Routine-Biased Technological Change and Endogenous Skill Investments*

Danyelle Branco Bladimir Carrillo Wilman Iglesias

July 28, 2022

Abstract

Although an extensive and influential literature examines the effects of labor-displacing technologies on labor markets, very little is known about whether and how individuals alter their human capital investments in response to these technological innovations. This paper provides detailed empirical evidence on this question by examining the consequences of the unprecedented advance in robotics technology in the United States. Our research design exploits variation in the penetration of robots across locations and in the timing of exposure across birth cohorts in a cross-cohort identification strategy. Our results show that cohorts differentially exposed to robots before or during the typical college-going ages are significantly more likely to complete a Bachelor's degree and experience an increase (or a smaller decline) in their labor market earnings. Empirical tests suggest that changes in the college premium and opportunity costs are the key mechanisms generating these effects. We then estimate a structural model of human capital investments to evaluate mechanisms and the importance of these effects for the earnings inequality. Mapping this model to the data, we find that the skill premium is the single most important component of our results, accounting for approximately two-thirds of the overall effect. Further simulations from the estimated model indicate that the effect of robots on earnings inequality declines substantially over time as younger worker generations with different educational choices enter the economy. These findings have important implications for the role of skill investments for the adjustment of the economy to technology in models of skill-biased technological change.

JEL codes: I21, J23, J24

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

*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: 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, 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

A long-standing question dating back to the times of Adam Smith concerns the effects of industrial technologies on the economy. Despite the extensive and influential literature on the consequences of labor-saving technologies for the demand for labor and wage rates, very little is known about whether and how individuals alter their skill investments in response to these technological innovations. Answers to these questions have significant implications for how we model skill-biased technological change, how changes in technology affect the structure of earnings over the long run, and how the government should guide the target of policies intended to reduce technology-driven inequalities by promoting education.

Skill-biased technologies potentially alter the incentives to make human capital investments. If high-skill professional occupations requiring critical thinking, management, or other non-routine skills are shielded from automation, then the premium to higher education should rise as technology advances. This premium effect, together with declines in the opportunity cost of time, should increase the incentives for individuals to invest in high-skill professions. In this paper, we offer evidence on the effects of a major automation technology on individuals' educational decisions, document in detail the mechanisms at work, and estimate a structural model of educational choice to evaluate the importance of this endogenous response for the dynamic of earnings inequality in the United States.

Though a large literature examines the effects of economic shocks on human capital, existing research has largely focused on variations in economic conditions that primarily affect individuals at the bottom of the skill distribution, such as changes in the construction sector (Charles et al., 2018), trade-induced changes in low-skill labor demand (Atkin, 2016; Greenland and Lopresti, 2016) or shocks to the natural resources and agricultural industries (Shah and Steinberg, 2017; Cascio and Narayan, 2015; Carrillo, 2020).[1] 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). The extent to which these technological changes induce skill acquisition in advanced economies depends on whether individuals on this margin effectively transition to Bachelor-level college. It is not obvious that this will be the case: college is an expensive, long-duration training investment that requires a more complex set of skills. 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:

> *"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."*

This paper provides detailed empirical evidence on this important question. Our focus is 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

[1] Within the natural resource sector, the literature on agricultural shocks in developing countries is voluminous. For a comprehensive review of these studies, see Ferreira and Schady (2009).

of sophisticated 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 cohort exposure in a difference-in-differences empirical strategy. We construct a measure of robot exposure intensity based on the interaction between the industry-specific robot penetration and initial employment composition in each location, following the neat approach developed by Acemoglu and Restrepo (2020). 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. 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 similar trends among older cohorts exposed after the typical college-going ages, but begin to diverge over time when new birth cohorts exposed before or during the critical timing of college decisions enter the economy. Our estimates imply that early exposed cohorts from the state experiencing the average robot penetration are 1.8 percent more likely to obtain a Bachelor's degree compared to late exposed cohorts. 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 driven by other factors, or differences in trends related to a diverse set of initial socioeconomic characteristics. We also show that our estimated effects are not confounded by other major shocks to the labor market, such as increased import competition and the recession of the early 1980s, or by major social programs, such as school finance reforms, war-on-poverty programs, and Medicaid. Together, these results demonstrate that the adoption of routine-biased technologies such as industrial robots has important implications for the accumulation of human capital.

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)

[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.

in their incomes *relative* to late exposed cohorts. The driving force behind this effect is education. When we account for the relationship between robots and college attainment, the relative income effect disappears entirely and if anything becomes slightly negative. 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 on the mechanisms underpinning our findings. We estimate that labor markets with greater exposure to robots saw a sizeable rise in the premium from having a Bachelor's degree relative to a high school diploma or two-year college 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. Counterbalancing these market incentive mechanisms, we find a significant and negative effect on parental income. These results suggest that the market incentive effects combined are large enough to outweigh the negative family income effect. 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 also examine changes in the sorting pattern to the major fields of study to provide further evidence on the skill premium channel. We exploit the fact that some individuals with a Bachelor's level training could still end up mismatched into routine-intensive jobs directly affected by robots, and the extent to which this occurs differs substantially across majors. The premium channel predicts that exposed cohorts should be less likely to major in subjects where this occupation-education mismatch is more common, as the returns to such occupations and thus majors became relatively lower. This is exactly what we find. Cohorts exposed to robots in childhood are more likely to sort into fields with a lower prevalence of routine-intensive occupations. They are more likely to avoid fields with a higher prevalence of machine operators, assemblers, inspectors, and precision-production occupations, exactly the occupations that experienced the bulk of the displacement consequences created by robots.

In the final part of the paper, we estimate a structural model of human capital investments to evaluate mechanisms and quantify the importance of these effects for the dynamics of earnings inequality. The model is parsimonious enough to be tractably estimated yet rich enough to capture the mechanisms discussed above. The model generates estimates of key parameters, such as the coefficient of risk aversion, within the range of existing estimates in the literature. Moreover, the model-based and reduced-form estimates of the effects of robots on college attainment are fairly comparable, and the model performs quite well in replicating other features of the data such as the age-specific pattern in enrollment rates. Remarkably, the model produces an elasticity of Bachelor's attainment with respect to parental income

[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 be not identified by our cross-location empirical strategy. But this does not necessarily affect the interpretation of our results and parameter of interest.

that is extremely similar to the one estimated by recent work using exogenous variation in family income generated by lottery wins (Bulman et al., 2021).

Counterfactual simulations suggest that the college premium channel is the single most important component behind the human capital response to robots. It accounts for approximately two-thirds of the overall effect, with most of the rest accounted for by the opportunity cost mechanism. Using the structural estimates of the model, we also estimate key elasticities. The elasticity of college with respect to the college premium is 0.60, whereas that with respect to the opportunity cost of college going is -0.18. The fact that the premium elasticity is larger than the opportunity cost elasticity is consistent with early work suggesting that lifetime future earnings is more important than initial earnings for individuals deciding whether or not to go to college (Berger, 1988). We perform further simulation exercises to explore the role of policy, finding that a government subsidy that covers 50 percent of the tuition costs for all students would enhance the college response to robots significantly.

To explore the implications of these results for the dynamics of earnings inequality, we extend the model to include an aggregate version of the production function introduced in Acemoglu and Autor (2011) where robots compete against labor in the production of tasks. Our analysis suggests that the effects of robots on high-skill earnings become increasingly negative over time as younger generations enter the market and generate an outward shift in the supply of high-skilled workers. By contrast, the effect on low-skill earnings declines over time due to the mechanical inward-shift effect on the supply of low-skilled workers. As a consequence, the technological-induced increase in earnings inequality becomes lower. These effects are quantitatively important. In the absence of the market incentive mechanisms, the long-run effect of robots on earnings inequality would be as much as 75 percent larger, holding all else equal.

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 increase in the college premium. We close the paper with a discussion of this question and believe that the most plausible explanation is that there have been changes in other important factors offsetting the premium effect. 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 the market incentives generated by changes in technology.

Our findings contribute to a vast literature on the impacts of technology. This literature has documented extensively that the routine-biased technologies affect the demand for labor, displacing workers specialized in routine tasks and increasing earnings inequality (see Jaimovich and Siu (2019) for an overview of the literature).[6]We contribute by providing evidence on whether, how, and why these tech-

[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). Subse-

nologies induce skill acquisition. While changes in the relative supply of high-skill workers have often been emphasized as an important force stabilizing the technological-induced gap between rich and poor (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]

Our study is also related to the recent work by Dauth et al. (2021), who show that the share of young workers without college education declines in firms potentially adopting industrial robots in Germany. This result could be driven by younger cohorts altering their skill investments, but also by employers altering their hiring decisions. Particularly related to our study is Berger and Engzell (2020), who examine the effects of robots on social mobility across US commuting zones. They document that robots have negative effects on social mobility in high-exposure areas. Our study differs from this paper since we study the impact of robots on *individuals* rather than on *places*. These effects are not necessarily the same, as the effects on places may reflect in part compositional effects due to endogenous migration responses of displaced workers. This is important in view of the evidence of workers sorting across local labor markets in response to the implementation of industrial robots (Acemoglu and Restrepo, 2020).[8]

This paper is organized as follows. Section 2 presents the data and empirical strategy. Section 3 reports the basic findings, documents their robustness, and provides empirical evidence on mechanisms. Section 4 identifies and structurally estimates a model of human capital investments to evaluate the quantitative importance of mechanisms and explore the implications of our findings for the long-run evolution of earnings inequality. Section 5 discusses the implications of our findings for the aggregate trends in college attainment, and section 6 concludes.

2 Data and Research Design

In this section, we provide an overview of the data sources, present the robot exposure variable, and describe the baseline specification.

2.1 Data Sources, Samples and 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.

quent 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 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).

[8]Since we assign individuals to robot adoption intensities based on their place of birth (determined before the robot shock), biases from selective migration are not an issue. Importantly, the data reveal that almost 80 percent of individuals reside in their state of residence during the critical ages when college decisions are formed (i.e., around 18). This high correspondence between birth and childhood residence state reduces the risk of measurement error in our assignment of childhood exposure due to migration, which is most likely common in more granular geographic divisions such as counties or commuting zones. We discuss our assignment of cohort exposure in Section 2.2.A.

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.[9] 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 whether an individual has a Bachelor's degree or higher, the level of education that is less prone to experience the displacement consequences created by robots.

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 later in Section 2.2. 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 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 $\lambda_j M_{jb}/L_{jb}$). For consistency with their conceptual framework and ease of comparison, we keep this adjustment term.[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 the "exposed" cohorts in our sample.

A concern with using realized penetration of robots in the United States is potential reverse causation[11] and unobservable shocks that simultaneously affect industrial robot usage and human capital investments. Following Acemoglu and Restrepo (2020), we construct the robot penetration variable in equation (1) using data of average robot adoption in the top 5 non-US countries with greater advances

[9] 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 A1). 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.

[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] In particular, local skills may themselves influence firm decisions on robot adoption across regions. For example, the adoption of robotics technologies complementing high-skill workers may be more likely if there is an abundant supply of high-skill workers when such technologies become available.

in robotics (Denmark, Finland, France, Italy, and Sweden).[12] Much of the robotics advances occurred first in these countries, so their industrial robot usage trends are a powerful predictor of that in the US, with a correlation coefficient of 0.95 (see Figure 2). At the same time, the extent of robot adoption in these countries is unlikely to be directly related to individuals' education across regions within 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 where the observed US robot penetration.

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 1966-83 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.[13] We exclude individuals residing in institutional group quarters to increase consistency between the different rounds of the ACS, a restriction that results in dropping about 3 percent of the sample.[14] We pool all of the ACS rounds into a single file to increase the precision of our estimated results.[15] Our basic sample consists of approximately 7.1 million observations.

2.2 Research Design

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.

In this section, we discuss the unit of analysis and present the definition of cohort exposure and the

[12] 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.

[13] 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.

[14] 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.

[15] 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 A3). In Section 3.1, 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.

basic estimating equation.

2.2.A The Geography of the Exposure to Robots

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.[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]

Importantly, there is a great deal of variation in the intensity of robot exposure across states, as shown in Figures A7 and A6. 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.

2.2.B Timing of Exposure

The research design compares cohorts exposed before or during the robot shock to those who were "too" old when the same shock became particularly salient. Figure 1 provides compelling evidence that robot adoption rose sharply and discontinuously in the early 1990s. In this subsection, we show that this sudden and large increase in robot adoption had immediate and first-order consequences on labor markets. Armed with this evidence, we then define the approximate date from which the effects of robots became particularly salient to implement our cross-cohort identification strategy.

To implement 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

[16] 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.

[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. We find similar figures when we track cohorts using the Current Population Survey, a representative sample that is collected annually since the 1960s.

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

$$(\text{emp/pop})_{s,t} - (\text{emp/pop})_{s,1989} = \alpha_t \ + \ \gamma_t \text{Robot penetration}_s \ + \mathbf{Z}_s'\Omega \ + \ \xi_{s,t} \tag{2}$$

where emp/pop is the employment-to-population ratio in each state s at time t $\in \{1981, 1984\ldots, 2007\}$. The regression includes a basic set of baseline characteristics computed from the 1990 census to account for differences in trends related to baseline conditions. The parameter of interest is γ_t, which measures the impact of robots in different moments in time. The path of these year-specific coefficients provides a detailed depiction of the dynamic effects of robots on employment. Standard errors are adjusted to account for arbitrary heteroskedasticity.

In Figure 3, Panel A plots the set of coefficients γ_t from equation (2) 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. Panel B of Figure 3 presents a trend-break analysis on the estimated coefficients and confirms this visual impression: the location of the structural break occurs around 1995. This is consistent with the view that the adoption of industrial robots spread significantly by the mid-1990s with the technological advances in robotics.

Definition of post-robot cohorts. These results help to motivate our basic specification measuring the impact of robots on educational attainment. The employment effects of robots become particularly salient around 1995, so we choose this date as an approximate benchmark. Individuals born after 1972 were exposed to this major episode before and during their transition ages to college. If there exists a causal link between robots and human capital, one would expect robot penetration to make a difference for many individuals in these birth cohorts. On the other hand, the cohorts born before 1972 were above age 23 by the mid-1990s, and thus they had largely completed their schooling decisions before the onset of rapid advances in robotics.[19]To provide a convenient means of summarizing magnitudes in tables and subjecting the estimates to sensitivity checks, we define the post-robot cohort group as those individuals born in 1977 or later. This definition does not consider the partially exposed cohorts born between 1973 and 1976 to account for the possibility that the effects of robot penetration on human capital may not occur instantaneously.

[19]For the constant sample of birth cohorts born between 1966 and 1972, about 80 percent of individuals had completed their schooling decisions by age 23. See Appendix Figure A4.

2.2.C Basic Specification

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

$$S_{ist} = \alpha \ + \ \beta \text{ Robot penetration}_s \times \text{Post}_t + \ \mathbf{X}'_{ist}\Omega \ + \ \sum_{z \in \mathbf{Z}} \Phi_z(z \times \mathbf{FE}_t) + \ \mathbf{FE}_s + \ \mathbf{FE}_t + \ \xi_{ist} \tag{3}$$

where S_{ist} is the outcome of interest for individual i born in state s and birth cohort t. The term Robot penetration$_s$ is our time-invariant measure of robot exposure intensity defined in equation (1). $Post_t$ is an indicator for the post-robot cohorts born after 1976. The key independent variable of interest is given by the interaction between these two variables. 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 fixed effects. The 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 initial state characteristics $\mathbf{Z}$, including baseline levels of educational attainment. These controls help account for potential mean-reverting dynamics in college attainment spuriously correlated with the exposure to robots. The residual term, ξ_{ist}, is clustered at the state-of-birth level to allow for serial correlation across birth cohorts.

