\}$, where 1 indicates attendance and 0 otherwise. Periods are represented by $t = \{0, 1, \ldots, 11\}$.

Preferences. The utility is intertemporally separable and depends on consumption (c) and preferences for college education. The discount factor is $\beta = 1/(1 + \rho)$, with a discount rate of ρ. The instantaneous utility function over consumption is CRRA:

$$u(c_{st}) = \frac{c_{st}^{1-\gamma}}{1-\gamma} \tag{7}$$

The parameter γ is the inverse of the intertemporal elasticity of substitution of consumption. A higher value of γ implies that individuals place a higher value on immediate consumption and are thus less willing to delay consumption today to pursue college. Individuals face a heterogeneous disutility shock from attending college, denoted by η. One can interpret this heterogeneity as differences in the psych cost of learning or capacity of individuals to adapt to the academic and college social environment. For example, students who struggle with academic coursework may find it difficult to keep up with the demands of college-level work, leading to stress, and anxiety, resulting in a greater disutility associated with attending college. Similarly, some students may have difficulty adapting to the social environment of college, such as making new friends, joining clubs or organizations, or participating in extracurricular activities. These students may feel socially isolated or excluded from the college community, increasing the disutility from attending college. We assume that each individual draws an idiosyncratic disutility shock from a normal distribution with mean μ_η and standard deviation σ_η.

[35] An interpretation of this assumption is that there are important psych costs of learning that become growingly important with age. This is consistent with data showing that educational attainment changes little with age once an individual completes her formal education decisions.

[36] The assumption of no saving is unlikely to significantly affect our quantitative analysis, as we are studying an "once-for-all" decision that is made at the beginning of the life cycle, when individuals rarely save. This assumption is fairly standard in the literature whose focus is to understand educational decisions, including the prominent studies by Arcidiacono et al. (2012) and Wiswall and Zafar (2015).

Budget constraint. The total costs of attending college include both tuition and fees as well as living expenses such as housing, food, transportation, and personal items. These costs can be covered by a combination of government grants and transfers from parents. Grants are distributed based on predetermined rules that consider the recipient's parental income, $G(I)$. Similarly, parental transfers are a function of family income and college attendance, such that $T_s(I) = \lambda_s I$ with $\lambda_s \in (0,1)$. Students who are unable to cover their full college expenses through grants and parental transfers can take out loans A up to make up the difference to an interest rate r. Individuals repay the loan during the two first periods after graduation under a plan with equal payments. Then, consumption in each period for $s = 1$ can be written as follows:

$$c_{1t} = \begin{cases} \lambda_1 I - (e - g(I)) + c_u + A & \text{if } t = 0 \\ w_{1t} - \text{debt} & \text{if } t = \{1,2\} \\ w_{1t} & \text{if } t \geq 3 \end{cases}$$

where $e - g(I)$ represents tuition and fees net of grants, c_u living expenditures while in college, A college loans, *dept* payment of the loan, and w_{1t} labor market income of college-educated people. Since there are borrowing constraints, students can only borrow up to limit A:

$$(e - g(I)) + c_u - \lambda_1 I \leq A$$

in practical terms, consumption of students who require loans to finance college attendance will be equal to c_u. For individuals who do not enroll in college, consumption is equal to income in each period, except in the first period, where they also receive transfers from their parents.

Earnings process. The earnings w_{st} for an individual at period t with education s is given by:

$$\ln w_{st} = g_s(\mathbf{Z}) + \kappa_{sa} R + \xi_{st} \tag{8}$$

where $g_s(\cdot)$ is an education-specific function measuring the importance of demographic and background characteristics, ξ_{st} a stochastic shock, and R is the robot penetration in the market. The semi-elasticity of wages with respect to robots, denoted as κ_{sa}, is allowed to vary between young and old workers, as well as between college-educated and non-college-educated people, with $a \in \{young, old\}$. This captures the possibility that skilled workers who are less prone to engage in routine-related jobs or those with more experience are less susceptible to the displacement consequences of industrial robots, as suggested by the evidence in Section 6.

Decision problem. An individual will attend college if the utility gains in consumption due to college are greater or equal to the disutility of attending college η:

$$\sum_{t=0}^{11} \beta^t \cdot \mathbb{E}\Big[u(c_{1t}) - u(c_{0t})\Big] \geq \eta \tag{9}$$

under this framework, the proportion of individuals who attend college is simply the proportion of individuals who draws a sufficiently low value of η. Therefore, the normal cumulative distribution function, $F(\cdot)$, determines aggregate college attainment of a given cohort. Increased robot penetration makes college education more attractive by reducing the labor market income a person would receive in the first period and by leading to a larger decline in the income of non-college-educated people.

7.1.B Production

Final goods are produced by a representative firm combining high-skill labor (ℓ_H), low-skill labor of younger workers (ℓ_{Ly}), low-skill labor of older workers (ℓ_{Lo}), and robots (R). By high- and low-skill labor, we mean bachelor'-level college and non-bachelor's-level college labor. The production function takes the following constant elasticity of substitution form:

$$Q = \left[\sum_{j \in \chi} (a_j \ell_j)^{\frac{\varepsilon-1}{\varepsilon}} + (a_R R)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{10}$$

where $\chi = \{H, Ly, Lo\}$ is the set of labor categories, $a's$ are effective share parameters, and $\varepsilon \in (0, \infty)$ is the elasticity of substitution. This production function can be viewed as a reduced-form version of the task-based model developed by Acemoglu and Autor (2011) and extended in Acemoglu and Restrepo (2018), where capital competes against labor in the production of tasks. Increases in a_R are interpreted as a task-replacing technological change that expands the range of tasks that capital can perform. This expansion in turn reduces the effective share parameters of labor.

Workers are paid their marginal products, so the automation-induced log-change in the price of labor $j \in \chi$ can be expressed as follows:

$$d\ln w_j = \underbrace{\frac{\varepsilon - 1}{\varepsilon} d\ln a_j + \frac{1}{\varepsilon} d\ln Q}_{\substack{\text{net demand} \\ \text{effect}}} - \underbrace{\frac{1}{\varepsilon} d\ln \ell_j}_{\substack{\text{supply} \\ \text{effect}}} \tag{11}$$

Equation (11) determines the equilibrium path of earnings after a technological shock. The first two terms on the right-hand side represent a net demand effect. A skill-replacing technological change will reduce the demand for labor via a displacement effect, as captured by a decline in a_j. But this technological change also creates a productivity effect ($d\ln Q$) that increases the demand for labor, so the overall demand effect ultimately depends on the relative importance of these forces. The last term on the right-hand side represents a labor supply effect. As more individuals invest in college education and become high-skill workers, the supply of low-skill labor declines, and this will drive up labor earnings of low-skill workers.

Definition 7.1 (Aggregate earnings)**.** *Aggregate labor market earnings w in a given moment in time is:*

$$w = \sum_{j \in \chi} \theta_j w_j, \quad \textit{with} \quad \theta_j = \frac{\ell_j}{\sum_{j \in \chi} \ell_j}$$

In words, overall labor market earnings are calculated as a weighted average of earnings categories. In our simulation exercises below, we use this definition to track aggregate earnings over time.

7.2 Identification and Estimation

We implement a three-step procedure to take the model to the data. First, we estimate the key earnings parameters $\{\kappa_{1old}, \kappa_{0old}, \kappa_{young}\}$ separately from the rest of the structure of the model, following the first-differences model used in Section 6. Second, we set some parameters externally. The curvature coefficient of the utility function, γ, is fixed to 1.5, following Abbott et al. (2019). We assume that the discount rate ρ is equal to the interest r, and set the latter to 5 percent per year, as in Heckman et al. (1998). As each period corresponds to 4 years, an annual interest rate of 5 percent is equivalent to an interest rate of 20 percent per period. This implies a discount factor β of approximately 0.83 ($= 1/1.2$). The parameters that determine parental transfers in the first period, λ_s, are computed separately for enrollers and non-enrollers based on estimates reported by Abbott et al. (2019).[37] The costs of college attendance, including tuition and fees as well as living expenses while in college, come from the National Center for Education Statistics (NCES, 2004) for the academic years of 1989-90. Individuals make college decisions based on "expected" grants, calculated as $G(I) = \pi_q \phi_q (e + c_u)$ for a youth from quartile q of the family income distribution, where π_q and ϕ_q are respectively probability of receiving grants and the share of college costs covered by grants. We calculate π_q and ϕ_q using information on college grants by family income reported in NCES (2004).

We estimate expected income for a worker with college attendance status s at period t of the working life as the average earnings of workers in that group observed in the 1990 census. These expected earnings profiles are allowed to differ across states. The elasticity of substitution coefficient ε is set to 1.4, based on the evidence in Katz and Murphy (1992). Table 8 summarizes the earnings coefficient estimated in the first step and the parameters set externally.

In the third step, we use the simulated method of moments to estimate the remaining parameters, $\boldsymbol{\psi} = \{\mu_\eta, \sigma_\eta\}$, by exploiting the variation provided by the robot-induced shock. Our target moments are the baseline bachelor's degree rate and the reduced-form coefficient of the effects of robots on bachelor attainment. The estimation sample consists of all individuals aged 18 in the 1990 census, the pivotal age when young adults make their college decisions. Recall that these individuals vary in terms of parental income, preferences, and expected earnings profile, with the latter determined by an individual's state of residence. We first solve the college decision of each individual before the arrival of robots for any given candidate parameters $\boldsymbol{\psi}$ and compute the simulated bachelor's degree attainment rate $\hat{y}_0(\boldsymbol{\psi})$.[38] We then solve the model after the introduction of robots and obtain the corresponding post-robot bachelor's degree rate $\hat{y}_1(\boldsymbol{\psi})$. Under this scenario,

[37] In Appendix Table G.1, Abbott et al. (2019) report yearly transfers by college enrollment status using data from the National Longitudinal Survey of Youth 1997. We divide these amounts by total family income to compute the corresponding share parameters, λ_s.

[38] Specifically, individuals randomly draw a disutility shock η_i from a normal distribution with a given value for parameters μ_η and σ_η. Using these draws, the college decision of each individual is solved using equation (9).

the adoption of robots represents a shock to the income profile that individuals expect to receive in their state. Since the penetration of robots varies across states, individuals from different states were exposed to this shock with varying degrees of intensity.[39] The average effect of robots is then calculated as $\Delta\hat{y}(\boldsymbol{\psi})/R$, where $\Delta\hat{y}(\boldsymbol{\psi}) = \hat{y}_1(\boldsymbol{\psi}) - \hat{y}_0(\boldsymbol{\psi})$ and R is the average robot penetration observed in the data. We repeat this process iteratively for different values of $\boldsymbol{\psi}$ until the difference between the simulated and empirical moments is as small as possible. Formally, $\boldsymbol{\psi}$ is estimated as the vector that minimizes the criterion function:

$$[\mathsf{m}(\boldsymbol{\psi}) - \hat{\mathrm{m}}]' \cdot \Omega \cdot [\mathrm{m}(\boldsymbol{\psi}) - \hat{\mathrm{m}}]$$

where $\mathrm{m}(\boldsymbol{\psi})$ is the vector of moments simulated from the model evaluated at $\boldsymbol{\psi}$, $\hat{\mathrm{m}}$ the vector of empirical moments to be matched, and Ω a diagonal matrix containing the inverse of squared moments. In practice, this definition of Ω implies that the estimation strategy minimizes the sum of squared percentage deviations between the set of moments in the model and in the data. The standard errors for the estimated parameter vector are computed using bootstrap. We first randomly draw 500 samples stratified by state and then estimate the parameters in each of these samples. The standard errors are derived from the distribution across these 500 estimations.

Identification. While the parameters are estimated simultaneously, it is useful to discuss the moments that best identify each of them. The mean of the disutility parameter, μ_η, is pinned down by seeking the values consistent with the baseline college attainment rate. Holding all else equal, the simulated college attainment rate declines monotonically as μ_η increases. On the other hand, the response of college attainment to robots provides information to identify the parameter σ_η. For values of σ_η close to zero, college attendance is either a rare or universal phenomenon for low or high values of μ_η, scenarios that are inconsistent with the data.[40] For a less extreme value of μ_η, such that the model matches the baseline bachelor attainment perfectly, values of σ_η close to zero imply an unrealistically large college response to robots.[41] It follows that there needs to be sufficient dispersion in the disutility from attending college for the model to rationalize the data. For values of σ_η that are not very close to zero, the predicted college response is monotonically decreasing in σ_η, with a unique value matching the model's moments to the empirical ones. Figure 4 illustrates this visually. Appendix Figure F.1 provides further evidence that the targeted moments effectively inform the estimated parameters by plotting how the criterion function changes around the estimated value of the parameters.

[39] Specifically, for an individual residing in state g with college attainment s and age group a, we approximate the expected earnings after the robot shock as $w_{gsa} = \bar{w}_{sag} \cdot \exp\{\hat{\kappa}_{sa} R_g\}$. In this expression, R_g is the robot penetration observed in state g, computed as discussed in Section 2, and $\bar{w}_{gsa}$ is pre-robot earnings.

[40] Intuitively, for low values of μ_η, college attendance becomes nearly universal for a value of σ_η close to zero, as few individuals receive disutility shocks sufficiently high to counterbalance the benefits of going to college. The opposite holds for high values of μ_η.

[41] In particular, when $\mu = 0.0097$ and σ_η ranges between 0.0001 and 0.0002, the model closely replicates the baseline bachelor's degree completion rate. However, the predicted reduced-form coefficient of the effects of robots on college attainment is 30 times as large as the one observed empirically.

7.3 Parameter Estimates and Model Fit

Table 9 presents the estimated parameters along with their targeted moments. The model nearly replicates the reduced-form effect of robots on bachelor's degree attainment as well as the bachelor's degree attainment rate, with the simulated moments well within 1 percent of the observed moments. This tight fit is unsurprising given that the model is exactly identified. The mean and standard deviation of the disutility from attending college are estimated at 0.038 and 0.064 respectively and statistically distinguishable from zero at less than 1 percent level.

To interpret the magnitude of these parameters, we can simulate how bachelor attainment changes in response to a shock in labor market earnings. Our model predicts that a 10-percentage point increase in the lifetime college premium would generate a 5-percent increase in bachelor attainment. This corresponds to a semi-elasticity of 0.5 and an elasticity of 0.7. The magnitude of this result is comparable to the findings reported in Abramitzky et al. (Forthcoming), perhaps the best quasi-experimental evidence on the effects of changes in skill returns on bachelor attainment. They show that a pay reform in Israel that increased the returns to bachelor's degree from 0 to 50 percent led to a 40-percent increase in bachelor attainment —a semi-elasticity of 0.8. We also simulate the effects of an increase in the opportunity cost of college attendance and find an elasticity of bachelor attainment with respect to initial earnings of -0.21. This smaller elasticity (in absolute value) is plausible and consistent with early work suggesting that future lifetime earnings is more important than initial earnings for human capital investment decisions (Berger, 1988).

In Appendix F.2, we explore the robustness of the results to alternative assumptions about the model. Our benchmark model sets the curvature coefficient of the utility function to 1.5. Appendix Table F.1 documents that the results are robust to assuming different curvatures for the utility of consumption, including values of γ equal to 1 (log utility) or 2. The estimated parameters are larger for smaller values of γ because the marginal utility of consumption is decreasing in γ and thus the disutility from attending college must be larger to counteract the benefits of attending college. However, the fit of the model remains tight. Our baseline analysis uses an annual interest rate of 5 percent, which is somewhat larger than the free-risk interest rate used in some studies analyzing the college decisions of recent cohorts (Lawson, 2017). Using interest rates that range between 1 and 4 percent (or 4 and 16 percent per period), we find estimates of the disutility parameters that are similar to the baseline (Appendix Table F.2). In our analysis, we have assumed that the life of the loan is 2 periods, which corresponds to 8 years. As a robustness check, we re-estimate the model exploring different repayment periods, ranging from 16 to 32 years (or 4 to 8 periods). This alters the value of the payments per period but has little impact on the estimated parameters (Appendix Table F.3).

7.4 Simulations

Having estimated the model, we can proceed to perform simulations to gain insights into the long-run evolution of earnings and the role of policy. We carry out two counterfactual exercises. First, we trace the effect of a technological shock on earnings in a world where people do not alter their human capital investments. In the second counterfactual, we simulate the dynamics of earnings and supply of skills in a scenario where the government subsidizes college attendance more generously.

