

Chuan, Amanda; Zhang, Weilong

Working Paper

Non-college Occupations, Workplace Routinization, and the Gender Gap in College Enrollment

IZA Discussion Papers, No. 16089

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Chuan, Amanda; Zhang, Weilong (2023) : Non-college Occupations, Workplace Routinization, and the Gender Gap in College Enrollment, IZA Discussion Papers, No. 16089, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/272716>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 16089

**Non-college Occupations, Workplace
Routinization, and the Gender Gap in
College Enrollment**

Amanda Chuan
Weilong Zhang

APRIL 2023

DISCUSSION PAPER SERIES

IZA DP No. 16089

Non-college Occupations, Workplace Routinization, and the Gender Gap in College Enrollment

Amanda Chuan

Michigan State University and IZA

Weilong Zhang

University of Cambridge

APRIL 2023

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Non-college Occupations, Workplace Routinization, and the Gender Gap in College Enrollment*

This paper explores how non-college occupations contributed to the gender gap in college enrollment, where women overtook men in college-going. Using instrumental variation from routinization, we show that the decline of routine-intensive occupations displaced the non-college occupations of women, raising female enrollment. Embedding this instrumental variation into a dynamic Roy model, we find that routinization decreased returns to the non-college occupations of women, increasing their college premium. In contrast, men's non-college occupations were less susceptible to routinization. Our model estimates that workplace routinization accounted for 44% of the growth in female enrollment during 1980-2000.

JEL Classification: I23, I24, I26, J16, J24, J23

Keywords: human capital, college enrollment, gender, occupations, automation

Corresponding author:

Amanda Chuan
Michigan State University
220 Trowbridge Rd
East Lansing, MI 48824
USA

E-mail: achuan@msu.edu

* We are grateful for excellent feedback from Eduardo Azevedo, Iwan Barankay, Jeff Biddle, Sandra Black, Kerwin Charles, Matthew Davis, Emma Dean, Gilles Duranton, Ryan Fackler, Clayton Featherstone, Fernando Ferreira, Eric French, Marina Gertsberg, Pavitra Govindan, Maia Guell, Joseph Gyourko, Steven Haider, Daniel Hamermesh, Robert Jensen, Judd Kessler, Ben Keys, Ilyana Kuziemko, Kai Liu, Corinne Low, Ioana Marinescu, Katherine Milkman, Olivia Mitchell, Jose Montalban, Karthik Nagarajan, Muriel Niederle, Eva Ranehill, Christopher Rauh, Alex Rees-Jones, Paul Sangrey, Katja Seim, David Schindler, Kailing Shen, Todd Sinai, Jeremy Tobacman, Petra Todd, Maisy Wong, and seminar participants at Cambridge University, Tufts University, Michigan State University, the Berlin Applied Micro Seminar Series, the American Economic Association Annual Meetings, IZA Junior/Senior Symposium, H2D2 Research Workshop, and WSAWBA.

“Out of high school, men are more willing than women to enter a trade. For example, there are jobs open to become electricians, carpenters, plumbers and more...Many of my male peers entered a career right out of high school and they are successful and happy.”

-Laura Thomas, Quinnipiac University, “Why the Future at U.S. Colleges is Female” (2021)

1 Introduction

Women used to lag behind men in college enrollment. As their work outcomes improved, it was predicted that the college gender gap would eventually close, and that men and women would attend college at equivalent rates thereafter. Women indeed closed the gap in 1970-1980, as shown in Figure 1. Contrary to expectations, the gap then reversed: women began attending college at increasingly higher rates relative to men. Female enrollment now exceeds male enrollment by over 10%, and the gap shows no signs of closing

It remains an open puzzle as to why women exceed men in college enrollment, when male college graduates work longer and earn higher median salaries than female college graduates (Becker et al., 2010). To reconcile this apparent contradiction, prior work has posited a greater share of women prepared for college than men. It argues that men face greater obstacles to formal human capital investment because more of them struggle to pay attention, stay disciplined, and persevere through school (Becker et al., 2010; Bertrand and Pan, 2013; Goldin et al., 2006). Other work posits that women experience a marriage market premium to attending college (see Chiappori et al., 2009; Chiappori et al., 2017; Ge, 2011; and Zhang, 2021). A separate literature examines gender differences in labor market returns to a college degree, but has mostly focused on the wages of college graduates (Charles and Luoh, 2003; Dougherty, 2005; Huang, 2014; Jacob, 2002; Olivieri, 2014; Pitt et al., 2012; Rendall, 2017).

Our paper distinguishes itself from existing research by concentrating on the non-college labor market. We argue that men have a comparative advantage in non-college work and center our analysis around non-college occupations, defined as occupations in which the majority of employed workers hold solely a high school diploma. If non-college occupations provide more gainful employment opportunities to men than women, men’s lower college enrollment rates may reflect their greater array of outside options, leading to a persistent gender gap in college enrollment as long as the non-college labor market favors men.¹

¹Our goal is not to explain the entirety of the college gender gap. To narrow down on how the *choice* to attend college differs between men and women, we examine the enrollment decision among high school graduates. We purposefully abstract away from marriage market considerations, parenthood, and gender differences in preparation prior to college entry, while acknowledging that these existing explanations operate

Examining the role of non-college occupations changes our current understanding of the college gender gap in important ways. If men lag behind women primarily because they struggle more in school, the appropriate policy response would be to remedy this “underinvestment” in education. Current suggestions include improving instruction for boys, such as through gender-segregated classrooms and teaching books that appeal to boys’ interests (Marcus, 2017; “Men Adrift”, 2015; Rosin, 2010; Sommers, 2013). However, if fewer men than women *choose* to enroll in college because they have more lucrative outside options, men may not be underinvesting in education at all. Instead, their lower enrollment could be a rational response to their more appealing non-college opportunities. Consequently, unequal college enrollment may be an efficient reaction to disparities in non-college labor opportunities.

This paper starts by observing that the non-college labor market is severely polarized by gender, in that the majority of occupations are either male- or female-dominated. From this observation emerge two stylized facts: first, non-college occupations dominated by women tend to pay less than those dominated by men; and second, many female-dominated occupations disappeared from the non-college labor market between 1970 and 2000. Together, these facts suggest that outside options to college-going were worse for women, and deteriorated even further over time. We posit that the widening disparity in non-college job prospects contributed to the widening of the reverse college gender gap.

To assess this hypothesis, we leverage *routinization* – the displacement of routine tasks by automation – as a shifter of non-college job prospects. A burgeoning literature on routine-biased technical change has established that over time, automated devices such as answering machines and computers increasingly substituted for human labor in performing routine tasks, eroding demand for workers in routine-intensive occupations (Acemoglu and Autor, 2011; Autor and Dorn, 2013; Autor et al., 2003; Cortes et al., 2014; Cortes et al., 2017; Goos et al., 2009, 2014; Jaimovich and Siu, 2012; Spitz-Oener, 2006). A few papers have noted that routinization has had especially severe impacts for the job prospects of women (Autor and Wasserman, 2013; Beaudry and Lewis, 2014; Black and Spitz-Oener, 2010; Dillender and Forsythe, 2019). Our analysis of census data reveals that *non-college* women were the most vulnerable to displacement, with over 70% of non-college young female workers employed in “routinizable” occupations in 1970. From 1970 to 2000, routinization lowered aggregate labor share among 18-30 year olds *only* for non-college women, but not college men, non-college men, and college women.

Following Autor and Dorn (2013), we measure local susceptibility to routinization using *routine task intensive (RTI) share*, the share of occupations that involve a high number

alongside the particular channel we investigate.

of routine tasks relative to other tasks. We use instrumental variation in routinization to address two challenges to causal inference. One is that RTI share in a local labor market could depend on the share of college and non-college workers, reflecting reverse causation. Another is that both RTI share and college enrollment rates could be correlated with unobserved factors, such as social norms regarding women’s education, the ease of graduating high school, or opportunities to finance a college education.

Our instrument predicts a local labor market’s routinizability using job posting data on administrative activity from Atalay et al. (2020). The intuition is that labor markets with high shares of administrative industries would experience more displacement in routine-intensive work. Time-series variation stems from within-occupation changes at the national level, which should not depend on changes in any particular commuting zone. Cross-sectional variation stems from 1950 industry composition, which pre-dates labor market and educational changes that occur during our analysis period of 1960-2000. Our identifying assumption is that within-occupation changes in administrative activity at the national level should only influence college enrollment in a commuting zone in ways reflected by routinization. We test these identifying assumptions in robustness checks, which verify that our results are not driven by changes in labor force participation, other changes to the share of non-college workers, or local shocks to markets from which the job postings originated. We also validate our results using alternative instruments, such as historic routine labor share and clerical requirements within occupations. All instruments exploit different sources of identifying variation but arrive at equivalent results.

Our first set of results comes from the two stage least squares (2SLS) regressions. The first stage regressions indicate that labor markets with higher shares of routine industries in 1950 experienced greater routinization in 1960-2000. The second stage results demonstrate that routinization increased college enrollment among young women. We find that a 1 percentage point rise in routinization would increase the proportion of 18-25 year old women who attend college by 0.58-0.61 percentage points. Equivalently, moving from a commuting zone that experienced the 25th percentile of routinization to one that experienced the 75th percentile of routinization (a difference of 5.51 percentage points) leads to a 3.20-3.36 percentage point rise in female college-going. For men, who experienced less displacement in their non-college job prospects, coefficient estimates are directionally smaller and not systematically significant. We thus use routinization to argue that the deteriorating availability of non-college jobs increased college enrollment. Since women’s non-college jobs were more vulnerable to routinization, female college enrollment grew especially quickly.

In order to further explore the mechanisms underlying our 2SLS results, we develop a Roy model that examines educational and occupational choices. We consider men and

women with varying levels of cognitive, manual, and administrative skills, as measured by the Armed Services Vocational Aptitude Battery (ASVAB) from the National Longitudinal Survey of Youth 1979 (NLSY79). These gender differences in skill endowments result in distinct comparative advantages for men and women, leading to gender polarization in non-college occupations. Due to this polarization, routinization had disparate effects on occupational returns for men and women.

Our model differs from the 2SLS approach in that it allows us to simulate the average treatment effect for all individuals in the sample. Because there is heterogeneity in treatment effects, our model predicts a larger gender gap in college enrollment due to routinization than estimated by a simple back-of-the-envelope calculation based on 2SLS estimates. Simulations show that the change in occupational returns led to a 4 percentage point increase in female enrollment, or 44% of the observed 9 percentage point growth in female enrollment from 1980 to 2000. In contrast, it can only account for a 1.6 percentage point increase in male enrollment. The change in female enrollment from declining non-college returns is over twice the change attributable to rising female labor force participation during this period. Our counterfactual analyses suggest that the decline in women’s non-college returns due to routinization is one major factor behind the growth in female enrollment relative to male enrollment.

Contributions to the literature. To our knowledge, this is the first paper that uses automation as a source of variation to investigate how the non-college labor market shaped the college gender gap over time. We use a new instrument to exploit the impact of automation on the demand for non-college workers in routine-intensive jobs. Prior work focusing on labor market returns has mostly relied on cross-sectional comparisons (Charles and Luoh, 2003; Dougherty, 2005; Jacob, 2002), occupational choice models (Olivieri, 2014), or general equilibrium models (Huang, 2014; Pitt et al., 2012; Rendall, 2017). Relative to these approaches, our paper better accounts for potential sources of endogeneity, such as supply-side factors which could influence both non-college occupation share and college enrollment (e.g., social norms regarding women’s work, ease of graduating high school, financial resources for pursuing college).

Second, we contribute to the literature on routine-biased technical change by quantifying automation’s impact on the rise of female college-going. To our knowledge, this is the first paper to estimate the causal impact of automation on the college gender gap. Most prior studies focus on the gender asymmetric impact of technological change on the labor market outcomes (Autor and Wasserman, 2013; Black and Spitz-Oener, 2010; Borghans et al., 2014; Cortes et al., 2021; Dillender and Forsythe, 2019; Juhn et al., 2014; Ngai and Petrongolo,

2017; Olivetti and Petrongolo, 2014, 2016; Yamaguchi, 2018). Our paper demonstrates substantial impacts on human capital acquisition, and therefore the skills that men versus women bring to the future workforce. Specifically, we show that routinization, a gender-neutral process, generates gender-asymmetric changes in college enrollment due to differences in skill endowments and skill prices. Our findings illuminate the role of technological change in shaping gender disparities in human capital. According to our simulation, changes in occupational returns from routinization can explain 4 of the 9 percentage point growth in female enrollment and 1.6 of the 2 percentage point growth in male enrollment from 1980 to 2000.

Third, our paper uses a model-based approach to link the gender polarization of the non-college labor market with the college gender gap. Since most prior papers use job task requirements to indirectly infer gender differences in skill levels, they cannot disentangle the skill endowments of individuals from the skill returns of jobs (Duran-Franch, 2020; Ngai and Petrongolo, 2017; Olivetti and Petrongolo, 2014; Rendall, 2017; Yamaguchi, 2018). We overcome this limitation by separately measuring skill endowments and task requirements, which is necessary to determine how routinization changed the value of different skills. The closest frameworks are Prada and Urzúa (2017) and Roys and Taber (2019), but our model deviates from them in two ways. We introduce instrumental variation from routinization to shift skill prices, following the spirit of Eisenhauer, Heckman, and Vytlačil (2015) and Heckman et al. (2018). Furthermore, we study both male and female workers and focus on gender inequality as it pertains to college enrollment choices, whereas the other two papers only analyze male workers.

The paper is organized as follows. Section 2 describes the data, stylized facts regarding non-college occupations, and descriptive evidence on routinization. Sections 3 and 4 describe our 2SLS methodology and results. Sections 5 and 6 describe our structural model methodology and results. We conclude in Section 7.

2 Data and Descriptive Evidence on Routinization

This section begins with a data overview. We then discuss the descriptive evidence that motivates our analytical approach. First, we present two stylized facts regarding gender polarization among non-college occupations. Second, we describe our measure of routinization, followed by descriptive evidence that links routinization with the widening gender gap in non-college job prospects.

2.1 Data

Our analysis starts with data from the U.S. decadal census for 1950-2000, which are collected by the U.S. Census Bureau and publicly provided by the Integrated Public Use Microdata Series (IPUMS; Ruggles et al., 2021). The census data for 1950, 1960, and 1970 include 1% of the entire U.S. population, while the census data for 1980, 1990, and 2000 include 5% of the population. The dependent variable is the share of 18-25 year olds who have ever enrolled in college. Our analysis focuses on employment patterns among youth, since older adults' education and labor decisions were influenced by factors that occurred prior to our analysis period of 1960-2000. Since our paper investigates the decision to attend college among qualified individuals, we limit our analysis to those with a high school diploma or GED. We focus on college enrollment rather than college completion since our goal is to understand how non-college job prospects impact the *choice* to pursue higher education.