Identification. The validity of any estimate of the impact of robots based on equation (3) 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 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 across birth cohorts. By conditioning on state and birth-year fixed effects, the parameter of interest is identified from within-state differences between cohorts that were exposed earlier and later to robotics advances 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 cross-cohort trends in educational and other socioeconomic outcomes prior to the technological advances in automation. After presenting the basic findings, we discuss specific threats to internal validity and provide a variety of evidence suggesting that the identification condition is likely to hold in our setting.

3 Results

3.1 Impacts on College Attainment

Visual evidence. We begin by presenting results from estimating a fully flexible version of equation (3) that replaces the $Post_t$ dummy with birth-cohort indicators. This specification allows the estimates for β to vary over time across birth cohorts and examine timing of impacts. Figure 4 plots the estimated coefficients of β_t and corresponding 95 percent confidence intervals for each birth cohort. We separate the birth cohorts into three groups, denoted by the vertical bars: "pre-robot cohorts"= birth-year $\leq$ 1972, "partially exposed" =birth-year $\in \{1973, 1976\}$ and "post-robot cohorts"= $\geq$ 1977. Since individuals born prior to 1972 were over age 23 in the mid-1990s and thus had largely completed their schooling decisions before the dramatic advance in robot adoption, there is no reason to expect higher- versus lower-exposure states to have differential trends before the 1972 cohort. Thus, the path of the pre-robot coefficients allows us to inspect the plausibility of the identifying assumption.

Figure 4 shows that the robot exposure intensity is unrelated to changes in college attainment before the 1972 robot cohort. The pre-robot coefficients are very close to zero and statistically insignificant. Notably, there is not any clear tendency toward improving or deteriorating college attainment during the pre-robot period. Indeed, the pre-robot coefficients display a flat trend and fluctuate randomly around zero. These results provide strong evidence in support of the identifying assumption that higher-versus-lower-exposed areas would have experienced similar trends in the absence of the robot shock.

After the 1972 cohort, more-versus-less exposed areas begin to diverge gradually, a pattern that remains persistently increasingly from the 1976 cohort. Post-robot cohorts from states with greater exposure to robots are significantly more likely to have at least a bachelor's degree. The persistent increase in the effects of robots after the 1976 cohort is natural as later cohorts spent more childhood years exposed to robots, and given the increasing adoption of robots over time. The estimated coefficient for the youngest post-cohort is approximately 0.006. It implies that cohorts from states experiencing the average exposure to robots see 1.2 percentage points larger increase in the likelihood of obtaining a bachelor's degree or higher ($0.006 \times 2 = 1.2$). This represents an increase of about 3.3 percent relative to the sample mean outcome.

Baseline estimates. The flexible specification provides 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 (3) to summarize magnitudes in tables and perform specification checks.

These results are reported in Table 1. Column (1) presents results from a specification that incorporate only birth-cohort and state-of-birth fixed effects as well as indicators for gender, race, and survey year. The coefficient of interest is estimated at 0.0027 with a standard error of 0.0015 and significant at the 10 percent level. Column (2) adds interactions between 1990 college attainment levels and birth-cohort fixed effects, which control for any possible mean reversion in college attainment spuriously correlated with the exposure to robots. While the inclusion of these interactions leaves the coefficient β virtually unchanged at 0.0027, it makes the point estimate much more precise. Indeed, the standard error falls by a factor of 2.5, from 0.0015 to 0.0006, and the estimated coefficient becomes highly significant at less than 1 percent. This suggests that the mean-reversion controls make a good job in reducing noise and increasing precision, and that our results are not simply an artifact of some possible convergence effect across more and less exposed states.

Column (3) presents results from our preferred specification, which incorporates the full set of additional baseline socioeconomic characteristics, as measured in 1990, interacted with birth-cohort fixed effects. The relationship is quite similar and somewhat larger with these controls. The coefficient β

now stands at 0.0034 and is slightly more precise, with a standard error of 0.0005. Quantitatively, this coefficient estimate implies that early exposed cohorts from the states such as Connecticut, which experienced the average penetration of robots, are 0.7 percentage points more likely to obtain a Bachelor's degree or higher relative to later exposed cohorts. Relative to the sample mean, this effect represents an increase of approximately 1.8 percent. While this effect might seem small, note that it is based on the average exposure to robots and hence this effect is significantly larger in states experiencing a greater exposure to robots. The same calculations for states at the 75 and 90th percentiles of the exposure to robots reveal effects of 2.1 and 3.7 percent respectively.

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.[20] However, by construction, younger cohorts are disproportionately underrepresented in the estimation sample.[21] As a robustness check, we rerun the baseline specification but restrict the estimation sample to the ACS conducted between 2015 and 2019 where all the birth cohorts are observed with nearly equal likelihood. This restriction reduces sample size by approximately 60 percent, yet both the point estimate and standard error are not appreciably affected (column 4, Table 1).

As shown in Appendix Table A2, the penetration of robots was, in general, much higher in manufacturing than in nonmanufacturing industries. This raises the concern that our results may be capturing the post-1990 decline in manufacturing employment due to factors unrelated to robotics technology. We address this issue in Appendix Table A3 by including interactions between birth-cohort fixed effects and 1990 baseline manufacturing share in each state, which proxies for the post-1990 declining trends in manufacturing employment. This is a very demanding specification because these controls may mechanically absorb part of the underlying variation we use to identify the effects of robots —recall that the robot exposure variable is a function of the baseline employment share of each industry. Still, the coefficient of interest remains very similar to the baseline and highly significant. As a further check, we directly control for interactions between birth-cohort fixed effects and 1990 level of employment rates or average wages. Once again, this has very little impact on the coefficient of interest.

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 A4. 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. The 2LS estimate is comparable to our baseline 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

[20]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.

[21]This is illustrated in Figure A1, which plots the share of each birth cohort in the estimation sample.

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 A5); *ii*) excluding outlier observations based on regression residuals as well as on leverage and Cook's distance measures (Appendix Table A6); *iii*) excluding industries with the largest Rotemberg weights (Appendix Table A7), as recommended by Goldsmith-Pinkham et al. (2020); and *iv*) alternative inference procedures (Appendix Table A8), 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 source of the gains in college completion. Since individuals in the middle of the skill distribution were the ones displaced by industrial robots, one would expect that the increase in college attainment comes primarily from this part of the education distribution if changes in the college premium and opportunity costs are the key mechanisms generating this relationship. To shed light on this hypothesis, we estimate the baseline specification (3) for different education categories: less than high school, high school, two-year college degree, and Bachelor's degree or higher. As shown in Figure 5, the robot-induced increase in Bachelor's degree attainment is driven by reductions in the probabilities of having high school and a two-year college degree. The effect on the likelihood of having less than high school, though significant, is much smaller in magnitude and of the opposite sign.[22] These results confirm that the improvement in Bachelor's degree attainment comes from individuals in the middle of the skill distribution, for whom the adoption of robots differentially altered the incentives to invest in Bachelor's-level college. 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. To save space, we focus on four-year attainment in the remaining analyses.

Overall, the results of this section show that growing up in robot-exposed markets leads to a significant improvement in Bachelor's degree attainment. This interpretation of the results depends critically on the assumption that there were no major differential trends in college attainment across more- and less-exposed areas driven by other factors. We critically evaluate assess potential threats to the identification strategy.

3.2 Threats to Internal Validity

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

3.2.A 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 the comparison of pre-robot cohort

[22]This result suggests heterogeneity among individuals within the middle of the education distribution. If the adoption of robots reduces the returns to middle education, then forward-looking individuals who are unable to go to college or lack college-ready skills could find optimal to reduce their educational investments and dropping out of the school at earlier ages.

trends does not suggest significant differences, one might still be worried about the possibility that our estimates are capturing some pre-existing convergence or catchup effect in human capital across states. We perform several additional exercises to address this issue in Table 2.

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 dynamics in college attainment taking place around the onset of recent advances in robotics technology. As an additional check, we also include interactions between birth-cohort fixed effects and 1960, 1970 and 1980 college attainment levels. Columns (2) of Table 2 shows that the coefficient of interest remains almost unchanged with these additional controls, going from 0.34 to 0.36. And although the relationship is estimated with less precision, it continues to be highly significant. Column (3) goes a step further and controls rather for the 1960-1990 change in college attainment interacted with birth-cohort fixed effects. The results remain extremely similar to the baseline, providing further evidence that mean reversion is unlikely to explain the post-1972 cohort changes in college attainment.

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 intensity in robot exposure across states.[23]As shown in column (4) of Table 2, the inclusion of these pretrends has virtually no impact on our results, with both the coefficients and standard errors nearly identical to the baseline.

Within-region variation. While the results above are very reassuring, one could still be concerned that our results 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 2). 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 dynamics or differences in trends across regions. While the inclusion of this detailed set of fixed effects reduces the variation in the data, which is natural as there are fewer

[23]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 cohorts (Lee and Solon, 2011; Wolfers, 2006).

states within each region, the results are strikingly similar to the baseline. The coefficient of interest is somewhat larger in magnitude, and while its standard error increases, the estimated relationship remains significant at the conventional levels of significance.

3.2.B 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.

Chinese import competition. The first obvious source of bias is the unprecedented rising Chinese import competition since the 1990s, which occurred in parallel with the recent advances in robotics technologies and had important implications for manufacturing employment (Autor et al., 2013; Acemoglu et al., 2016). It is important to note, however, that the states more affected by increased Chinese imports are far from being the same as those states housing industries with greater adoption of robots. Indeed, the correlation between the exposure to robots and Chinese import penetration, as measured in (Autor et al., 2013), is only 0.06 and statistically insignificant. Consistent with this lack of correlation, controlling for a full set of interactions between the intensity in Chinese import competition across states and birth-cohort fixed effects does not appreciably change the point estimate (column 2, Table 3).

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 (0.08) 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 3).

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.[24]During this period, Medi-

[24] Johnson and Jackson (2019) and Hoynes et al. (2016) document that Head Start and Food Stamp respectively lead to

caid 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,[25] Appendix Table A9 documents that if anything the post-robot cohorts from states with greater robot penetration are *less* likely to have been exposed to these programs in childhood. This suggests that these programs cannot explain the gains in college attainment we report in Table 1. Consistent with this notion, controlling for the childhood-exposure probability to these programs has very little impact on our estimates (columns 4-6, Table 3). In Appendix Table A10, 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.

3.3 Impacts 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. We examine this question systematically in Table 4 using several income measures as dependent variables. Note that since robots have overall negative impacts on earnings in areas with greater robot penetration, positive values of β in our estimation of (3) imply that cohorts exposed to robots in childhood experienced a *smaller negative* impact on their labor market income *relative* to those cohorts exposed later in the life-cycle when their educational decisions had finalized.

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. The odd-numbered columns include the baseline controls used previously. 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. Early-exposed cohorts from the state witnessing the average robot penetration experience a relative increase of 0.44 to 0.62 percent depending on the income measure being considered.

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 in driving the income effects, we perform a mediation-style analysis by controlling for a Bachelor's degree indicator in the income regressions and establishing the extent 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.

[25] The shares of the population exposed in the relevant childhood years of individuals in our sample are 30, 81, 91, 97 and 38 percent for community health centers, head start, food stamp, Medicaid and school finance reforms, respectively.

which the estimated coefficient of interest is reduced. The results from this exercise are presented in the even-numbered columns of Table 4. 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. In fact, the estimated coefficient of β becomes opposite signed in some cases. 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.

Are these effects driven by scarring effects? A complication when interpreting these results is that increased use of industrial robots also affected older adult workers. 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 towards finding an improvement in the labor market incomes of younger cohorts even in the absence of a causal relationship. To mitigate this concern, we repeat the baseline specification but include a detailed set of state-of-residence $\times$ birth-cohort fixed effects. Now the coefficient 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. This is a very narrow source of variation as most individuals work in the same place where they lived during their childhood years. Still, we continue to observe meaningful and highly significant income effects using this more demanding specification (see Appendix Table A11). This suggests that biases from scarring effects are unlikely to be a major issue.

In summary, cohorts exposed to robots at the beginning of the life-cycle experienced an increase in their incomes relative to late exposed cohorts. We reiterate that these results *do 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.

3.4 Mechanisms

This section provides evidence on the likely mechanisms generating the basic picture documented so far. Our analysis suggests that changes in the college premium and opportunity cost of college are the key drivers of our findings.

3.4.A 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 skill-replacing 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 data simply as 2008. We follow Charles et al. (2018) and assume that individuals attending college in a given year forgo immediate income gains equivalent to the average earnings of individuals aged 18-25 without any college training. We measure the college premium as the earnings gap between older working adults (ages 25-65) 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 ($\times 52$).[26] For a given outcome y of demographic group g in state s, we estimate the following first-difference specification:

$$\Delta ln y_{gs,90-08} = \alpha \; + \; \gamma \text{Robot penetration}_s \; + \mathbf{Z}_s' \Omega \; + \; \alpha_g \; + \; \xi_{gs} \tag{5}$$

where $\Delta ln y_{gs,90-08}$ 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 (5) are presented in Table 5. 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. The point estimate of -0.030 implies that the average increase in the stock of robots is associated with a decline of 6 percent in the labor market income a young adult worker receives.

Columns (2) to (4) also document a decline in the average earnings of older adult workers across all education groups, but these effects are much smaller for individuals with a Bachelor's degree. Columns (5) and (6) show that these differences in the estimated effects of robots between individuals with and without a Bachelor's degree are highly significant. In the state experiencing the average penetration of robots, the earnings premium from having a Bachelor's degree rose by 2.3 percent relative to a high school or less and by 1.7 relative to a two-year college degree. On the other hand, there are no statistically meaningful effects on the premium from having a two-year college degree relative to high school or less was (column 7). These findings show that investing in Bachelor's-level type of college became more appealing than two-year college attendance. This pattern in the data is broadly consistent with our findings in Figure 5 showing that part of the increase in Bachelor's degree attainment stems from individuals who would rather have gone to a tow-year college in the absence of the robot shock. This suggests that the skill premium is likely the single most important driver of our findings: if the opportunity cost channel were the dominant mechanism, we would have observed an *increase* in a two-year degree attainment, as the average earnings an individual must forego to acquire a two-year college degree became significantly lower.

Taken together, the evidence strongly supports the hypothesis that the adoption of industrial robots altered the market incentives to invest in Bachelor's-level training by rising its relative returns and reducing its opportunity cost. 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

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

by robots. This could occur because there exists education-job mismatching and even individuals with Bachelor's-level training could end up in less-skilled jobs being directly affected by robots. We will return to this point in Section 3.5 and exploit heterogeneity in the education-job mismatching across majors to develop an additional test of the college premium channel.

3.4.B 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 (5) using the log-change in parental income as the dependent variable. We define parental income as the household heads' and their spouses' incomes from all sources. We focus on household heads over age 40, as they are more likely to have children of college-going ages. 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 (8) of Table 5. As one can see, there is a highly significant decline in parental income in areas housing industries with greater exposure to robots. While precisely estimated, the magnitude of this result is smaller than the effects on the opportunity cost and college premium measures. The point estimate implies that the average increase in the stock of robots is associated with a 0.57 percent decline in parental income. This result suggests that we most likely would have observed a larger positive college response to the penetration of robots in the absence of this income effect.

3.4.C Supply-Side Responses

We next explore the possibility that our results may reflect supply-side responses of colleges and universities. For example, institutions that award 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. State administrations may increase their investments in education and training programs or directly provide 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 Integrated Postsecondary Education Data System. 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 model (5). As can be seen from Appendix Table A12, 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.