We start in a steady state where approximately 30 percent of workers have a bachelor's degree, the pre-robot shock period. We then assume that each state experiences a permanent shock in robotics technology equivalent to the robot penetration observed in the data. This generates an initial "net demand effect" on the earnings of workers with college attendance status s and in age category a, as captured by the first term on the right-side of equation (11). For simplicity, we approximate these net demand effects as the reduced-form coefficient of the effect of robots on earnings, $\kappa = \{\kappa_{1old}, \kappa_{0old}, \kappa_{young}\}$, multiplied by the robot penetration observed in each state.[42] These changes in earnings alter the incentives to invest in college education, and individuals respond accordingly. This alters the supply of high- and low-skill workers. These changes in the supply of skills in turn influence the dynamic of earnings, as captured by the last term on the right-side of equation (11). We assume that these supply effects on earnings are materialized after the initial generation of workers retires and is replaced by the incoming generation of workers who made different educational choices.[43] In each moment in time, overall earnings are computed using definition 7.1. We perform these simulations separately for each state and then compute the population-weighted mean value across all states.

Figure 5 shows the long-run evolution of bachelor attainment and earnings, which are normalized to zero in the initial period. As one can infer from the figure, the introduction of robots leads to an increase in college attainment among the first generation of affected workers. However, as the supply of high-skill workers increases, the returns to skill fall and incoming cohorts adjust their educational investments accordingly. Over the long term, the effect of robots on college attendance is approximately two-thirds the effect observed among the first affected generation. Another key insight we gain is that this endogenous educational response is not of the first order of magnitude to fully offset the negative demand effect induced by the adoption of robots. The increase in college attendance reduces the effect of technology on earnings by 53 percent in the short run. But over the long run, as earnings adjusts to changes in the supply of skills, this number falls to 33 percent.

Policy experiments. We next turn to our counterfactual policy analysis. We consider policies that alter the subsidies for students in the first and second quartiles of the income distribution, denoted as the target population. We first consider a policy that increases the value of grants such that college costs for the target population are zero, while keeping fixed the coverage of students who receive such grants. The second counterfactual policy we consider is an increase in the coverage

[42] This is imperfect because these reduced-form coefficients likely incorporate some of the effects of the endogenous increased supply of skills on earnings. Under the assumption that the earnings adjustment to an increased supply of skills takes a long time to materialize, then our approach provides a reasonable characterization of the net demand effects.

[43] We adopt this assumption for simplicity, but we believe that it is reasonable to assume that a single birth cohort who just completed their educational investments is unlikely to significantly impact market earnings. Rather these effects materialize slowly over time after many cohorts with different educational choices have entered the market.

of existing grants to 100 percent among the target population. Analogously, the share of college costs covered by grants remains fixed under this alternative intervention. Put differently, the first policy affects the intensive margin of subsidies by increasing the recipient's grants, whereas the second policy experiment affects the extensive margin by increasing the number of students who can receive grants.

Panel A of Figure 6 shows that these policies encourage a large response of bachelor attainment to robots. Comparing both types of policies, increasing the share of the population receiving existing grants is significantly more effective than increasing the value of grants. The key reason for this is that the value of grants is already high for students at the bottom of the income distribution, so the scope for gains is more limited. By contrast, the baseline share of the population covered by grants is far from universal and thus the scope of impact is larger. When both policies are implemented combinedly, the baseline increase in college attainment is more than doubled over the long term.

Panel B documents that the effect of endogenous human capital accumulation on earnings is magnified by both policies. Unsurprisingly, given the pattern observed in Panel A, the policy increasing the coverage of grants is more effective in dampening the net demand effect brought by automation. According to our calculations, this policy reform would dampen the robot-induced decline in earnings by approximately 60 percent. When both policies are implemented simultaneously, the long-run effects of robots on earnings are substantially reduced by 92 percent. Taken in its entirety, these exercises suggest that endogenous human capital accumulation cannot undo most of the earnings effects of automation unless there are sufficiently generous educational subsidies.

While these findings are striking, it is important to stress that they could miss important unmodeled features of the economy. A such aspect is migration. Acemoglu and Restrepo (2020) provide evidence that workers move away from commuting zones differentially exposed to robots. This endogenous migration response is likely to be less important in our state-level analysis. But to the extent to which it is important, our counterfactual analysis could underestimate or overestimate the aggregate effects of endogenous human capital accumulation. Another aspect is that we assume that workers supply labor inelastically in the second part of the life cycle, but if workers adjust their labor supply endogenously, then this could affect the path of earnings over the long run. Finally, our model does not incorporate behavioral responses of parents to policy. If parents respond by reducing financial transfers, then the impact of increased subsidies would be smaller. Therefore, our counterfactual exercises correspond to scenarios where these forces are absent. Developing and estimating a richer model that incorporates these and other forms of general equilibrium responses is beyond the scope of this paper. Nevertheless, we believe the quantitative analysis presented above is useful because it provides an intuitive way to interpret the magnitude of the reduced-form findings while highlighting avenues for future research.

8 Concluding Remarks

The last few decades have seen an intense debate on the impacts of automation technologies on workers. While a vast literature has studied this question both empirically and theoretically, much less evidence is available on whether and how individuals respond to automation. In this paper, we consider one of the most natural margins of adjustment —human capital. We investigate the extent to which the adoption of industrial robots affected individuals' college decisions in the United States. By exploiting variation in the baseline industrial mix of each state interacted with plausibly exogenous changes in sector-specific robot penetration rates, we find strong evidence that growing up in labor markets heavily exposed to industrial robots leads to greater investments in bachelor's-level education. This effect is large enough to have consequences for the labor market income. Our estimates suggest that cohorts exposed to robots in childhood experienced an increase (or a smaller decline) in their labor market income relative to those cohorts exposed later in the life cycle who could not alter their educational decisions. The data suggest that changes in the college premium and opportunity costs of college-going are the key drivers of these results.

To interpret the magnitude of the findings, we develop and estimate a simple model of human capital investments where changes in the demand and supply of skills shape the evolution of earnings. Mapping this model to the data, we find that the endogenous educational response is not of the right order of magnitude to fully counterbalance the decline in earnings induced changes in technology. Our simulations suggest that this mechanism mitigate these automation-induced effects by only 33 percent over the long run. We also conduct policy counterfactuals to explore the role of subsidies to college attendance. We find that a reform that increases the coverage and value of grants can offset the earnings effects by approximately 92 percent. In sum, these exercises suggest that endogenous human capital accumulation has little effect unless the government significantly subsidizes college attendance. Our quantitative analysis is exploratory in nature, and we believe that incorporating more complex general equilibrium effects represents an interesting avenue for further investigation.

Bibliography

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L Violante, "Education policy and intergenerational transfers in equilibrium," *Journal of Political Economy*, 2019, *127* (6), 2569–2624.

Abramitzky, Ran, Victor Lavy, and Maayan Segev, "The effect of changes in the skill premium on college degree attainment and the choice of major," *Journal of Labor Economics*, Forthcoming.

Acemoglu, Daron and David Autor, "Skills, tasks and technologies: Implications for employment and earnings," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

_ and Pascual Restrepo, "Low-skill and high-skill automation," *Journal of Human Capital*, 2018, *12* (2), 204–232.

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

_ , David Autor, Jonathon Hazell, and Pascual Restrepo, "AI and Jobs: Evidence from Online Vacancies," Technical Report, National Bureau of Economic Research 2020.

Adão, Rodrigo, Martin Beraja, and Nitya Pandalai-Nayar, "Technological transitions with skill heterogeneity across generations," Technical Report, National Bureau of Economic Research 2020.

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad, "The skill complementarity of broadband internet," *The Quarterly Journal of Economics*, 2015, *130* (4), 1781–1824.

Angrist, Joshua, David Autor, and Amanda Pallais, "Marginal effects of merit aid for low-income students," *Quarterly Journal of Economics*, Forthcoming.

Arcidiacono, Peter, V Joseph Hotz, and Songman Kang, "Modeling college major choices using elicited measures of expectations and counterfactuals," *Journal of Econometrics*, 2012, *166* (1), 3–16.

Athreya, Kartik and Janice Eberly, "Risk, the College Premium, and Aggregate Human Capital Investment," *American Economic Journal: Macroeconomics*, April 2021, *13* (2), 168–213.

Atkin, David, "Endogenous Skill Acquisition and Export Manufacturing in Mexico," *American Economic Review*, August 2016, *106* (8), 2046–85.

Attanasio, Orazio P and Guglielmo Weber, "Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey," *Journal of political Economy*, 1995, *103* (6), 1121–1157.

Autor, David and David Dorn, "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, 2013, *103* (5), 1553–97.

_ , _ , and Gordon H Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Autor, David H, Lawrence F Katz, and Alan B Krueger, "Computing inequality: have computers changed the labor market?," *The Quarterly journal of economics*, 1998, *113* (4), 1169–1213.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, "Artificial Intelligence, Firm Growth, and Industry Concentration," 2020.

Bailey, Martha J and Andrew Goodman-Bacon, "The War on Poverty's experiment in public medicine: Community health centers and the mortality of older Americans," *American Economic Review*, 2015, *105* (3), 1067–1104.

_ , Shuqiao Sun, and Brenden Timpe, "Prep School for poor kids: The long-run impacts of Head Start on Human capital and economic self-sufficiency," *American Economic Review*, 2021, *111* (12), 3963–4001.

Berger, Mark C, "Predicted future earnings and choice of college major," *ILR Review*, 1988, *41* (3), 418–429.

Blom, Erica, Brian C Cadena, and Benjamin J Keys, "Investment over the business cycle: Insights from college major choice," *Journal of Labor Economics*, 2021, *39* (4), 1043–1082.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental shift-share research designs," Technical Report 1 2022.

Burstein, Ariel, Eduardo Morales, and Jonathan Vogel, "Changes in between-group inequality: computers, occupations, and international trade," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 348–400.

Cadena, Brian C and Benjamin J Keys, "Human capital and the lifetime costs of impatience," *American Economic Journal: Economic Policy*, 2015, *7* (3), 126–53.

Carrillo, Bladimir, "Present Bias and Underinvestment in Education? Long-run Effects of Childhood Exposure to Booms in Colombia," *Journal of Labor Economics*, 2020, *38* (4), 1127–1265.

Caselli, Francesco and Alan Manning, "Robot arithmetic: new technology and wages," *American Economic Review: Insights*, 2019, *1* (1), 1–12.

Castro, Rui and Daniele Coen-Pirani, "Explaining the evolution of educational attainment in the United States," *American Economic Journal: Macroeconomics*, 2016, *8* (3), 77–112.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo, "Housing Booms and Busts, Labor Market Opportunities, and College Attendance," *American Economic Review*, October 2018, *108* (10), 2947–94.

Chodorow-Reich, Gabriel and Johannes Wieland, "Secular labor reallocation and business cycles," *Journal of Political Economy*, 2020, *128* (6), 2245–2287.

Cragg, John G and Stephen G Donald, "Testing identifiability and specification in instrumental variable models," *Econometric Theory*, 1993, pp. 222–240.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner, "The adjustment of labor markets to robots," *Journal of the European Economic Association*, 2021, *19* (6), 3104–3153.

Donovan, Kevin and Christopher Herrington, "Factors affecting college attainment and student ability in the us since 1900," *Review of Economic Dynamics*, 2019, *31*, 224–244.

Ersoy, Fulya Y, "The effects of the great recession on college majors," *Economics of Education Review*, 2020, *77*, 102018.

Faber, Marius, Andrés P Sarto, and Marco Tabellini, "Local shocks and internal migration: the disparate effects of robots and Chinese imports in the US," Technical Report, National Bureau of Economic Research 2022.

Galor, Oded and Omer Moav, "Ability-biased technological transition, wage inequality, and economic growth," *The Quarterly Journal of Economics*, 2000, *115* (2), 469–497.

Giacomo, Giuseppe Di and Benjamin Lerch, "Automation and Human Capital Adjustment: The Effect of Robots on College Enrollment," 2022.

Goldin, Claudia Dale and Lawrence F Katz, *The race between education and technology*, harvard university press, 2009.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," *American Economic Review*, 2020, *110* (8), 2586–2624.

Goodman-Bacon, Andrew, "The long-run effects of childhood insurance coverage: Medicaid implementation, adult health, and labor market outcomes," *American Economic Review*, 2021, *111* (8), 2550–2593.

Goos, Maarten and Alan Manning, "Lousy and lovely jobs: The rising polarization of work in Britain," *The review of economics and statistics*, 2007, *89* (1), 118–133.

Graetz, Georg and Guy Michaels, "Robots at work," *Review of Economics and Statistics*, 2018, *100* (5), 753–768.

Greenland, Andrew and John Lopresti, "Import exposure and human capital adjustment: Evidence from the US," *Journal of International economics*, 2016, *100*, 50–60.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles, "Should robots be taxed?," *The Review of Economic Studies*, 2022, *89* (1), 279–311.

Hakobyan, Shushanik and John McLaren, "Looking for local labor market effects of NAFTA," *Review of Economics and Statistics*, 2016, *98* (4), 728–741.

Han, Luyi and John V Winters, "Industry fluctuations and college major choices: Evidence from an energy boom and bust," *Economics of Education Review*, 2020, *77*, 101996.

Hansen, Bruce E, "The new econometrics of structural change: Dating breaks in US labor productivity," *Journal of Economic perspectives*, 2001, *15* (4), 117–128.

Heckman, James J, Lance Lochner, and Christopher Taber, "Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents," *Review of economic dynamics*, 1998, *1* (1), 1–58.

Hémous, David and Morten Olsen, "The rise of the machines: Automation, horizontal innovation, and income inequality," *American Economic Journal: Macroeconomics*, 2022, *14* (1), 179–223.

Hickman, Daniel C and William W Olney, "Globalization and investment in human capital," *ILR Review*, 2011, *64* (4), 654–672.

Hjort, Jonas and Jonas Poulsen, "The arrival of fast internet and employment in Africa," *American Economic Review*, 2019, *109* (3), 1032–79.

Hoynes, Hilary, Diane Whitmore Schanzenbach, and Douglas Almond, "Long-run impacts of childhood access to the safety net," *American Economic Review*, 2016, *106* (4), 903–34.

Hummels, David, Jakob R Munch, and Chong Xiang, "Offshoring and labor markets," *Journal of Economic Literature*, 2018, *56* (3), 981–1028.

International Federation of Robotics, "World Robotics Industrial Robots," 2006.

Jackson, C Kirabo, Rucker C Johnson, and Claudia Persico, "The effects of school spending on educational and economic outcomes: Evidence from school finance reforms," *The Quarterly Journal of Economics*, 2016, *131* (1), 157–218.

Jäger, Kirsten, "EU KLEMS growth and productivity accounts 2017 release-description of methodology and general notes," in "The Conference Board Europe. Available at: http://www. euklems. net/TCB/2016/Metholology_EU% 20KLEMS_2016. pdf" 2016.

Jaimovich, Nir and Henry E Siu, "How automation and other forms of IT affect the middle class: Assessing the estimates," *Brookings Economic Studies, Report*, 2019.

Johnson, Rucker C and C Kirabo Jackson, "Reducing inequality through dynamic complementarity: Evidence from Head Start and public school spending," *American Economic Journal: Economic Policy*, 2019, *11* (4), 310–49.

Jones, John Bailey and Fang Yang, "Skill-biased technical change and the cost of higher education," *Journal of Labor Economics*, 2016, *34* (3), 621–662.

Katz, Lawrence F and Kevin M Murphy, "Changes in relative wages, 1963–1987: supply and demand factors," *The quarterly journal of economics*, 1992, *107* (1), 35–78.

Krueger, Alan B, "How computers have changed the wage structure: evidence from microdata, 1984–1989," *The Quarterly Journal of Economics*, 1993, *108* (1), 33–60.

Lavecchia, Adam M, Heidi Liu, and Philip Oreopoulos, "Behavioral economics of education: Progress and possibilities," in "Handbook of the Economics of Education," Vol. 5, Elsevier, 2016, pp. 1–74.

Lawson, Nicholas, "Liquidity constraints, fiscal externalities, and optimal tuition subsidies," *American Economic Journal: Economic Policy*, 2017, *9* (4), 313–43.

Lee, David S, Justin McCrary, Marcelo J Moreira, and Jack Porter, "Valid t-ratio Inference for IV," *American Economic Review*, 2022, *112* (10), 3260–90.