To measure the impact of routinization, we use the labor share of routine-intensive occupations at the commuting zone level. Measures of routine-intensity and other occupational characteristics come from Autor and Dorn (2013). Our main instrumental variable is the share of occupations that involve high levels of administrative activity, as measured by job posting data constructed by Atalay et al. (2020).

We follow Autor and Dorn (2013) in defining a local labor market as a commuting zone, which captures commuting patterns for work across counties. Commuting zones are defined across the entire contiguous United States, in contrast to other geographic constructs that may under-represent certain industries (e.g., metropolitan statistical areas may underrepresent industries in rural areas such as agriculture or mining).²

Our structural model uses individual level data from the geocoded National Longitudinal Survey of Youth 1979 Cohort (NLSY79), which follows 12,686 respondents annually from 1979-1994 and every two years from 1996 until present day (Bureau of Labor Statistics, 2019). We construct a binary college attendance decision, based on years of education exceeding 12, and designate the individual's occupation choice as the modal occupation between ages 25 to 35. The occupation's monetary return is set to the individual's average annual earnings in that occupation. The final sample contains 8,540 individuals, with 4,217 men and 4,323 women. The NLSY79 provides insight into individual skill heterogeneity through ability measures obtained from the Armed Services Vocational Aptitude Battery (ASVAB). It contains detailed location data for each respondent in each year, allowing us to account for potential composition effects due to migration. We provide further details in Appendix A and summary statistics in Appendix Table A.2.

²Following Acemoglu and Autor (2011), we calculate labor supply weights by adjusting the sampling weight using the number of hours worked per week and the number of weeks worked per year.

2.2 Gender polarization among non-college occupations

Our empirical approach is motivated by two stylized facts from the census data. To describe them, we classify occupations according to gender and education. “Male-dominated” occupations are those with less than 30% women; “female-dominated” occupations comprise of more than 70% women; and “gender-equal” occupations comprise of 30-70% women. We define “non-college occupations” as those with at least 50% high school graduates, and “college occupations” as those consisting of at least 50% college enrollees.

The first stylized fact is that **female-dominated non-college occupations tend to earn lower pay than male-dominated occupations do**. Figure 2a shows a *missing quadrant* in the non-college labor market. Plenty of male-dominated occupations pay above the median income of all workers, including college graduates. For example, miners, machinists, and truck drivers are over 90% male and earn between the 40th to the 80th percentile of annual earnings. In contrast, female-dominated occupations, such as cashier, housekeeper, and cosmetologist, pay below the 20th percentile. While the evidence is merely correlational, it suggests that male high school graduates have the potential for high earnings, while their female counterparts are less likely to sort into high-paying occupations. Experimental evidence squares with this interpretation, as male high school students cite their lucrative non-college job prospects as a reason why interventions to increase college-going were less successful for them compared to female high school students (Carrell and Sacerdote, 2017).

College occupations display the opposite missing quadrant, as shown in Figure 2b. There is a dearth of low-paying occupations that are male-dominated, but plenty of low-paying female-dominated occupations. The evidence in Figure 2 is consistent with an underlying sorting mechanism for college enrollment, where few men enter low-paying college occupations given the availability of high-paying non-college occupations. In contrast, many women would accept low-paying college jobs if their non-college job prospects were not particularly appealing.³

The second stylized fact is that **many female-dominated occupations disappeared from the non-college labor market over time**. Figure 3a displays how non-college occupations vary by gender composition in 1970. Non-college occupations exhibited substantial gender polarization. 34% were female-dominated, 53% male-dominated, and only 13% gender-equal. By 2000, female-dominated occupations plummeted from 34% to 13%, male-dominated occupations rose even higher to 76%, and gender-equal occupations remained low at 12%. College occupations demonstrate the opposite trend, exhibiting gender convergence. Figure 3b shows that the share of gender-equal occupations tripled from 17% to

³In Section 5, we will show this sorting mechanism can arise naturally given different comparative advantages to non-college work among men and women.

50%, while the share of male-dominated occupations dropped from 72% to 21%. The share of female-dominated occupations doubled from only 12% to 29%. The descriptive evidence suggests that over time, men and women became more substitutable in college work, but non-college occupations remained polarized by gender.

Upon further investigation of the disappearance of female-dominated occupations, we find that those experiencing the greatest decline were intensive in *routine* tasks. We split Figure 3a into routine and nonroutine occupations. The resulting histograms, displayed in Appendix Figure A.1, reveal that the decline in female-dominated occupations occurred *exclusively* in routine occupations. In fact, among routine occupations, those that were female-dominated virtually all disappeared from the non-college labor market.

What happened to these occupations? Some became obsolete, while others transitioned to the college labor market. For example, typists became virtually eliminated, experiencing a 95% decline in labor share from 1970 to 2000. In contrast, secretarial and administrative assistant roles experienced an *across occupation* decline of 66% in 1970-2000, as well as a *within occupation* rise in the value of a college degree. They adopted more cognitive tasks over time, such as writing original documents and using spreadsheets (Atalay et al., 2018). A college degree became increasingly valuable for applicants who had to compete for a shrinking number of vacancies in these positions. It is important to note that our descriptive evidence deliberately conflates these across-occupation and within-occupation changes to motivate that women’s jobs have changed more than men’s along this dimension. It does not isolate causal effects, but illustrates the variation used in our instrumental variable approach.

2.3 Routinization and occupational composition

2.3.1 Measuring routinization

We use routinization to extract plausibly exogenous changes in non-college job opportunities. Examining all forms of routinization would make it difficult to assess the validity of our instrumental approach. Since routinization is challenging to measure, using broad measures of routinization could impact college enrollment through channels other than labor market returns (Autor, 2015).⁴ Instead, we chose to focus on a particular type of routine-biased technical change: the *routinization of the office* during 1960-2000 (Autor et al., 2003; Black

⁴Beyond the scope of this paper are other forms of automation which could impact college-going through different channels. For example, beginning in the 1990s, industrial roboticization substituted for manually intensive work and displaced the job prospects of men (Acemoglu et al., 2020; Acemoglu and Restrepo, 2019, 2020). Prior to the 1960s, improvements in household production technologies contributed to the mass entry of women into the labor force (Greenwood et al., 2005). Other early automation forms include machinery that substituted for manual-intensive labor in agriculture and manufacturing prior to the 1950s (Adams, 2019; Atack et al., 2019; Autor, 2015).

and Spitz-Oener, 2010). Examples include the electric typewriter, fax machine, word processor, and personal computer, which substituted for human labor in executing routine tasks (Atlassian, 2022).

We define “routinizable” occupations as those where the majority of tasks can be easily executed using automated devices. To more precisely capture routinizability, we employ the “routine task intensity” (RTI) measure from Autor and Dorn (2013). The RTI of occupation k is calculated using the logged index of its routine, manual, and abstract tasks:

$$RTI_k = \ln(routine_{k,1980}) - \ln(manual_{k,1980}) - \ln(abstract_{k,1980})$$

The RTI measure captures an occupation’s routine content net of its manual and abstract content. “Routine,” “manual,” and “abstract” task content are compiled from census data and the Dictionary of Occupational Titles.⁵ “Routine” tasks are codifiable tasks that can be executed following an explicit set of rules. As technology progressed, automating devices replaced human labor in executing routine tasks, decreasing employer demand for workers who specialize in these tasks. For example, electric typewriters and carbon paper obviated the need for clerical workers to fill out forms one by one using pen and paper (Decker, 2016).

“Manual” tasks are defined as tasks requiring in-person execution, which tend to be physical or service-oriented. These tasks were difficult to automate since they often involved the handling of objects across space, such as lifting materials or moving in a constantly changing environment (Autor, 2015; Autor and Dorn, 2013). Consequently, occupations that involved more manual tasks, such as operative and production occupations, were shielded from automation compared to clerical occupations. Lastly, “abstract” tasks involve complex mental processes that are not easily programmable, such as problem-solving, management, and complex communication.

Prior work found that automation directly substituted for routine tasks while complementing abstract and manual tasks.⁶ Greater routine content made an occupation more susceptible to routinization, while more manual or abstract tasks made an occupation less susceptible to routinization. Occupations high in RTI are therefore vulnerable to routinization. In fact, the decline in female-dominated non-college occupations shown in Figure A.1 was driven by occupations in the top third of RTI. Therefore, to measure the impact of routinization, we focus on “RTI share”, or the share of high RTI occupations:

⁵We fix them to 1980 levels, which nets out within-occupation changes over time so that any change in RTI across labor markets will stem only from changes in occupational composition.

⁶Brynjolfsson and Hitt (2000) and Bresnahan et al. (2002) demonstrate that computers and routine tasks functioned as substitutes in production. On the other hand, by increasing the marginal productivity of abstract tasks, computers and similar automating devices raised labor demand for workers with abstract skills (Autor et al., 2003; Bresnahan et al., 2002; Brynjolfsson and Hitt, 2000; Spitz-Oener, 2006).

$$\text{RTI share}_{ct} = \frac{\sum_{k=1}^K \mathbf{1}(\text{RTI}_k > \text{RTI}_{1980}^{P66}) L_{ckt}}{\sum_{k=1}^L L_{ckt}}$$

where L_{ckt} is the total number of workers 16-64 years of age in commuting zone c , occupation k , and year t . Occupation k is designated high-RTI if falls in the top third of routine task intensity for all occupations in 1980: $\text{RTI}_k > \text{RTI}_{1980}^{P66}$.⁷

We then define our routinization measure as the change in RTI share for commuting zone c in year t :

$$\text{routinization}_{ct} = \text{RTI share}_{c,1950} - \text{RTI share}_{ct}$$

where t ranges from 1960 to 2000.

2.3.2 Linking routinization, job polarization, and college enrollment

Related work has established that the routine content of jobs declined over time because automation substituted for human labor in executing routine tasks (see Autor and Dorn, 2013; Goos et al., 2009). We find that among youth, these changes predominantly affected women, as a large share of the female workforce sorted into routinizable occupations. Table 1 shows that as a share of the 18-25 year old female workforce, workers in office and administrative support occupations declined by 15 percentage points, from 42% in 1970 to 27% by 2000. In comparison, the largest occupational decline for comparable men was about half that amount. This occurred in the production occupations, where declines were comparable between men and women.⁸

Columns (4) and (8) compare the decline in routine task intensity (RTI) for women and men, measured in standard deviations. They show that this decline was largest among office and administrative support occupations, while other non-college occupations (including production) did not experience a decline in RTI. Furthermore, this decline was almost three times larger for women than men among office and administrative support occupations.⁹

⁷The impact of automation is better captured by the share of high RTI occupations than other measures, such as average RTI level, which would not capture the full extent of each local labor market that is vulnerable to automation. We set the threshold of RTI to be the 66th percentile of the 1980 occupational distribution following Autor and Dorn (2013). In robustness checks, we designate occupations as high-RTI if they are in the top half of routine task intensity, rather than the top third. This alternative definition does not appreciably change our results (results available upon request).

⁸The evidence in Table 1 is consistent with Autor et al. (2015), who find a significant decline in the routine jobs of women but not men, and gender-symmetric impacts of automation on production occupations.

⁹We decompose the decline in routine task intensity (RTI) into its component parts of routine content, manual content, and abstract content. We find that the decline in RTI among office and administrative support occupations is driven by the decline in routine tasks. In contrast, the change in RTI among production occupations is driven by the decline in manual tasks. We conclude that the decline in production occupations documented by prior work may be due to other forms of technical change, but not routinization

Overall, Table 1 illustrates that if non-college occupations helped shape the college gender gap, the effect would primarily occur within the office and administrative support occupations. First, among non-college occupations, these occupations alone experienced declines in RTI. Second, declines in other occupations, such as production, were not gender-biased. The descriptive evidence suggests that the tasks most vulnerable to routinization were administrative. Therefore, our main instrument will exploit use changes in administrative activity within all occupations over time.

How could these changes have impacted college-going by gender? Our main descriptive evidence compares the labor share of routinizable occupations (high RTI) with non-routinizable occupations (low RTI). Figure 4 shows a decline in the routinizable occupations of non-college women, but not college women, non-college men, or college men. Panel (a) shows stark differences in how labor share changed over time in routinizable versus non-routinizable jobs for non-college women. First, the share of routinizable occupations reached almost one-third (31.8%) of all non-college working women 18-30 years old in 1970. However, it then plummeted to less than half this level by 2000, from 31.8% to 14.1%. In contrast, the share of non-routinizable jobs was quite small and constant at 5.4-7.3% over the same period. College jobs did not exhibit this differential trend. For female college enrollees, the labor share of both routinizable and non-routinizable jobs followed parallel trajectories, increasing by 10-12 percentage points from 1970 to 2000.

These patterns are not observed for men. Panel (b) shows that among young men, the labor shares of routinizable and non-routinizable jobs follow parallel trajectories. The exception is that non-routinizable jobs among college men grew by 8 percentage points from 1970 to 2000, which closely aligns with the growth rate experienced by college women in non-routinizable jobs. This growth among college workers is to be expected if automation complemented abstract skills and increased employer demand for college-educated workers.

This descriptive evidence aligns with Black and Spitz-Oener (2010), who report a “strong decline in routine tasks experienced by women and almost not at all by men” for Western Germany (pg. 188). Similarly, Autor et al. (2015) find that routinization impacts female employment, with no significant effects for male employment. We extend on this work by further isolating the decline to non-college women, which suggests women’s outside options to college-going were disproportionately vulnerable to routinization. The gender asymmetric impacts are expected, given differential sorting into non-college occupations. Men’s “blue-collar” occupations were highly manual, making them difficult to displace even if some

specifically (Autor and Wasserman, 2013; Autor et al., 2015; Cerina et al., 2021). Our results are consistent with Dillender and Forsythe (2019), who find that routinization was concentrated among office occupations in the 2010s.

involved routine tasks. Women’s “pink-collar” occupations were less manual, as they involved operations such as bookkeeping and calculating, which were easy to routinize (Autor et al., 2003; Black and Spitz-Oener, 2010).

3 Two Stage Least Squares Approach

The descriptive evidence we have shown so far does not establish the causal effect of routinization on college enrollment. Similarly, ordinary least squares (OLS) regressions may not be sufficient to isolate causal impacts. Table 2 presents positive OLS estimates between routinization and college enrollment, which suggest that labor markets which underwent greater routinization experienced growth in female enrollment ($p < 0.01$) and marginal growth in male enrollment ($p < 0.10$). Alternative explanations could be driving these estimates. For example, if more students enrolled in college for other reasons, the college workforce would rise relative to the non-college workforce, mechanically decreasing the non-college labor share.