3.5 The Premium Channel and Sorting Pattern to Fields of Study

Most high-skill jobs that cannot be performed by industrial robots require college education. However, some individuals with college training may even end up in routine-intensive occupations being directly replaced by robots, and this education-to-job mismatching likely explains in part why we observe a decline in the earnings of college-educated workers. This mismatching could happen either because some majors are oriented toward routine-intensive occupations in some way or due to the existence of frictions in the labor market. Remarkably, the prevalence of this mismatching varies substantially across fields. The likelihood of engaging in routine-related occupations such as machinists, assemblers and material handlers is particularly high in fields such as electrical and mechanic repair technologies (25 percent) and mechanical engineering-related technologies (10 percent).[27] By contrast, this rate is less than 1 percent in fields such as actuarial science, elementary education, and environmental engineering.

These differences suggest a simple test of the college premium channel. This channel implies that exposed cohorts should be less likely to major in subjects with a greater prevalence of routine-related occupations, as the returns to such occupations (and thus fields) became relatively lower with the widespread adoption of robots. Other mechanisms, in particular the opportunity cost of college attendance, are unlikely to generate such sharp predictions on the specific major field of study. Therefore, estimating how the exposure to robots affected the sorting pattern to fields of study serves as a cleaner examination of the skill premium channel.

To explore this hypothesis, we construct measures of the extent to which a field of study is susceptible to the displacement effects of robots. We do this in a two-step procedure. In the first step, we match occupational-task scores from Autor and Dorn (2013) to all 330 time-consistent occupation categories in the ACS data. These scores indicate the extent to which a given occupation is intensive in routine, manual and abstract tasks, measured on a zero to ten scale in a non-mutually exclusive manner.[28]For each occupation k, we then compute the routinization index developed in the seminal work of Autor et al. (2003):

$$\text{Routine Task Share}_k = 100 \times \frac{r_k}{r_k + m_k + a_k} \tag{6}$$

where the r, m and a are respectively the routine, manual, and abstract task scores standardized with equal mean and variance. This index varies between 0 and 100 and captures the extent to which a given occupation is intensive in routine task inputs relative to other task inputs. In the second step, we use data from the ACS on all individuals with a Bachelor's degree to match this routinization index to each major-occupation pair. The resulting index is then aggregated at the major level by taking the employment-weighted average across all occupations within each major.

Column (1) of Table 6 presents the results from estimating the baseline model (3) using the average routinization index in each major as the dependent variable.[29] In line with the college premium channel,

[27] We create this prevalence of routine-intensive occupations in each field of study using the 2010-19 rounds of the ACS. Information on the field of study is available in the ACS since 2009. However, the codes for the field of degree changed between 2009 and 2010, so we use the ACS conducted since 2010 to maintain consistency.

[28] Indeed, these scores allow occupations to be intensive at different degrees in multiple tasks. The correlation between these task-intensity measures is far from perfect. The raw correlation between these task scores ranges from 0.002 to -0.32. The partial correlation between these indexes after controlling for the measure of manual tasks is -0.33.

[29] Since the exposure to robots is associated with significant gains in college attainment, any estimates of the exposure

we find that exposed cohorts sort into fields with lower shares of routine-intensive occupations. The coefficient estimate is negative and highly significant. In columns (2) to (6), we consider the prevalence of specific occupations in each field of study rather than the routinization index (6), focusing on those directly affected by robots such as machine operators, assemblers, inspectors, mechanics, and repairers. Once again, the evidence is broadly consistent with the premium hypothesis. We estimate negative and highly significant coefficients in all cases. The advance in robotics technology reduced the returns to routine-related occupations, and individuals responded by pursuing fields where such occupations are less commonplace, or equivalently, those fields where the returns became relatively higher.

4 Structural Analysis

The results presented so far are consistent with changes in market incentives driving the college response to robots. Without more structure, however, it is not possible to say what part of the effect is due to the college premium versus opportunity cost channels or whether these market incentive effects are large or small relative the parental income effects. To investigate these questions, we propose and estimate a simple partial equilibrium model of college choice. The model is parsimonious enough to be tractably estimated yet rich enough to capture the mechanisms discussed above.

4.1 A Model of College Choice

4.1.A Setup

There is a continuum of individuals born in cohort t that live for J periods. The first period ($j = 0$) corresponds to the first year of adult life when individuals are aged 19 and have just completed high school. Each period in the model corresponds to a year and there are 46 periods as a whole, matching the typical working life cycle. In the first period, individuals choose either to attend college for four years or enter the labor force. They are allowed to delay entry to college every year up to the age of 23 (the timing of enrollment in the data), a point from which their educational decisions become irreversible. While delaying their entry into college, individuals are participating in the labor market. An agent who never enrolls in college enters the labor force since age 19 uninterruptedly. Individuals are credit-constrained and can only borrow to finance college costs. We also assume that there is no saving, so consumption is equal to income in each period.[30]

Preferences. The utility is intertemporally separable and depends on consumption and preferences for college education. Individuals discount the future at rate $\beta = 1/(1 + \rho)$, with a discount factor of ρ.

effects on changes in college majors can be driven by individuals who otherwise would have not enrolled in college (extensive margin), by changes in the type of fields of individuals who would have enrolled in college independently of the robot shock (intensive margin), or by a combination of both margins. While we cannot identify the importance of both margins separately, they most likely go in the same direction.

[30] With savings, the model becomes greatly complicated and intractable to be estimated. Yet, 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).

During the study stage, the utility of being in college is ψ_k, which varies with the period k of enrollment. Without loss of generality, the disutility of work is normalized to zero. Let s equal to one if an individual attends college and zero otherwise. The instantaneous utility in period j takes the form:

$$U(c_j) = u(c_j) - \mathbb{I}_e \psi_k = \frac{c_t^{1-\sigma}}{1-\sigma} - \mathbb{I}_e \psi_k s \tag{7}$$

where c_j is consumption, and $\mathbb{I}_e$ an indicator function that takes the value of one during the study phase and zero otherwise. Remember that an individual who delayed college entry until age 23 completes her degree at age 26, so the study phase goes from period 0 to 8. The utility function over consumption is CRRA and σ is the curvature parameter that determines the degree of relative risk aversion. Attending college generates utility or disutility. To the extent that individuals face a psych cost of learning to complete a college degree, ψ_k will be positive and attending college generates disutility. If individuals enjoy the social life during college (e.g., meeting friends, participating in clubs and sports), or do not want to disappoint parents, ψ_k will be less positive. Note that the subscript k is equivalent to an individual's age at enrollment, so delaying entry alters the utility of college attendance. If the costs of learning become significantly more important as people age and remain out of a regular study routine, then ψ_k will be increasing in the age of enrollment.

Earnings process. The earnings y_j^s for an individual at period j with education s is the product between the price of an effective unit of labor w^s and her accumulated human capital stock $h_j(s, \xi^s)$, where:

$$\log h_j(s, \xi^s) = g^s(\mathbf{Z}) + \zeta_j^s + \xi_{jt}^s \tag{8}$$

$$\log w^s = \omega_a^s R_j \tag{9}$$

where $g^s(\cdot)$ is an education-specific function measuring the importance of demographic and background characteristics, ζ_j^s an education-specific age component, ξ_{jt}^s a cohort-specific stochastic shock, and R is the demand for robots in the market. The subscript $a \in \{young, old\}$ denotes young and old workers, corresponding to ages 19 to 25 and over 25 respectively. Therefore, we allow the semi-elasticity of wages with respect to robots, ω_a^s, to differ between college- and non-college persons and between young and old workers. 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.

Government subsidy and education costs. The government pays a subsidy G for each period spent in college. Therefore, the net price of college p is equal to the difference between the tuition costs F and government grants. Both government grants and tuition costs are exogenously determined.

Students loans. Since there are borrowing constraints, students can only borrow up to a fraction α_b of the annual college cost net of grants. We assume that the constraint is binding, and all students borrow $\alpha_b p$ with $\alpha_b < 1$. The interest rate r is constant during the entire life of the loan. Interest accumulates during college, so the total debt owed at the end of college is $D = \alpha_b p \sum_{l=1}^{4} (1+r)^l$. An individual starts repaying the loan from the first period after graduation under a plan with equal payments that last N periods. The amount paid each year n after graduation until period N is:

$$\text{Payment}_n = D\frac{r(1+r)^N}{(1+r)^N - 1} \tag{10}$$

Parental support. The parents cover the remaining fraction $(1-\alpha_b)$ of the out-pocket costs of college that is not financed with borrowing. In addition, while attending college, students receive an income transfer from the parents that is equal to a fraction λ of the parental income I net of the out-of-pocket costs of college. Thus, the consumption for college-goers during the study stage will be:

$$c_j = \overbrace{\lambda(I-(1-\alpha_b)p)}^{\substack{\text{income}\\\text{transfer}}} + \overbrace{\alpha_b p}^{\substack{\text{amount}\\\text{borrowed}}} + \overbrace{(1-\alpha_b)p}^{\substack{\text{parental}\\\text{contribution}}} - \overbrace{p}^{\substack{\text{net price}\\\text{of attendance}}}$$

or,

$$c_j = \lambda(I-(1-\alpha_b)p)$$

This implies that parental income, borrowing constraints, and college costs affect an individual's college decision by altering their consumption during the study phase. Individuals who choose not to attend college and enter the labor force do not receive any income transfer. In this case, $c_j = y_j^0$

Parental income. Parental income is the parents' and their spouses incomes from labor ($I \equiv \sum w^s h_j(s, \xi^s)$). Parental income depends on the household head's human capital and demand for robots:

$$\log I = \log h(s, \xi^{p,s}) + \kappa R_j \tag{11}$$

The parameter κ measures the overall impact of robots on parents' income. To the extent that parents have, on average, higher levels of educational attainment and more work experience, or are inherently less prone to engage in routine-related jobs, the effects of robots on parental income will be less negative when compared to the mean worker.

4.1.B Decision Problem

An agent chooses whether or not to attend college by maximizing the present value of lifetime utility. This implies a threshold disutility level ψ_k^* such that an individual at period j will decide to go to college if $\psi_k \leq \psi_k^*$. This threshold is known to the individual at the time of the college decision. If she decides to delay the college decision for the next period $j+1$, she does not observe the disutility threshold in advance but does have some expectation over it. Let π_k be the perceived probability of $\psi_k \leq \psi_k^*$ in each period k, and $V_{j|J}^s$ the expected utility over periods j to J. At period j of the study phase, an individual will decide to attend college if the net present value of this choice is greater or equal to zero:

$$\Omega_{j|J} \equiv V_{j|J}^1 - V_{j|J}^0 \geq 0$$

Since the value of entering the labor force rather than attending college, $V_{j|J}^0$, depends on an individual's college decision in the next period and on whether she is facing the last period of college decision, we have that:

$$V^0_{j|J} \equiv \begin{cases} u_j(c_j) + \beta(\pi_{j+1}V^1_{j+1|J} + (1-\pi_{j+1})V^0_{j+1|J})) & \text{if } j < 4 \\ \sum_{l=j}^{J} \beta^{l-j} u_l(c_l) & \text{if } j = 4 \end{cases}$$

Recall that period 4 corresponds to the last period of college decision and delaying college is not possible anymore. Now suppose that Ω is normally distributed with density $\phi(\cdot)$ and distribution $\Phi(\cdot)$. Then, the fraction S of individuals in cohort t who complete a college degree during their life cycle is:

$$S_t = \sum_{j=0}^{4} \Phi(\Omega_{j|J})q_j, \text{with } \; q_0 = 1 \tag{12}$$

where q_j is the fraction of individuals that make their college decisions in period j or later. Hence, the product $\Phi(\Omega_{j|J})q_j$ represents the probability that an individual enrolls in college at age $19 + j$. When there is an increase in the penetration of robots in the market, it has different impacts on different cohorts because they were exposed in different moments in the study stage. On the one hand, exposure after age 23 (or period 4) makes no difference because the timing of college decisions has finalized for everyone, and individuals cannot alter past educational choices. On the other hand, increased robot penetration should have larger impacts when the exposure occurs at the onset of the study phase where the pool of individuals potentially altering their educational decisions is the entire population in a cohort.

Effects of robots. We now characterize the effects of a robot shock on the college attainment of different cohorts. Suppose that there is an advance in robotics technology, such that $R_{j+1} - R_0 = R_j - R_0$. Differentiating (12), we obtain the average effect of robots on cohort $t = 0, 1, \ldots T$ whose earliest period of exposure is $\tau = t$:

$$\begin{aligned} dS_{t,\tau} = \mathbb{I}_{\tau<5} \Bigg[& \overbrace{\sum_{j=\tau}^{4} \left[\frac{d\Omega_{j+4|J}}{dy^1} y^1 \omega^1_{old} - \frac{d\Omega_{j+4|J}}{dy^0} y^0 \omega^0_{old}\right] q_j \phi_j}^{\text{premium effect}} \quad + \quad \overbrace{\sum_{j=\tau}^{4} \frac{d\Omega_{j|j+3}}{dy^0} q_j \phi_j y^0 \omega^0_{young}}^{\text{opportunity cost effect}} \\ & + \quad \underbrace{\sum_{j=\tau}^{4} \frac{d\Omega_{j|j+3}}{dI} q_j \phi_j \lambda I \kappa}_{\text{parental income effect}} \Bigg] dR \end{aligned} \tag{13}$$

where $\mathbb{I}_{\tau<5}$ is an indicator function that is equal to one if the earliest period of exposure is lower than 5 and zero otherwise. The expression illustrates transparently the importance of the three mechanisms discussed above: college premium, opportunity costs, and parental income. Notably, the strength of these mechanisms depends crucially on how the exposure to robots affects labor market earnings and how heterogeneous these effects are with respect to an individual's educational attainment. In the case that college- and non-college workers' lifetime earnings are identically affected, the premium channel plays no role. We exploit equation (13) to form the basis of our estimation strategy and quantitative analysis below.

4.2 Identification and Estimation

We implement a three-step procedure to take the model to the data. First, we estimate the key earnings parameters $\{\omega^1_{old}, \omega^0_{old}, \omega^0_{young}, \kappa\}$ separately from the rest of the structure of the model. In practice, the values of these coefficients are taken directly from the estimates in Section 3.4. Second, we set some parameters and initial conditions externally. We assume that the discount factor ρ is equal to the interest r, and set the latter to 5 percent, following Heckman et al. (1998). This implies that the discount rate β is approximately 0.95 ($\approx 1/1.05$). We choose the initial earnings by age and Bachelor's degree status as well as initial parental income to match the average values observed in the 1990 census. We set $\lambda = 0.07$ based on Kalenkoski and Pabilonia (2010), who study similar cohorts and report data on the transfers that four-year college students receive from their parents. The direct tuition costs of college, grants, and student loans are parameterized using the average values for the academic years of 1989-90 as published by the National Center for Education Statistics (NCES, 2004). With these data, it is straightforward to calculate α_p. We calibrate the perceived probabilities of enrollment into college π_k and proportion of individuals making college decisions q_j to be consistent with the data on actual enrollment and entry into college rates by age. In doing so, we use the 1990 census and use the questions on college attendance and on whether the highest degree completed is "some college but less than a year" to construct attendance rates and proxies for the rate of persons entering into college by age. Finally, we set the life of the loan N to 25 years, as in Ionescu and Simpson (2016).