Lee, Jin Young and Gary Solon, "The Fragility of Estimated Effects of Unilateral Divorce Laws on Divorce Rates," Technical Report, National Bureau of Economic Research 2011.

Lleras-Muney, Adriana, "The relationship between education and adult mortality in the United States," *The Review of Economic Studies*, 2005, *72* (1), 189–221.

Lovenheim, Michael F, "The effect of liquid housing wealth on college enrollment," *Journal of Labor Economics*, 2011, *29* (4), 741–771.

Michaels, Guy, Ashwini Natraj, and John Van Reenen, "Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years," *Review of Economics and Statistics*, 2014, *96* (1), 60–77.

National Center for Education Statistics, "Waiting to Attend College," 2005. Report (Table C-1).

NCES, "Paying for College: Changes Between 1990 and 2000 for Full-Time Dependent Undergraduates. Findings from The Condition of Education 2004 NCES.," *National Center for Education Statistics*, 2004.

Shah, Manisha and Bryce Millett Steinberg, "Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital," *Journal of Political Economy*, 2017, *125* (2), 527–561.

Shapiro, Doug, Afet Dundar, Faye Huie, Phoebe Khasiala Wakhungu, Xin Yuan, Angel Nathan, and Youngsik Hwang, "Tracking Transfer: Measures of Effectiveness in Helping Community College Students to Complete Bachelor's Degrees.(Signature Report No. 13).," *National Student Clearinghouse*, 2017.

Solon, Gary, Steven J Haider, and Jeffrey M Wooldridge, "What are we weighting for?," *Journal of Human resources*, 2015, *50* (2), 301–316.

Stuart, Bryan A, "The long-run effects of recessions on education and income," *American Economic Journal: Applied Economics*, 2022, *14* (1), 42–74.

Weinstein, Russell, "Local labor markets and human capital investments," *Journal of Human Resources*, 2022, *57* (5), 1498–1525.

Wiswall, Matthew and Basit Zafar, "Determinants of college major choice: Identification using an information experiment," *Review of Economic Studies*, 2015, *82* (2), 791–824.

Wolfers, Justin, "Did unilateral divorce laws raise divorce rates? A reconciliation and new results," *American Economic Review*, 2006, *96* (5), 1802–1820.

World Robotics, "World Robotics," 2001, pp. xiii, 348 p.:.

Young, Alwyn, "Improved, nearly exact, statistical inference with robust and clustered covariance matrices using effective degrees of freedom corrections," *Manuscript, London School of Economics*, 2016.

Figure 1: Trends in the robot market

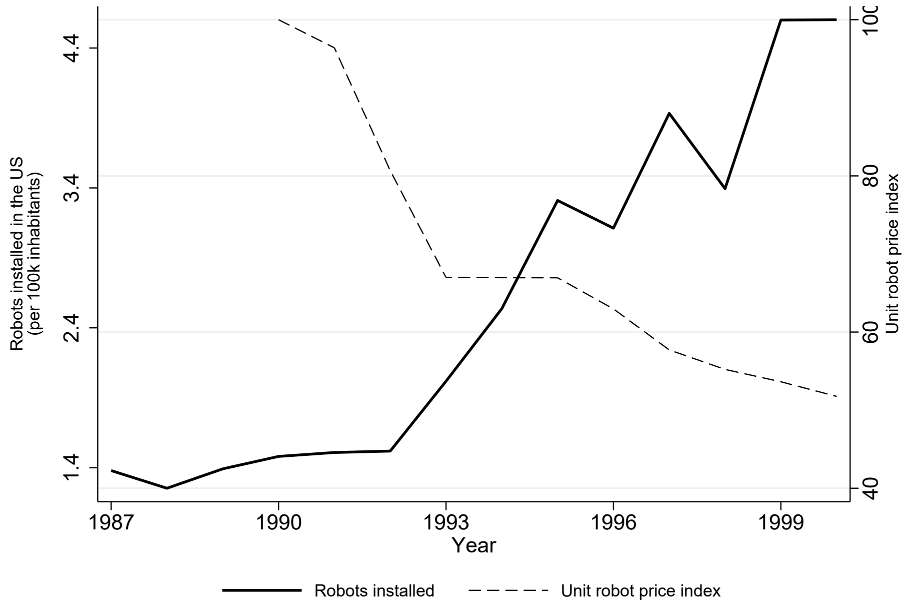

Notes. Data on newly installed robots come from the World Robotics (2001), whereas the robot price index is from the International Federation of Robotics (2006). The robot price index is calculated as an unweighted arithmetic average price index across the countries with available annual price data: United States, Germany, France, Italy, United Kingdom, and Sweden.

Figure 2: Adjusted Penetration of Robots across States

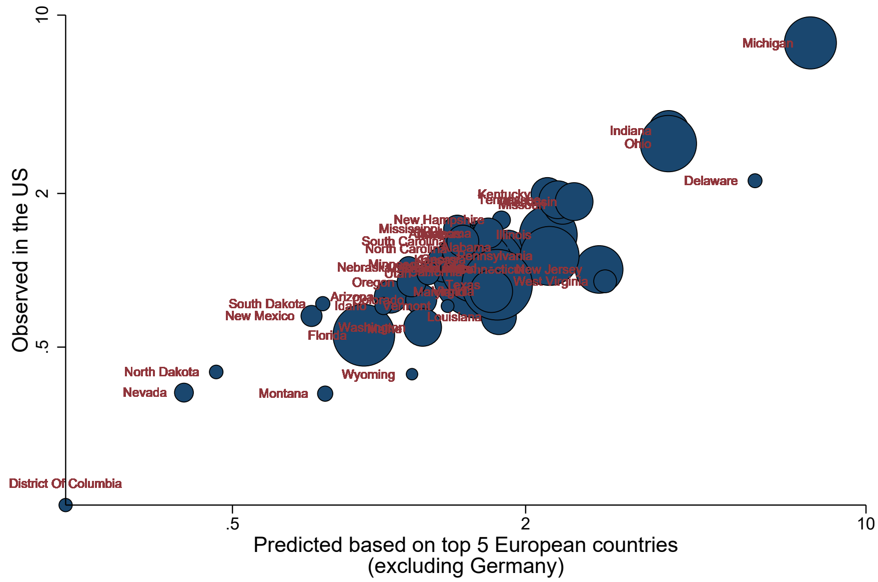

Notes. This figure plots the exposure to robots across states based on the adjusted penetration of robots in the United States and top 5 countries (excluding Germany). The adjusted penetration of robots is measured for the 2004-2007 period (rescaled to a 14-year equivalent change) for the United States, and for the 1993-2007 period for the United States.

Figure 3: Flexible Estimates

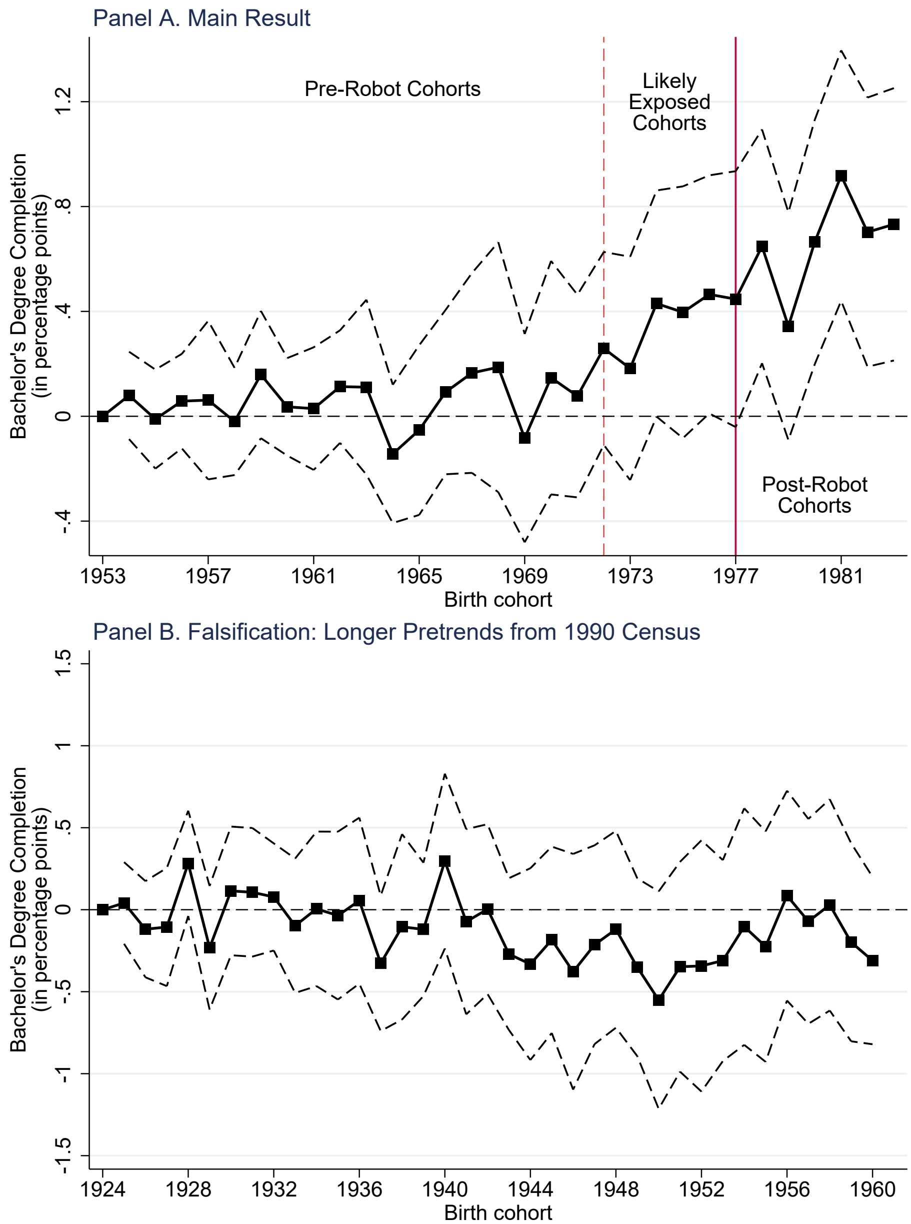

Notes. This figure plots estimates of the interaction between the robot exposure variable and indicators for birth years, using the flexible model (3). Panel A presents the main results, whereas Panel B a falsification exercise using data from the 1990 Census. The covariates include those from column 5 of Table 1. For Panel B, baseline state characteristics are constructed using the 1970 Census. See notes to column (5) of Table 1 for details on sample and specification. The dashed lines represent 95 percent confidence intervals based on standard errors clustered at the state-of-birth level.

Figure 4: Identification of the Structural Parameters

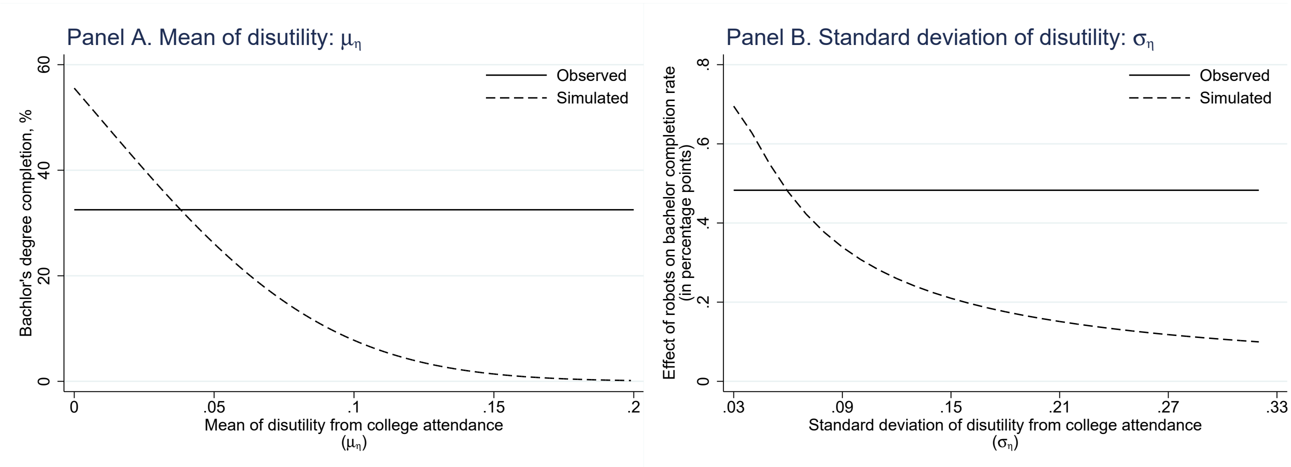

Notes. These figures illustrate the identification of the structural parameters with respect to the observed and simulated moments. The $x-$axis is the respective parameter of interest, which we vary while fixing the other parameter to its estimated value. The $y-$axis represents the corresponding observed and simulated moments. The solid lines indicate the observed value in the data, while the dashed ones indicate the value simulated from the estimated model. Each parameter is identified by finding the point on the $x-$axis where the solid and dashed lines intersect.

Figure 5: Long-Run Evolution of Bachelor Attainment and Earnings

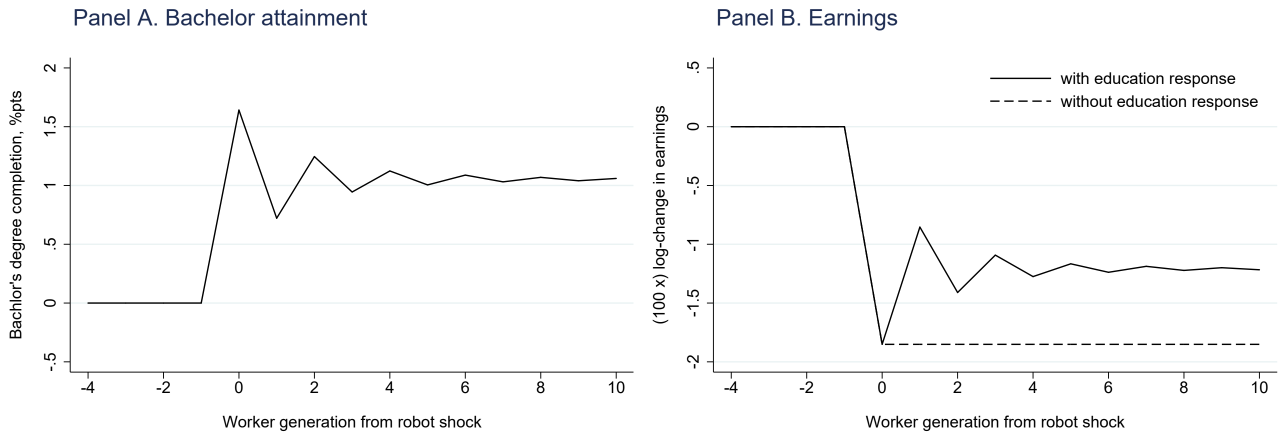

Notes. These figures shows the long-run evolution of bachelor attainment and earnings. Initial earnings and bachelor attainment are normalized to zero. We start in a steady state where approximately 30 percent of workers have a bachelor's degree, the pre-robot shock period. We then assume that each state experiences a permanent shock in robotics technology equivalent to the robot penetration observed in the data. In each moment in time, overall earnings are computed using definition 7.1. We perform these simulations separately for each state and then compute the population-weighted mean value across all states.