We therefore use an instrument to predict *displacement due to routinization* using newspaper job posting data. Atalay et al. (2020) extract information the skills and activities involved in an occupation based on job postings in *the Boston Globe*, *the New York Times*, and *the Wall Street Journal* during 1940-2000. They use textual machine learning approaches to map each job title in a posting to a code in the Census 2000 Occupation Index (Census Bureau, 2021). We then use their measures of job characteristics at the occupation code level to construct our instrument.

Specifically, we use their measure of administrative activity, measured by the frequency with which postings for an occupation mention the following keywords: “administrative,” “paperwork,” “filing,” and “typing”. These routine office tasks were the most vulnerable to automation, whereas operative and production occupations were shielded from routinization because they involved manual tasks that were more difficult to automate (Autor and Dorn, 2013; Dillender and Forsythe, 2019; Spitz-Oener, 2006). Our measure focuses on the subset of tasks that are most likely to be routinized, but may not fully capture the range of routine tasks affected. We prefer this conservative measure over measures that incorporate a larger range of tasks, which would risk violating the exclusion restriction by attributing too many task changes to automation.¹⁰

To predict routinization at the commuting zone level, we fix industry shares in 1950 and interact them with “administrative share”, the share of occupations with high administrative

¹⁰The risk of using too conservative a measure is that the relevance condition may not be satisfied. However, the first stage F-statistics exceed 200 (Table 3), indicating a strong relationship between routinization and our instrument.

activity in each industry. The intuition is that commuting zones with high 1950 shares of administrative industries should experience greater routinization as these industries automate over time. Therefore, the administrative share IV is:

$$\text{admin share IV}_{ct} = \sum_{i=1}^I E_{i,c,1950} \frac{\sum_k L_{ikt} \mathbf{1}(\text{admin}_{kt} > \text{admin}_{1950}^{P66})}{\sum_k L_{ikt}}$$

where i indexes industry, k indexes occupation, t indexes year from 1960 to 2000, and c indexes commuting zone. $E_{i,c,1950}$ represents the share of industry i in commuting zone c in 1950. The expression $\frac{\sum_k L_{ikt} \mathbf{1}(\text{admin}_{kt} > \text{admin}_{1950}^{P66})}{\sum_k L_{ikt}}$ is the administrative share in industry i in year t . It is constructed using L_{ikt} , which represents the number of workers in occupation k , industry i , year t . The indicator $\mathbf{1}(\text{admin}_{kt} > \text{admin}_{1950}^{P66})$ equals 1 if occupation k in year t is in the top third of administrative activity, based on the occupation distribution in 1950.¹¹ While $E_{i,c,1950}$ and L_{ikt} are constructed from the census, the indicator $\mathbf{1}(\text{admin}_{kt} > \text{admin}_{1950}^{P66})$ is constructed using job posting data from Atalay et al. (2020).

One concern regarding this conservative measure is that it captures changes in female-dominated occupations more than male-dominated occupations. However, the occupational sorting shown in Figure A.1 indicates that any gender-neutral shock that is routine-biased will inevitably have gender-asymmetric impacts. For young workers, routinization hit administrative occupations hardest, as shown in Table 1. Since women were more likely to hold these occupations, we argue that a stronger first stage relationship should exist for women than men. Indeed, when we use a more broad-based, less conservative instrumental variable – predicted routine share among all occupations – we find very similar results (see Table 5). This is because any attempt to find a “gender-neutral” shock will lead to a gender-asymmetric response, given that women sorted into the jobs that were disproportionately displaced.

We perform the following two stage least squares regression. The first stage regression captures the relationship between routinization and the instrument within commuting zone c and year t :

$$\text{routinization}_{ct} = \alpha_0 + \alpha_1 \text{admin share IV}_{ct} + \alpha_2 W_{ct} + \theta_c + \phi_t + u_{ct} \quad (1)$$

In our regression approach, our measure of routinization focuses on the routinization among *non-college workers between 25 to 65 years old*. Focusing on routinization among

¹¹Following the logic of Autor and Dorn (2013), we define “highly administrative occupation” based on whether the occupation is in the top third of the 1950 distribution. Fixing the occupational distribution to 1950 allows us to compare how administrative share changes over time for industries that traditionally involved intensive administrative activity.

this group enables us to directly measure how routinization changed the outside options to college-going.¹² Focusing on jobs held by 25-65 year olds fits our underlying premise that 18-25 year olds make their college-going decisions based on the job prospects of *those currently working*. Furthermore, excluding 18-25 year olds also avoids simultaneity concerns: if enrollment among 18-25 year olds rose for other reasons during this time, fewer workers would take routine task intensive jobs, and RTI share would mechanically decline.

We control for commuting zone-year level controls W_{ct} , commuting zone, census region, and year. The matrix of control variables W_{ct} includes the proportion of female, Black, and Hispanic residents. It also includes the proportion of people by 10-year age bin. Additional controls are discussed later in this section.

The intuition behind our first stage regression is that commuting zones starting out with high levels of administrative work should have undergone greater routinization. For example, the commuting zone around Republic city in the state of Washington had high 1950 shares of the legal services industry, which used to comprise of many administrative jobs that involved completing and filing forms. As the legal industry automated, the extent of routinization would be especially severe in Republic city compared to other labor markets. This would lead to a positive first stage coefficient between predicted administrative share, our instrument, and routinization.

The second stage regression then uses the first stage linear prediction of routinization to isolate the impact on college enrollment in commuting zone c , year t for gender g :

$$\text{college enrollment}_{ct}^g = \beta_0 + \beta_1 \text{routinization}_{ct} + \beta_2 W_{ct} + \eta_c + \psi_t + \epsilon_{ct}^g \quad (2)$$

As with the first stage regression in Equation 1, the second stage regression controls for commuting zone-year characteristics W_{ct} , commuting zone dummies, and year dummies.¹³

Under the frameworks of Adao et al. (2019) and Borusyak et al. (2018), the shift-share approach is equivalent to a weighted instrumental variable regression in which industry-level shocks are the instrument and industry shares are the weights. The exclusion restriction is therefore that the administrative share at the national industry level can only affect college enrollment in ways reflected by routinization at the commuting zone level. This restriction is met if no commuting zone plays a large role in determining administrative share in an

¹²However, this measure is endogenous to supply-side considerations that influence educational choices, such as social norms regarding education or the ease of graduating from high school. In Section 4.1, we apply our 2SLS specification to routinization among both college and non-college workers and find no change in our results.

¹³Note that while the reduced form and second stage effects on enrollment are gender-specific, we pool gender in estimating the first stage effect. This avoids the assumption that men and women operate in isolated markets and allows for correlation between impacts for men and women.

industry. Since our job posting data come from newspapers located in New York City and Boston, in robustness checks we exclude the commuting zones containing these cities to determine whether our 2SLS results are driven by local omitted variables correlated with both college enrollment and administrative work.

The general threat to the exclusion restriction is that industry-level changes in routine activity, measured by administrative occupation share, could be correlated with enrollment in ways not captured by commuting zone-level changes in RTI share. Using commuting zone dummies accounts for time-invariant omitted factors, but not changes across time correlated with both enrollment and labor market prospects. We next discuss plausible time-varying confounders that could generate the gender differences in college-going we report in Section 4. These confounders motivate the inclusion of certain controls into the W_{ct} matrix.¹⁴

One possibility is that non-automation factors could drive industry level changes correlated with both enrollment and routinization in a commuting zone. For instance, the decline in manufacturing over this period could change both college enrollment and the proportion of high-RTI occupations within an industry (see Autor et al., 2013). We therefore include in W_{ct} lagged shares of the largest industries: manufacturing, mining, and retail trade.¹⁵ We also control for lagged service sector shares, given the Autor and Dorn (2013) finding that automation raised service sector employment. Using the lagged shares is preferable to current shares, since current shares may directly depend on college enrollment rates.

Supply-side factors could influence enrollment in ways correlated with the instrument. For example, high female labor force participation in a commuting zone may raise the share of industries that employ female high school graduates in 1950. More non-college jobs may be available to women in this commuting zone than in others, which would then increase their outside options to college-going, leading to lower growth in female enrollment in 1960-2000. We therefore control for both female and male labor force participation among 25-65 year olds. Since 25-65 year olds are beyond typical college age, their labor force participation should not directly depend on the college enrollment of 18-25 year olds.

Related concerns are serial correlation in RTI share, as well as persistence in other unobservable factors that could influence women’s labor market prospects. For instance, commuting zones with more routine jobs in 1950 may have more favorable social norms regarding women’s schooling in 1960-2000. In some specifications, we control for lagged RTI share

¹⁴In Section 4.1, we check that our results hold even when we do not use the control variables discussed below (Table 5).

¹⁵A trade-off exists between controlling for some industries versus all industries. Our identification relies on industry-level shocks, so controlling for all industries would lead the industry dummies to absorb valuable identifying variation. We therefore only control for major industries that compose a large share of the overall labor force.

to capture the effects of these and related social norms. Finally, as mentioned above, routinization changed both the returns to non-college work and college work. To separate the pull factor of rising college earnings from the push factor of declining non-college job opportunities, we control for median earnings in abstract-intensive occupations in certain specifications.

We use the standard error correction procedure of Adao et al. (AKM, 2019). AKM (2019) demonstrate that shift-share instruments introduce correlation across labor markets with similar industry shares, and that clustering standard errors at the local labor market level is insufficient to account for such correlation. To report the results of our weak instrument tests, we calculate Montiel Olea-Pflueger F-statistics, which are preferable to Kleibergen-Paap F-statistics in assessing instrument strength (Andrews et al., 2018; Andrews and Stock, 2018; Olea and Pflueger, 2013). Given recent literature on the limitations of using t-ratio based inference and first stage F-statistics to assess instrument strength (Lee et al., 2020), we report Anderson-Rubin weak instrument-robust confidence intervals.

4 Two Stage Least Squares Results

We begin by investigating the first stage relationship between the instruments and routinization, presented in Table 3. As discussed in Section 3, we use various sets of controls to account for potential confounds. Column (1) controls for demographic characteristics at the commuting zone level, male and female labor force participation, the ten-year lagged service sector share, and the ten-year lagged shares of the industries with the highest labor shares in our data: manufacturing, retail, and mining. Adding to these controls, columns (2) and (4) include the median annual log earnings of occupations in the top third of abstract intensity. Columns (3) and (4) include the ten-year lag of RTI share.

Throughout this paper, we measure routinization as the percentage point reduction in RTI share from 1950 levels. We find that on average, a commuting zone with a 1 percentage point higher share of administrative industries in 1950 experienced 0.38-0.39 percentage points more routinization in 1960-2000 ($p < 0.01$). Coefficient estimates remain constant even when we control for median earnings in abstract-intensive work in columns (2) and (4), suggesting that the decline in RTI share is driven by declining routine task demand rather than growing returns to abstract-intensive work. Similarly, our estimates do not change when we control for lagged RTI share in columns (3) and (4), indicating that serial correlation in unobservables is unlikely to explain these relationships. Across all specifications, Montiel Olea-Pflueger F-statistics hover at 201.45-214.57. To visually assess fit, Appendix Figure A.2a plots the raw data against the linear prediction. The raw data exhibit a clear

positive relationship between routinization and the administrative share IV, indicating that commuting zones with higher historical administrative industry shares experienced greater routinization. Appendix Figure A.2b then shows that this positive relationship arises from automation. Commuting zones with greater 1950 shares of administrative industries, as predicted by our instrument, adopt more office personal computers per worker in 1980-1990.

Next, Table 4 reports the reduced form results for female enrollment (panel a) and male enrollment (panel b). Across all regressions, we find greater female enrollment rates among commuting zones with higher instrument values. This finding is consistent with the premise that women’s non-college job opportunities diminished in labor markets more vulnerable to routinization. Commuting zones predicted to undergo 1 percentage point more routinization exhibit on average a 0.22-0.23 percentage point rise in female enrollment ($p < 0.01$). The coefficient for men is about 75% of the estimate for women and marginally significant at 0.17 percentage points ($p < 0.10$).

We next turn to the two stage least squares results in panels c-d of Table 4. By isolating variation in routinization based on changes in administrative activity over time, we aim to capture declines in employer demand for routine-intensive occupations. This then translates into fewer job options for high school graduates, since most routine-intensive occupations provided opportunities for non-college workers. Consistent with this story, panel c demonstrates that commuting zones that underwent more routinization experienced higher female enrollment rates. Our estimates indicate that a 1 percentage point rise in routinization led to a 0.58-0.61 percentage point rise in the proportion of 18-25 year old women enrolled in college ($p < 0.01$). Panel d shows that the corresponding estimate for male enrollment is 0.44 percentage points ($p < 0.10$). Equivalently, moving from a commuting zone which experienced the 25th percentile of routinization to one in the 75th percentile of routinization would have increased female enrollment by 3.18-3.33 percentage points and male enrollment by 2.40-2.44 percentage points.

Our Anderson-Rubin weak instrument-robust 95% confidence intervals exclude 0 for female enrollment but cannot reject the null hypothesis of no effect for male enrollment. The results establish a consistently significant negative relationship for women, but not for men. However, the coefficient estimates on male and female enrollment do not statistically differ. It is possible that the erosion of routine jobs also impacted male college-going, since some men worked in occupations vulnerable to automation.¹⁶ Yet, even if male and female enrollment were equally elastic to routinization, far more women worked in routinizable jobs

¹⁶In addition to the few men who worked in secretarial and clerical occupations, routinizable occupations that were dominated by men include shipping clerks, meter readers, security guards, machinists, and machinery repairers.

than men – around 70% of non-college women compared to 40% of non-college men during 1960-2000. The aggregate change in non-college job prospects for women would still exceed that for men. We explore the implications of these estimates on aggregate trends in the college gender gap over time in Sections 5 and 6.

4.1 Additional specifications

We next address potential concerns regarding our main regression specification from Table 4 panels c and d. Table 5 summarizes results from our additional specifications. Column (1) excludes controls which may be correlated with omitted variables that influence routinization: contemporaneous labor force participation, lagged service sector share, lagged major industry shares, and lagged routine share. Column (2) excludes the commuting zones that contain Boston and New York City. This checks whether our job posting variables are influenced by local supply shocks, which could impact both education and employment through channels other than routinization.

Column (3) controls for abstract occupation share, which our main specifications allow to vary freely with routine content. This additional control places severe restrictions on the variation we use, but better nets out the impact of non-automation forces that shift routine and abstract content simultaneously. Column (4) examines the RTI share among both non-college and college workers, rather than only non-college workers. This accounts for the possibility that endogenous supply-side forces, such as greater high school completion rates, changed the composition of the non-college workforce in ways that also correlate with routine labor share.