In the third step, we solve analytically the derivatives within the brackets in equation (13) and estimate the remaining utility parameters $\theta = \{(\psi_0, \dots, \psi_4), \gamma\}$ using the general method of moments, given the earnings parameters estimated in the first step and externally calibrated parameters and initial conditions from the second step. To implement this estimator, we first collapse the data by year-of-birth and state-of-birth cells, the level of variation we exploit in our reduced-form analysis, and then define $dS_{it} = S_{it} - S_{i1966}$ for each cohort born in year t and state i. Cohorts born in 1976 or later were age 19 or younger by the mid-1990s. We aggregate these cohorts into one group to facilitate our estimation procedure. We then augment equation (13) to include an additional error term v_{it} that reflects measurement error and unobserved differences across cohorts and regions that influence human capital investments. In the end, the effects of robots for each cohort can be written as follows:

$$\begin{aligned} S_{it} - S_{i1966} = \mathbb{I}_{\tau<5} \Bigg[& \overbrace{\sum_{j=\tau}^{4} \Big[\frac{d\Omega_{j+4|J}}{dy^1} y^1 \hat{\omega}^1_{old} - \frac{d\Omega_{j+4|J}}{dy^0} y^0 \hat{\omega}^0_{old} \Big] q_j \phi_j}^{\text{premium effect}} \quad + \quad \overbrace{\sum_{j=\tau}^{4} \frac{d\Omega_{j|j+3}}{dy^0} q_j \phi_j y^0 \hat{\omega}^0_{young}}^{\text{opportunity cost effect}} \\ & + \quad \underbrace{\sum_{j=\tau}^{4} \frac{d\Omega_{j|j+3}}{dI} q_j \phi_j \lambda I \hat{\kappa}}_{\text{parental income effect}} \Bigg] dR_i + v_{it} \text{ for all } t = 1967, \dots, 1976 \text{ and } \tau = 1995 - (t+19) \end{aligned} \tag{14}$$

The moment conditions used to jointly estimate the parameters are based on the system (14) and the pattern of college attendance by age. Specifically, our model has multiple moment conditions of the form:

$$\mathbb{E}\Big[dR_i \cdot \mathbf{v}_i(\theta), \Phi_j(\theta) - \Phi_j \Big] = \mathbf{0} \tag{15}$$

where $\mathbf{v}_i = \{v_{i1967}, \ldots, v_{i1976}\}$ is the vector of residuals and Φ_j are the observed entry into college probabilities by age. The vector θ of parameters is estimated as the vector that minimizes: $m(\theta)'Vm(\theta)$, where $m(\cdot)$ corresponds to the moments conditions mentioned above and V is a weighting matrix. As a whole, there are 15 moments conditions and 6 parameters to be estimated.

While the parameters are estimated simultaneously, it is useful to discuss the moments that best identify each of them. The disutility parameters are largely pinned down by seeking the values of the disutility that are consistent with the pattern of enrollment rates and exposure effects across cohorts. Heterogeneity in the disutility of college attendance is important to rationalize the differences in enrollment rates by age. The curvature parameter γ of utility over consumption is identified by the magnitude of the effects of robots on college attainment. Intuitively, conditional on all the other parameters, a large college response means that individuals place a relatively high weight on future consumption, implying a low value of γ.

4.3 Parameter Estimates

The point estimates and standard errors are displayed in Panel A of Table 7. The disutility parameters are all positive and as one could expect, the disutility of college is increasing with the age at enrollment. They range from 0.35 to 0.63, all of which are statistically distinguishable from zero. The coefficient of risk aversion γ is estimated at 1.37 and statistically significant at less than 1 percent. This estimate is in the lower range of existing estimates in the literature using US microdata. In a series of seminal studies, Attanasio et al. (1999) and Attanasio and Weber (1995) estimate this coefficient to be 1.35, and 1.5 respectively for the US households.[31]

Sensitivity analyses. Appendix Table A13 explores the robustness of our results to alternative assumptions regarding some of the calibrated parameters. Our baseline analysis uses the 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 an interest rate of 3 percent, we estimate the disutility coefficients that range between 0.41 and 0.68 and a coefficient of risk aversion of 1.40, which are extremely similar to the baseline estimates. In our analysis, we have assumed that the life of the loan is 25 years. As a robustness check, we re-estimate the model exploring different repayment periods, ranging from 10 to 30 years. This alters the value of the payments per period but leaves the coefficient estimates unchanged. Finally, we re-estimate the model assuming that the perceived probabilities of enrollment in subsequent periods, π_{j+z}, are equal to zero, rather than matching them to the actual patterns of enrollment observed in the data. The parameter estimates remain nearly the same. The reason why this does not materially alter our results is that the bulk of college entry occurs at age 19, and then falls rapidly to zero afterward. Thus, setting $\pi_{j+z} = 0$ does not seem to be unreasonable.

Comparison to reduced-form estimates. Panel B of Table 7 compares the structural to reduced-form average effects of robots on Bachelor's degree attainment. To compute the average effect implied

[31] Other important studies in the literature are Blundell et al. (1994) and Banks et al. (2001), who use data from United Kingdom and estimate a coefficient of risk aversion of 1.37 and 1.96 respectively.

by the structural model, we first calculate the marginal effect of an additional robot for each cohort and then take the weighted average across all cohorts, weighting by cohort size. In Figure A7, we in addition compare the cohort-specific effects generated by the structural and reduced-form estimates. As one can infer, the model tracks the reduced form estimates quite well. The average estimated effect implied by the structural model is 0.0039, which is comparable to the 0.0034 estimate from the reduced-form analysis. In Appendix Figure A9, we show that the model performs well in replicating other key moments, such as the average and age-specific probabilities of enrollment in a Bachelor's-level degree college.

4.4 Understanding the Relative Importance of Mechanisms

With the estimated parameters from the structural model at hand, we next evaluate the importance of the mechanisms for generating the college response to robots. To do so, we simulate the counterfactual effects of robots by eliminating each channel one at time while keeping the other mechanisms of impact unchanged (see Panel C of Table 7). We eliminate the opportunity cost effect by imposing $\omega^0_{young} = 0$, the college premium channel by imposing $\omega^1_{old} = \omega^0_{old} = 0$, and the parental income mechanism by imposing $\kappa = 0$. These results are illustrated in Figure 7. Eliminating the premium channel yields an effect of robots that is only one-third the actual observed effect, while eliminating the opportunity cost effect cuts in half the effect of robots. This shows that the college premium effect is more important than the opportunity cost in explaining an individual's response to the widespread adoption of industrial robots. When we eliminate both market incentive effects, the simulated effect of robots becomes negative. This should not come as a surprise given that it reflects the negative impact of robots on parental income. Indeed, shutting off the parental income channel leads to an effect of robots that is slightly larger than the observed baseline. Therefore, in the absence of the market incentive effects, the rapid advances in robotics that took place since the 1990s would likely have caused a decline in college attainment.

4.5 Recovering Key Elasticities

In panel D of Table 7, we present the implied elasticities of college with respect to parental income, opportunity cost and college premium. These elasticities are of independent interest because they could serve as inputs when evaluating the adjustment response of other technologies and labor market shocks. We obtain these elasticities by simulating the percentage change in college attainment generated by the robot-induced percentage change in one of these factors while keeping constant the others. Our calculations suggest an elasticity with respect to the college premium of 0.60. This estimate is slightly smaller than that documented in Long et al. (2015), who estimate an elasticity of 0.67 studying the effect of local labor market wages on major choices, but larger than the around 0.10 elasticity based on lab experimental variation among students of New York University (Wiswall and Zafar, 2015) and on business cycle variation in France (Beffy et al., 2012). The elasticity with respect to the opportunity cost is -0.18. Although it is difficult to find directly comparable estimates from the literature, the fact that it is noticeably smaller in magnitude than the premium elasticity is consistent with the evidence in Berger (1988) that the lifetime future earnings is more important than initial earnings for individuals deciding whether or not to attend college. This is why the premium channel plays a relatively more important role for generating the college response to robots despite that the decline in unskilled earnings among young adults, and thus in the opportunity cost, is disproportionately and remarkably larger than the increase in the college premium.

With respect to parental income, we observe an elasticity of 0.28. Existing research exploiting experimental variation in parental income suggests that this figure is reasonable. In a recent study, Bulman et al. (2021) use rich administrative datasets covering the US population and leverage variation in lottery wins to examine the impact of parental income shocks on children's college attendance. They show that the likelihood of having a four-year college degree increases by about 0.22 percent for each 1 percent increase in parental income around the year of high school. This estimate is only slightly lower than that obtained from our simple structural model.

4.6 Policy Counterfactuals: Robot Subsidies

From a policy perspective, it is important to understand whether and by how much government subsidies magnify or dampen the endogenous college response to robots. To investigate this question, we carry out simulations of alternative policies regarding the government college grants. In this exercise, we replace the actual system of subsidies by a government subsidy which covers 50 percent of the tuition costs for all students. For the state experiencing the average penetration of robots, this policy change would have increased the fraction of people with a Bachelor's degree by 1.3 percentage points (or 3.7 percent) rather than by 0.7 percentage points (or 1.8 percent). These findings suggests that an endogenous government subsidy to college can significantly enhance the magnitude of the college response to robots.

4.7 Earnings and Inequality Dynamics

We now turn to the question of what happens with the dynamics of earnings and inequality as younger generations enter the labor force and replace older ones. Understanding these dynamics is important to understand the long-run adjustment of the economy to changes in technology. To explore this question, we extend our partial equilibrium model to include the production sector. Consider an aggregate production function Q that combines high-skill labor (S_H), low-skill labor (S_L), and robots (R) to produce output. By high- and low-skill labor, we mean college and non-college labor. This production function takes the following constant elasticity of substitution form:

$$Q = \left[(a_H S_H)^{\frac{\eta-1}{\eta}} + (a_L S_L)^{\frac{\eta-1}{\eta}} + (a_R R)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \tag{16}$$

The $a's$ terms represent effective share parameters and $\eta \in (0, \infty)$ 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. Under this framework, 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 a_H and a_L with implications for the dynamic of earnings and inequality. More formally, let us assume that workers are paid their marginal products.

In this case, the automation-induced log-change in earnings can be expressed as follows:

$$d\log w_H = \frac{\eta-1}{\eta} d\log a_H + \frac{1}{\eta} d\log Q - \frac{1}{\eta} d\log S_H$$
$$d\log w_L = \frac{\eta-1}{\eta} d\log a_L + \frac{1}{\eta} d\log Q - \frac{1}{\eta} d\log S_L$$

The first two terms on the right-hand side of this system represent the effects of robots on earnings in the absence of a change in the supply of high- and low-skill labor. A skill-replacing technological change will reduce the demand for labor via a displacement effect, as captured by a decline in a_H and a_L. But this technological change also creates a productivity effect that increases the demand for both high- and low-skill labor, so the overall demand effect ultimately depends on the relative importance of these forces. The overall demand effect is not directly observed, but we can recover it for a given change in the observed earnings and supply of labor and given an estimate of η. Knowing the overall demand effect, we then can simulate the dynamics of earnings and inequality as new generations with different skills enter the workforce.

To obtain an approximate estimate of the overall demand shift we restrict the parameter η based on the evidence in Katz and Murphy (1992), who estimate the elasticity of substitution to be 1.4. Our findings in Section 3.4.A suggest that the earnings of low- and high-skill workers declined by 5.6 and 2.4 percent in the state experiencing the average penetration of robots between 1990 and 2008. The results in section 3.1 indicate the fraction of individuals with a Bachelor's degree increased by 1.8 percent in the same state. Since the cohorts exposed to the robot shock during or before the period of college decisions represent 25 percent of the workforce in 2008, these figures indicate that the overall demand effect is -5.6 and -2 percent for low- and high-skill workers respectively.

With these calculations, we now simulate the dynamic effects on earnings by altering the fraction of exposed cohorts that are part of the workforce and simulating the subsequent education response implied by the change in earnings and college premium. These results are shown in Figure 8. Panel A of this figure documents that the effect on high-skill earnings becomes more negative over time as younger generations enter the labor force, going from -2 to -3 percent. This is a direct consequence of the outward shift in the supply of high-skill labor as more people go to college. At the same time, there is a mechanical inward shift in the supply of low-skill labor and as a result, the decline in low-skill earnings becomes less negative over time. The magnitude of this adjustment response is relatively small. Even when the replacement is complete, the effect of robots on low-skill earnings over the long run would be only 10 percent smaller (in absolute value) than the baseline effect. This suggests that in the absence of policies that effectively push more people through college, workers at the end of the skill distribution are likely to continue experiencing significant earnings declines as robotic technology advances.

Note that since the earnings effects on high-skill workers become larger over time, earnings inequality falls. The automation-induced change in inequality can be expressed as follows:

$$d\log\left(\frac{w_H}{w_L}\right) = \frac{\eta-1}{\eta} d\log\left(\frac{a_H}{a_L}\right) - \frac{1}{\eta} d\log\left(\frac{S_H}{S_L}\right)$$

With this expression, it is straightforward to calculate the dynamics of inequality. Figure 8, panel B shows that the decline in earnings inequality is sizeable. When the generational replacement is complete, the automation-induced earnings gap between college- and non-college-educated persons in the mean state falls from 3.5 to 2 percent. This result implies that in the absence of the endogenous educational response, the long-run effect of robots on earnings inequality would be about 75 percent larger.

The role of policy. To explore the role of policy, panels C and D of Figure 8 repeat the same simulation exercise but introduce a government subsidy that covers 50 percent of the tuition costs for all students. With this policy change, the dynamics of earnings and inequality become steeper, with the effects on high-skill earnings becoming larger in magnitude with the replacement of old generations of workers. This leads to a long-run earnings gap induced by industrial robots of 1 percent, a reduction of a factor of 2 relative to the scenario with no policy change. Overall, these results suggest that policies that enhance college education have the potential to reduce the effects of automation technologies such as industrial robots on earnings inequality.

5 Implications and Discussion

The previous sections provide detailed empirical evidence that the unprecedented advances in robotic technology taking place since the 1990s have induced skill acquisition via an increase in the net returns to skill investments. These results raise the following question: if the skill premium is an important force inducing skill acquisition, why has there been little progress in the aggregate trends of college attainment for cohorts entering the labor market after the 1980s despite the rapid increase in the college premium? Most of the progress took place earlier, but college completion rates have remained roughly constant after the 1950s. A possible explanation is that college attendance decisions across states are substitutes. Specifically, if individuals systematically move from higher-to lower-exposed areas, they could crowd out the college education of other individuals in those areas. Consequently, as more people born in exposed regions complete college, we see less completion in other areas so that the aggregate college response to robotics technology is negligible. While we believe that such negative spillover effects are plausible, they are likely to play only a minor role in our state-level analysis. Most moves occur between counties within the same state and since our analysis is at the state-of-birth$\times$birth-cohort level, any spillover effect within the state is built into the estimate. Consistent with the limited scope for spillover effects between states, the data indicate that most individuals attend college in their state of birth —approximately 80 percent. Therefore, the few ones moving out would need to displace the college completion rates of a disproportionately large share of individuals in other states to annihilate the aggregate gain in college attainment. This seems implausible.

Another possibility, which we find more appealing, is that there have been changes in other factors offsetting the endogenous skill response to technology. This possibility has been already raised in Goldin and Katz (2009), who provide a comprehensive discussion on possible factors. They highlight the large baby boom, reduced college readiness, increased neighborhood segregation, and the sharp rise in public and private college tuition since 1980. Castro and Coen-Pirani (2016) evaluate the role of tuition costs, skill prices, and education quality in explaining the evolution of college attainment over the 20th century in a model of human capital investments. They provide evidence that rising tuition costs and declining learning ability account for almost all of the slowdown in college attainment of recent cohorts. Other studies employing an analogous approach reach similar conclusions (Jones and Yang, 2016; Donovan and Herrington, 2019). 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).

It is important to note that these insights do not imply that the skill premium has not been important. Our findings suggest that aggregate college attainment would have increased at a slower rate or even declined in the absence of the market incentive mechanisms generated by changes in automation technology. The parameters we estimate can serve as inputs to discipline models of skill-biased technological transitions with endogenous skill choice, as in Caselli and Manning (2019) and Guerreiro et al. (2022).

6 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 college 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.

Our exploration of mechanisms suggests that changes in the college premium and opportunity costs of college-going appear to be a major driver of the improvements in Bachelor's degree attainment. Areas with greater robot penetration witnessed a meaningful rise in the premium from having a Bachelor's degree and a decline in the opportunity cost of time as measured by the average earnings a young unskilled receives. We estimate a structural model of college decisions and find that the premium channel is the most important component generating the college response to robots, accounting for around two-thirds of the overall effect. Counterfactual simulations suggest that expansions in government grants have the potential to enhance significantly this endogenous human capital response to automation.