Figure 6: Counterfactual Policies

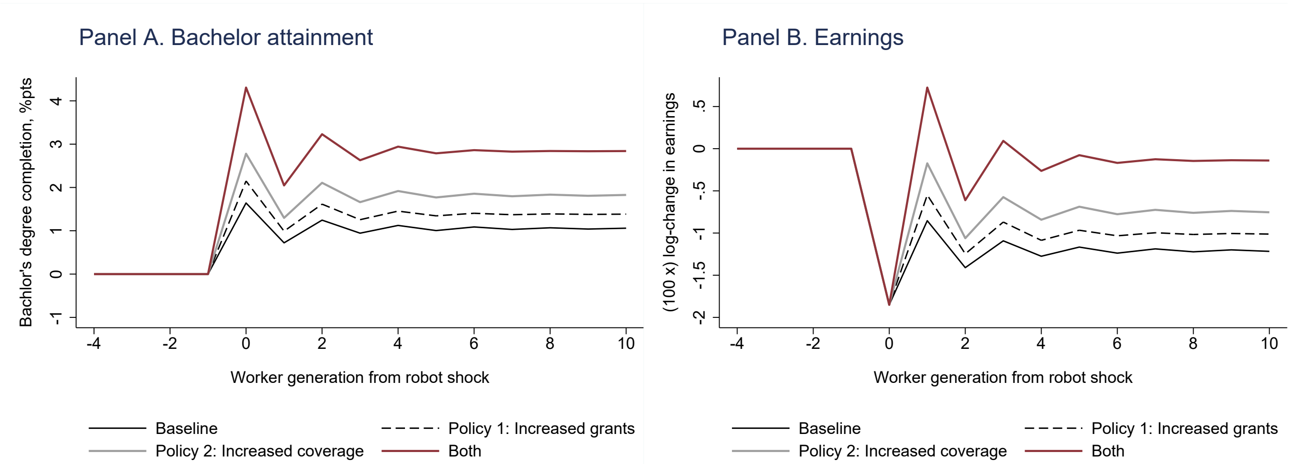

Notes. These figures shows the long-run evolution of bachelor attainment and earnings under different counterfactual policies. Initial earnings and bachelor attainment are normalized to zero. We start in a steady state where approximately 30 percent of workers have a bachelor's degree, the pre-robot shock period. We then assume that each state experiences a permanent shock in robotics technology equivalent to the robot penetration observed in the data. In each moment in time, overall earnings are computed using definition 7.1. We perform these simulations separately for each state and then compute the population-weighted mean value across all states.

Table 1: Childhood Exposure Effects on Bachelor Completion (in % pts)

	Dependent variable is Bachelor's degree completion					
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration × post	0.320	0.3632	0.4172	0.3994	0.4828	0.5053
	[0.0683]	[0.1569]	[0.2009]	[0.1800]	[0.1338]	[0.1428]
R^2	0.669	0.670	0.670	0.671	0.673	0.689
Mean Dep. Variable	32.51	32.51	32.51	32.51	32.51	34.39
Observations	15372069	15372069	15372069	15372069	15372069	5110175
1990 state demographics×birth-year FE	✓	✓	✓	✓	✓	✓
1990 manufacturing shares×birth-year FE		✓	✓	✓	✓	✓
Exposure to trade×birth-year FE			✓	✓	✓	✓
1990 share of routine jobs×birth-year FE				✓	✓	✓
1990 state college level×birth-year FE					✓	✓
Keep ACS 2015-19						✓
Birth-year FE	✓	✓	✓	✓	✓	✓
State-of-birth FE	✓	✓	✓	✓	✓	✓

Notes. This table reports estimates of β in equation (2). Coefficients shown in percentage points for ease of reading. The sample is limited to individuals who are over age 30 at survey time and born in one of the States covering the mainland of the United States. Post is an indicator for individuals born in 1974 onward. Robot penetration is the intensity of exposure to robots in one's state of birth, as described in Section 2. All regressions control for race, gender, state-of-birth, and survey-year×birth-year fixed effects. Column (1) includes interactions between birth-year fixed effects and 1990 state demographics: log of population, the share of the population over 65 years old, the share of the population under 5 years of age, the share of blacks, and the share of population that is urban. Column (2) includes interactions between birth-year fixed effects and 1990 state industry shares: the share of manufacturing employment, and the share of light manufacturing employment (textile industry and the paper, publishing, and printing industry). Column (3) includes interactions between birth-year fixed effects and exposure to Chinese imports. Column (4) includes interactions between birth-year fixed effects and the share of employment in routine jobs. The share of employment in routine jobs is defined as in Autor and Dorn (2013): routine occupations that are in the top employment-weighted third of routine task-intensity. Column (5) includes interactions between birth-year fixed effects and state college level in 1990. Column (6) restricts the sample to the ACS conducted between 2015 and 2019. Robust standard errors in brackets are clustered at the state-of-birth level.

Table 2: Childhood Exposure Effects on Different Education Groups (in % pts)

	Dependent variable is					
		High school degree				
	Less than high school	High school	High school+ less than 1 yr. college	High school + more than 1 yr. college but no degree	Associate's degree	Bachelor's degree completion
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration × post	0.0625	-0.2695	-0.0196	0.1239	-0.38	0.4828
	[0.1164]	[0.1640]	[0.0386]	[0.0647]	[0.0973]	[0.1338]
R^2	0.458	0.517	0.120	0.267	0.246	0.673
Mean Dep. Variable	7.34	27.82	7.51	14.98	9.85	32.51
Observations	15372069	15372069	15372069	15372069	15372069	15372069
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes. This table reports estimates of β in equation (2). Coefficient shown in percentage points for ease of reading. The sample is limited to individuals who are over age 30 at survey time and born in one of the States covering the mainland of the United States. Baseline covariates correspond to those reported in column (5) of Table 1. See notes to Table 1 for details on the sample and baseline covariates. Robust standard errors in brackets are clustered at the state-of-birth level.

Table 3: Childhood Exposure Effects on Bachelor Attainment (in % pts)
(Mean Reversion, Pre-cohort Trends and Within-Region Variation)

	Dependent variable is Bachelor's degree completion				
	(1)	(2)	(3)	(4)	(5)
Robot penetration×post	0.4828	0.5022	0.5022	0.4807	0.4701
	[0.1338]	[0.1430]	[0.1430]	[0.1344]	[0.1996]
R^2	0.673	0.673	0.673	0.509†	0.674
Mean Dep. Variable	32.51	32.51	32.51	32.51	32.51
Observations	15372069	15372069	15372069	15372069	15372069
1970 state college level (× birth-year FE)		✓			
1970-1990 change in state college level (× birth-year FE)			✓		
State-specfic pre-cohort trends				✓	
Region-of-birth×birth-year FE					✓
Baseline covariates	✓	✓	✓	✓	✓

Notes. This table explores the robustness of the baseline estimates to additional controls for mean reversion (columns 2-3), pre-cohort linear trends (column 4), and region-of-birth×birth-year fixed effects (column 5). The sample is limited to individuals who are over age 30 at survey time and born in one of the States covering the mainland of the United States. Robot penetration is the intensity of exposure to robots in one's state of birth, as described in Section 2. All regressions include the baseline controls included in column (3) of Table 1 (see footnotes to that table for details). The regions are defined by the US Census Bureau: the Northeast, the Midwest, the South, and the West. Robust standard errors in brackets are clustered at the state-of-birth level.
† note that the R^2 in column (4) is not comparable with that of the other columns given the two-step procedure described in Section 4.1.

Table 4: Childhood Exposure Effects on College Attainment (in % pts)
(Controlling for Other Labor Market Shocks, and Social Reforms)

	Dependent variable is Bachelor's degree completion									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Robot penetration×post	0.4828	0.4826	0.4844	0.5116	0.484	0.5047	0.4512	0.4839	0.4456	0.6123
	[0.1338]	[0.1483]	[0.1310]	[0.1393]	[0.1374]	[0.1222]	[0.1346]	[0.1341]	[0.1197]	[0.1418]
R^2	0.673	0.673	0.673	0.673	0.673	0.673	0.673	0.673	0.673	0.674
Mean Dep. Variable	32.51	32.43	32.51	32.51	32.51	32.51	32.51	32.51	32.51	32.43
Observations	15372069	15274531	15372069	15372069	15372069	15372069	15372069	15372069	15372069	15274531
Adding controls for:										
Mexican import competition		✓								✓
Offshoring			✓							✓
IT capital				✓						✓
Computer technology					✓					✓
1980-82 recession						✓				✓
War on poverty programs							✓			✓
Medicaid								✓		✓
School finance reforms									✓	✓
Baseline covariates	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes. This table reports estimates that evaluate the robustness of our baseline results to controlling for other labor market shocks and social reforms. Column (1) repeats the baseline specification reported in column (5) of Table 1 (see footnotes to that table for details). Column 2 includes interactions between birth-year fixed effects and the exposure in Mexican import competition in the state of birth. Column (3) controls for interactions between birth-year fixed effects and a Bartik measure of exposure to offshoring in the state of birth. Column (4) controls for interactions between birth-year fixed effects and the exposure to information technology capital in the state of birth. Column (5) controls for interactions between birth-year fixed effects and the exposure to computer technology in the state of birth. Column (6) controls for interactions between birth-year fixed effects and the measure of exposure to the 1980-82 recession in the state of birth (Stuart, 2022) . Column (7) includes the fraction of childhood years exposed to war-on-poverty programs for each birth cohort in the state of birth: Head Start, Food Stamp, and Community Health Centers. Column (8) includes the fraction of childhood years exposed to Medicaid in the state of birth. Column (9) adds the fraction of school-going ages (5 to 17) exposed to a school finance reform in the state of birth. The measures of exposure to Mexican imports, offshoring, IT capital, and computer techonology are measured as in Acemoglu and Restrepo (2020). Robust standard errors in brackets are clustered at the state-of-birth level.

Table 5: Effects on Bachelor's Degree Enrollment (in % pts)
(Evidence from IPEDS)

	Dependent variable is change in bachelor's degree enrollment rate	
	1990-2008	1987-1990 (placebo)
	(1)	(2)
Robot penetration	1.681	-0.187
	[0.755]	[0.278]
R^2	0.263	0.706
Mean enrollment rate in 1990	23.98	23.98
Observations	48	48
Baseline covariates	✓	✓

Notes. This table reports estimates of γ in equation (5). Coefficient shown in percentage points for ease of reading. The sample is limited to the States covering the mainland of the United States. Baseline covariates include: 1990 state college enrollment, 1990 state demographics (log of population, the share of the population over 65 years old, the share of the population under 5 years of age, the share of blacks, and the share of population that is urban), 1990 state industry shares (the share of manufacturing employment, and the share of light manufacturing employment), exposure to Chinese imports, and the share of employment in routine jobs. Standard errors in brackets are robust to arbitrary forms of heterokedasticity.

Table 6: Childhood Exposure Effects on Income (in logs)

	Dependent variable is					
	Log total income		Log earned income		Log income wages	
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration×post	0.0085	0.0041	0.0076	0.0037	0.0079	0.004
	[0.0033]	[0.0027]	[0.0030]	[0.0023]	[0.0030]	[0.0023]
Bachelor's degree completion		0.7672		0.7072		0.7021
		[0.0101]		[0.0091]		[0.0095]
R^2	0.416	0.479	0.418	0.477	0.43	0.49
Observations	14156782	14156782	12576491	12576491	11800877	11800877
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes. This table reports estimates of β in equation (2) for different income measures as outcomes. Sample sizes vary across outcomes because of missing observations. All regressions control for the baseline demographic and socioeconomic state characteristics described in Table 1. In addition, all regressions control for state-of-residence×birth-year fixed effects. Robust standard errors in brackets are clustered at the state-of-birth level.

Table 7: Effects on Market Earnings and Family Income
(in logs)

	Long differences, 1990-2008					
	ages 18-21	ages 22-65				
	No college attendance	High school	Associate's degree	High school + Associate's degree	Bachelor's degree	Log family income
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration	-0.07682	-0.01688	-0.0309	-0.01898	-0.01497	-0.01922
	[0.01676]	[0.00486]	[0.00813]	[0.00468]	[0.00429]	[0.01992]
R^2	0.250	0.210	0.248	0.218	0.127	0.715
Observations	6483	130064	39511	139487	101517	24295
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes. This table reports the results from estimating equation (6). The dependent variable in columns (1) to (5) is the log-change in market earnings for different subgroups. The dependent variable in column (6) is the log-change in family income. In column (6), the sample is limited to families where parents co-reside with children aged 17-18. The outcomes are computed within cells defined by demographic × state groups, where the demographic groups are gender, age, race, and place-of-birth. In column (6), these cells are defined by the demographic characteristcis of the household head and also include education level categories. All regressions control for the baseline demographic and socioeconomic state characteristics described in Table 1. In addition, all regressions control for the full set of demographic cell fixed effects. All regressions are weighted by the 1990 cell size. Robust standard errors in brackets are clustered at the state level.

Table 8: Earnings Coefficients and Preset Parameters

Description	Parameter	Value	Source
Panel A: Estimated earnings coefficients			
Young workers without college education: aged 18-21	κ_{young}	-0.07682	Table 7, column 1
Adult workers without a bachelor's degree: aged 22-65	κ_{0old}	-0.01898	Table 7, column 4
Adult workers with a bachelor's degree: aged 22-65	κ_{1old}	-0.01497	Table 7, column 5
Panel B: Preset parameters			
Curvature parameter of utility function	γ	1.5	Abbott et al. (2019), Attanasio and Weber (1995)
Interest and discount rate per period	$r = \rho$	0.2	Heckman et al. (1998)
Parental transfers shares	$\{\lambda_0, \lambda_1\}$	$\{0.048,\ 0.098\}$	Abbott et al. (2019)
Annual tuition and fees (1990 dollars)	e	4014.65	NCES (2004)
Estimated annual living expenses while in college (1990 dollars)	c_u	4725.42	NCES (2004)
Probability of receiving grants by family income:			
Lowest quarter	π_1	0.556	NCES (2004)
Lower middle quarter	π_2	0.453	NCES (2004)
Upper middle quarter	π_3	0.381	NCES (2004)
Highest quarter	π_4	0.329	NCES (2004)
Share of college costs covered by grants by family income:			
Lowest quarter	ϕ_1	0.811	NCES (2004)
Lower middle quarter	ϕ_2	0.583	NCES (2004)
Upper middle quarter	ϕ_3	0.495	NCES (2004)
Highest quarter	ϕ_4	0.294	NCES (2004)
Elasticity of substitution in production function	ε	1.4	Katz and Murphy (1992)

Notes. This table summarizes the earnings coefficient estimated in the first step and the parameters set externally. The earnings coefficients are estimated using the first-difference model (6). The model is estimated for each worker group separately.

Table 9: Estimated Parameters for the Structural Model

Description	Parameter	Estimate	Target moment	Data	Model
Mean of disutility from college attendance	μ_η	0.0379	Bachelor completion (% pts.)	32.506	32.547
		[0.0023]			
Standard deviation of disutility from college attendance	σ_η	0.064	Reduced-form effect of robots	0.482	0.486
		[0.0024]	on bachelor completion (% pts.)		

Notes. This table shows parameter estimates obtained using the simulated method of moments. The estimation sample consists of all youths aged 18 in the 1990 census, when they are ready to make college decisions. For a vector of possible values of structural parameters ψ, we simulate a value η_i for each individual i based on the cumulative normal distribution function $F(\cdot)$ and then solve for their college decision under a scenario with and without robots. For the latter scenario, the corresponding updated earnings profiles are simulated using the estimates of κ_{sa} obtained in the first step. The parameters are estimated by minimizing the distance between the target empirical moments and simulated moments as predicted by the model for a given vector of free parameters. The target moments are the bachelor attainment rate and the reduced-form effect of robots on bachelor attainment reported in column (5) of Table 1. The estimated coefficients and respective standard errors are reported in the third column. The empirical and simulated moments are reported in the last two columns. Standard errors reported in brackets are obtained through a bootstrap of the structural estimation.