Across all specifications, we find remarkably similar results. The displacement of routine occupations corresponds to significant growth in female enrollment, while the impact on male enrollment is insignificant or marginally significant. Anderson-Rubin confidence intervals reject the null hypothesis of no effect for women, but not for men.

Alternate Instruments Lastly, we use four alternate instruments. Our first instrument predicts routine share across all occupations. It is constructed by interacting the share of occupations in the top third of routine intensity in 1950 with the 1950 industry composition of a commuting zone. This instrument is a modified version of the instrument from Autor and Dorn (2013). Rather than focusing on administrative activity, the routine share instrument examines how routine work has changed in the broader labor market. Its advantage over the administrative activity instrument is that it encompasses the full range of routine tasks conducted by both men and women. This also is its disadvantage: because this instrument incorporates a wider range of activities, it is more likely to violate the exclusion restriction

since it could incorporate changes that do not arise from automation, such as social norms influencing changes in education and routine labor within a commuting zone.

Our second instrument re-defines administrative share based on the share of occupations in the top half of administrative activity, rather than the top third. Our third instrument takes the average level of administrative activity across all occupations in an industry as the predictor of routinization. Our fourth instrument uses job posting data on “clerical requirements”, based on the number of times a clerical knowledge requirement is specified per job posting for an occupation. Despite using different sources of variation, all instruments yield the same set of quantitative results as our main specification in Table 4 panels c and d. Online Appendix B discusses our instrument construction and justification in greater detail.

5 Structural Model Approach

Our 2SLS results show that commuting zones with greater predicted routinization experienced higher female college enrollment. We next propose a mechanism to explain these findings. An augmented Roy model with latent skills describes how individual choices respond to non-college job prospects. We explicitly model sequential education and occupation decisions based on the dynamic discrete choice literature (Eisenhauer, Heckman, and Mosso, 2015; Keane and Wolpin, 1997; Roys and Taber, 2019; Todd and Zhang, 2020). Changes in routinization shift occupational-specific skill prices, which alter occupational returns and impact the college premium.

Identifying the effect of routinization on skill prices is challenging because both are endogenous outcomes determined jointly by supply and demand. Our innovation involves incorporating instrumental variation in routinization into occupation-specific skill prices. By doing so, we leverage plausibly exogenous shifts in skill prices to identify the causal effects of routinization on well-defined margins of the education and occupation decisions.¹⁷

The 2SLS estimate is commonly interpreted as a convex combination of conditional local average treatment effects (LATEs, Angrist and Imbens, 1995.) However, recent studies argue that conditioning the 2SLS estimate on covariates yields negative weights on certain subgroups (Blandhol et al., 2022; Słoczyński, 2020). Our model tackles this concern by simulating the treatment effect for every individual in the entire sample. The average treatment effect derived from our simulation is a simple average for the entire sample with equal weights, which offers a more transparent method to link the 2SLS estimates to the average treatment effect of the treated sample (ATT).

¹⁷Eisenhauer, Heckman, and Vytlačil (2015) and Heckman et al. (2018) have also incorporated instruments into discrete choice models. However, the decision rules in their models are not fully dynamic.

Since there is heterogeneity in treatment effects, the average treatment effect on treated individuals could significantly differ from the linear extrapolation derived from the 2SLS approach. We employ our ATT estimates to simulate the impact of routinization on aggregate college enrollment trends between 1980 and 2000. Our simulations enable us to compare the importance of non-college returns with another contributing factor to the college gender gap: the gender difference in labor force participation.

5.1 Basic environment

The model has two periods with transitions and nodes shown in Figure 5. Individuals are forward looking and sequentially choose their education D_i in period 1 and their occupation O_i in period 2. In the first period, individuals choose whether to attend college based on the flow utility of schooling and expected values from the second period. Initial skill endowments are unobserved by the econometrician but fully observed by each individual. Following Heckman et al. (2006) and Prada and Urzúa (2017), we identify workers' unobserved skills by constructing a measurement system based on individuals' test scores from the NLSY79. We use $\theta_i = [\theta_{ci}, \theta_{mi}, \theta_{ai}]$ to represent a vector of skill sets for individual i , where subscripts c , m , and a are used to denote cognitive, mechanical, and administrative skills, respectively. We allow for gender differences in skill distributions.

We demarcate three different occupation choices $O_i \in \{\text{White collar, Blue collar, Pink collar}\}$. White collar occupations ($O_i = 1$) refer to occupations dominated by college workers; blue collar occupations ($O_i = 2$) refer to occupations dominated by male high school graduates; and pink collar occupations ($O_i = 3$) refer to occupations dominated by female high school graduates. This classification is derived from Figure 3, which shows gender polarization in non-college occupations and gender convergence among college occupations. Lastly, we allow for home-staying as an outside option to working ($O_i = 4$), which endogenizes the labor force participation decision.

Our specification is intentionally parsimonious compared to prior life-cycle dynamic discrete choice models (Keane and Wolpin, 1997, 2001; Roys and Taber, 2019; Todd and Zhang, 2020). It assumes that all occupation and education choices are made once and permanent. This simplicity enables us to focus on how routinization changes the price of different skills, creating heterogeneous impacts on occupational returns.

5.2 Sequential schooling and occupation choices

The model is solved through backwards induction. In the second period, individual i with gender $g \in \{m, f\}$ chooses an occupation depending on perceived expected values across

alternatives. Ex post, individual i who chooses occupation O_i given an education level D_i receives utility $U(O_i|D_i)$:

$$U(O_i|D_i) = Y(O_i|D_i) + P(O_i|D_i) + \epsilon_{O,D,i} \quad (3)$$

where $Y(O_i|D_i)$ denotes the monetary return from occupation O_i given an education level D_i , while $P(O_i|D_i)$ is the non-pecuniary utility of working in occupation O_i (e.g., from job amenities or potential discrimination costs). The term $\epsilon_{O,D,i}$ is an idiosyncratic preference shock that follows the type I extreme value distribution. Note that we can only identify differences among options, as opposed to their levels. We therefore normalize the value of the home-staying option to be 0 for identification purposes.

Earnings in occupation O_i are expressed as

$$Y(O_i|D_i) = X_i^Y \beta_{O,X}^g + D_i \beta_{O,D}^g + \theta_i \beta_{O,\theta}^g(c, t) + \theta_i D_i \beta_{O,D,\theta}^g(c, t) + u_{O,D,i}^g \quad (4)$$

where X_i^Y is a vector of relevant observed variables, including cohort, region, and urban dummies. The subscript $g \in \{m, f\}$ denotes male and female, respectively. The college premium comes from both $D_i \beta_{O,D}^g$ and $\theta_i D_i \beta_{O,D,\theta}^g$, in which $\beta_{O,D}^g$ captures the common return to education while $\beta_{O,D,\theta}^g$ captures the component that varies by skill level θ_i . Lastly, $u_{O,D,i}^g$ is the random component, realized only after occupation O_i has been chosen. Analogously, the non-pecuniary utility $P(O_i|D_i)$ from entering occupation O_i has the following expression

$$P(O_i|D_i) = X_i^Y \alpha_{O,X}^g + D_i \alpha_{O,D}^g + \theta_i \alpha_{O,\theta}^g(c, t) + \theta_i D_i \alpha_{O,D,\theta}^g(c, t) \quad (5)$$

where $\alpha_{O,D}^g$ represents the non-pecuniary return to education shared by all workers and $\alpha_{O,D,\theta}^g$ captures the extra non-pecuniary education premium that varies by worker's skill level θ .

In the first period, individuals decide whether to attend college based on a cost function $C(D_i)$, which includes both monetary and psychic costs. Since we can only identify the difference in the cost of attending versus not attending college, we normalize the cost of not attending college to zero: $C(D_i = 0) = 0$. The cost of attending college is

$$C(D_i = 1) = X_i^D \lambda_X^g + \theta_i \lambda_\theta^g + \xi_{D,i}^g$$

where X_i^D captures a vector of characteristics commonly believed to be relevant factors for education choice.¹⁸ The term $\theta_i \lambda_\theta^g$ captures the heterogeneous cost of attendance for individual i with skill θ_i and gender g . The preference shock on education $\xi_{D,i}^g$ is assumed

¹⁸Following Eisenhauer, Heckman, and Mosso (2015) and Prada and Urzúa (2017), X_i^D includes parental education, the number of siblings, an indicator variable for broken home, and family income at age 14.

to be orthogonal to X_i^D and θ_i .

The continuation value of education choice D_i , given the available information at D_i , has the expression:

$$EV(D_i) = \beta E \left\{ \max_{O_i} (U(O_i|D_i)|\zeta(D)) \right\}$$

where $\zeta(D)$ captures all the information that is realized when individuals make their education decision D_i . Therefore, the individual makes her education decision according to the following rule:

$$D_i = \begin{cases} 1 & \text{if } C(D_i = 1) + EV(D_i = 1) > EV(D_i = 0) \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

5.3 Incorporating routinization

To capture the effects of routinization on skill prices, a straightforward approach is to assume that routinization affects both the pecuniary and non-pecuniary returns for each occupation O , with magnitudes that depend on education. However, using routinization directly would yield estimates that depend on unobserved supply-side confounders. We therefore adopt the instrument defined in Section 3. In particular, the functions $\{\beta_{O,\theta}^g(c,t), \beta_{O,D,\theta}^g(c,t), \alpha_{O,\theta}^g(c,t), \alpha_{O,D,\theta}^g(c,t)\}$ all depend on $\widehat{routinization}_{c(i),t}^g$, which is the first-stage predicted level of routinization in commuting zone $c(i)$ and year t . Equations (4) and (5), which specify pecuniary and non-pecuniary returns that individual i receives from occupation O , would then be:

$$Y(O_i|D_i) = X_i^Y \beta_{O,X}^g + D_i \beta_{O,\theta}^g + \theta_i \beta_{O,\theta}^{g,0} + \theta_i D_i \beta_{O,D,\theta}^{g,0} + \theta_i \beta_{O,D,\theta}^{g,1} \widehat{routinization}_{c(i),t}^g + \theta_i D_i \beta_{O,D,\theta}^{g,1} \widehat{routinization}_{c(i),t}^g + u_{O,D,i}^g \quad (7)$$

$$P(O_i|D_i) = X_i^Y \alpha_{O,X}^g + D_i \alpha_{O,\theta}^g + \theta_i \alpha_{O,\theta}^{g,0} + \theta_i D_i \alpha_{O,D,\theta}^{g,0} + \theta_i \alpha_{O,D,\theta}^{g,1} \widehat{routinization}_{c(i),t}^g + \theta_i D_i \alpha_{O,D,\theta}^{g,1} \widehat{routinization}_{c(i),t}^g \quad (8)$$

Based on the above equations, occupational returns depend on both individual characteristics (e.g., gender, education, and skills), as well as predicted routinization in the resident commuting zone. Therefore, identical workers in the same occupation may have different returns if they live in areas that experienced different levels of routinization.

5.4 Structural model estimation strategy

5.4.1 Latent abilities

To construct multi-dimensional skill profiles at the individual level, we use the NLSY79’s ASVAB test scores.¹⁹ We follow the Exploratory Factor Analysis (EFA) methods of Prada and Urzúa (2017). While prior work focuses only on men (Eisenhauer, Heckman, and Mosso, 2015; Heckman et al., 2018; Prada and Urzúa, 2017; Roys and Taber, 2019), we conduct EFA separately for men and women to allow for gender differences in how latent skills explain test performance.

The factor loadings, depicted in Figure 6, indicate that two separate skills (“factors”) explain the variation in ASVAB scores. For both men and women, the first factor has the highest loadings for subtests designed to assess cognitive skill. However, we discover a gender difference in the second factor, which suggests that men and women have distinct skills that explain their ASVAB scores. Loadings on the second factor are significant for only the mechanical ability subtests for men and only the administrative ability subtests for women. Our EFA results match with Prada and Urzúa (2017) regarding mechanical skills for men, but are novel regarding administrative skills for women.²⁰

Based on our factor loadings, we characterize each individual’s skill set θ_i by three dimensions: cognitive skill $\theta_{c,i}$, mechanical skill $\theta_{m,i}$, and administrative skill $\theta_{a,i}$. This skill structure sheds light on how men and women can have different comparative advantages, leading to the occupational sorting shown in Figure 2. Men tend to have higher mechanical skill, which provides a comparative advantage in manually intensive tasks. Women tend to have higher administrative skill, and therefore a comparative advantage in routine office work. Online Appendix A.4 describes the EFA implementation in greater detail.

Guided by the exploratory factor analysis, we specify the measurement equations for an

¹⁹The Armed Services Vocational Aptitude Battery (ASVAB) is the battery of tests commonly used to measure IQ (see, among many others, Neal and Johnson, 1996; Cameron and Heckman, 1998; Heckman et al., 2006), but it also measures non-cognitive ability. Subtests designed to measure cognitive ability and IQ are: arithmetic reasoning, word knowledge, mathematics knowledge, and paragraph comprehension. Subtests designed to measure mechanical ability are: automotive and shop information, electronics information, and mechanical comprehension. Subtests designed to measure administrative ability are: coding speed and numerical operations (ASVAB Prep Tests, 2022). See Online Appendix A.4.

²⁰We are agnostic as to whether these gender differences in skills arise from nature, nurture, or both. Men may be born with greater mechanical skill than women, or may have invested more in mechanical skill during their youth. The origins of these skill differences do not affect our model results, since our model conditions on these gender differences at college-going age.

individual i with latent skill vector $\theta_i = [\theta_{c,i}, \theta_{m,i}, \theta_{a,i}]$ as follows:

$$\begin{aligned} C_{j,i} &= \lambda_j^c \theta_{c,i} + e_{j,i}^c, j = 1, 2, \dots, 4 \\ M_{j,i} &= \lambda_j^c \theta_{c,i} + \lambda_j^m \theta_{m,i} + e_{j,i}^m, j = 5, 6, 7 \\ A_{j,i} &= \lambda_j^c \theta_{c,i} + \lambda_j^a \theta_{a,i} + e_{j,i}^a, j = 8, 9 \end{aligned} \tag{9}$$

where $C_{j,i}$ denotes the four subtests for the cognitive ability measure, $M_{j,i}$ denotes the three mechanical subtests, and $A_{j,i}$ denotes the two administrative subtests.²¹ We restrict the loading coefficients $\{\lambda_j^c, \lambda_j^m, \lambda_j^a\}$ to be gender neutral so that any gender differences in test scores reflect only gender differences in latent abilities. Lastly, to identify the system, we assume that all error terms $\{e_{1,i}^c, \dots, e_{4,i}^c, e_{5,i}^m, e_{6,i}^m, e_{7,i}^m, e_{8,i}^a, e_{9,i}^a\}$ are mutually independent and uncorrelated with the skill vector θ_i .