These findings have important implications for the long-run structure of earnings. Our model-based simulations suggest that the effect of robots on earnings inequality declines over time as younger generations of workers with different educational choices enter the market. This is driven by an increase in the relative supply of workers since entering worker generations are more likely to go to college. Ignoring this endogenous supply effect would lead to effects on inequality that are remarkably larger. A direct implication of these findings is that policies fostering access to higher education have the potential to mitigate the disruptive effects of labor-displacing technologies on inequality.

Bibliography

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, David Dorn, Gordon H Hanson, and Brendan Price**, "Import competition and the great US employment sag of the 2000s," *Journal of Labor Economics*, 2016, *34* (S1), S141–S198.

__ **,** __ **, 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.

__ **, James Banks, Costas Meghir, and Guglielmo Weber**, "Humps and bumps in lifetime consumption," *Journal of Business & Economic Statistics*, 1999, *17* (1), 22–35.

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, Frank Levy, and Richard J Murnane, "The skill content of recent technological change: An empirical exploration," *The Quarterly journal of economics*, 2003, *118* (4), 1279–1333.

__ **, 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.

Bai, Jushan, "Estimating multiple breaks one at a time," *Econometric theory*, 1997, pp. 315–352.

__ **and Pierre Perron**, "Estimating and testing linear models with multiple structural changes," *Econometrica*, 1998, pp. 47–78.

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.

Banks, James, Richard Blundell, and Agar Brugiavini, "Risk pooling, precautionary saving and consumption growth," *The Review of Economic Studies*, 2001, *68* (4), 757–779.

Beffy, Magali, Denis Fougere, and Arnaud Maurel, "Choosing the field of study in postsecondary education: Do expected earnings matter?," *Review of Economics and Statistics*, 2012, *94* (1), 334–347.

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

Berger, Thor and Per Engzell, "Intergenerational mobility in the fourth industrial revolution," 2020.

Blundell, Richard, Martin Browning, and Costas Meghir, "Consumer demand and the life-cycle allocation of household expenditures," *The Review of Economic Studies*, 1994, *61* (1), 57–80.

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

Bulman, George, Robert Fairlie, Sarena Goodman, and Adam Isen, "Parental resources and college attendance: Evidence from lottery wins," *American Economic Review*, 2021, *111* (4), 1201–40.

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.

Cascio, Elizabeth U and Ayushi Narayan, "Who Needs a Fracking Education? The Educational Response to Low-Skill-Biased Technological Change," *ILR Review*, 2015, p. 0019793920947422.

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.

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.

Ferreira, Francisco HG and Norbert Schady, "Aggregate economic shocks, child schooling, and child health," *The World Bank Research Observer*, 2009, *24* (2), 147–181.

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

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), 2550–93.

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.

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.

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.

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

Ionescu, Felicia and Nicole Simpson, "Default risk and private student loans: Implications for higher education policies," *Journal of Economic Dynamics and Control*, 2016, *64*, 119–147.

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.

Kalenkoski, Charlene Marie and Sabrina Wulff Pabilonia, "Parental transfers, student achievement, and the labor supply of college students," *Journal of Population Economics*, 2010, *23* (2), 469–496.

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, 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.

Long, Mark C, Dan Goldhaber, and Nick Huntington-Klein, "Do completed college majors respond to changes in wages?," *Economics of Education Review*, 2015, *49*, 1–14.

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.

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.

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

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

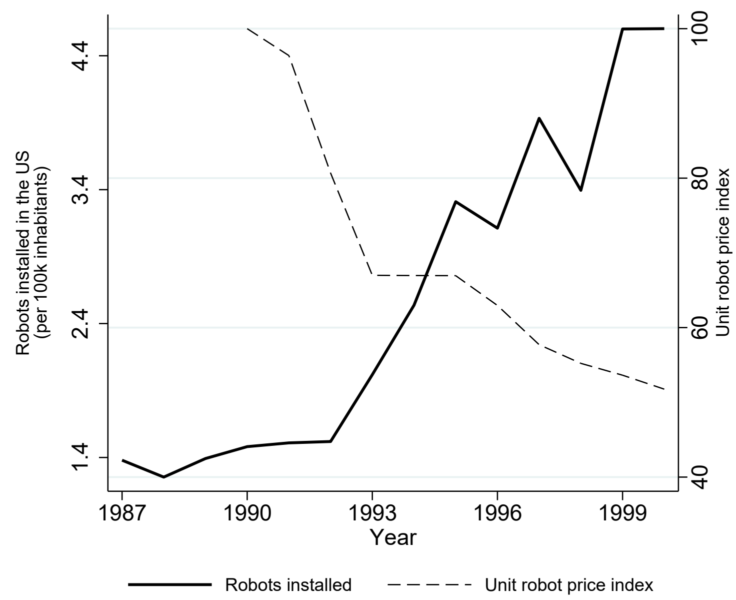

Notes: Data on newly installed robots come from the World Robotics (2001), whereas the robot price index is from the International Federation of Robotics (2006). The robot price index is calculated as an un-weighted arithmetic average price index across the countries with available annual price data: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

Figure 2: Adjusted Penetration of Robots across States

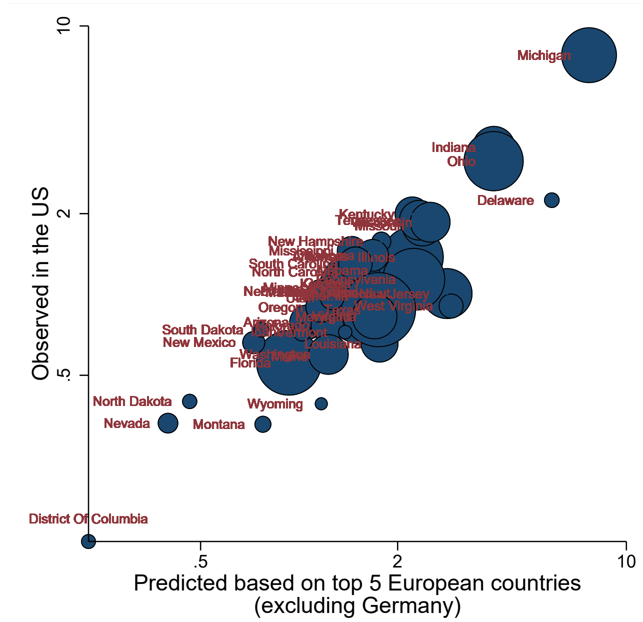

Notes: This figure plots the exposure to robots across based on the adjusted penetration of robots in the United States and top 5 countries (excluding Germany). The adjusted penetration of robots is measured for the 2004-2007 period (rescaled to a 14-year equivalent change) for the United States, and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

Figure 3: Timing of Robot Impacts on Labor Markets

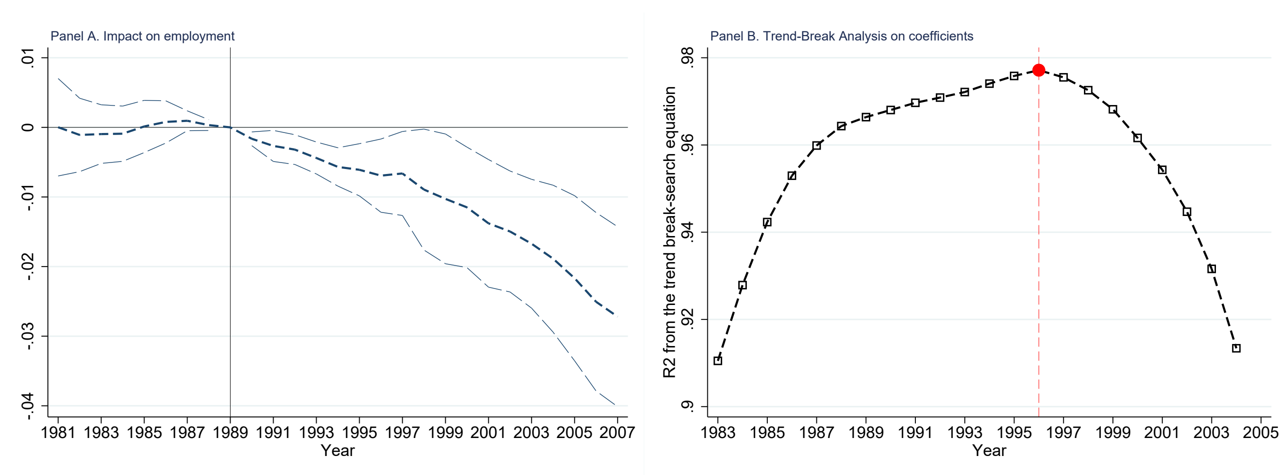

Notes: Panel A presents estimates of γ_t from estimating equation (2) for different years. Control variables include a set of 1990 state characteristics: og population, share of population that is under five, share of population that is over 65, share of population that is black, share of population that is urban, and share of population that is part of the labor force. Standard errors are robust to arbitrary forms of heteroskedasticity. Panel B reports the R^2 from the following trend-break equation on the coefficients obtained in panel A: $\hat{\gamma}_t = \alpha + \tau \cdot t + \lambda \cdot (t - t^*)\mathbf{1}\{t > t^*\} + \epsilon_t$. In this expression, t^* is the year of the potential structural break. Following the literature on structural break (Bai, 1997; Bai and Perron, 1998), the date that maximizes the R^2 represents the location of the structural break.

Figure 4: Flexible Estimates

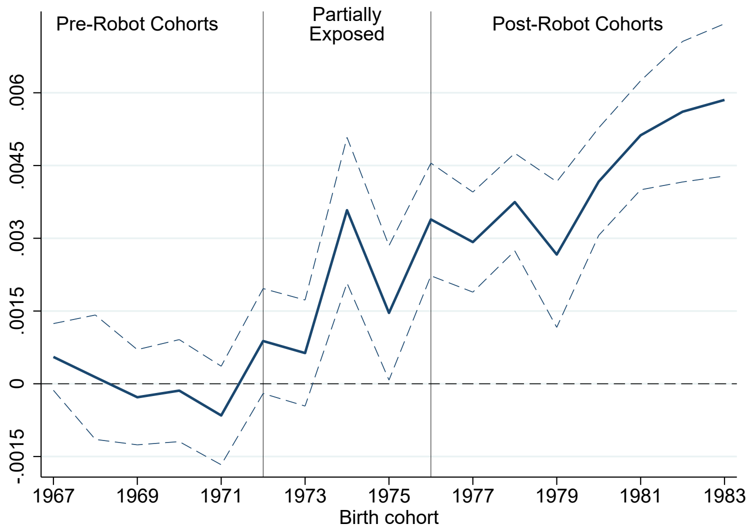

Notes: This figure plots estimates of the interaction between the robot exposure variable and indicators for birth years, using the flexible version of model (3). The interaction term for individuals born in 1966 is normalized to zero. The dashed lines represent 95 percent confidence intervals based on standard errors clustered at the state-of-birth level. See notes to column (3) of Table 1 for details on sample and specification.

Figure 5: Effects on Education Groups

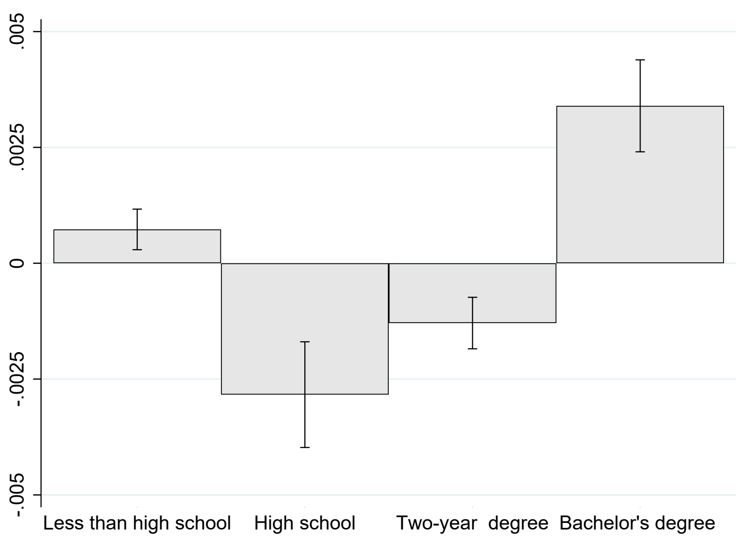

Notes: This figure shows estimates of the effects of robots on different education categories, based on specification (3). The confidence intervals are based on standard errors clustered at the state-of-birth level. See notes to column (3) of Table 1 for details on sample and specification.

Figure 6: Comparing Model-Based Estimates to Empirical Results

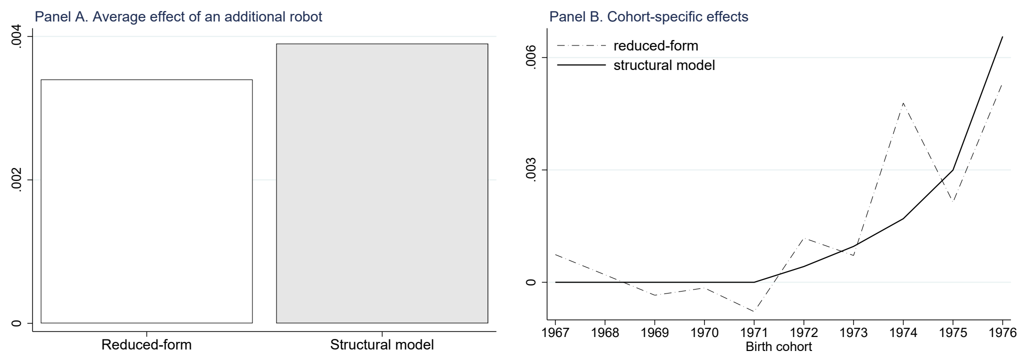

Notes: This figure compares the structural to reduced-form effects of robots on bachelor attainment. To compute the average effect implied by the structural model in Panel A, we first calculate the marginal effect of an additional robot for each cohort and then take the weighted average across all cohorts, weighting by cohort size. Panel B compares the cohort-specific effects generated by the structural and reduced-form estimates.

Figure 7: Simulated Effects on Bachelor Attainment by Components

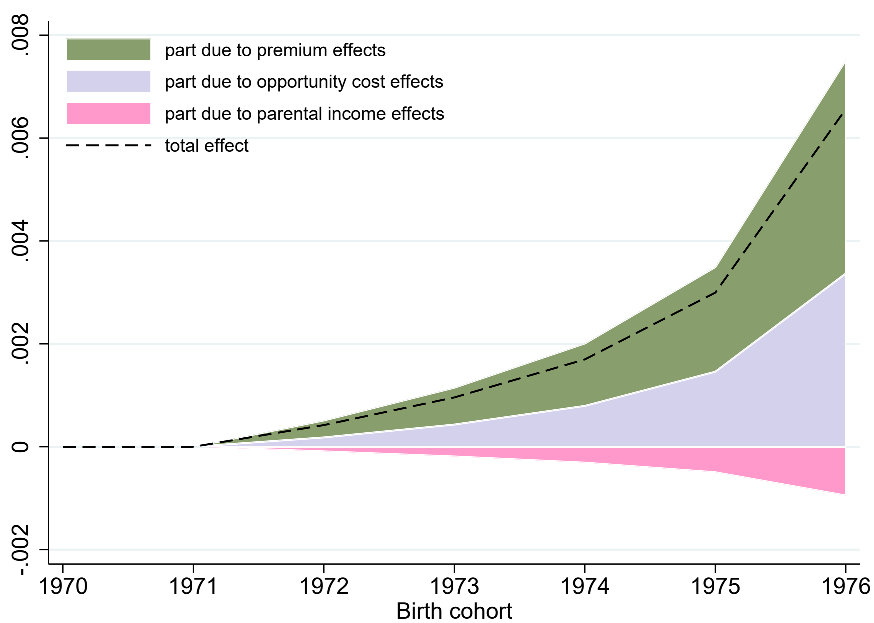

Notes: This figure decomposes the simulated effects of robots on bachelor attainment into the premium, opportunity cost, and parental income effects.