Online Appendix to "Routine-Biased Technological Change and Endogenous Skill Investments"

Danyelle Branco Bladimir Carrillo Wilman Iglesias

May 9, 2023

List of Figures 2

List of Tables 2

A Data 4

A.1 Details on ACS and Variable Definitions . . . 4

A.2 Construction of Robot Exposure . . . 4

B Timing of Effects on Labor Markets 15

C Long-Run Labor Market Pre-Trends 17

D Defining Post-Robot Birth Years 18

E Additional Results and Robustness Checks 19

E.1 Gender Differences . . . 19

E.2 Weighted versus Unweighhted Regressions . . . 20

E.3 2SLS Estimates . . . 21

E.4 Alternative Constructions of Exposure to Robots . . . 22

E.5 Outlier Analysis . . . 23

E.6 Decomposing Variation: Rotemberg Weights . . . 24

E.7 Additional Covariates . . . 26

E.8 Robust Inference . . . 27

E.9 Robots and Exposure to Other Programs . . . 28

E.10 Supply-Side Responses . . . 30

Contact information: Branco: Department of Economics, Universidade Federal de Pernambuco, Av. Marielle Franco, Caruaru- PE, 55014-900 (e-mail: danyelle.branco@ufpe.br). Carrillo: Department of Economics, Universidade Federal de Pernambuco, AV. Prof. Moraes Rego, 1235 - Cidade Universitaria, Recife - PE, 50670-420 (e-mail: bladimir.carrillo@ufpe.br). Iglesias (e-mail: wilman.iglesias@unl.huskers.edu). The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

F Model Appendix **31**
F.1 Local Identification . . . 31
F.2 Sensitivity Tests . . . 32

References **35**

List of Figures

A.1 Birth Cohorts in the ACS . . . 7
A.2 Observations by Survey Year . . . 7
A.3 Composition of Birth Cohorts by Survey Year . . . 8
A.4 Robot Exposure Intensity by State . . . 9
A.5 Baseline Industrial Shares across States . . . 10
A.6 Robot Exposure in Selected States and Industries . . . 11
A.7 Robot Exposure across Commuting Zones after Removing State Effects . . . 12
B.1 Timing of Robot Impacts on Labor Markets . . . 16
D.1 R^2 on Different Definitions of Post_t . . . 18
F.1 Minimum Distance, Model . . . 31

List of Tables

A.1 Crosswalks between 1990 Census Bureau industrial classification and IFR Industries 13
A.2 Robot Exposure . . . 14
C.1 Labor Market Pre-Trends 1970-1990: State Level . . . 17
E.1 Childhood Exposure Effects on Bachelor Completion (in % pts) (Gender Heterogeneity) 19
E.2 Childhood Exposure Effects on Bachelor Completion (in % pts) (Weighted vs. Unweighted Estimates) . . . 20
E.3 Exposure Effects on Bachelor's Degree (in % pts) (2SLS Estimates) . . . 21
E.4 Exposure Effects on Bachelor's Degree (in % pts) (Alternative Definitions of Robot Exposure) . . . 22
E.5 Exposure Effects on Bachelor's Degree (in % pts) (Outlier Analaysis) . . . 23
E.6 Exposure Effects on Bachelor's Degree (Rotemberg Weights) . . . 25
E.7 Exposure Effects on Bachelor's Degree (in % pts) (Additional Covariates) . . . 26
E.8 Exposure Effects on Bachelor's Degree (in % pts) (Robustness to Alternative Inference Procedures) . . . 27
E.9 Robots and Exposure to Safety Net Programs and School Finance Reforms . . . 28
E.10 Exposure Effects on Bachelor's Degree (in % pts) (Controlling Flexibly for Other Social Reforms) . . . 29

E.11 Robots and Supply-Side Responses 30
F.1 Estimated Parameters for the Structural Model: Alternative Curvature Coefficients 32
F.2 Estimated Parameters for the Structural Model: Alternative Interest Rates 33
F.3 Estimated Parameters for the Structural Model: Alternative Repayment Peridos . . 34

A Data

A.1 Details on ACS and Variable Definitions

Our basic sample uses data from all the available rounds of the annual American Community Survey (ACS), ranging from 2001 to 2019. These data are publicly available from the Integrated Public Use Microdata Series (IPUMS). The samples are limited to native-born in the 1953-83 birth cohorts who are above age 30 at the time of the survey. This restriction excludes individuals from Hawaii and Alaska, so the resulting sample includes all individuals born in one of the remaining 48 states or the District of Columbia. In the ACS, the District of Columbia is considered a separate state. This sample restriction also excludes immigrants (about 10 percent of the observations), as it is not possible to infer whether or not they were exposed to automation technologies in the United States. In addition, we exclude individuals residing in institutional group quarters to increase consistency between the different rounds of the ACS. The basic sample consists of approximately 15.3 million records.

In terms of labor market outcomes, we consider total personal income, earned income, and income wages. Total personal income (INCTOT) refers to pre-tax personal income or losses from all sources for the previous year. Earned income (INCEARN) is the income earned from wages or a person's own business or farm for the previous year. Income wages (INCWAGE) represent the pre-tax wage received as an employee for the previous year. Income observations at the top of the distribution (typically 99 percentile) are top coded, with the top code value often defined as the state means of values above a given income cutoff. Following Acemoglu and Autor (2011), we replace the top code values for 1.5 times the value of the respective top code values. To render the income variables comparable across time, we convert them to constant 1999 dollars applying the CPI-U to the relevant year.

A.2 Construction of Robot Exposure

Our main analysis relies on the measure of robot exposure developed by Acemoglu and Restrepo (2020):

$$\text{Robot penetration}_s = \sum_{j \in \chi} \overbrace{\ell_{sj}}^{\text{Industry share}} \underbrace{\left(\frac{\Delta M_j}{L_{jb}} - g_j \frac{M_{jb}}{L_{jb}} \right)}_{\text{Robot Penetration}} \tag{A.1}$$

where ℓ_{sj} is the initial employment share of industry j in state s, which we calculate using the census conducted in 1970 to capture the long-term industrial composition that was prevailing before the major advance in automation. The variable $\Delta M_j = M_{j\tau} - M_{jb}$ is the change in the number of robots in each industry between the base year b and final year τ, normalized by the number of workers L_{jb}. In the model of automation developed by Acemoglu and Restrepo (2020), the labor market effects are related to the change in the number of robots per thousand workers after adjusting for the growth rate of output g_j of each industry (captured by the expression $g_j M_{jb}/L_{jb}$). For consistency with their conceptual framework and ease of comparison, we keep this adjustment term in equation (1).

Data on robots come from the International Federation of Robotics (IFR), which are available since 1993 based on yearly surveys of robot suppliers. These data cover 50 countries, including the United States, and are consistently available for 13 manufacturing and 6 non-manufacturing industry categories. The manufacturing sector is disaggregated into 13 categories (automotive, plastics and chemicals, metal products, industrial machinery, food and beverages, basic metals, electronics, miscellaneous manufacturing, minerals, wood and furniture, shipbuilding and aerospace, textiles, and paper and printing), while the remaining non-manufacturing corresponds to six broad groups (mining, education and research, agriculture, utilities, construction, and services). We use the 1993 to 2007 period to measure the adjusted penetration of robots, using data on average robot adoption in the top 5 non-US countries with greater advances in robotics (Denmark, Finland, France, Italy, and Sweden). This group excludes Germany, which is well ahead of the United States and thus is less relevant for robot adoption trends in the latter. In robustness exercises, we use measures of robot penetration expanding the top 5 to include Germany and other European countries with available data on robots. We also present results using a measure of robot penetration in the United States.

To compute the measure of adjusted penetration of robots by industry, we use industry-level data on employment from the European Union–level analysis of capital, labor, energy, materials, and service inputs (EUKLEMS) Growth and Productivity Accounts (Jäger, 2016). We use a "crosswalk" between the US industry codes in the census and IFR industry codes to match the robot penetration variable to the baseline employment shares in each state. We collapse the 199 detailed industry categories in the census into the 19 IFR industries, as detailed in Table A.1.

To sum up, we construct the overall measure of robot exposure in each state using the following step-by-step procedure in which we:

- Step 0: collapse the 199 detailed industry codes in the census to the 19 IFR industries.
- Step 1: construct the initial employment share of each industry in state s using the 1970 census.
- Step 2: compute the adjusted penetration of robots for each industry using data from the IFR and EUKLEMS database.
- Step 3: combine the results in steps 1 and 2 using equation (A.1) to generate the measure of robot exposure.

Table A.2 provides descriptive statistics for the main measure of exposure to robots, displaying the substantial variation in the adjusted penetration of robots across industries.

Cross-sectional Variation in Robot Exposure Intensity. Figure A.4 shows that there is substantial variation in the data, with a standard deviation of about 1.35 robots per thousand workers (relative to the mean of 2 robots per thousand workers). This variation stems not only from differences in robot adoption rates across industries but from substantial differences in the baseline industrial composition of employment across states. Figure A.5 shows this substantial variation in initial employment share across states.

The labor market analysis of Acemoglu and Restrepo (2020) relies on data at the commuting-zone level. Because we have no information on an individual's birthplace detailed at the commuting zone level, our analysis focuses on state-level data. While this comes at a cost in terms of loss of variation, much of the variation in the commuting-zone level data in fact stems from differences between (rather than within) states. Figure A.7 illustrates this visually. Remarkably, state fixed effects account for about 75 percent of the overall cross-commuting zone variation in robot exposure intensity. This suggests that our state-level analysis captures a substantial portion of the relevant identifying variation.

Figure A.1: Birth Cohorts in the ACS

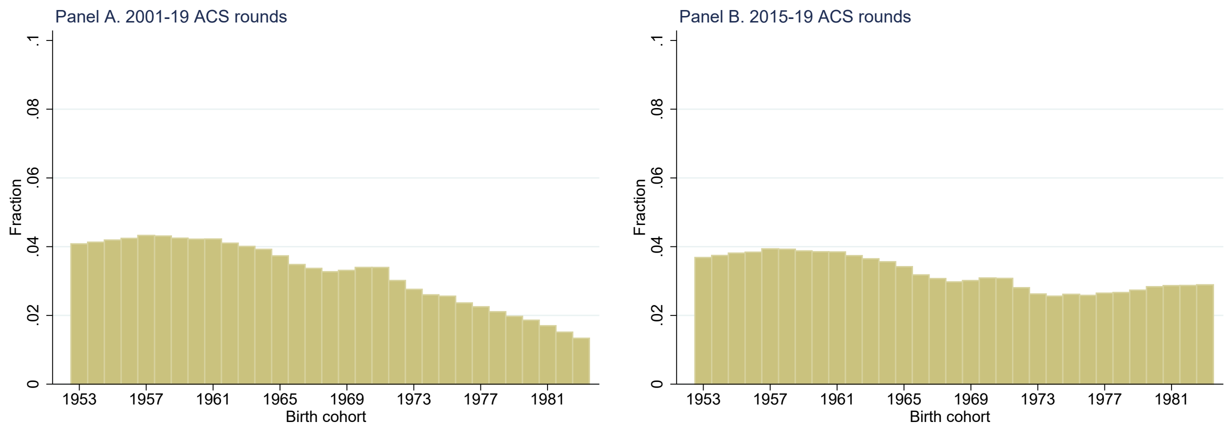

Figure A.2: Observations by Survey Year

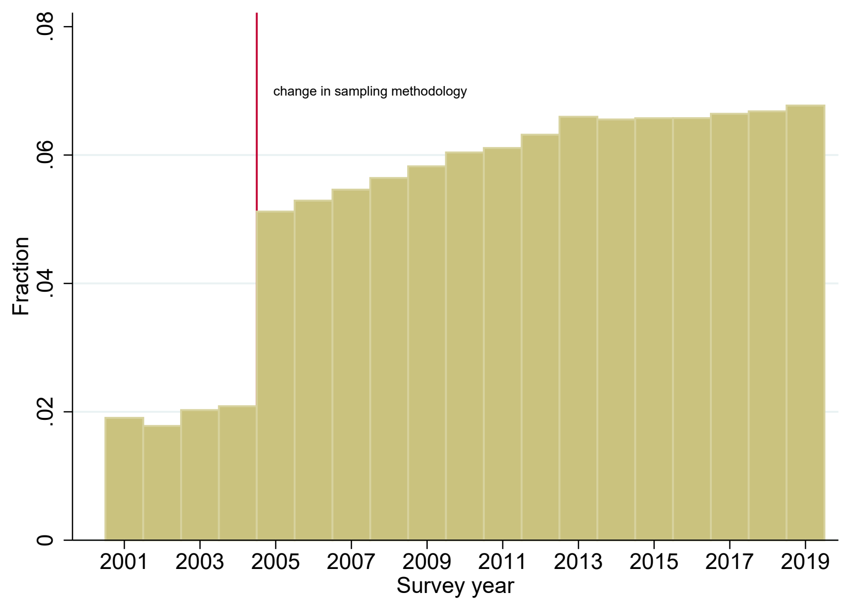

Figure A.3: Composition of Birth Cohorts by Survey Year

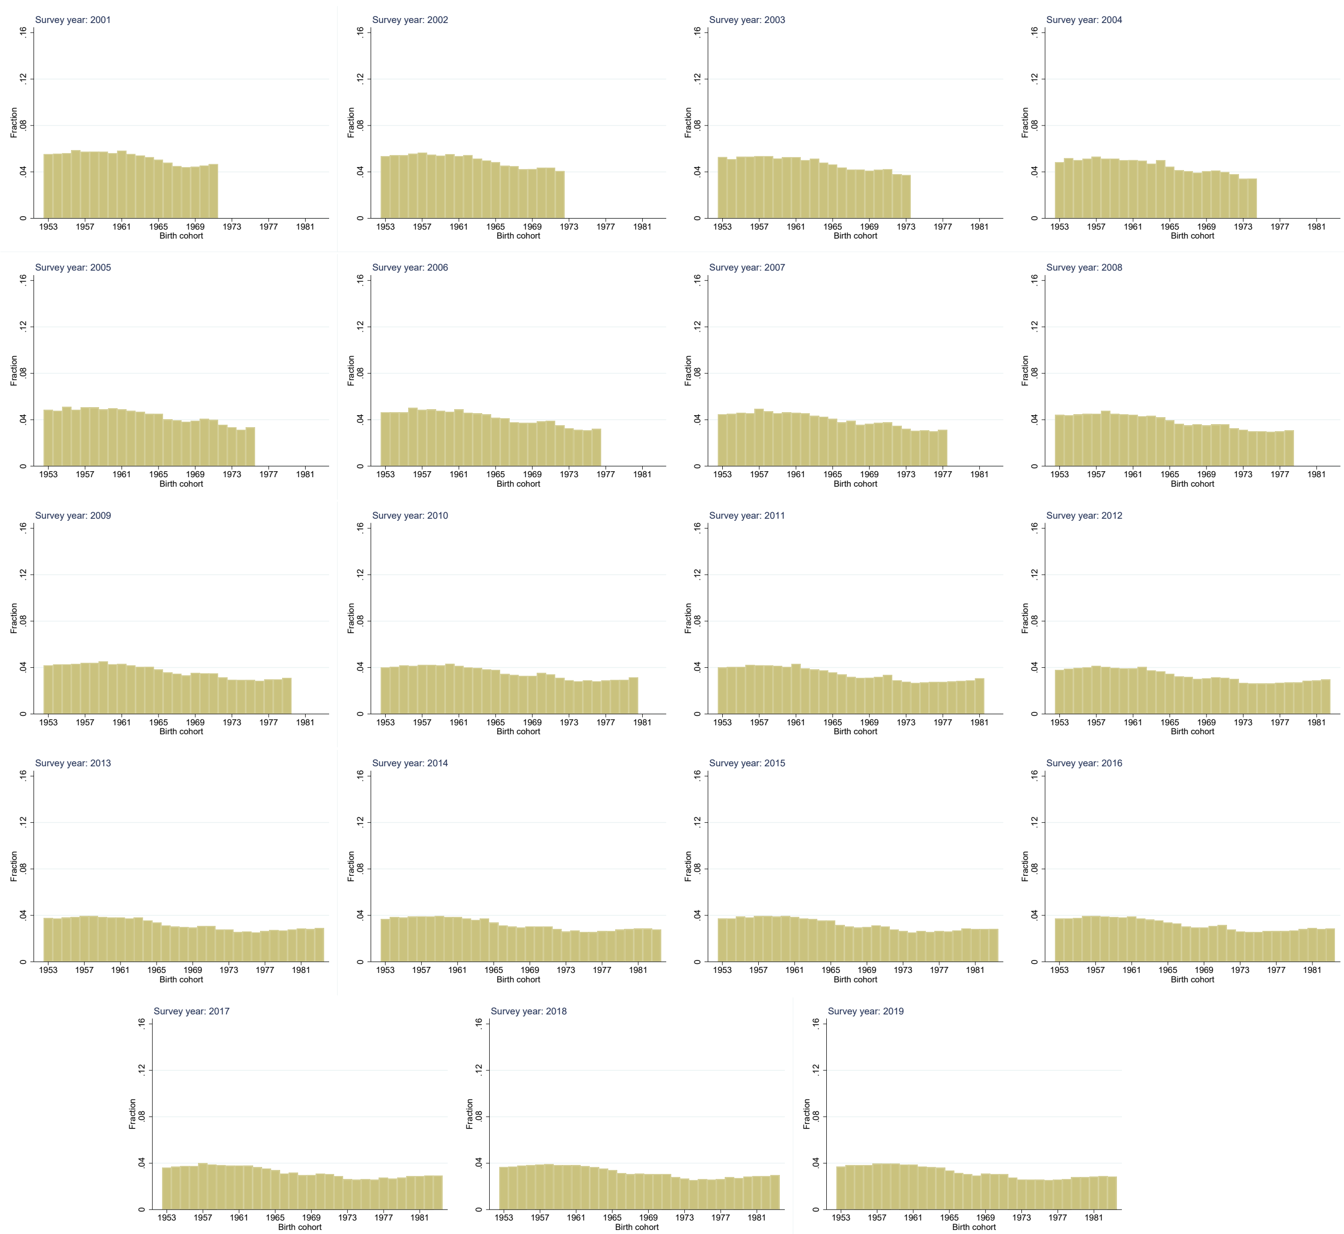

Notes. This figure shows the birth cohorts by survey year.