It is worth noting that we allow latent abilities to be correlated with each other, as several test scores are relevant for multiple abilities. To identify the system, we follow prior work and assume that at least one measure in $M_{j,i}$ is exclusively driven by mechanical skill, one measure in $A_{j,i}$ is exclusively driven by administrative skill, and a set of standard normalizations (Carneiro et al., 2003; Eisenhauer, Heckman, and Mosso, 2015; Heckman et al., 2006; Prada and Urzúa, 2017).²² We refer interested readers to the aforementioned papers or Robustness Appendix B for further details on identification.

5.4.2 The maximum likelihood function

The measurement equations are jointly estimated with the model using maximum likelihood. Let $\psi \in \Psi$ denote a vector of structural parameters and $\Omega_i = \{X_i, T_i, O_i, Y_i, D_i\}$ be the vector of observable characteristics of individual i , including exogenous control variables X_i , a college dummy D_i , occupations O_i , and annual earnings Y_i . Test scores T_i include cognitive test scores $C_{j,i}$, mechanical test scores $M_{j,i}$, and administrative test scores $A_{j,i}$.

²¹In particular, $C_{j,i} \in \{\text{arithmetic reasoning, word knowledge, paragraph comprehension, mathematics knowledge}\}$, $M_{j,i} \in \{\text{automotive and shop information, electronics information, and mechanical comprehension}\}$ and $A_{j,i} \in \{\text{coding speed and numerical operations}\}$.

²²In practice, we assume the factor loadings of cognitive skill on automotive shop information test (λ_5^c) and on coding speed test (λ_9^c) are equal to 0. The loading of cognitive skill on mathematics knowledge (λ_2^c), the loading of mechanical skill on mathematics knowledge (λ_7^m) and the loading of administrative skill on numerical operations (λ_9^a) are standardized to 1.

The likelihood function for individual i is given by

$$\begin{aligned}
\ell_i(\Omega_i|\psi) = & \int_{\theta} \underbrace{\Pi_{j=1}^4 f_j(C_{j,i}|\theta_i; \psi) \Pi_{j=5}^7 f_j(M_{j,i}|\theta_i; \psi) \Pi_{j=8}^9 f_j(A_{j,i}|X_i, \theta_i; \psi)}_{\text{skill measurements}} \\
& \underbrace{(f_Y(Y_i|D_i, O_i, X_i, \theta_i; \psi))^{I(O_i \neq 4)}}_{\text{wage outcomes}} \underbrace{\Pi_{k=1}^4 (\Pr(O_i|D_i, X_i, \theta_i; \psi))^{I(O_i=k)}}_{\text{occupations}} \\
& \underbrace{\Pi_{l=0}^1 (\Pr(D_i|X_i, \theta_i; \psi))^{I(D_i=l)}}_{\text{college}} dF_{\theta}(\theta; \psi)
\end{aligned} \tag{10}$$

where $\Pr(\cdot)$ represents the probability of occupation choice O_i or education choice D_i defined in Equations 3 and 6, $f_j(\cdot)$ is the probability density function for test j defined by Equation 9, $f_Y(\cdot)$ is the probability density function of earnings Y_i in Equation 4, and $F_{\theta}(\cdot)$ is the joint cumulative distribution of the latent skill vector $\theta \in \Theta$.

Lastly, we impose distributional assumptions. In particular, $\epsilon_{O,D,i}$ follows the standard Gumbel distribution while other error terms follow the normal distribution. For latent skills, we use mixtures of normal distributions, which provide minimal restrictions on the underlying distributions of $[\theta_c, \theta_m, \theta_a]$.²³ Following Prada and Urzúa (2017), we use mixtures of two normal distributions and assume $E[\theta_c] = E[\theta_m] = E[s] = 0$.²⁴ After plugging the distribution assumptions into Equation (10), $\Pr(O_i)$ will be a multinomial logit function and $\Pr(D_i)$ will be a probit function.

We obtain the sample log likelihood after taking the logarithm of Equation (10) and summing across all individuals. We can then obtain the estimates $\hat{\psi}$ by maximizing the total likelihood function

$$\hat{\psi} = \operatorname{argmax}_{\psi} \sum_{i=1}^N \log \ell_i(\Omega_i|\psi).$$

The standard errors are computed using the BHHH algorithm (Berndt et al., 1974).

6 Structural model results

To be brief, we report the parameter estimates for our equations in Appendix Tables A.3-A.6. Appendix Table A.7 reveals that moments from the model simulation are close to the real data. We focus this section on summarizing how these model estimates inform occupation and education choices.

²³Ferguson (1983) shows that any probability distribution can be approximated arbitrarily well by a finite mixture of normal densities. Therefore, this distributional assumption should provide sufficient flexibility while imposing a minimal number of restrictions on the underlying distributions.

²⁴However, the mean values for men and women may differ and do not necessarily equal 0.

6.1 The relationship between skills, occupational sorting, and education decisions

Our model estimates reveal notable gender differences in skill profiles, depicted in Figure 7. First, Figure 7a demonstrates similar distributions of cognitive skill for men and women, although the variance is lower for women than men.²⁵ This provides further evidence that men and women are substitutable in white collar work, and can explain the gender convergence among college occupations from 1970 to 2000 (see Figure 3b).

In contrast, there are substantial gender differences in mechanical and administrative skills. Figure 7b shows that the mechanical skill distribution for men is higher in mean and variance than for women, and that mechanical skills for women max out near the male mean. Figure 7c shows that women on average have higher administrative skills than men do. These differences provide a basis for the gender polarization among non-college occupations shown in Figure 3. They also help substantiate related research claiming that gender-based occupational segregation arose from higher mechanical skill among men (Huang, 2014; Pitt et al., 2012; Rendall, 2017; Welch, 2000).

The gender difference in skill endowments leads to comparative advantages at different occupations. This then drives occupational sorting patterns, as shown in Figure 8. Cognitive skill is positively correlated with white-collar work for both men and women. As cognitive skill increases, men shift from blue-collar occupations to white-collar occupations, while women shift from pink-collar occupations and home-staying to white-collar occupations. Mechanical skill is positively correlated with blue-collar occupations only for men. When moving up the quintiles of the mechanical skill distribution, men increasingly sort into blue-collar occupations and out of white-collar occupations. Lastly, administrative skill raises the likelihood of pink-collar work for women but barely influences occupational sorting for men. This occupational sorting matches the trends in occupational returns for women versus men, shown in Appendix Figure A.3.

We then examine the correlation between skill endowments and college attendance in Figure 9. While cognitive skill predicts college-going for both men and women, it explains more of the variation in men’s college-going. Among individuals with low cognitive skill, women attend college at higher rates than men. The disparity in college attendance diminishes as cognitive skill increases. The patterns are consistent with the idea that women have worse outside options to attending college than men, and therefore lower opportunity costs.

²⁵This result is consistent with Becker et al. (2010), who argue that the lower variance in skills among women contributes to why more women than men are prepared to attend college. Our paper argues that independent of any differences in the *supply* of students prepared for college, demand for a college degree is also higher among women than men.

College is worthwhile only for men whose white collar returns compensate them for giving up blue-collar work – that is, men whose cognitive skill is sufficiently high relative to their mechanical skill. In contrast, women’s non-college work options tend to be less lucrative, making it worthwhile to attend college even if their cognitive ability was relatively low.

Figure 9 indicates a sharper trade-off between college enrollment and non-cognitive skills for men than women. As mechanical skill increases, enrollment declines for men but stays flat for women. Administrative skill does not drive college-going for men or women. These patterns are consistent with mechanical skill raising blue-collar returns and presenting lucrative outside options to attending college. Mechanical skill therefore raises the opportunity cost of college enrollment for men, whereas administrative skill does not impose sharp costs on college enrollment for women since pink-collar returns tend to be lower than blue-collar returns.

The interactions between college attendance and skill endowments imply different levels of occupational polarization in the college and non-college labor markets, shown in Figure 10. The non-college labor market exhibits severe gender polarization. Non-college men specialize in blue collar jobs given their higher mechanical skills, while non-college women specialize in pink collar jobs. In contrast, both male and female college graduates tend to hold white collar jobs due to strong complementarities between their cognitive skills and white-collar work. Together, these results recreate the gender polarization among non-college occupations *and* the gender convergence among college occupations observed in the census data (Figure 3).

6.2 The effect of routinization on occupation choice and college enrollment

To understand how the 2SLS estimates map to the aggregate enrollment trends in Figure 1, our model simulates effects for each individual and averages the estimates over the entire sample. Figure 11 depicts the ATT estimates alongside the 2SLS estimates. It shows that treatment effects are more heterogeneous for female workers than for male workers, subject to the same level of routinization. The 2SLS for female workers is 0.578, while the ATT is 1.026. The results indicate that using a simple linear extrapolation of the 2SLS estimates would considerably underestimate the average effect for female workers. In contrast, the ATT and 2SLS estimates are similar for male workers, since their treatment effects are more homogeneous.

Next, we employ the ATT estimates to quantify how much routinization contributed to the college gender gap. Specifically, we simulate counterfactual college enrollment choices and occupational choices based on routinization levels in 1980 and 2000, while keeping other

primitive parameters constant. This approach holds constant alternative explanations such as shifts in female labor force participation, social norms regarding women’s education, or marriage market concerns. The results are presented in Table 6. Although routinization increases college attendance for both genders, enrollment growth for women is more than double that of men. Female enrollment rises by 4.0 percentage points from 1980 to 2000, while male enrollment increases by only 1.6 percentage points. This growth among women is driven by the decreasing returns to pink-collar occupations relative to white-collar occupations. We observe a 20.6 percentage point increase in white-collar jobs between 1980 and 2000, primarily due to the shift from pink-collar jobs. The simulated decline in pink-collar returns is consistent with the empirical fact that many female-dominated occupations vanished from the non-college labor market over time. In comparison, the share of men transitioning from blue-collar to white-collar occupations was only one-sixth as large, at 3.2 and 4.2 percentage points, respectively.

Lastly, we compare the decline in non-college occupations with another secular trend that could have contributed to women’s college-going: their rising labor force participation. What would female college enrollment have been in 1980 if women worked as much as men? In Table 7 columns (2)-(3), we set female labor force participation equal to male labor force participation. This raises the share of women in white- and pink-collar occupations. Since more women work in white-collar occupations than before, female enrollment increases by 1.9 percent. Together, Tables 6 and 7 indicate that the decline in women’s non-college returns due to routinization increased female enrollment by 4.0 percentage points, accounting for 44% of the 9 percentage point change in female enrollment from 1980 to 2000. This is over double the growth in female enrollment if we raised women’s labor force participation to the level of men’s. The model simulation indicates that the decline in non-college occupations is one major reason behind the growth in female college-going relative to male college-going.

7 Conclusion

The college gender gap reversed when women exceeded men in college enrollment. This came as a surprise, since many scholars anticipated that male and female enrollment rates would eventually converge. We argue that women’s greater enrollment is partly attributable to their worse outside job options. We establish two stylized facts based on the premise that the non-college labor market is highly polarized by gender, in that most occupations were male- or female-dominated. First, non-college occupations dominated by men tend to pay better than those dominated by women, suggesting that job opportunities may be worse for high school graduates if they are female. Second, this discrepancy grew over time as

automation displaced routine-intensive occupations, which employed the majority of young, working non-college women.

Informed by these stylized facts, we instrument for routinization. Our instrument predicts the share of occupations intensive in administrative activity based on job posting data from major newspapers in 1950-2000. The intuition behind our instrument is that industries with higher administrative activity involve more routine tasks, and local labor markets with greater historic shares of these industries would experience more routinization over time. Consistent with this intuition, our first stage regressions show that local labor markets with higher predicted administrative shares in 1950 experienced greater routinization as workplaces automated. This decline led to significant enrollment growth among 18-25 year old women, but effects for men are not systematically significant. We estimate that moving from a commuting zone which experienced the 25th percentile of routinization to one which experienced the 75th percentile of routinization corresponds to a 3.34 rise in female enrollment.

To investigate the mechanisms that explain these results at the individual level, we develop a discrete choice Roy model. The model embeds instrumental variation from the job posting data to examine how routinization affects the value of different skills. Using a maximum likelihood procedure, we find that gender differences in skills lead non-college men to sort into manual-intensive work and non-college women into routine-intensive work. The resulting gender polarization among non-college occupations translates to a comparative advantage for men in non-college work, given the greater pay in manual occupations relative to administrative occupations. Over time, automation decreased the value of administrative skill in routine-intensive work, lowering the opportunities for non-college women and exacerbating their comparative disadvantage in non-college work.

Our focus on the non-college labor market addresses the apparent contradiction regarding why women exceed men in college-going yet lag behind men in college earnings. Due to the difference in outside options, women have lower opportunity costs to attend college. This means that on average, they are willing to accept lower wages than men to pursue college. Women's lower non-college job opportunities can thus explain both why the college gender gap reversed fifty years ago, yet the gender wage gap has continued to persist.²⁶ Exploring this linkage in greater detail will be a fruitful avenue for future work.

²⁶Relatedly, women's lower opportunity costs mean that they can make college pay off even if they pursue non-STEM majors, which tend to be less lucrative. Indeed, Card and Payne (2021) find that most of the gender gap in STEM majors can be explained by the fact that women attend college at greater rates to begin with.

References

- Acemoglu, D., and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. C. Ashenfelter and D. Card (Eds.), *Handbook of labor economics* (pp. 1043–1171). Elsevier.
- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2020). AI and jobs: Evidence from online vacancies. *NBER Working Paper 28257*.
- Acemoglu, D., and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Acemoglu, D., and Restrepo, P. (2020). Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy*, 128(6), 2188–2244.
- Adams, B. T. (2019). Farm machinery automation for tillage, planting cultivation, and harvesting. *Handbook of farm, dairy and food machinery engineering* (pp. 115–131). Elsevier.
- Adao, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4), 1949–2010.
- Andrews, I., Stock, J., and Sun, L. (2018). Weak instruments and what to do about them.
- Andrews, I., and Stock, J. H. (2018). Weak instruments and what to do about them.
- Angrist, J. D., and Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American statistical Association*, 90(430), 431–442.
- Atack, J., Margo, R. A., and Rhode, P. W. (2019). “automation” of manufacturing in the late nineteenth century: The hand and machine labor study. *Journal of Economic Perspectives*, 33(2), 51–70.
- Atalay, E., Phongthientham, P., Sotelo, S., and Tannenbaum, D. (2018). New technologies and the labor market. *Journal of Monetary Economics*, 97, 48–67.
- Atalay, E., Phongthientham, P., Sotelo, S., and Tannenbaum, D. (2020). The evolution of work in the united states. *American Economic Journal: Applied Economics*, 12(2), 1–34.
- Autor, D., and Wasserman, M. (2013). *Wayward sons: The emerging gender gap in labor markets and education* (tech. rep.). Third Way.
- Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives*, 29(3), 3–30.
- Autor, D. H., and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. *American Economic Review*, 103(5), 1553–97. <http://doi.org/10.3886/E112652V1>
- Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6), 2121–68.
- Autor, D. H., Dorn, D., and Hanson, G. H. (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125(584), 621–646.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4), 1279–1333.