Figure 8: Simulated Dynamics of Earnings and Inequality

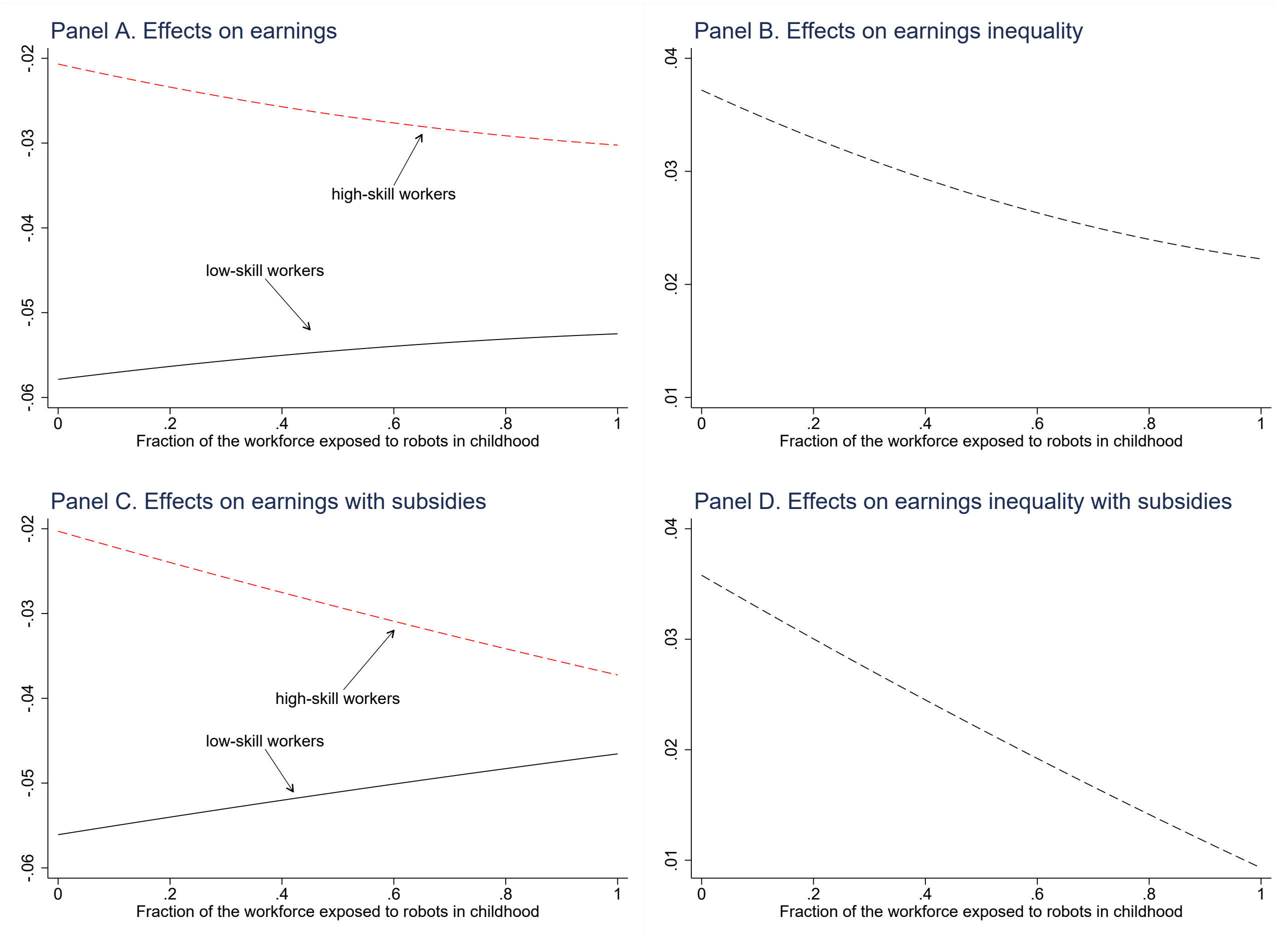

Notes: This figure simulates the dynamics of earnings and inequality as younger generations replace older ones in the labor market.

Table 1: Childhood Exposure Effects on College Attainment

	Dependent variable is Bachelor's degree or higher			
	(1)	(2)	(3)	(4)
Robot penetration × post	0.0027	0.0027	0.0034	0.003
	[0.0016]	[0.0006]	[0.0005]	[0.0005]
Mean Dep. Variable	0.3617	0.3617	0.3617	0.3833
Observations	7111535	7111535	7111535	2611948
1990 college level × birth-year FE		✓	✓	✓
1990 demographics × birth-year FE			✓	✓
Sample restricted to 2015-19 ACSs				✓
State-of-birth FE	✓	✓	✓	✓
Birth-year FE	✓	✓	✓	✓

Notes: This table reports estimates of β in equation (3). 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 after 1976. 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, and survey-year fixed effects. The 1990 demographics controls include: log population, the share of population that is under five, the share of population that is over 65, the share of population that is black, the share of population that is urban, and the share of population that is part of the labor force. Column (4) 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 College Attainment
(Mean Reversion, Pre-cohort Trends and Within-Region Variation)

	Dependent variable is Bachelor's degree or higher				
	(1)	(2)	(3)	(4)	(5)
Robot penetration × post	0.0034	0.0036	0.0032	0.0034	0.0037
	[0.0005]	[0.0009]	[0.0007]	[0.0006]	[0.0010]
Mean Dep. Variable	0.3617	0.3617	0.3617	0.3617	0.3617
Observations	7111535	7111535	7111535	7111535	7111535
1960, 1970, and 1980 college level (× birth-year FE)		✓			
1960-1990 change in college level (× birth-year FE)			✓		
State-specfic pre-cohort linear 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. Post is an indicator for individuals born after 1976. 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.

Table 3: Childhood Exposure Effects on College Attainment
(Controlling for Other Labor Market Shocks, and Social Reforms)

	Dependent variable is Bachelor's degree or higher					
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration × post	0.0034	0.0031	0.0028	0.0028	0.0028	0.0027
	[0.0005]	[0.0005]	[0.0008]	[0.0007]	[0.0007]	[0.0007]
Mean Dep. Variable	0.3617	0.3617	0.3617	0.3617	0.3617	0.3617
Observations	7111535	7111535	7111535	7111535	7111535	7111535
Adding controls for:						
Chinese import competition		✓				
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 3 of Table 1. Column 2 includes interactions between birth-year fixed effects and the intensity in Chinese import competition (as measured in Autor et al. (2013)) in the state of birth. Column (3) controls for interactions between birth-year fixed effects and the measure of 1980-82 recession severity constructed by Stuart (2022) in the state of birth. Column (4) 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 (5) includes the fraction of childhood years exposed to Medicaid in the state of birth. Column (6) adds the fraction of school-going ages (5 to 17) exposed to a school finance reform in the state of birth. Robust standard errors in brackets are clustered at the state-of-birth level.

Table 4: Childhood Exposure Effects on Income

	Dependent variable is					
	Log total income		Log earned income		Log income wages	
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration × post	0.0031	0.0006	0.0021	-0.0003	0.0022	-0.0002
	[0.0008]	[0.0008]	[0.0008]	[0.0008]	[0.0008]	[0.0008]
Bachelor's degree indicator		0.717		0.6742		0.6651
		[0.0082]		[0.0071]		[0.0068]
Observations	6498308	6498308	6054011	6054011	5738488	5738488
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes: This table reports estimates of β in equation (3) for different income measures as outcomes. Sample sizes vary across outcomes because of missing observations. See notes to Table 1 for details on sample and baseline covariates. Robust standard errors in brackets are clustered at the state-of-birth level.

Table 5: Robots, Opportunity Costs, Skill Premium, and Parental Income

	Long differences, 1990-2008							
	Log earnings				Log skill premium			
	ages 18-25	ages 25-65			ages 25-65			
	No college	No college	Two-year degree	Bachelor's degree	Bachelor vs no college	Bachelor vs two-year degree	Two-year degree vs No college	Log parental income
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Robot penetration	-0.0309	-0.0282	-0.0219	-0.0124	0.0117	0.0087	0.0028	-0.0057
	[0.0049]	[0.0027]	[0.0017]	[0.0018]	[0.0023]	[0.0021]	[0.0017]	[0.0018]
Observations	18736	90912	89427	96302	62190	64610	62328	193411
Baseline covariates	✓	✓	✓	✓	✓	✓	✓	✓

Notes: This table reports the results from estimating equation (5). 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 (8), the demographic groups include in addition education level categories. All regressions control for the baseline demographic and socioeconomic state characteristics described in Table 1. In addition, all the 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 6: Childhood Exposure Effects on Occupational Characteristics of Majors

		100 × Share of				
	Avg. routine task share	Operators, assemblers, inspectors, and production	Operators	Assemblers	Inspectors	Precision production
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration × post	-0.09165	-0.00176	-0.00098	-0.00029	-0.00049	-0.0024
	[0.01356]	[0.00037]	[0.00017]	[0.00009]	[0.00021]	[0.00051]
Mean Dep. Variable	27.01	0.78	0.39	0.15	0.24	1.11
Observations	1987977	1987977	1987977	1987977	1987977	1987977
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes: This table reports estimates of β in equation (3) for occupational characteristics of majors. Column (1) considers the average routine task share in each major, as defined by equation (6). Columns (2)-(6) repeat the same logic as in column (1) but consider specific occupations. The sample is limited to individuals with a Bachelor's degree. See notes to Table 1 for details on the sample and basic controls. Robust standard errors in brackets are clustered at the state-of-birth level.

Table 7: Estimates of Model Parameters

Panel A: Parameter Estimates		Estimate	SE	*Panel B: Comparison of effects*	
Coefficient of relative risk aversion	σ	1.3781	[0.0106]	Reduced-form, %	Structural model, %
Disutility of college by age at entry :				1.89	2.167
19	ψ_0	0.3549	[0.0073]		
20	ψ_1	0.5239	[0.0064]		
21	ψ_2	0.6147	[0.0077]		
22	ψ_3	0.6301	[0.0072]		
23	ψ_4	0.6300	[0.0059]		
Panel C: Simulated Effects, %				*Panel D: Implied Elasticities*	
Eliminating premium effect	0.778			Premium elasticity	0.5962
Eliminating opp. cos effect	1.056			Opportunity cost elasticity	-0.1837
Eliminating income effect	2.444			Parental income elasticity	0.2842

Notes: Panel A reports results from estimating the model via the general method of moments. Panel B compares the model-generated relative effects to the reduced-form estimates. These relative effects are calculated for the state experiencing the average penetration of robots. Panel C simulates the effects of robots by eliminating each channel one at time while keeping the other mechanisms unchanged. Panel D reports the implied elasticities by the model.

Online Appendix

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 1966-83 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. We also exclude individuals residing in institutional group quarters to increase consistency between the different rounds of the ACS. The basic sample consists of approximately 7.1 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 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 by 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 of average robot adoption in the top 5 non-US countries with greater advances in robotics (Denmark, Finland, France, Italy, and Sweden). 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. In some results, we also present results using a measure of robot penetration for 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 A1.

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 A2 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 A5 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 A6 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 A8 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.

B Details on Some Additional Results and Robustness Checks

B.1 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 *Robot penetration*$_s$ $\times$ *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 $\ell_{sj}\cdot$ *Robot penetration*$_j$ $\times$*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:

$$\beta = \textstyle\sum_j \alpha_j \tilde{\beta}_j$$

$$\text{where} \qquad \tilde{\beta}_j \equiv \big(\textstyle\sum_i z_{ijst}\cdot x_{ist}\big)^{-1}\sum_i z_{ijst}\cdot \tilde{y}_{ist}$$

$$\alpha_j \equiv \big(\textstyle\sum_j \mu_j \sum_i z_{ijst}\cdot x_{ist}\big)^{-1}\mu_j \sum_i z_{ijst}\cdot x_{ist}$$

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 A7 reports estimates of $\tilde{\beta}_j$ and α_j for the top five industries with the highest Rotemberg weights. We find that the automotive industry has the largest share of the overall weight, about 89 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 coefficients β are in fact somewhat larger when one excludes the automotive industry. Indeed, the coefficient for bachelor attainment goes from 0.0034 in the baseline to 0.0060 when we exclude the automotive industry.

B.2 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 A8, the results are in general very similar to our baseline.

B.3 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 (3) using these measures of program exposure as dependent variables. Table A9 documents that if anything the post-robot cohorts from states with greater robot penetration are *less* likely to have been exposed to these programs in childhood. This suggests that these programs cannot explain the gains in college attainment we report in Table 1.