Figure A.4: Robot Exposure Intensity by State

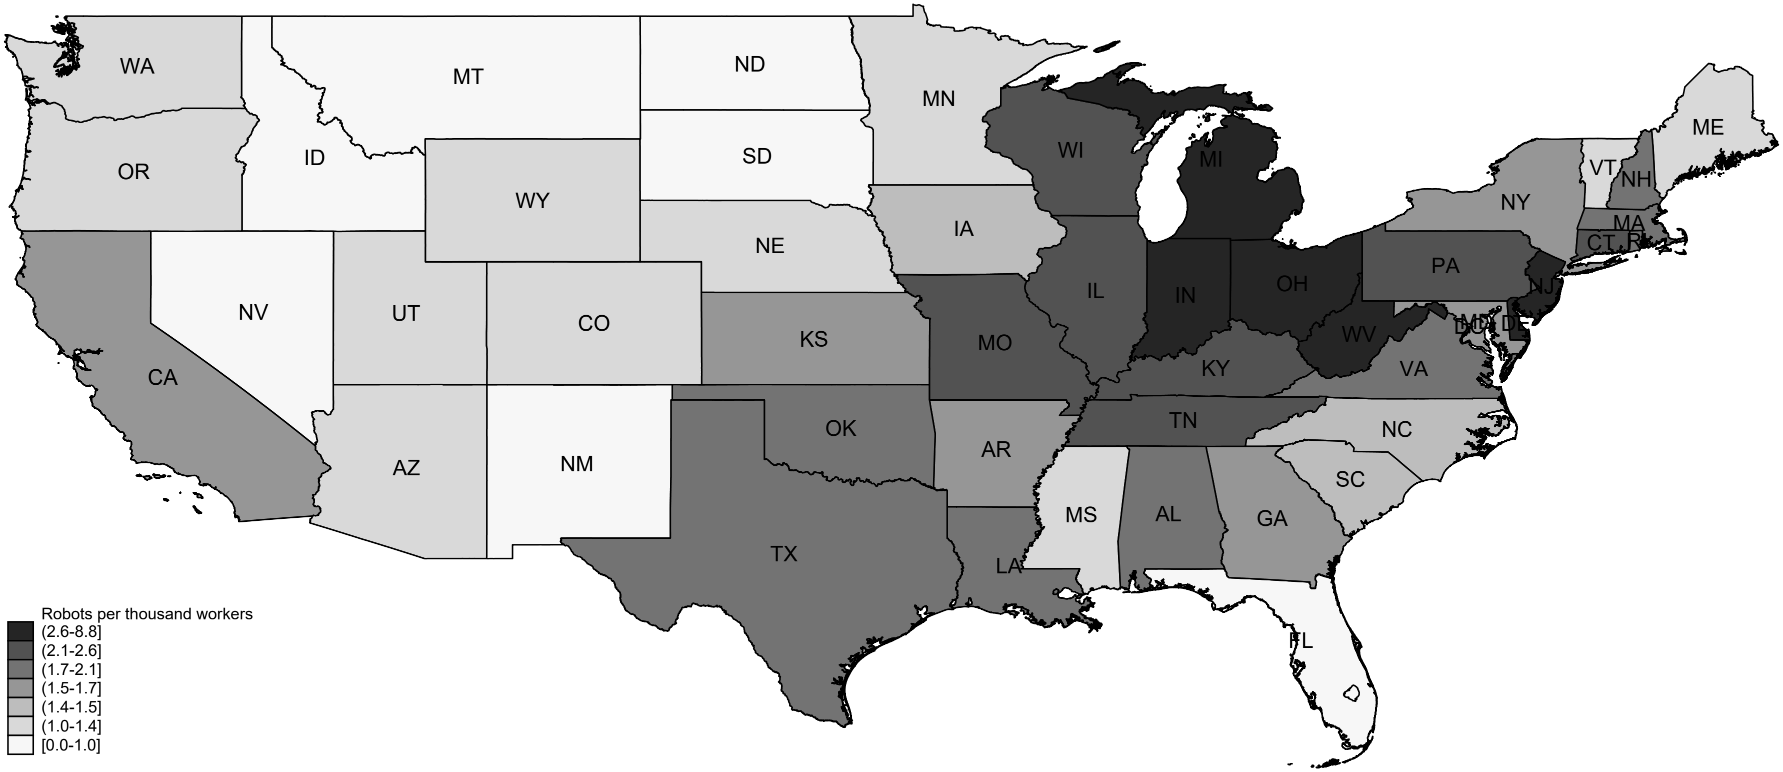

Notes. This map displays the intensity of robot exposure across states.

For Online Publication

Figure A.5: Baseline Industrial Shares across States

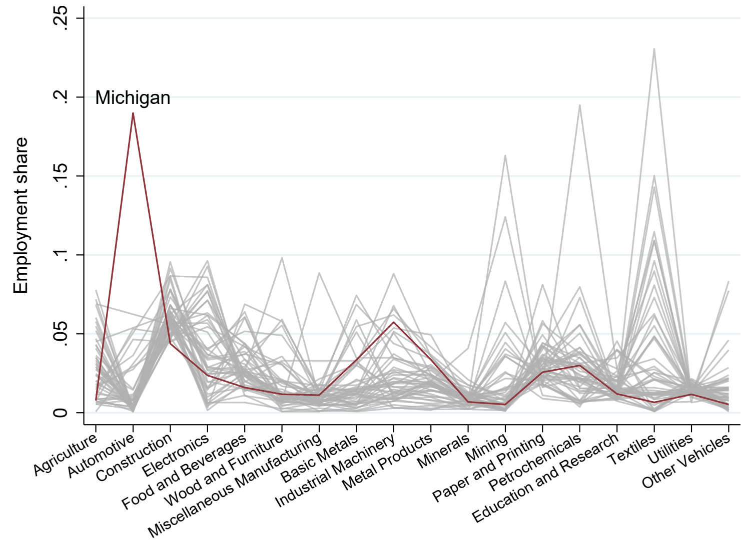

Notes. This figure shows the variation in initial industrial shares by state. This figure is constructed using data from the 1970 census.

Figure A.6: Robot Exposure in Selected States and Industries

Notes. These figures show the variation underlying the robot exposure variable. Panel A shows the initial industry's share of employment in Louisiana and Michigan. Panel B shows the adjusted penetration of robots in each industry.

Figure A.7: Robot Exposure across Commuting Zones after Removing State Effects

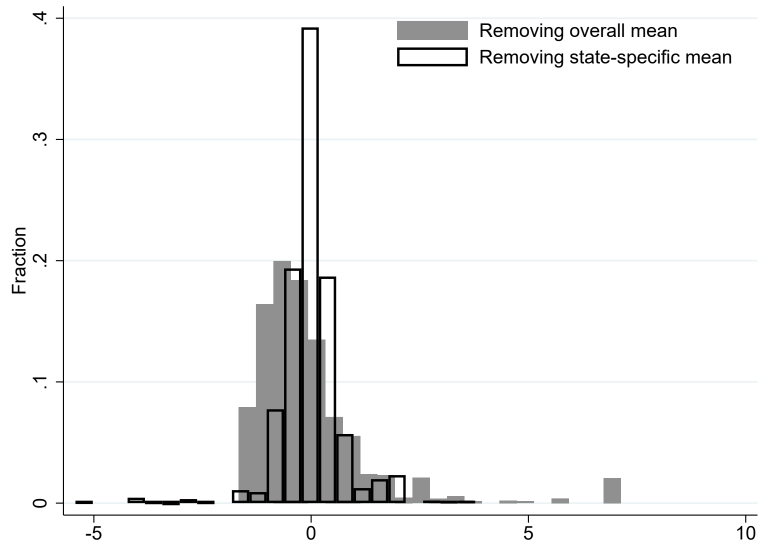

Notes. This figure shows the variation in the robot exposure across commuting zones before and after removing state fixed effects.

Table A.1: Crosswalks between 1990 Census Bureau industrial classification and IFR Industries

IFR Industry	Census Industry Code	Number of Groups
Manufacturing:		
Food and Beverages	100-130	10
Textiles	132-152, 220-222, and 450-472	6
Paper and Printing	160-172	5
Petrochemicals	180-192 and 200-212	10
Wood and Furniture	231, 241, and 242	3
Minerals	250-262	5
Basic Metals	270-272, 280, and 301	5
Metal Products	281-300	6
Industrial Machinery	310-312, 320, 331, and 332	6
Electronics	321-350 and 371-381	10
Automotive	351	1
Miscellaneous Manufacturing	391 and 392	2
Nonmanufacturing:		
Agriculture	10-32 and 2030	6
Mining	40-42 and 50	4
Construction	60	1
Shipbuilding and Aerospace	352-370	4
Services	400-442, 500-842, 870-890, and 892	101
Utilities	450-452 and 470-472	6
Education and Research	850-860 and 891	4

Notes. This table shows the crosswalks between the industry codes in the 1970 census and that in the IFR data.

Table A.2: Robot Exposure

	Mean	Standard Deviation	Observations N	Observations Aggreg. Level
Robots per thousand workers	2.09	1.35	49	States
Adjusted penetration of robots per thousand workers (overall)	4.77	8.46	19	
Adjusted penetration of robots per thousand workers...				
Manufacturing				
Automotive	32.94			
Petrochemicals	21.46			
Metal Products	8.01			
Industrial Machinery	1.01			
Food and Beverages	5.20			
Basic Metals	5.70			
Electronics	3.46			
Miscellaneous Manufacturing	-1.20			IFR industries
Minerals	2.66			
Wood and Furniture	3.65			
Shipbuilding and Aerospace	2.83			
Textiles	1.06			
Paper and Printing	0.61			
Nonmanufacturing				
Mining	2.69			
Education and Research	0.30			
Agriculture	0.16			
Utilities	0.02			
Construction	0.07			
Services	0.00			

Notes. This table provides descriptive statistics for the main measure of exposure to robots, displaying the variation in the adjusted penetration of robots across industries.

B Timing of Effects on Labor Markets

Figure 1 provides compelling evidence that robot adoption rose sharply and discontinuously in the early 1990s. In this section, we show that this sudden and large increase in robot adoption had immediate and first-order consequences on labor markets.

To estimate the dynamic effects of robots on labor markets, we use high-precision data on employment from the Bureau of Labor Statistics Quarterly Census of Wages and Employment (QCEW) at the state-year level. These data are derived from administrative tax reports submitted to state employment security agencies by all employers covered by unemployment insurance laws, accounting for about 95 percent of total administrative employment records. These data are collected since 1975. Starting in 1976, unemployment insurance laws were extended to cover a greater number of industries and establishments. This resulted in a staggered expansion of the coverage of employment in the QCEW across states between 1976 and 1980, introducing significant measurement challenges during this period (see Chodorow-Reich and Wieland (2020) for further discussion). Therefore, we exclude the 1975-1980 period from our analysis.

With these data, we estimate the following first-difference equation:

$$(\text{emp/pop})_{s,t} - (\text{emp/pop})_{s,1989} = \alpha_t \ + \ \gamma_t \text{Robot penetration}_s \ + \mathbf{Z}_s'\Omega \ + \ \xi_{s,t} \qquad \text{(B.1)}$$

where emp/pop is the employment-to-population ratio in each state s at time $t \in \{1981, 1984 \ldots, 2007\}$. The parameter of interest is γ_t, which measures the impact of robots at different moments in time. The path of these year-specific coefficients provides a detailed depiction of the dynamic effects of robots on employment. The regression control for baseline state characteristics and standard errors are adjusted to account for arbitrary heteroskedasticity.

Appendix Figure B.1 plots the set of coefficients γ_t from equation (B.1) for each year, along with 95 percent confidence intervals. The figure shows that the intensity in robot exposure is not associated with statistically meaningful changes in employment prior to 1990. These estimated coefficients are small in magnitude and statistically indistinguishable from zero. After 1990, the coefficients begin to be negative and statistically significant. By the mid-1990s, the estimated relationship becomes sizeable and rapidly increasing in magnitude.

Figure B.1: Timing of Robot Impacts on Labor Markets

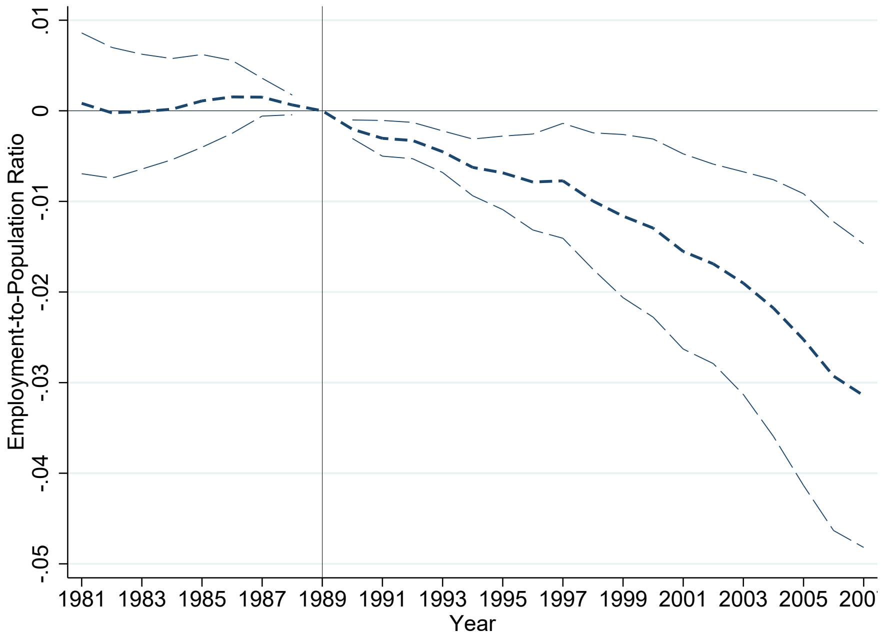

Notes: This figure presents estimates of γ_t from estimating equation (B.1) for different years. Each coefficient is emanates from a separate regression. All regressions control for baseline state characteristics. Standard errors are robust to arbitrary forms of heteroskedasticity.

C Long-Run Labor Market Pre-Trends

Table C.1: Labor Market Pre-Trends 1970-1990:
State Level

	Long differences, 1970-1990					
	Employment to population ratio	Manufacturing employment to population ratio	Employment to population ratio (include public sector and self-emp.)	Non-participation rate	Unemployment rate	Log wages
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration	0.001	0.0013	0.0003	-0.0014	0.0009	-0.0061
	[0.0016]	[0.0012]	[0.0013]	[0.0016]	[0.0013]	[0.0048]
R^2	0.371	0.797	0.508	0.544	0.584	0.717
Observations	48	48	48	48	48	7070
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes. This table reports the results from estimating the change in labor market outcomes between 1970 and 1990 on the robot penetration. The unit analysis is a state. Baseline covariates include: 1970 state college enrollment, 1970 state demographics (log of population, the share of the population over 65 years of age, the share of the population under 5 years of age, the share of blacks, and the share of population that is urban), 1970 state industry shares (the share of manufacturing employment, and the share of light manufacturing employment), exposure to Chinese imports, and the share of employment in routine jobs. The specification in column (6) is estimated at the demographic cell× state level, where demographic cells are defined by age, gender, education, and race. All regressions are weighted by population in 1970. Robust standard errors in brackets are clustered at the state level.