- Beaudry, P., and Lewis, E. (2014). Do male-female wage differentials reflect differences in the return to skill? cross-city evidence from 1980-2000. *American Economic Journal: Applied Economics*, 6(2), 178–94.
- Becker, G. S., Hubbard, W. H., and Murphy, K. M. (2010). Explaining the worldwide boom in higher education of women. *Journal of Human Capital*, 4(3), 203–241.
- Berndt, E. R., Hall, B. H., Hall, R. E., and Hausman, J. A. (1974). Estimation and inference in nonlinear structural models. *Annals of economic and social measurement, volume 3, number 4* (pp. 653–665). NBER.
- Bertrand, M., and Pan, J. (2013). The trouble with boys: Social influences and the gender gap in disruptive behavior. *American economic journal: applied economics*, 5(1), 32–64.
- Black, S. E., and Spitz-Oener, A. (2010). Explaining women’s success: Technological change and the skill content of women’s work. *The Review of Economics and Statistics*, 92(1), 187–194.
- Blandhol, C., Bonney, J., Mogstad, M., and Torgovitsky, A. (2022). When is tsls actually late? *National Bureau of Economic Research Working Paper*.
- Borghans, L., Weel, B. T., and Weinberg, B. A. (2014). People skills and the labor-market outcomes of underrepresented groups. *ILR Review*, 67(2), 287–334. <https://doi.org/10.1177/001979391406700202>
- Borusyak, K., Hull, P., and Jaravel, X. (2018). Quasi-experimental shift-share research designs. *NBER Working Paper*, (w24997).
- Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, 117(1), 339–376. <https://doi.org/10.1162/003355302753399526>
- Brynjolfsson, E., and Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic perspectives*, 14(4), 23–48.
- Bureau of Labor Statistics. (2019). *National longitudinal survey of youth 1979 cohort, 1979-2016 (rounds 1-27)* (tech. rep.). Center for Human Resource Research (CHRR), The Ohio State University. Columbus, OH.
- Cameron, S. V., and Heckman, J. J. (1998). Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of American males. *Journal of Political economy*, 106(2), 262–333.
- Card, D., and Payne, A. A. (2021). High school choices and the gender gap in stem. *Economic Inquiry*, 59(1), 9–28.
- Carneiro, P., Hansen, K. T., and Heckman, J. J. (2003). Estimating distributions of treatment effects with an application to the returns to schooling and measurement of the effects of uncertainty on college choice. *International Economic Review*, 44(2), 361–422.
- Carrell, S., and Sacerdote, B. (2017). Why do college-going interventions work? *American Economic Journal: Applied Economics*, 9(3), 124–51.
- Census Bureau. (2021). *Industry and occupation indexes*. Retrieved December 30, 2021, from <https://www.census.gov/topics/employment/industry-occupation/guidance/indexes.html>
- Cerina, F., Moro, A., and Rendall, M. (2021). The role of gender in employment polarization. *International Economic Review*, 62(4), 1655–1691.

- Charles, K. K., and Luoh, M.-C. (2003). Gender differences in completed schooling. *Review of Economics and statistics*, 85(3), 559–577.
- Chiappori, P.-A., Iyigun, M., and Weiss, Y. (2009). Investment in schooling and the marriage market. *American Economic Review*, 99(5), 1689–1713.
- Chiappori, P.-A., Salanié, B., and Weiss, Y. (2017). Partner choice, investment in children, and the marital college premium. *American Economic Review*, 107(8), 2109–2167.
- Cortes, G. M., Jaimovich, N., Nekarda, C. J., and Siu, H. E. (2014). The micro and macro of disappearing routine jobs: A flows approach. *NBER Working Paper 20307*.
- Cortes, G. M., Jaimovich, N., and Siu, H. E. (2017). Disappearing routine jobs: Who, how, and why? *Journal of Monetary Economics*, 91, 69–87.
- Cortes, G. M., Jaimovich, N., and Siu, H. E. (2021). The growing importance of social tasks in high-paying occupations: Implications for sorting. *Journal of Human Resources*, 0121–11455R1.
- Decker, K. D. (2016). Why the office needs a typewriter revolution. *Low-Tech Magazine*. <https://www.lowtechmagazine.com/2016/11/why-the-office-needs-a-typewriter-revolution.html>
- Diekhoff, G. (1992). *Statistics for the social and behavioral sciences: Univariate, bivariate, multivariate*.
- Dillender, M., and Forsythe, E. (2019). Computerization of white collar jobs. *Upjohn Institute Working Paper*, (19-310).
- Dougherty, C. (2005). Why are the returns to schooling higher for women than for men? *Journal of Human Resources*, 40(4), 969–988.
- Duran-Franch, J. (2020). Oh, man! What happened to women? The blurring of gender-based occupational segregation.
- Eisenhauer, P., Heckman, J. J., and Mosso, S. (2015). Estimation of dynamic discrete choice models by maximum likelihood and the simulated method of moments. *International Economic Review*, 56(2), 331–357.
- Eisenhauer, P., Heckman, J. J., and Vytlacil, E. (2015). The generalized Roy model and the cost-benefit analysis of social programs. *Journal of Political Economy*, 123(2), 413–443.
- Ferguson, T. S. (1983). Bayesian density estimation by mixtures of normal distributions, 287–302.
- Ge, S. (2011). Women’s college decisions: How much does marriage matter? *Journal of Labor Economics*, 29(4), 773–818.
- Goldin, C., Katz, L. F., and Kuziemko, I. (2006). The homecoming of American college women: The reversal of the college gender gap. *Journal of Economic perspectives*, 20(4), 133–156.
- Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in europe. *American economic review*, 99(2), 58–63.
- Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American economic review*, 104(8), 2509–26.
- Greenwood, J., Seshadri, A., and Yorukoglu, M. (2005). Engines of liberation. *The Review of Economic Studies*, 72(1), 109–133.

- Heckman, J. J., and Cameron, S. (2001). The dynamics of educational attainment for black, Hispanic, and white males. *Journal of Political Economy*, 109(3), 455–499.
- Heckman, J. J., Humphries, J. E., and Veramendi, G. (2018). Returns to education: The causal effects of education on earnings, health, and smoking. *Journal of Political Economy*, 126(S1), S197–S246.
- Heckman, J. J., Stixrud, J., and Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor economics*, 24(3), 411–482.
- Hershbein, B., and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7), 1737–72.
- Huang, L. (2014). *A revolution in education: Determinants of the gender gap reversal* (Doctoral dissertation). University of Pennsylvania.
- Jacob, B. A. (2002). Where the boys aren't: Non-cognitive skills, returns to school and the gender gap in higher education. *Economics of Education review*, 21(6), 589–598.
- Jaimovich, N., and Siu, H. (2012). Job polarization and jobless recoveries. *NBER Working Paper 18334*. <https://doi.org/10.3386/w18334>
- Juhn, C., Ujhelyi, G., and Villegas-Sanchez, C. (2014). Men, women, and machines: How trade impacts gender inequality. *Journal of Development Economics*, 106, 179–193.
- Keane, M. P., and Wolpin, K. I. (1997). The career decisions of young men. *Journal of political Economy*, 105(3), 473–522.
- Keane, M. P., and Wolpin, K. I. (2001). The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review*, 42(4), 1051–1103.
- Lee, D. L., McCrary, J., Moreira, M. J., and Porter, J. (2020). Valid t-ratio inference for IV. *arXiv preprint arXiv:2010.05058*.
- Marcus, J. (2017). Why men are the new college minority.
- Men adrift. (2015). *The Economist*. Available at <https://www.economist.com/essay/2015/05/28/men-adrift> (2023/02/14).
- Neal, D. A., and Johnson, W. R. (1996). The role of premarket factors in black-white wage differences. *Journal of political Economy*, 104(5), 869–895.
- Ngai, L. R., and Petrongolo, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4), 1–44.
- Olea, J. L. M., and Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business and Economic Statistics*, 31(3), 358–369.
- Olivetti, C., and Petrongolo, B. (2014). Gender gaps across countries and skills: Demand, supply and the industry structure. *Review of Economic Dynamics*, 17(4), 842–859.
- Olivetti, C., and Petrongolo, B. (2016). The evolution of gender gaps in industrialized countries. *Annual review of Economics*, 8, 405–434.
- Olivieri, E. (2014). *Occupational choice and the college gender gap* (Doctoral dissertation). University of Chicago.
- Pitt, M. M., Rosenzweig, M. R., and Hassan, M. N. (2012). Human capital investment and the gender division of labor in a brawn-based economy. *American economic review*, 102(7), 3531–60.

- Prada, M. F., and Urzúa, S. (2017). One size does not fit all: Multiple dimensions of ability, college attendance, and earnings. *Journal of Labor Economics*, 35(4), 953–991.
- Rendall, M. (2017). Brain versus brawn: The realization of women’s comparative advantage. *University of Zurich, Institute for Empirical Research in Economics, Working Paper*, (491).
- Rosin, H. (2010). The end of men. *The Atlantic*. Available at <https://www.theatlantic.com/magazine/archive/2010/07/the-end-of-men/308135/> (2023/02/14).
- Roys, N. A., and Taber, C. R. (2019). Skill prices, occupations, and changes in the wage structure for low skilled men.
- Ruggles, S., Flood, S., Foster, S., Goeken, R., Pacas, J., Schouweiler, M., and Sobek, M. (2021). IPUMS USA: Version 11.0. <https://doi.org/10.18128/D010.V11.0>
- Sheskin, D. J. (2004). *Test 28: the Pearson Product-Moment Correlation Coefficient. Handbook of Parametric and Nonparametric Statistical Procedures* (3rd ed.). CRC Press.
- Słoczyński, T. (2020). When should we (not) interpret linear iv estimands as late? *arXiv preprint arXiv:2011.06695*.
- Sommers, C. H. (2013). How to make school better for boys. *The Atlantic*. Available at <https://www.theatlantic.com/education/archive/2013/09/how-to-make-school-better-for-boys/279635/> (2023/02/14).
- Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of labor economics*, 24(2), 235–270.
- Todd, P. E., and Zhang, W. (2020). A dynamic model of personality, schooling, and occupational choice. *Quantitative Economics*, 11(1), 231–275. <https://doi.org/10.3982/qe890>
- Welch, F. (2000). Growth in women’s relative wages and in inequality among men: One phenomenon or two? *American Economic Review*, 90(2), 444–449.
- Welsh, J. R., Kucinkas, S. K., and Curran, L. T. (1990). Armed services vocational battery (ASVAB): Integrative review of validity studies.
- Why the future at U.S. colleges is female. (2021). *Wall Street Journal*. <https://www.wsj.com/articles/college-campus-men-women-gender-gap-role-imbalance-feminism-woke-trade-school-11632257601>
- Yamaguchi, S. (2018). Changes in returns to task-specific skills and gender wage gap. *Journal of Human Resources*, 53(1), 32–70.
- Zhang, H. (2021). An investment-and-marriage model with differential fecundity: On the college gender gap. *Journal of Political Economy*, 129(5), 1464–1486.

Table 1: Change in Occupation Groups, 1970-2000

	Females				Males			
	Share of Workforce (%)			RTI	Share of Workforce (%)			RTI
	1970-2000	1970	2000	1970-2000	1970-2000	1970	2000	1970-2000
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Office and Administrative Support Occupations	-14.89	42.37	27.48	-0.42	0.65	9.61	10.25	-0.16
Production Occupations	-9.24	14.72	5.47	0.09	-8.89	20.92	12.03	0.05
Agriculture and Construction Occupations	-0.19	0.94	0.76	0.02	2.40	11.99	14.39	-0.03
Installation, Maintenance, and Repair Workers	0.13	0.26	0.39	0.10	-0.46	7.59	7.13	0.05
Transportation and Material Moving Occupations	0.75	1.36	2.10	0.02	-5.32	16.56	11.24	-0.08
Computer, Mathematical, Engineering, and Science Occupations	1.15	1.34	2.49	-0.13	-0.34	5.82	5.48	-0.29
Healthcare and Protective Occupations	2.09	7.54	9.62	0.02	1.23	2.80	4.03	0.22
Community, Social Services, Education, Arts, Media Occupations	2.89	6.43	9.32	0.26	-0.09	3.95	3.86	0.16
Management, Business, Science, Arts Occupations	5.52	2.44	7.96	-0.08	-0.85	7.11	6.26	0.14
Service (Food, Maintenance, Sales) Occupations	11.80	22.61	34.41	0.09	11.69	13.64	25.32	-0.02

Notes: Occupation groups as share of 18-25 year old civilian workforce. Columns (1-4) show summary statistics for women, while columns (5-8) show summary statistics for men. Columns (1-3) and (5-7) show shares as proportions of female or male workforce. Columns (4) and (8) show change in standard deviation of routine task intensity (RTI) measure from Autor and Dorn (2013). Data from Census and Autor and Dorn (2013).