Figure A1: Birth Cohorts in the ACS

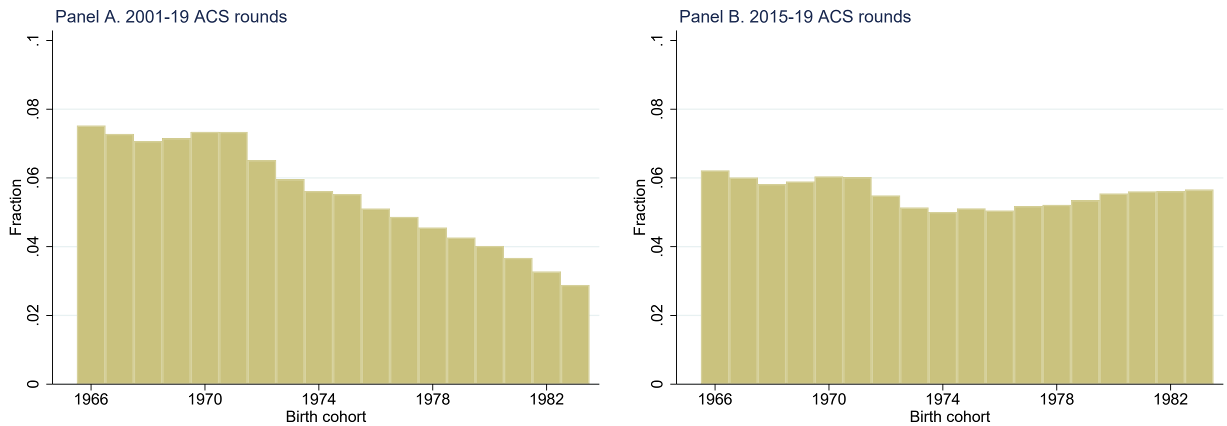

Figure A2: Observations by Survey Year

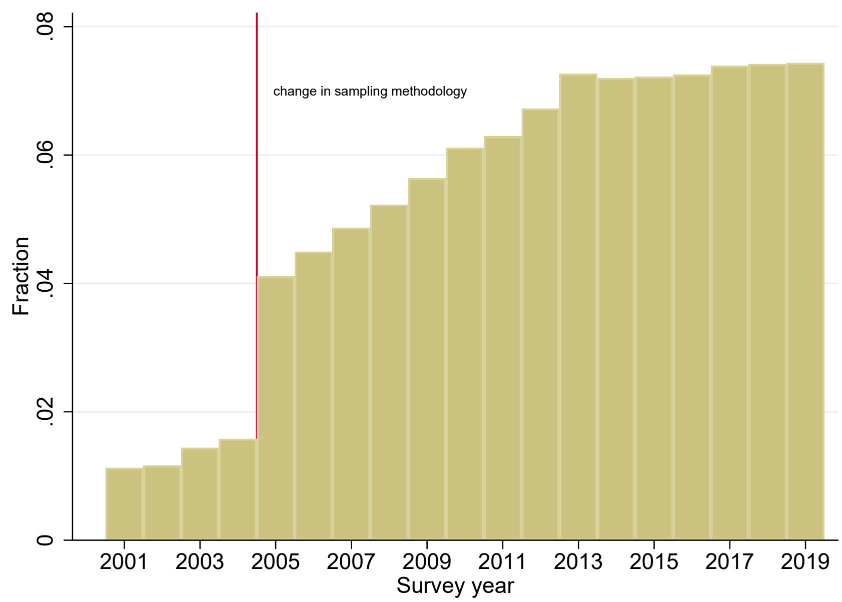

Figure A3: Composition of Birth Cohorts by Survey Year

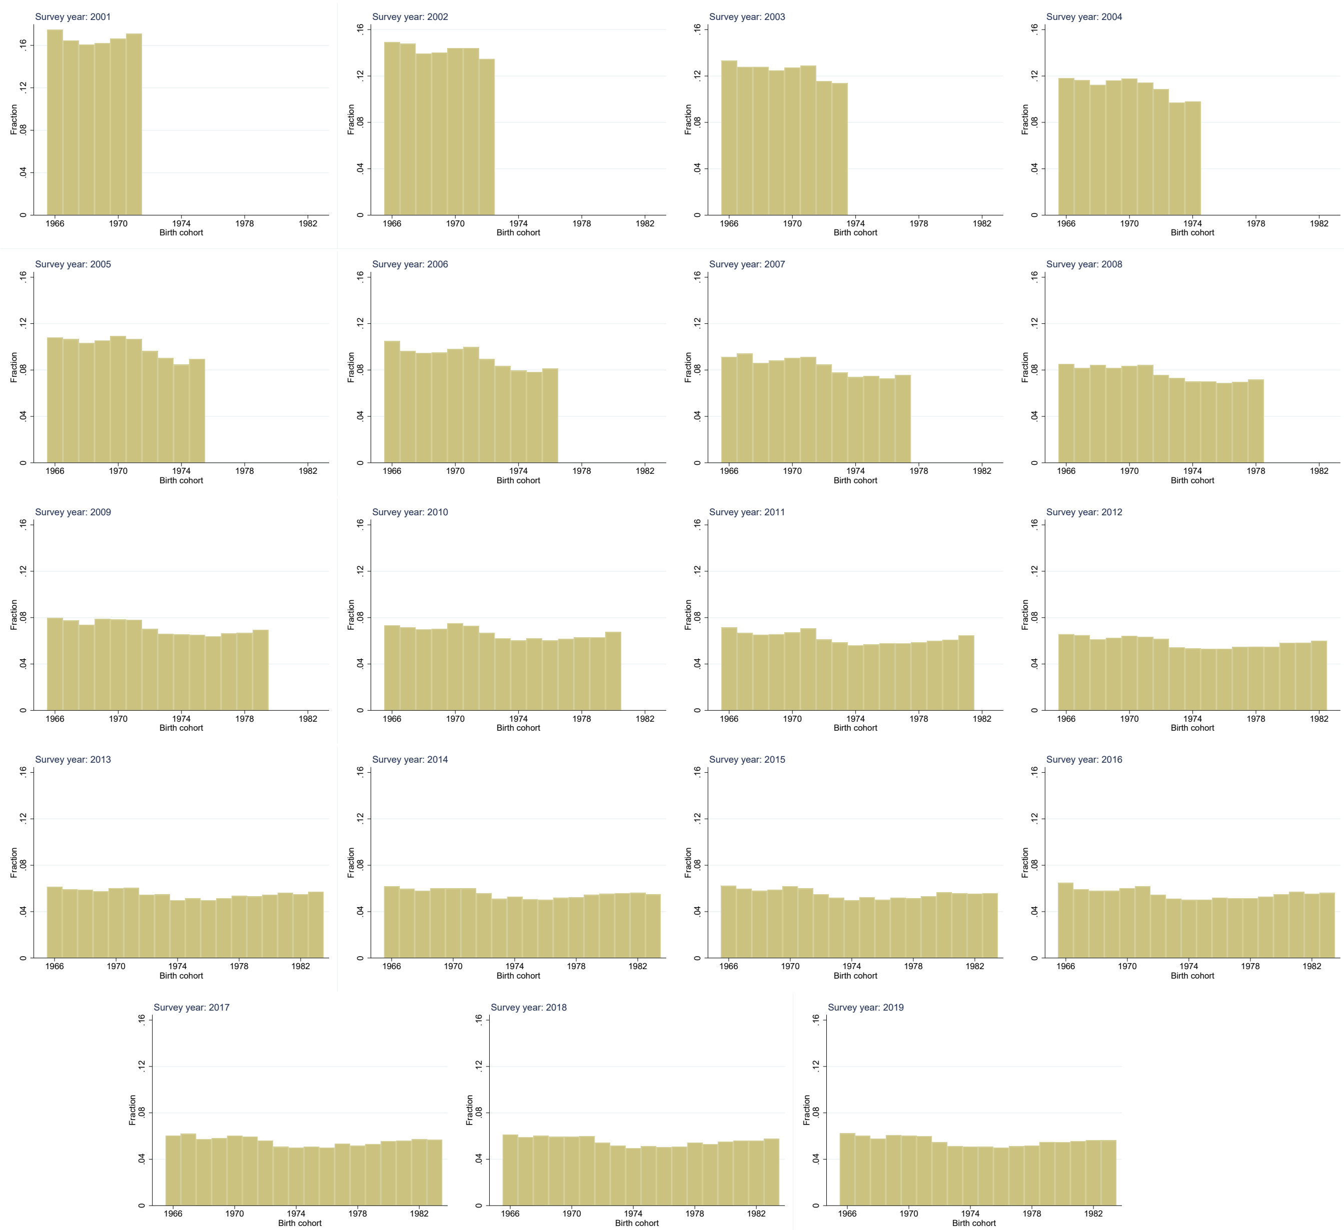

Notes: This figure shows the birth cohorts by survey year.

Figure A4: College Attendance by Age

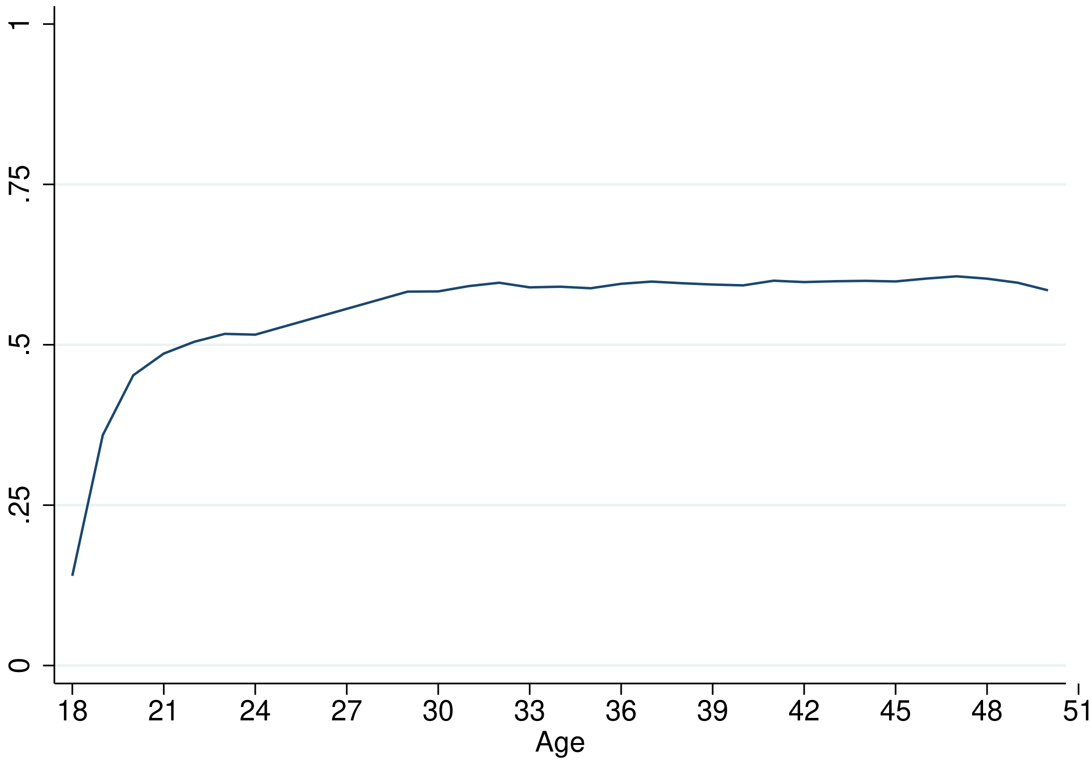

Notes: This figure shows the share of individuals with some college, for a constant sample born in the mainland of the United States from 1966 to 1972. We use the 1990 census and the 2001-2019 ACS to generate these figures.

Figure A5: Robot Exposure Intensity by State

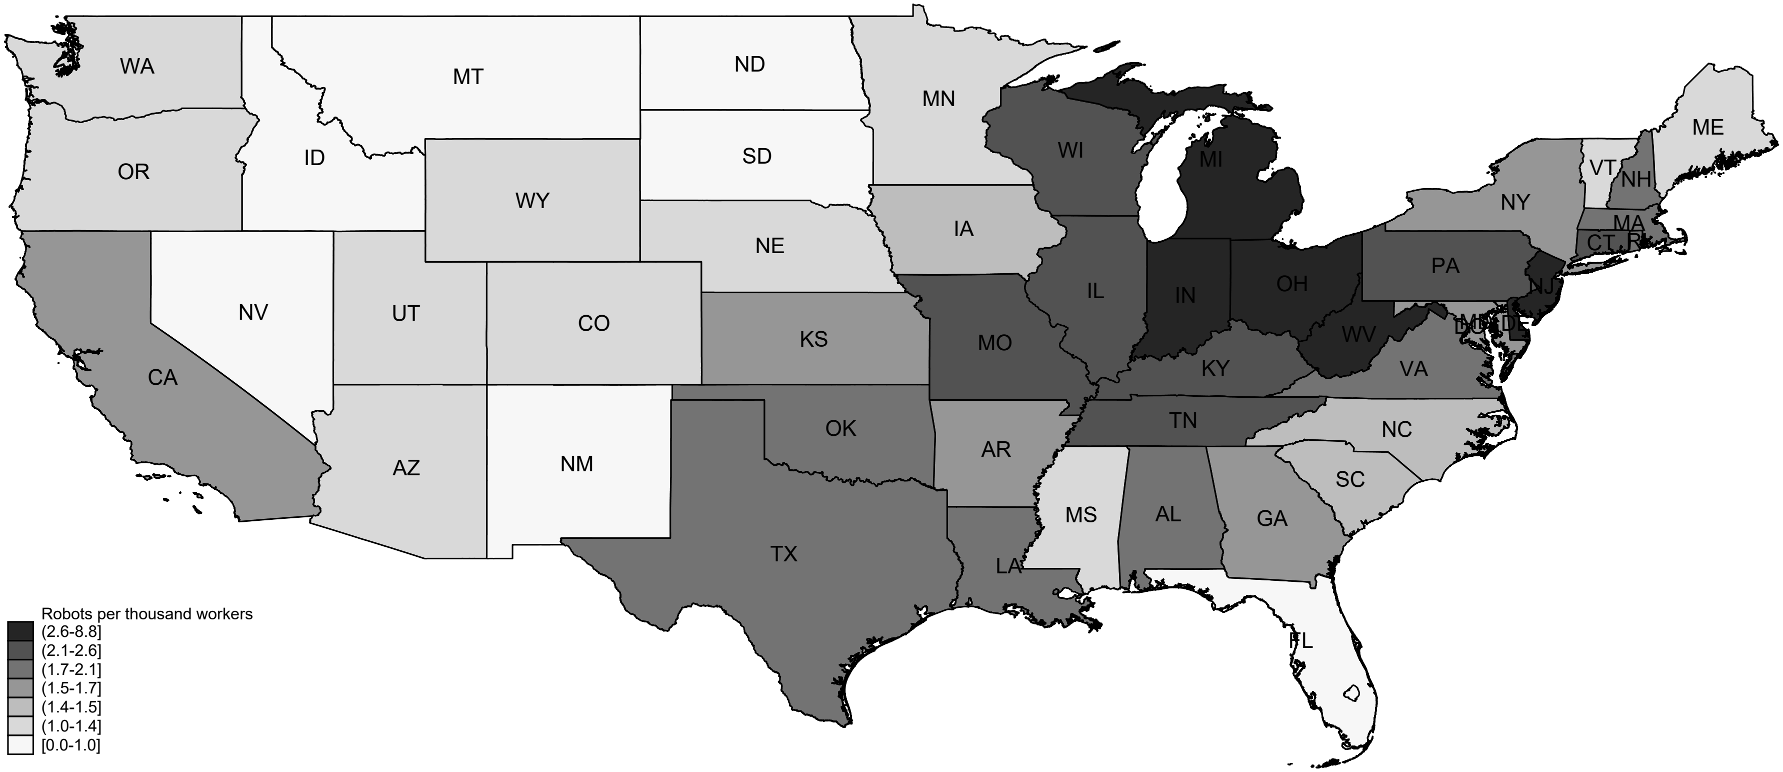

Notes: This map displays the intensity of robot exposure across states.

Figure A6: Baseline Industrial Shares across States

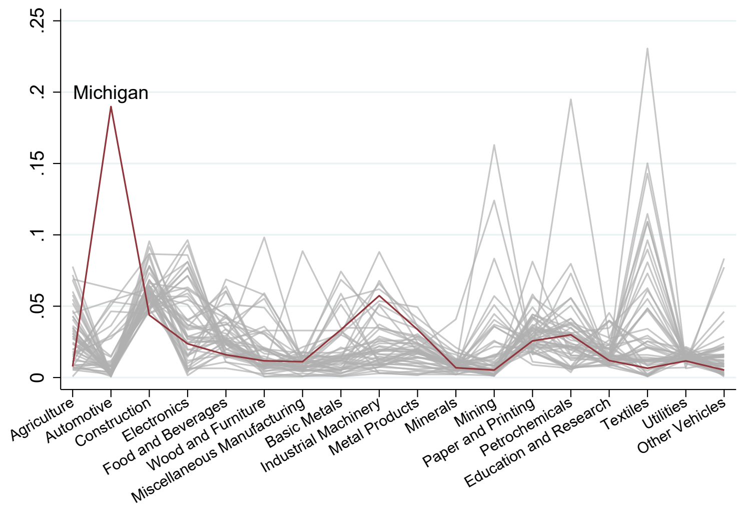

Notes: This figure shows the variation in initial industrial shares by state. This figure is constructed using data from the 1970 census.