D Defining Post-Robot Birth Years

Given the timing of robot penetration, one would expect the effects to begin to emerge around the mid-1970s birth cohorts and the evidence above is consistent with this prediction. However, to parsimoniously and precisely summarize our findings in table format, it is necessary to define Post_t in a way that best captures the basic picture presented in Figure 3. Theory suggests that this college response should naturally be stronger for younger cohorts, but there is no precise prediction about when exactly these effects could begin to manifest. To guide our definition of Post_t and summarize our findings in table format parsimoniously, we adopt a hands-off approach that is similar in spirit to Goodman-Bacon (2021). In particular, we estimate model (2) for all possible definitions of Post_t and choose the one that maximizes the R^2, following the idea of structural break tests (Hansen, 2001).

Figure D.1 presents these results. They confirm the visual inspection of the flexible estimates reported in Figure 3. These tests suggest that the breakpoint that best captures the pattern of college responses in the data is the 1974 birth cohort. Thus, we use this definition throughout the paper.

Figure D.1: R^2 on Different Definitions of Post_t

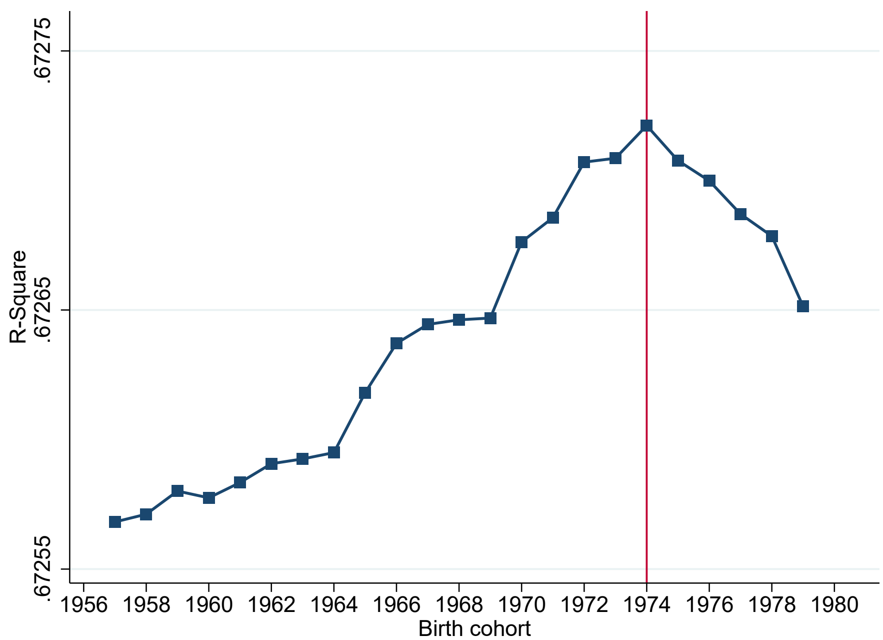

Notes. This figure presents the R^2 from estimating equation (2) for different definitions of Post_t. The covariates include those from column 5 of Table 1. See notes to column (5) of Table 1 for details on the sample and specification.

E Additional Results and Robustness Checks

E.1 Gender Differences

In Appendix Table E.1, we explore heterogeneous effects with respect to gender. We do not observe statistically meaningful differences in the estimated coefficients between males and females. At first glance, this lack of heterogeneous effects could seem surprising, given that the labor market impacts of robots on males tend to be somewhat larger (though not always statistically significant). However, a possible interpretation is that the educational responses to a same change in labor market conditions may be larger for females than males. This is in line with the recent evidence by (Charles et al., 2018) documenting that a labor market shock, which similarly affected male and female employment opportunities, had larger effects on female education. This interpretation is also consistent with the rapid growth in female schooling during our study period.

Table E.1: Childhood Exposure Effects on Bachelor Completion (in % pts) (Gender Heterogeneity)

	Dependent variable is Bachelor's degree completion	
	(1)	(2)
Robot penetration×post	0.4828	0.37717
	[0.1338]	[0.14950]
Robot penetration×post×female		0.07377
		[0.06107]
R^2	0.673	0.699
Mean Dep. Variable	32.51	32.51
Observations	15372069	15372069
Baseline covariates	✓	✓

Notes. This table tests for heterogeneity in the estimated effects by gender. We estimate model (2) interacted with a female indicator. Coefficient shown in percentage points for ease of reading. The sample is limited to individuals who are over age 30 at survey time and born in one of the States covering the mainland of the United States. Baseline covariates correspond to those reported in column (5) of Table 1. See notes to Table 1 for details on the sample and baseline covariates. Robust standard errors in brackets are clustered at the state-of-birth level.

E.2 Weighted versus Unweighhted Regressions

Because the underlying variation that we exploit is across states and birth cohorts, the use of individual-level data in the analysis effectively weights state-cohort cells by population size. As discussed by Solon et al. (2015), such weighting is appropriate in the presence of heteroskedasticity. Otherwise, it is thought to be counterproductive. We follow their practical recommendations and report both weighted and unweighted estimates. To estimate the unweighted version of model (2), we first estimate a regression of bachelor's degree attainment against individual-level characteristics (gender, race) and then collapse the residuals into state-of-birth and year-of-birth cells. We then estimate model (2) using these state-cohort level data. Appendix Table E.2 shows that the unweighted estimate is extremely similar to the weighted one (0.48 versus 0.44) but notably less precise, confirming the strengths of our main approach.

Table E.2: Childhood Exposure Effects on Bachelor Completion (in % pts) (Weighted vs. Unweighted Estimates)

	Dependent variable is Bachelor's degree completion		
		Data collapsed by state-birth-year cells	
	Baseline	Weighted	Unweighted
	(1)	(2)	(3)
Robot penetration×post	0.4828	0.4849	0.4407
	[0.1338]	[0.1568]	[0.2108]
R^2	0.673	0.978	0.962
Mean Dep. Variable	32.51	32.51	32.51
Observations	15372069	1519	1519
Baseline covariates	✓	✓	✓

Notes. This table presents weighted and unweighted estimates. Column (1) repeats the baseline results for ease of comparison. To obtain the results displayed in columns (2) and (3), we first estimate a regression of bachelor's degree completion on gender and race indicators. We then collapse the residuals into state-of-birth and year-of-birth cells. These collapsed residuals are used as the dependent variable of interest. Column (2) estimates equation (2) weighing the observations by cell group size. Column (3) estimate equation (2) without weighing the observations. Baseline covariates correspond to those reported in column (5) of Table 1. See notes to Table 1 for details. Robust standard errors in brackets are clustered at the state-of-birth level.

E.3 2SLS Estimates

We next present two-stage least squares (2SLS) estimates where our baseline, European-based robot penetration measured is used as an instrumental variable for the observed US robot penetration. These results are presented in Appendix Table E.3. Column (1) documents a powerful first-stage relationship, as had already been noted in Figure 2, with the F-statistics well above the conventional weak instrument threshold of 10. Notably, the F-statistics is also well above 100, which is important in view of the evidence in recent advances in instrumental variable estimation suggesting that it may be the relevant threshold for weak instruments (Lee et al., 2022). Turning to the second stage, we find that the 2SLS estimate is comparable to the reduced-form coefficient and corresponding OLS estimates, both in magnitude and statistical significance.

Table E.3: Exposure Effects on Bachelor's Degree (in % pts)
(2SLS Estimates)

	Dependent variable is:			
	US robot penetration × post (First stage)	Bachelor's degree completion (Reduced-form)	(OLS)	(2SLS)
	(1)	(2)	(3)	(4)
US robot penetration × post			0.373	0.4693
			[0.1506]	[0.1386]
Robot penetration × post	1.0288	0.4828		
	[0.0687]	[0.1338]		
Cragg and Donald (1993) F statistic	224.08			
Mean Dep. Variable	0.31	32.51	32.51	32.51
Observations	15372069	15372069	15372069	15372069
Baseline covariates	✓	✓	✓	✓

Notes. This table reports 2SLS estimates of the effect of exposure to robots on Bachelor's degree attainment. We instrument the US exposure to robots using exposure to robots from the top 5 European countries in terms of robot penetration. The sample is limited to individuals who are over age 30 at survey time and born in one of the States covering the mainland of the United States. All regressions include the baseline controls included in column (5) of Table 1 (see footnotes to that table for details). Robust standard errors in brackets are clustered at the state-of-birth level.

E.4 Alternative Constructions of Exposure to Robots

Table E.4: Exposure Effects on Bachelor's Degree (in % pts)
(Alternative Definitions of Robot Exposure)

		Alternative constructions of robot penetration			
	Baseline	Employment shares in 1990	Include Germany	Include all European countries with data	Unadjusted definition
	(1)	(2)	(3)	(4)	(5)
Robot penetration×post	0.4828	0.7817	0.3497	0.3225	0.3371
	[0.1338]	[0.2683]	[0.1035]	[0.0985]	[0.0947]
Rescaled coefficient	0.4828	0.4806	0.4559	0.4006	0.4969
R^2	0.673	0.673	0.673	0.673	0.673
Mean Dep. Variable	32.51	32.51	32.51	32.51	32.51
Observations	15372069	15372069	15372069	15372069	15372069
Baseline covariates	✓	✓	✓	✓	✓

Notes. This table presents results from alternative ways to construct the measure of exposure to robots. For ease of comparison, the table also reports rescaled coefficients. The rescaled coefficients are obtained by dividing the point estimates by the ratio of the standard deviations of the baseline to alternative measures of robot penetration. Column (1) repeats the baseline estimates reported in column (5) of Table 1. Column (2) uses the 1990 rather than the 1970 census to construct the initial industrial composition of employment in each state. Column (3) includes Germany to construct the adjusted penetration of robots. Column (4) uses data from all European countries to construct the adjusted penetration of robots. Column (5) uses the unadjusted penetration of robots to construct the overall measure of robot exposure. See notes to Table 1 for details on the sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

E.5 Outlier Analysis

Table E.5: Exposure Effects on Bachelor's Degree (in % pts)
(Outlier Analaysis)

	Baseline (1)	Exclude 3-sigma outliers (2)	Exclude 2-sigma outliers (3)	Exclude 1-sigma outliers (4)	Exclude 0.5-sigma outliers (5)	Exclude highly influential observations (Cook's distance) (6)
Robot penetration×post	0.4828	0.4861	0.4715	0.4717	0.5484	0.5017
	[0.1338]	[0.1347]	[0.1319]	[0.1257]	[0.1215]	[0.1074]
R^2	0.673	0.689	0.732	0.783	0.832	0.849
Mean Dep. Variable	32.51	32.49	32.45	32.39	32.47	32.47
Observations	15372069	15367892	15350905	15270926	14979839	14665860
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes. This table evaluates the robustness to outliers. Column (1) repeats the baseline results reported in column (2) of Table 1. Columns (2)-(5) exclude observations that are 3, 2, 1, and 0.5 standard deviations away from the residual mean respectively. Column (6) excludes observations that shift the baseline estimate at least to $4/N$ (Cook's distance). See notes to Table 1 for details on the sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

E.6 Decomposing Variation: Rotemberg Weights

We next investigate the relative importance of each industry for our results by computing the "Rotemberg" weights, as recommended by Goldsmith-Pinkham et al. (2020). Here the concern is that the positive effects on college attainment we find are completely driven by a particular industry, which would suggest that the results may be the product of unobservable shocks differentially affecting regions disproportionately specialized in certain types of industries. The Rotemberg weights decompose the Bartik difference-in-differences estimator into a weighted sum of estimates that use each industry share, along with the robot penetration in each industry, as a separate source of variation. Let x_{ist} and $\tilde{y}_{ist}$ denote $\textit{Robot penetration}_s \times \textit{Post}_t$ and the outcome variable after removing the basic sext of fixed effects and the rest of the control variables. Let also z_{ijst} denote $\ell_{sj} \cdot \textit{Robot penetration}_j \times \textit{Post}_t$, which is the robot exposure variable separately for each industry after filtering out the baseline covariates. Finally, let μ_j be the adjusted penetration of robots in each industry. In this case, the Rotemberg weights can be computed as follows:

$$
\begin{aligned}
& \beta = \textstyle\sum_j \alpha_j \tilde{\beta}_j \\
\text{where} \qquad & \tilde{\beta}_j \equiv \left(\textstyle\sum_i z_{ijst} \cdot x_{ist} \right)^{-1} \textstyle\sum_i z_{ijst} \cdot \tilde{y}_{ist} \\
& \alpha_j \equiv \left(\textstyle\sum_j \mu_j \textstyle\sum_i z_{ijst} \cdot x_{ist} \right)^{-1} \mu_j \textstyle\sum_i z_{ijst} \cdot x_{ist}
\end{aligned}
$$

Under this framework, $\tilde{\beta}_j$ is obtained in a 2SLS regression where the measure of robot exposure based only on industry j is used as an instrumental variable for the overall robot exposure variable. The weights $\{\alpha_j\}$ sum to one, but not all need to be positive.

Appendix Table E.6 reports thr Rotemberg weights. We find that the automotive industry has the largest share of the overall weight, with a weight above 80 percent. This is what one could expect given that the trends in robot adoption of this industry are almost of incomparable magnitude to that of any other industry. But most importantly, the automotive industry is not the only reason why we observe the positive effects of robots on college attainment. As shown in the table, the estimated coefficient β is in fact larger when we exclude the automotive industry.

Table E.6: Exposure Effects on Bachelor's Degree (Rotemberg Weights)

	Rotemberg weights		Estimate of β (in % pts) excluding each industry in row
	Raw data	Baseline covariates	(Baseline covariates)
	(1)	(2)	(3)
Automotive	0.8496	1.0125	1.1618
Plastics and Chemicals	0.1060	0.0826	0.5252
Basic Metals	0.0367	-0.0130	0.4550
Metal Products	0.0291	-0.0130	0.4625
Industrial Machinery	0.0080	-0.0015	0.4920
Electronics	0.0024	-0.0189	0.4817
Minerals	0.0021	-0.0011	0.4834
Paper and Printing	0.0001	0.0001	0.4901
Utilities	0.0000	0.0000	0.4820
Education and Research	-0.0003	0.0001	0.4838
Construction	-0.0003	-0.0001	0.4798
Miscellaneous Manufacturing	-0.0006	-0.0004	0.4831
Agriculture	-0.0008	-0.0007	0.4771
Mining	-0.0047	-0.0011	0.4963
Wood and Furniture	-0.0051	-0.0152	0.4866
Shipbuilding and Aerospace	-0.0061	-0.0174	0.4938
Textiles	-0.0070	0.0091	0.4844
Food and Beverages	-0.0092	-0.0221	0.4608

Notes. This table decomposes the baseline coefficient β into a weighted sum of estimates that use each industry share, along with the robot penetration in each industry, as a separate source of variation. Columns (1) and (2) presents Rotemberg weights for all industries, following Goldsmith-Pinkham et al. (2020). Column (3) shows the estimated coefficient when each industry is excluded from the overall measure of robot exposure. See notes to Table 1 for details on the sample and specification.

E.7 Additional Covariates

Table E.7: Exposure Effects on Bachelor's Degree (in % pts)
(Additional Covariates)

	Dependent variable is Bachelor's degree completion			
	(1)	(2)	(3)	(4)
Robot penetration×post	0.4828	0.5221	0.4152	0.4819
	[0.1338]	[0.1238]	[0.1628]	[0.1561]
R^2	0.673	0.673	0.674	0.674
Mean Dep. Variable	32.51	32.51	32.51	32.51
Observations	15372069	15372069	15372069	15372069
1990 employment-to-pop. ratio (× birth-year FE)		✓		✓
1990-2008 change in demographics (× birth-year FE)			✓	✓
Baseline covariates	✓	✓	✓	✓

Notes. This table explores the robustness of the baseline estimates to additional controls. The sample is limited to individuals over age 30 at survey time and born in one of the States covering the mainland of the United States. All regressions include the baseline controls included in column (5) of Table 1 (see footnotes to that table for details). Robust standard errors in brackets are clustered at the state-of-birth level.

E.8 Robust Inference

Our baseline analysis uses standard errors clustered at the state-of-birth level. In this section, we evaluate the robustness of our results to alternative inference approaches. First, we use standard errors clustered at the state level but adjust them by the effective sample size implied by the relative importance of each observation, as suggested by Young (2016). Second, because we are using a shift-share identification strategy, a particular concern is that standard procedures to inference may result in smaller standard errors if residuals are spatially correlated across areas with similar sectoral shares (Borusyak et al., 2022). Therefore, we evaluate the robustness of the results using the inference procedures proposed by Adao et al. (2019) and Borusyak et al. (2022) that address cross-region correlation in residuals in shift-share designs. Finally, we present results from a specification that uses standard errors two-way clustered at the state-of-birth and birth-year level, which account for possible serial and spatial correlation in a flexible manner. As shown in Table E.8, the results are in general very similar to our baseline.