Table 2: OLS Regression of College Enrollment on Routinization

	College enrollment			
	(1)	(2)	(3)	(4)
<i>A. Women</i>				
Routinization	0.416	0.448	0.431	0.467
	(0.096)***	(0.098)***	(0.094)***	(0.095)***
Observations	3610	3610	3610	3610
<i>B. Men</i>				
Routinization	0.215	0.221	0.235	0.243
	(0.120)*	(0.122)*	(0.128)*	(0.130)*
Observations	3610	3610	3610	3610
Commuting zone FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Median cognitive earnings		✓		✓
Lagged RTI share			✓	✓

OLS regressions of enrollment on instruments at the commuting zone-year level. All regressions include demographic controls for the proportion of female, Black, and Hispanic residents and by 10-year age bin. All regressions also control for U.S. census division, year, commuting zone, labor force participation, manual occupation share, and 10-year lagged major industry shares: services, manufacturing, retail, and mining. Columns (2) and (4) add median annual log earnings for occupations in the top third of abstract-intensive tasks. Columns (3) and (4) additionally control for the 10-year lag of RTI share. Standard errors clustered at commuting zone level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: First Stage Regression of Routinization on Instruments

	Routinization			
	(1)	(2)	(3)	(4)
Administrative share IV	0.387 (0.026)***	0.383 (0.027)***	0.388 (0.027)***	0.383 (0.027)***
F-statistic	214.572	204.993	204.654	201.452
Observations	3610	3610	3610	3610
Commuting zone FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Median cognitive earnings		✓		✓
Lagged RTI share			✓	✓

First stage regression of RTI share on instruments. All regressions include demographic controls for the proportion of female, Black, and Hispanic residents and by 10-year age bin. All regressions also control for U.S. census division, year, commuting zone, labor force participation, manual occupation share, and 10-year lagged major industry shares: services, manufacturing, retail, and mining. Columns (2) and (4) add median annual log earnings for occupations in the top third of abstract-intensive tasks. Columns (3) and (4) additionally control for the 10-year lag of the share of high-RTI occupations. Standard errors are clustered at the two-digit industry level and adjusted using the correction procedure of Adao et al. (2019). Olea-Pflueger F-statistics reported using AKM (2019) standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Reduced Form and Second Stage Regressions

	College enrollment			
	(1)	(2)	(3)	(4)
<i>A. Reduced form regression, women</i>				
Administrative share IV	0.224 (0.062)***	0.232 (0.062)***	0.224 (0.061)***	0.232 (0.061)***
Observations	3610	3610	3610	3610
<i>B. Reduced form regression, men</i>				
Administrative share IV	0.169 (0.092)*	0.170 (0.092)*	0.169 (0.091)*	0.170 (0.091)*
Observations	3610	3610	3610	3610
<i>C. Second stage regression, women</i>				
Routinization	0.578 (0.163)*** [0.258,0.898]	0.606 (0.166)*** [0.281,0.931]	0.578 (0.160)*** [0.265,0.891]	0.606 (0.161)*** [0.291,0.922]
First Stage F-statistic	214.572	204.993	204.654	201.452
Observations	3610	3610	3610	3610
<i>D. Second stage regression, men</i>				
Routinization	0.436 (0.236)* [-0.026,0.898]	0.444 (0.238)* [-0.022,0.910]	0.436 (0.232)* [-0.019,0.891]	0.444 (0.234)* [-0.015,0.904]
First Stage F-statistic	214.572	204.993	204.654	201.452
Observations	3610	3610	3610	3610
Commuting zone FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Median cognitive earnings		✓		✓
Lagged RTI share			✓	✓

This table presents the reduced form (panels A-B) and second stage (panels C-D) estimates. Panels A and C display the estimates for women, while panels B and D display the estimates for men. All regressions include demographic controls for the proportion of female, Black, and Hispanic residents and by 10-year age bin. All regressions also control for U.S. census division, year, commuting zone, labor force participation, manual occupation share, and 10-year lagged major industry shares: services, manufacturing, retail, and mining. Standard errors are clustered at the two-digit industry level and adjusted using the correction procedure of Adao et al. (2019). Montiel Olea-Pflueger first stage F-statistics reported using AKM (2019) standard errors. The second stage estimates include Anderson-Rubin (1949) weak instrument robust confidence intervals using the AKM (2019) correction procedure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Second Stage Regressions, Additional Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	<i>A: Female Enrollment</i>							
Routinization	0.495 (0.167)*** [0.167,0.822]	0.607 (0.158)*** [0.297,0.917]	0.628 (0.161)*** [0.312,0.944]	0.736 (0.242)*** [0.263,1.210]	0.504 (0.126)*** [0.256,0.752]	0.365 (0.145)** [0.080,0.650]	0.548 (0.145)*** [0.213,0.884]	0.784 (0.145)*** [0.350,1.219]
F-statistic	137.280	205.558	182.001	51.040	103.421	203.540	161.233	111.715
Observations	3610	3600	3610	3610	3610	3610	3610	3610
	<i>B: Male Enrollment</i>							
Routinization	0.315 (0.266) [-0.207,0.838]	0.503 (0.234)** [0.044,0.961]	0.441 (0.238)* [-0.025,0.907]	0.540 (0.315)* [-0.077,1.157]	0.218 (0.174) [-0.123,0.558]	0.254 (0.196) [-0.129,0.638]	0.432 (0.196) [-0.101,0.964]	0.616 (0.196)* [-0.048,1.280]
F-statistic	137.280	205.558	182.001	51.040	103.421	203.540	161.233	111.715
Observations	3610	3600	3610	3610	3610	3610	3610	3610
Minimum controls	✓							
Excluding Boston and NYC		✓						
Control for abstract occupation share			✓					
RTI share: non-college workers	✓	✓	✓		✓	✓	✓	✓
RTI share: college and non-college workers				✓				
IV: Administrative Share (top third)	✓	✓	✓	✓				
IV: Routine Share					✓			
IV: Administrative Share (top half)						✓		
IV: Administrative Activities							✓	
IV: Clerical Requirements								✓

Two stage least squares regressions, additional specifications. Column (1) uses a minimum set of controls: total commuting zone population, year dummies, census region dummies, commuting zone dummies, manual occupation share, proportion by gender, race, and ten-year age bin. Columns (2)-(7) start from the basic specification of Table 4 Column (1). Column (2) excludes commuting zones that contain Boston and New York City. Column (3) additionally controls for abstract occupation share. Column (4) uses the routinization of all workers, rather than only non-college workers used in the main specification. The IV in column (5) uses the share of occupations in the top half of administrative activity, rather than the top third. Column (6) uses the administrative activities IV, and column (7) the clerical requirements IV. Standard errors are clustered at the two-digit industry level and adjusted using the correction procedure of Adao et al. (2019). Montiel Olea-Pfueger F-statistics reported using AKM (2019) standard errors. Anderson-Rubin (1949) confidence intervals reported using AKM (2019) correction. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Simulated Changes in Occupation and Education due to Routinization

	Women			Men		
	Baseline: 1980 (1)	Routinization in 2000 (2)	Change (3)	Baseline: 1980 (4)	Routinization in 2000 (5)	Change (6)
<i>Occupation choices</i>						
White collar	0.404	0.610	0.206	0.392	0.434	0.042
Blue collar	0.038	0.049	0.011	0.485	0.452	-0.032
Pink collar	0.348	0.153	-0.195	0.048	0.050	0.002
Not working	0.209	0.188	-0.021	0.075	0.064	-0.012
<i>Education choices</i>						
High school	0.389	0.350	-0.040	0.481	0.465	-0.016
College	0.611	0.650	0.040	0.519	0.535	0.016

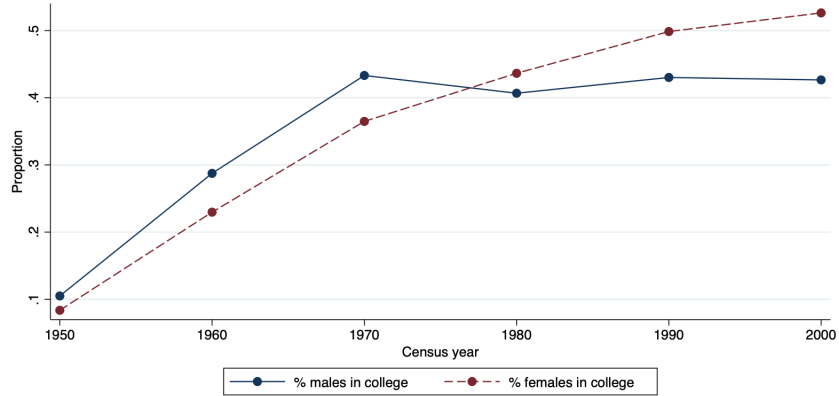
This table presents simulated education and occupation choices for the NLSY79 cohort. Columns (1) and (2) report simulated choices for women based on the routinization in 1980 and 2000, respectively. Column (3) reports the difference in simulated choices for women from 1980 to 2000. Columns (4) and (5) report simulated choices for men based on routinization in 1980 and 2000, respectively. Column (6) reports the difference in simulated choices for men from 1980 to 2000.

Table 7: Simulated Changes in Women's Outcomes from Equalizing Labor Force Participation

	Baseline	Equalizing LFP	
	1980 (1)	Level (2)	Change (3)
<i>Occupation choices</i>			
White collar	0.404	0.455	0.051
Blue collar	0.038	0.048	0.010
Pink collar	0.348	0.422	0.074
Not working	0.209	0.075	-0.134
<i>Education choices</i>			
High school	0.389	0.371	-0.019
College	0.611	0.629	0.019

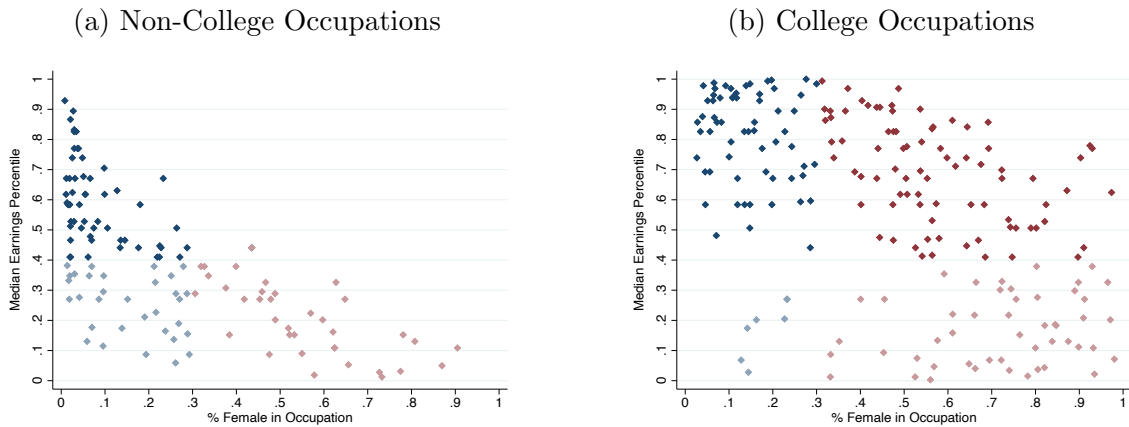
The first column displays simulated choices for women in the baseline model in 1980. Column (2) presents the alternative simulated choices for women after setting their labor force participation rates equal to men's. Column (3) reports the difference between columns (1) and (2).

Figure 1: College Enrollment by Gender, 1950-2000



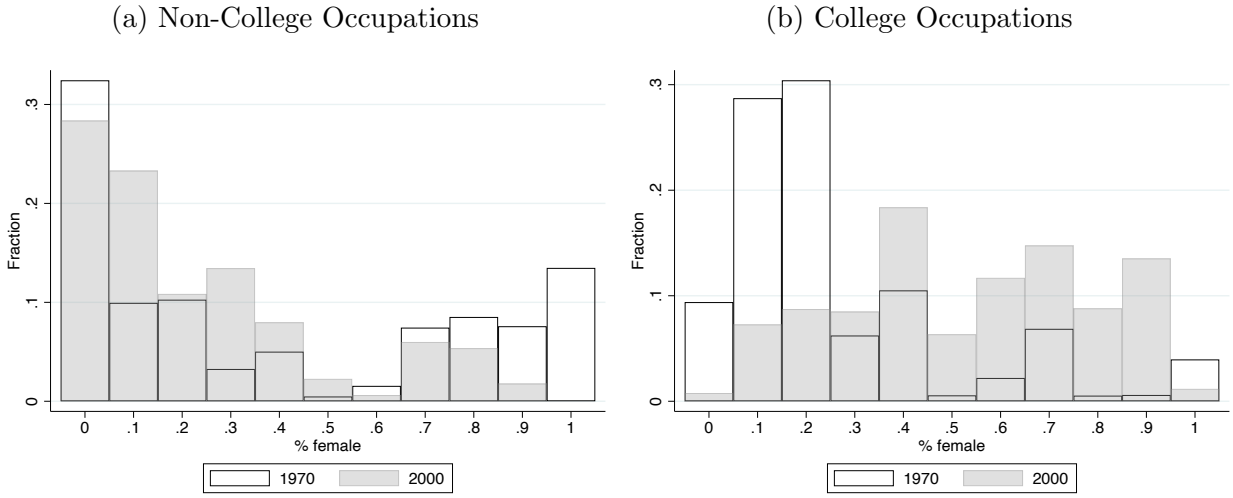
Proportion of 18-25 year olds ever enrolled in college. Solid lines represent male enrollment and dashed lines represent female enrollment. Data from the U.S. census.

Figure 2: Occupations by Gender Composition and Percentile Median Earnings, 2000



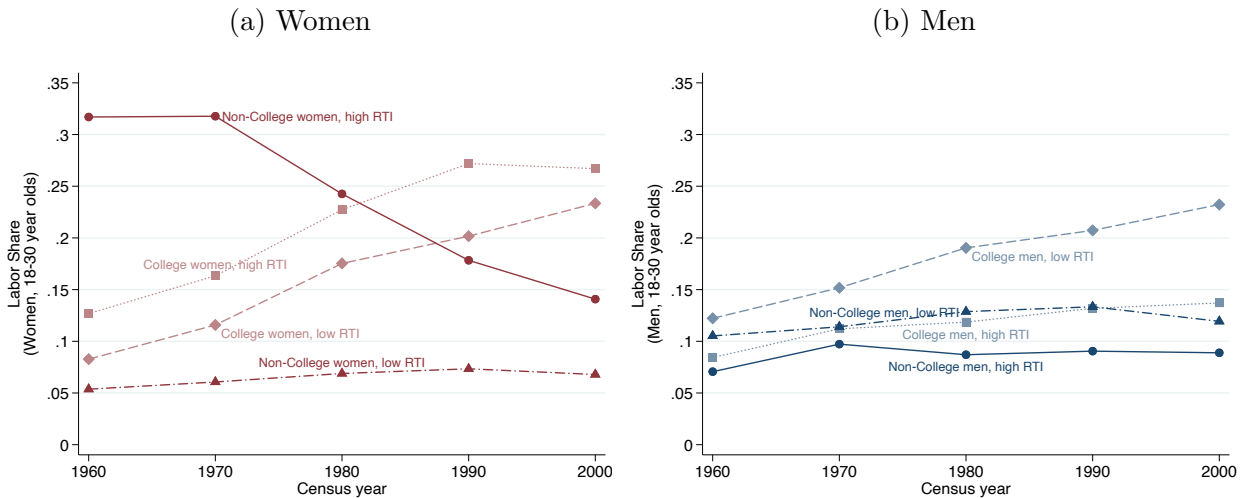
Occupations by proportion female and median annual earnings percentile in 2000. Panel a depicts occupations with 50% or fewer college graduates in 2000. Panel b depicts occupations with 50% or more college graduates in 2000. Navy markers indicate occupations where women comprise less than 30% of all workers, with dark navy markers representing occupations with earnings above the 40th percentile and light navy markers representing occupations with earnings below the 40th percentile. Maroon markers indicate occupations where women comprise 30% or more of all workers. Data from the U.S. census.