Figure A7: Robot Exposure in Selected States and Industries

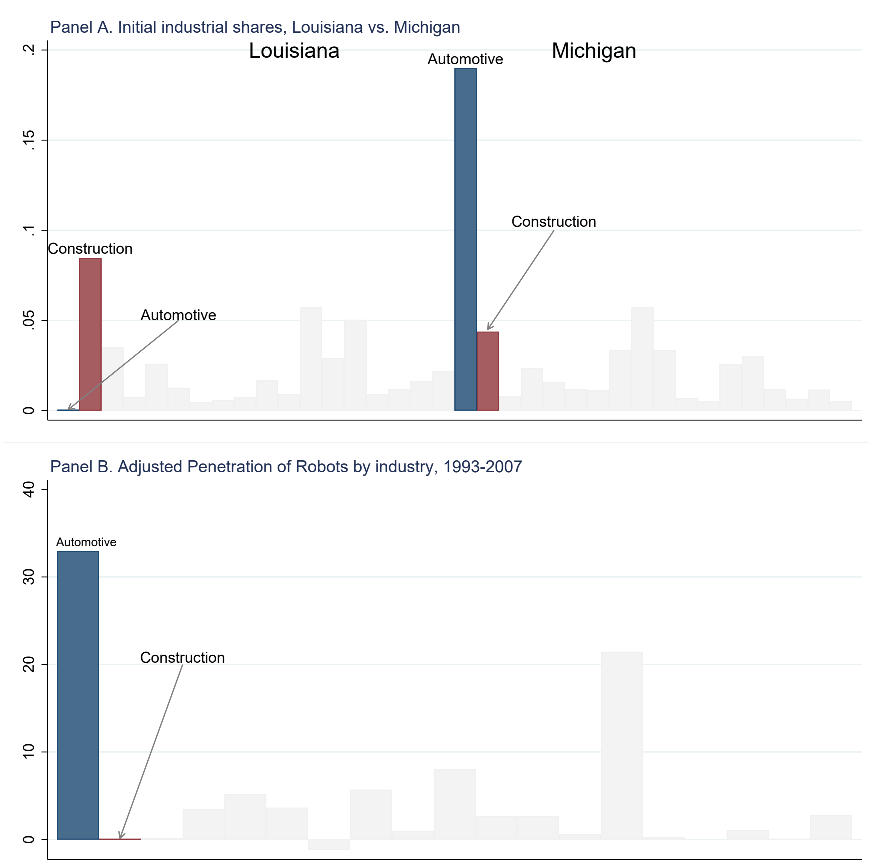

Notes: These figures show the variation underlying the robot exposure variable. Panel A shows the initial industry's share of employment in Louisiana and Michigan. Panel B shows the adjusted penetration of robots in each industry.

Figure A8: Robot Exposure across Commuting Zones after Removing State Effects

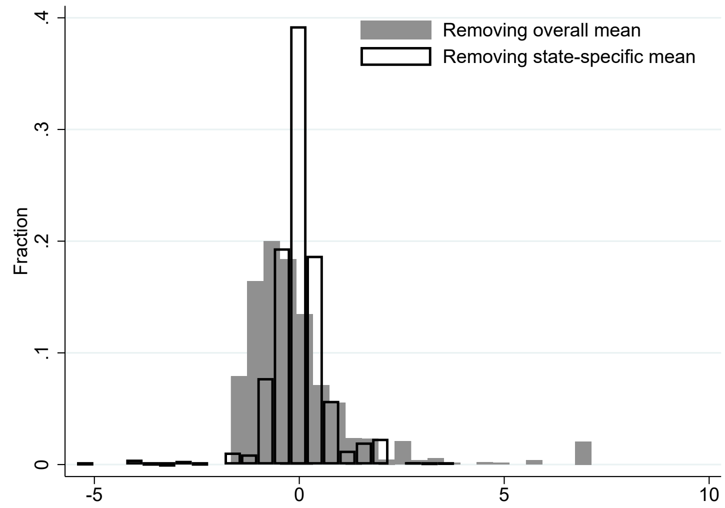

Notes: This figure show the variation the robot exposure variable across commuting zones.

Figure A9: Bachelor Attainment and Enrollment Probability in Model vs. Data

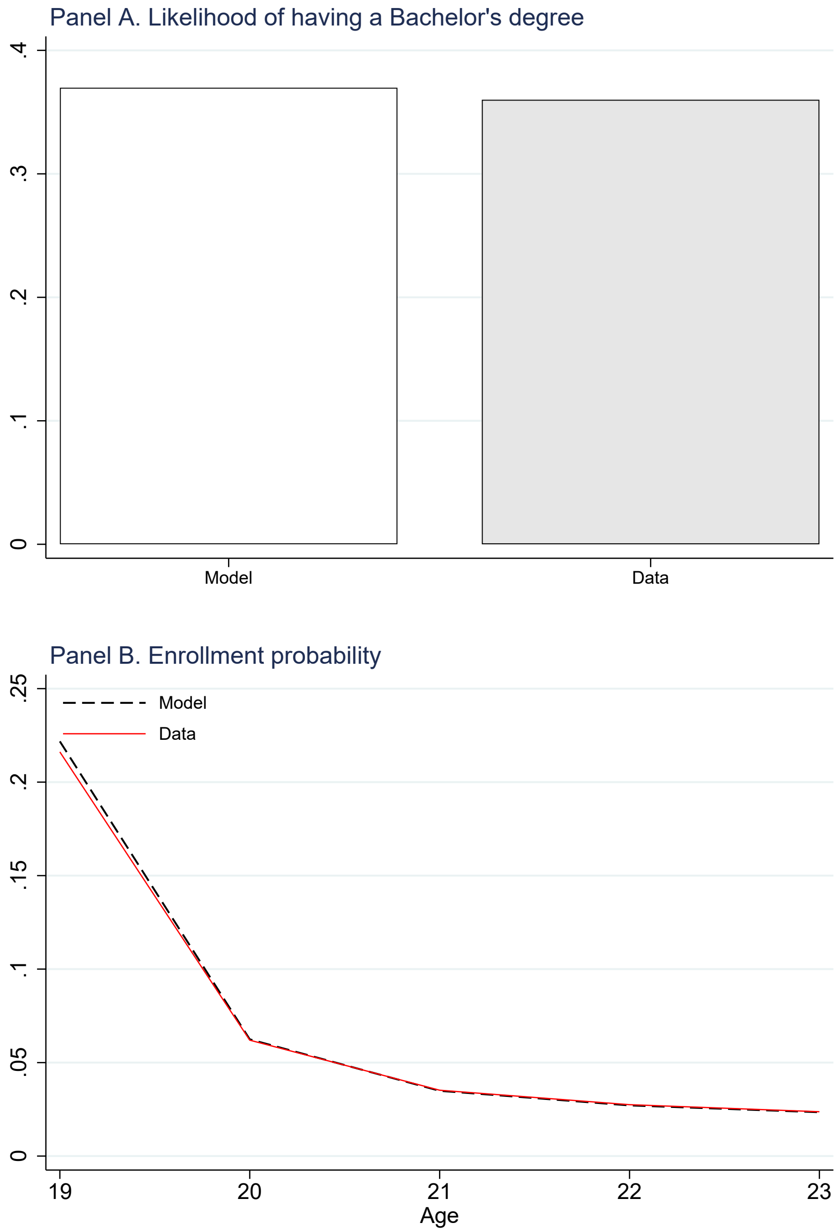

Notes: This figure compares actual and model-based moments.

Table A1: 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 A2: Robot Exposure

	Mean	Standard Deviation	Observations	
			N	Aggreg. Level
Robots per thousand workers	2.09	1.35	49	States
Adjusted penetration of robots per thousand workers (overall)	4.769	8.461	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.

Table A3: Exposure Effects on Bachelor's Degree (Controlling for Initial Market Conditions)

	Dependent variable is Bachelor's degree or higher				
	(1)	(2)	(3)	(4)	(5)
Robot penetration × post	0.0034	0.0029	0.0039	0.0028	0.0029
	[0.0005]	[0.0005]	[0.0007]	[0.0007]	[0.0009]
Mean Dep. Variable	0.3617	0.3617	0.3617	0.3617	0.3617
Observations	7111535	7111535	7111535	7111535	7111535
1990 manufacturing share (× birth-year FE)		✓			✓
1990 employment-to-pop. ratio (× birth-year FE)			✓		✓
1990 avg. wages (× birth-year FE)				✓	✓
Baseline covariates	✓	✓	✓	✓	✓

Notes: This table explores the robustness of the baseline estimates to additional controls for baseline labor market conditions. 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. Post is an indicator for individuals born after 1976. 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). Robust standard errors in brackets are clustered at the state-of-birth level.

Table A4: Exposure Effects on Bachelor's Degree
(2SLS Estimates)

	Dependent variable is:			
	US robot penetration × post	Bachelor's degree or higher		
	First stage	Reduced-form	OLS	2SLS
	(1)	(2)	(3)	(4)
US robot penetration × post	1.0158		0.0028	0.0033
	[0.0465]		[0.0003]	[0.0006]
Robot penetration × post		0.0034		
		[0.0005]		
Cragg and Donald (1993) F statistic	477.68			
Mean Dep. Variable	0.4168	0.3617	0.3617	0.3617
Observations	7111535	7111535	7111535	7111535
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. Post is an indicator for individuals born after 1976. Robots 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). Robust standard errors in brackets are clustered at the state-of-birth level.

Table A5: Exposure Effects on Bachelor's Degree
(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	0.0034	0.0059	0.0025	0.0023	0.0022
	[0.0005]	[0.0007]	[0.0003]	[0.0003]	[0.0003]
Rescaled coefficient	0.0034	0.0036	0.0033	0.0029	0.0032
Mean Dep. Variable	0.3617	0.3617	0.3617	0.3617	0.3617
Observations	7111535	7111535	7111535	7111535	7111535
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 (3) of Table 1. Column (2) uses the 1990 rather than 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 sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

Table A6: Exposure Effects on Bachelor's Degree (Outlier Analaysis)

	Baseline	Exclude 3-sigma outliers	Exclude 2-sigma outliers	Exclude 1-sigma outliers	Exclude 0.5-sigma outliers	Exclude high leverage observations	Exclude highly influential observations (Cook's distance)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Robot penetration × post	0.0034	0.0034	0.0034	0.0034	0.0036	0.0033	0.0036
	[0.0005]	[0.0005]	[0.0005]	[0.0005]	[0.0005]	[0.0004]	[0.0004]
Mean Dep. Variable	0.3617	0.3615	0.3611	0.3606	0.3614	0.3395	0.3615
Observations	7111535	7109758	7099378	7052943	6892806	6870213	6769715
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 with leverage superior to $2k/N$, with k being the number of predictors and N the number of observations. Column (7) excludes observations that shift the baseline estimate at least to $4/N$ (Cook's distance). See notes to Table 1 for details on sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

Table A7: Exposure Effects on Bachelor's Degree
(Top 5 Industries with Largest Rotemberg Weights)

	Rotemberg Weight (1)	Industry-specific Estimate of β (2)	Estimate of β without industry k (3)
Automotive	0.8947	0.0029	0.0060
Petrochemicals	0.0840	0.0058	0.0032
Basic Metals	0.0292	0.0078	0.0033
Metal products	0.0256	0.0044	0.0034
Industrial Machinery	0.0081	0.0041	0.0034

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. Column (1) table presents Rotemberg weights for the top 5 industries with the highest weights, following Goldsmith-Pinkham et al. (2020). The even-numbered columns report industry-specific coefficients, which are obtained in a 2SLS regression where the measure of robot exposure based only on industry k is used as an instrumental variable for the overall robot exposure variable. The remaining columns show the estimated coefficient when each industry is excluded from the overall measure of robot exposure. See notes to Table 1 for details on sample and specification.

Table A8: Exposure Effects on Bachelor's Degree
(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	0.0034	0.0034	0.0034	0.0034	0.0034
	[0.0005]	[0.0007]	[0.0004]	[0.0004]	[0.0006]
Mean Dep. Variable	0.3617	0.3617	0.3617	0.3617	0.3617
Observations	7111535	7111535	7111535	7111535	7111535
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 sample and specification.

Table A9: 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	-0.0018	-0.0027	-0.0168	-0.0085	-0.023
	[0.0019]	[0.0011]	[0.0062]	[0.0048]	[0.0175]
Mean Dep. Variable	0.3019	0.8182	0.9192	0.9752	0.3893
Observations	7111535	7111535	7111535	7111535	7111535
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 sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

Table A10: Exposure Effects on Bachelor's Degree
(Controlling Flexibly for Other Social Reforms)

		Controlling for birth-cohort FE × adoption year of...				
	Baseline	Community Health Centers	Head Start	Food Stamp	Medicaid	School Finance Reforms
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration × post	0.0034	0.0032	0.0034	0.0033	0.0035	0.0034
	[0.0005]	[0.0006]	[0.0005]	[0.0005]	[0.0005]	[0.0005]
Mean Dep. Variable	0.3617	0.3617	0.3617	0.3617	0.3617	0.3617
Observations	7111535	7111535	7111535	7111535	7111535	7111535

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 sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

Table A11: Exposure Effects on Income
(Controlling state-of-residence×Birth-cohort Fixed Effects)

	Dependent variable is					
	Log total income		Log earned income		Log income wages	
	(1)	(2)	(3)	(4)	(5)	(6)
Robot penetration × Post	0.0031	0.0082	0.0021	0.0077	0.0022	0.0082
	[0.0008]	[0.0016]	[0.0008]	[0.0022]	[0.0008]	[0.0023]
Observations	6498308	6498308	6054011	6054011	5738488	5738488
State-of-residence × birth-year FE		✓		✓		✓
Baseline covariates	✓	✓	✓	✓	✓	✓

Notes: This table demonstrates the robustness of the baseline estimates to controlling for state-of-residence× birth-cohort fixed effects. See notes to Table 1 for details on sample and specification. Robust standard errors in brackets are clustered at the state-of-birth level.

Table A12: Robots and Supply-Side Responses

	Log-differences 1990-2008:			
	Average net tuition and fees costs	Average revenue from state and local appropiations	Average revenue from state and local grants	Total Government transfers in education and training assistance
	(1)	(2)	(3)	(4)
Robot penetration	-0.013	-0.047	0.012	-0.025
	[0.012]	[0.048]	[0.019]	[0.030]
Mean Dep. Variable	1.132	1.605	0.367	1.062
Observations	49	49	49	49
Baseline covariates	✓	✓	✓	✓

Notes: This table reports results from estimating equation (5): $\Delta Y_{s,90-08} = \alpha + \gamma Robots_s + \mathbf{Z}_s'\Omega + \xi_s$. 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.

Table A13: Estimates of Model Parameters
(Sensitivity analyses)

		Baseline Estimate	Alternative interest rate	Alternative life of the loan 10 years	20 years	30 years	Perceived enrollment probability set to zero
		(1)	(2)	(3)	(4)	(5)	(6)
Coefficient of relative risk aversion	σ	1.3781	1.4089	1.3781	1.3781	1.3781	1.3867
		[0.0106]	[0.0107]	[0.0106]	[0.0106]	[0.0107]	[0.0106]
Disutility of college by age at entry:							
19	ψ_0	0.3549	0.3806	0.3546	0.3548	0.3549	0.2794
		[0.0073]	[0.0086]	[0.0073]	[0.0073]	[0.0073]	[0.0070]
20	ψ_1	0.5239	0.5465	0.5238	0.5238	0.5239	0.4976
		[0.0064]	[0.0067]	[0.0064]	[0.0064]	[0.0064]	[0.0063]
21	ψ_2	0.6147	0.6344	0.6146	0.6146	0.6147	0.597
		[0.0077]	[0.0088]	[0.0077]	[0.0077]	[0.0077]	[0.0072]
22	ψ_3	0.6301	0.6489	0.6299	0.63	0.63	0.6353
		[0.0072]	[0.0083]	[0.0072]	[0.0072]	[0.0072]	[0.0069]
23	ψ_4	0.63	0.6483	0.6297	0.6297	0.6298	0.6555
		[0.0059]	[0.0072]	[0.0059]	[0.0059]	[0.0059]	[0.0057]

Notes: This table explores the robustness of the structural results to alternative assumptions regarding some of the calibrated parameters. Column (1) repeats the baseline specification. Column (2) sets the interest rate to 3 percent (instead of 5 percent). Columns (3)-(5) explore different repayment periods, ranging from 10 to 30 years. Column (6) assumes that the perceived probability of enrollment in subsequent periods is zero.

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.

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.

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

Ferman, Bruno, "A simple way to assess inference methods," *arXiv preprint arXiv:1912.08772*, 2021.

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*, Forthcoming.

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.

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.