Table E.8: Exposure Effects on Bachelor's Degree (in % pts)
(Robustness to Alternative Inference Procedures)

		Alternative inference procedures			
	Baseline	Clustered by state + Young (2016) effective d.o.f.-adj.	Borusyak et al. (2022) robust SE	Adao et al. (2019) robust SE	Twoway clustering by state + birth year
	(1)	(2)	(3)	(4)	(5)
Robot penetration×post	0.4828	0.4828	0.4828	0.4828	0.4828
	[0.1338]	[0.1447]	[0.1216]	[0.1103]	[0.1341]
R^2	0.673	0.673	0.673	0.673	0.673
Mean Dep. Variable	32.51	32.51	32.51	32.51	32.51
Observations	15372069	15372069	15372069	15372069	15372069
Baseline covariates	✓	✓	✓	✓	✓

Notes. This table evaluates the robustness of the baseline results in Table 1 to alternative inference approaches: *i*) standard errors clustered at the state level but adjusted by the effective sample size implied by the relative importance of each observation, as suggested by Young (2016); *ii*) inference procedures proposed by Adao et al. (2019) and Borusyak et al. (2022) that address cross-region correlation in residuals in shift-share designs; *iii*) two-way clustering by state-of-birth and birth year. See notes to Table 1 for details on the sample and specification.

E.9 Robots and Exposure to Other Programs

We examine the extent to which cohort exposure to robots predicts the probability of childhood exposure to Community Health Centers, Head Start, Food Stamp, Medicaid and School Finance Reforms. As in Goodman-Bacon (2021), we consider ages 0 to 9 as the relevant window of exposure for Community Health Centers, Food Stamp and Medicaid, and ages 3 to 4 for Head Start. For the school finance reforms, we use the ages 5 to 17 as the relevant window of exposure (as in Jackson et al. (2016)). For each of these programs, we generate a variable measuring the fraction of the relevant years that a given cohort was exposed to the reform or program. We then estimate our baseline specification (2) using these measures of program exposure as dependent variables. Table E.9 documents that the post-robot cohorts from states with greater robot penetration are not significantly more likely to have been exposed to these programs in childhood.

Table E.9: Robots and Exposure to Safety Net Programs and School Finance Reforms

	Fraction of relevant years of exposure to...				
	Community Health Centers	Head Start	Food Stamp	Medicaid	School Finance Reforms
	(1)	(2)	(3)	(4)	(5)
Robot penetration×post	0.0246	0.0023	0.0028	-0.005	-0.0575
	[0.0181]	[0.0156]	[0.0138]	[0.0137]	[0.0371]
R^2	0.915	0.866	0.967	0.975	0.892
Mean Dep. Variable	0.210	0.150	0.520	0.580	0.300
Observations	15372069	15372069	15372069	15372069	15372069
Baseline covariates	✓	✓	✓	✓	✓

Notes. This table estimates the effects of robots on the share of childhood years exposed to each program and reform. As in Goodman-Bacon (2021), we consider ages 0 to 9 as the relevant window of exposure for Community Health Centers, Food Stamp, and Medicaid, and ages 3 to 4 for Head Start. For the school finance reforms, we use the ages 5 to 17 as the relevant window of exposure (as in Jackson et al. (2016)). See notes to Table 1 for details on the sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

Table E.10: Exposure Effects on Bachelor's Degree (in % pts)
(Controlling Flexibly for Other Social Reforms)

		Controlling for birth-cohort FE x adoption year of...					
	Baseline	Community Health Centers	Head Start	Food Stamp	Medicaid	School Finance Reforms	All simultaneously
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Robot penetration×post	0.4828	0.4819	0.554	0.51	0.4736	0.5678	0.5814
	[0.1338]	[0.1345]	[0.1417]	[0.1294]	[0.1316]	[0.1543]	[0.1558]
R^2	0.673	0.673	0.673	0.673	0.673	0.673	0.674
Mean Dep. Variable	32.51	32.51	32.51	32.51	32.51	32.51	32.51
Observations	15372069	15372069	15372069	15372069	15372069	15372069	15372069
Baseline covariates	✓	✓	✓	✓	✓	✓	✓

Notes. This table demonstrates the robustness of the baseline estimates to controlling flexibly for the timing of war-on-poverty programs, Medicaid, and School Finance reforms. Columns (2)-(6) repeat the baseline specification, but separately include birth-cohort fixed effects interacted with the year of each program or adoption across states. See notes to Table 1 for details on the sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

E.10 Supply-Side Responses

Table E.11: Robots and Supply-Side Responses

	Log-differences 1990-2008:			
	Average net tuition and fees costs	Percapita revenue from state and local appropiations	Per capita revenue from state and local grants	Per capita Government transfers in education and training assistance
	(1)	(2)	(3)	(4)
Robot penetration	0.0015	0.0014	0.0059	-0.017
	[0.0099]	[0.0288]	[0.0323]	[0.045]
R^2	0.0001	0.0001	0.0002	0.319
Mean Dep. Variable	1.13	0.35	1.63	1.06
Observations	49	49	49	49
Baseline covariates	✓	✓	✓	✓

Notes. This table reports results from estimating equation (6): $\Delta Y_{s,90-08} = \alpha + \gamma Robots_s + \mathbf{Z}'_s\Omega + \xi_s$. The unit of analysis is a state. All regressions control for the baseline demographic and socioeconomic state characteristics described in Table 1. Standard errors are robust to arbitrary forms of heteroskedasticity.

F Model Appendix

F.1 Local Identification

This subsection provides further evidence that the targeted moments effectively inform the estimated parameters. Figure F.1 plots the value of the criterion function for different values of a given parameter, holding the other fixed at its estimated value. As one can infer from the figure, the targeted moments provide relevant idenfitication information. The criterion function is strictly convex around the estimated parameters and consequently the coefficients are uniquely identified.

Figure F.1: Minimum Distance, Model

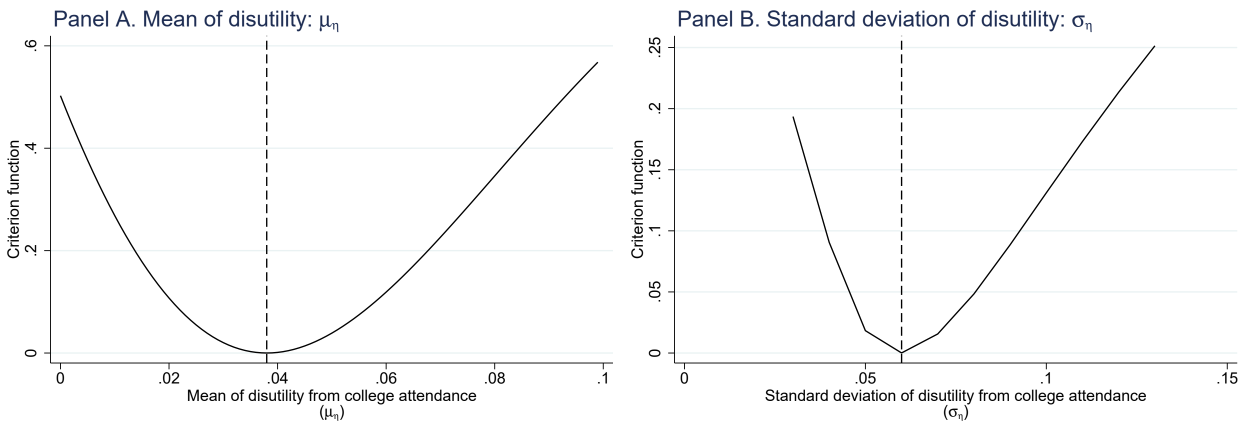

Notes. These figures show the impact of varying one parameter at time on the criterion function, holding the other parameter fixed at its estimated value. The $x-$axis is the respective parameter of interest, while the $y-$axis represents the value of the criterion function.

F.2 Sensitivity Tests

Table F.1: Estimated Parameters for the Structural Model:
Alternative Curvature Coefficients

Description	Parameter	Estimate	Target moment	Data	Model
			Panel A: $\gamma = 1$ *(log utility)*		
Mean of disutility from college attendance	μ_η	3.947 [0.288]	Bachelor completion (% pts.)	32.506	32.110
Standard deviation of disutility from college attendance	σ_η	4.630 [0.293]	Reduced-form effect of robots on bachelor completion (% pts.)	0.483	0.476
			Panel B: $\gamma = 2$		
Mean of disutility from college attendance	μ_η	0.00039 [0.00005]	Bachelor completion (% pts.)	32.506	32.322
Standard deviation of disutility from college attendance	σ_η	0.00078 [0.00001]	Reduced-form effect of robots on bachelor completion (% pts.)	0.483	0.482

Notes. This table shows the robustness of the results to assuming alternative values for the curvature parameter of the utility function. The estimates are obtained using the simulated method of moments. The estimation sample consists of all youths aged 18 in the 1990 census, when they are ready to make college decisions. For a vector of possible values of structural parameters ψ, we simulate a value η_i for each individual i based on the cumulative normal distribution function $F(\cdot)$ and then solve for their college decision under a scenario with and without robots. For the latter scenario, the corresponding updated earnings profiles are simulated using the estimates of κ_{sa} obtained in the first step. The parameters are estimated by minimizing the distance between the target empirical moments and simulated moments as predicted by the model for a given vector of free parameters. The target moments are the bachelor attainment rate and the reduced-form effect of robots on bachelor attainment reported in column (5) of Table 1. The estimated coefficients and respective standard errors are reported in the third column. The empirical and simulated moments are reported in the last two columns. Standard errors reported in brackets are obtained through a bootstrap of the structural estimation.

Table F.2: Estimated Parameters for the Structural Model: Alternative Interest Rates

Description	Parameter	Estimate	Target moment	Data	Model
			Panel A: annual interest rate=1%		
Mean of disutility from college attendance	μ_η	0.06376 [0.00254]	Bachelor completion (% pts.)	32.5062	32.555
Standard deviation of disutility from college attendance	σ_η	0.08164 [0.0025]	Reduced-form effect of robots on bachelor completion (% pts.)	0.4828	0.5032
			Panel B: annual interest rate=2%		
Mean of disutility from college attendance	μ_η	0.05768 [0.00252]	Bachelor completion (% pts.)	32.5062	32.5557
Standard deviation of disutility from college attendance	σ_η	0.0804 [0.00249]	Reduced-form effect of robots on bachelor completion (% pts.)	0.4828	0.4484
			Panel C: annual interest rate=3%		
Mean of disutility from college attendance	μ_η	0.04912 [0.00249]	Bachelor completion (% pts.)	32.5062	32.5925
Standard deviation of disutility from college attendance	σ_η	0.0729 [0.00256]	Reduced-form effect of robots on bachelor completion (% pts.)	0.4828	0.4535
			Panel D: annual interest rate=4%		
Mean of disutility from college attendance	μ_η	0.04185 [0.00248]	Bachelor completion (% pts.)	32.5062	32.6744
Standard deviation of disutility from college attendance	σ_η	0.0661 [0.00243]	Reduced-form effect of robots on bachelor completion (% pts.)	0.4828	0.4686

Notes. This table shows the robustness of the results to assuming alternative values for the annual interest rate. Note that the interest rate per period is equal 4(×) the annual interest rate, since each period corresponds to 4 years. The estimates are obtained using the simulated method of moments. The estimation sample consists of all youths aged 18 in the 1990 census, when they are ready to make college decisions. For a vector of possible values of structural parameters ψ, we simulate a value η_i for each individual i based on the cumulative normal distribution function $F(\cdot)$ and then solve for their college decision under a scenario with and without robots. For the latter scenario, the corresponding updated earnings profiles are simulated using the estimates of κ_{sa} obtained in the first step. The parameters are estimated by minimizing the distance between the target empirical moments and simulated moments as predicted by the model for a given vector of free parameters. The target moments are the bachelor attainment rate and the reduced-form effect of robots on bachelor attainment reported in column (5) of Table 1. The estimated coefficients and respective standard errors are reported in the third column. The empirical and simulated moments are reported in the last two columns. Standard errors reported in brackets are obtained through a bootstrap of the structural estimation.

Table F.3: Estimated Parameters for the Structural Model: Alternative Repayment Peridos

Description	Parameter	Estimate	Target moment	Data	Model
			Panel A: loans last 4 periods		
Mean of disutility from college attendance	μ_η	0.03925 [0.00251]	Bachelor completion (% pts.)	32.5062	33.1022
Standard deviation of disutility from college attendance	σ_η	0.06697 [0.00247]	Reduced-form effect of robots on bachelor completion (% pts.)	0.4828	0.4711
			Panel B: loans last 6 periods		
Mean of disutility from college attendance	μ_η	0.03954 [0.00239]	Bachelor completion (% pts.)	32.5062	32.5649
Standard deviation of disutility from college attendance	σ_η	0.06505 [0.00253]	Reduced-form effect of robots on bachelor completion (% pts.)	0.4828	0.478
			Panel C: loans last 8 periods		
Mean of disutility from college attendance	μ_η	0.04108 [0.00245]	Bachelor completion (% pts.)	32.5062	31.8347
Standard deviation of disutility from college attendance	σ_η	0.06582 [0.00248]	Reduced-form effect of robots on bachelor completion (% pts.)	0.4828	0.4909

Notes. This table shows the robustness of the results to assuming alternative repayment periods. The estimates are obtained using the simulated method of moments. The estimation sample consists of all youths aged 18 in the 1990 census, when they are ready to make college decisions. For a vector of possible values of structural parameters ψ, we simulate a value η_i for each individual i based on the cumulative normal distribution function $F(\cdot)$ and then solve for their college decision under a scenario with and without robots. For the latter scenario, the corresponding updated earnings profiles are simulated using the estimates of κ_{sa} obtained in the first step. The parameters are estimated by minimizing the distance between the target empirical moments and simulated moments as predicted by the model for a given vector of free parameters. The target moments are the bachelor attainment rate and the reduced-form effect of robots on bachelor attainment reported in column (5) of Table 1. The estimated coefficients and respective standard errors are reported in the third column. The empirical and simulated moments are reported in the last two columns. Standard errors reported in brackets are obtained through a bootstrap of the structural estimation.

References

Acemoglu, Daron and David Autor, "Skills, tasks and technologies: Implications for employment and earnings," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

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

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental shift-share research designs," Technical Report 1 2022.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo, "Housing Booms and Busts, Labor Market Opportunities, and College Attendance," *American Economic Review*, October 2018, *108* (10), 2947–94.

Chodorow-Reich, Gabriel and Johannes Wieland, "Secular labor reallocation and business cycles," *Journal of Political Economy*, 2020, *128* (6), 2245–2287.

Cragg, John G and Stephen G Donald, "Testing identifiability and specification in instrumental variable models," *Econometric Theory*, 1993, pp. 222–240.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," *American Economic Review*, 2020, *110* (8), 2586–2624.

Goodman-Bacon, Andrew, "The long-run effects of childhood insurance coverage: Medicaid implementation, adult health, and labor market outcomes," *American Economic Review*, 2021, *111* (8), 2550–2593.

Hansen, Bruce E, "The new econometrics of structural change: Dating breaks in US labor productivity," *Journal of Economic perspectives*, 2001, *15* (4), 117–128.

Jackson, C Kirabo, Rucker C Johnson, and Claudia Persico, "The effects of school spending on educational and economic outcomes: Evidence from school finance reforms," *The Quarterly Journal of Economics*, 2016, *131* (1), 157–218.

Jäger, Kirsten, "EU KLEMS growth and productivity accounts 2017 release-description of methodology and general notes," in "The Conference Board Europe. Available at: http://www. euklems. net/TCB/2016/Metholology_EU% 20KLEMS_2016. pdf" 2016.

Lee, David S, "Training, wages, and sample selection: Estimating sharp bounds on treatment effects," *The Review of Economic Studies*, 2009, *76* (3), 1071–1102.

Solon, Gary, Steven J Haider, and Jeffrey M Wooldridge, "What are we weighting for?," *Journal of Human resources*, 2015, *50* (2), 301–316.

Young, Alwyn, "Improved, nearly exact, statistical inference with robust and clustered covariance matrices using effective degrees of freedom corrections," *Manuscript, London School of Economics*, 2016.