Figure 3: Occupational Dispersion by Gender Composition



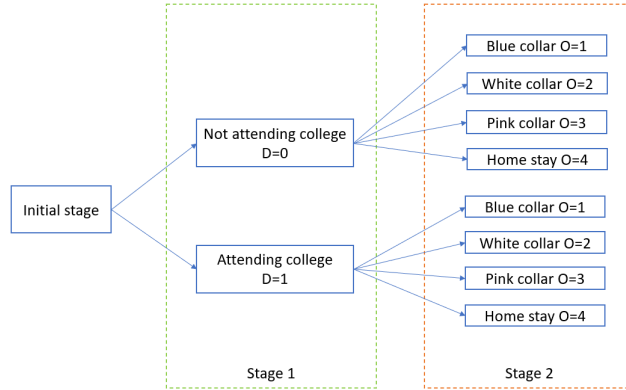
Distribution of occupations by proportion female in 1970 and 2000 for “non-college” occupations (a) and “college” occupations (b). “Non-college” occupations are those with 50% or fewer college graduates, while “college” occupations are those with over 50% college graduates. The designation of occupations as “college” or “non-college” changes each year based on the education composition of workers. Individuals aged 18-30 years old. Data from the U.S. census.

Figure 4: Low vs. High-RTI Labor Share, by Education



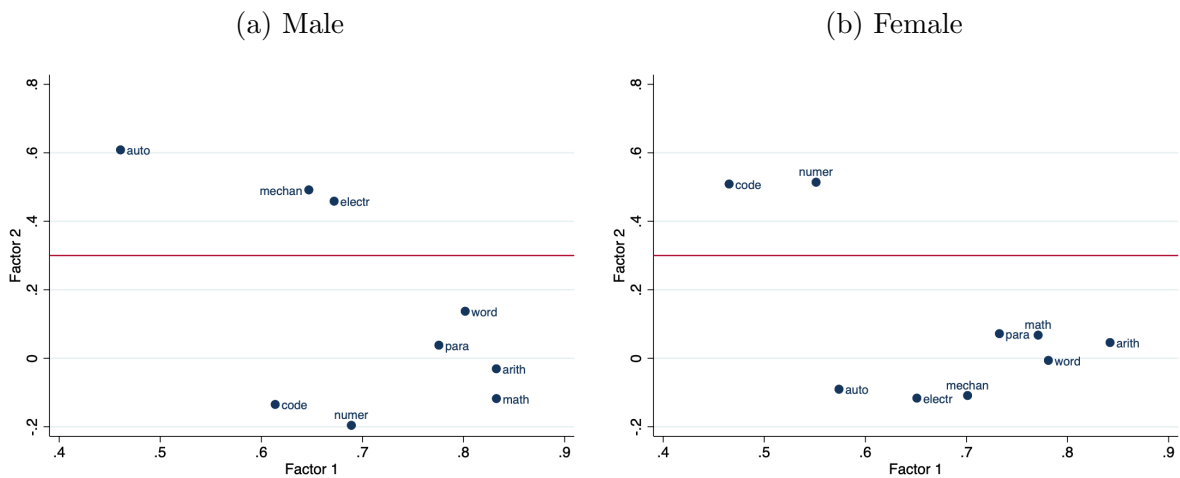
Share of occupations in top and bottom third of routine task intensity (low- and high-RTI share, respectively). Panel a plots this share for women, while panel b plots this share for men. Data from the U.S. census and Autor and Dorn (2013). 18-30 year olds.

Figure 5: Two Period Dynamic Discrete Choice Model



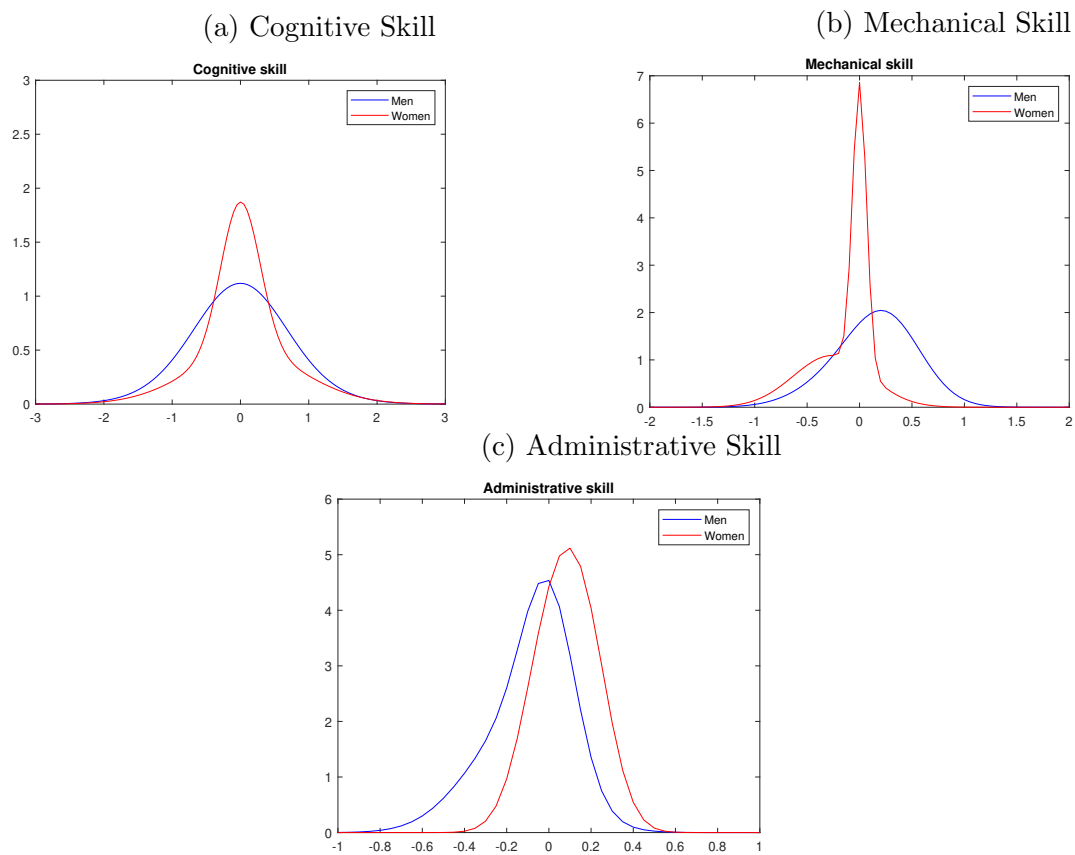
Description of structural discrete choice model. In Stage 1, individuals decide whether or not to attend college. In Stage 2, they choose their occupation from four choices: blue collar, white collar, pink collar, or home staying. The model is solved via backward induction.

Figure 6: Factor Analysis Loadings



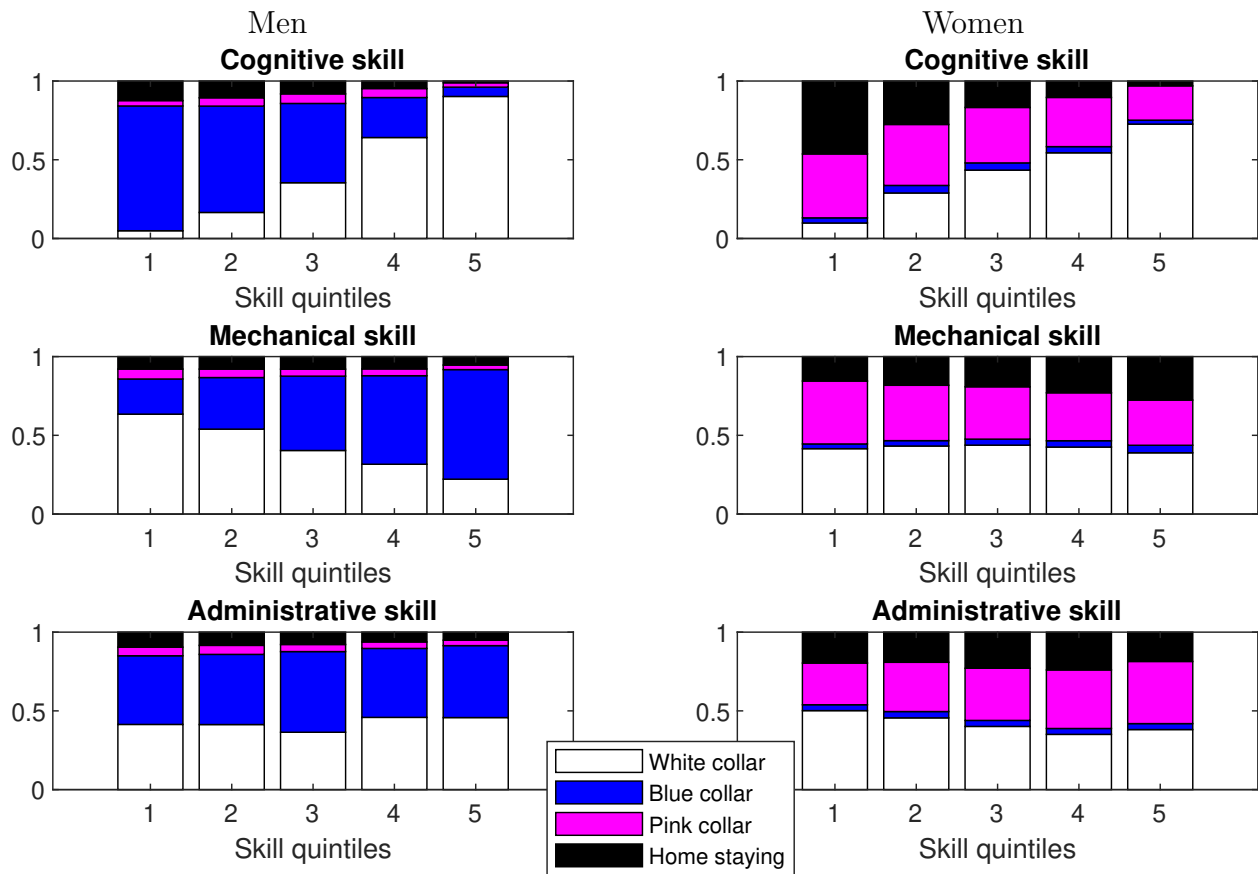
Loadings calculated from exploratory factor analysis (quartimax rotation). The red horizontal line marks the statistically significant threshold (see Diekhoff, 1992; Sheskin, 2004). arith = arithmetic reasoning; auto= automotive information and shop information; code = coding speed; electr = electronics information; math = mathematics knowledge; mechan = mechanical comprehension; numer = numerical operations; para = paragraph comprehension; word = word knowledge.

Figure 7: Distribution of Skills by Gender



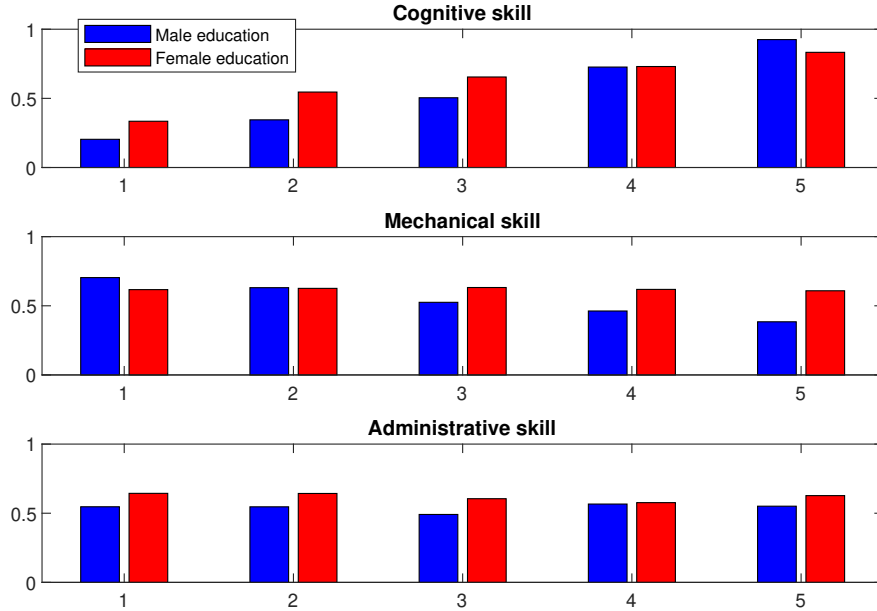
Distribution of skills by gender. The blue distribution is for men, and the red distribution is for women. Panel (a) presents the estimated distribution of cognitive skill, while panels (b) and (c) present analogous results for mechanical skill and administrative skill, respectively.

Figure 8: Occupation Choice Distribution by Skill Quintile and Gender



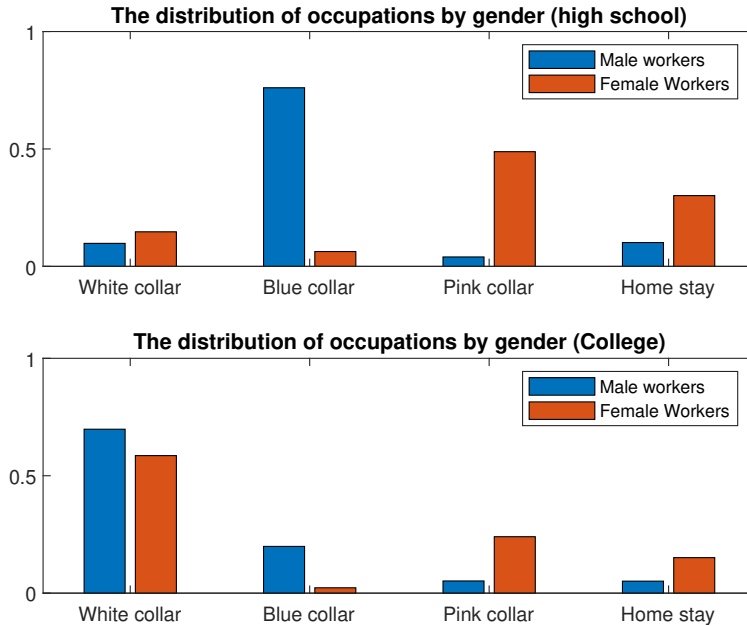
We simulate each individual 200 times based on the estimates of the model to calculate the distribution of occupation choices by skill quintiles and gender. The upper panels present the effect of cognitive skill by gender, integrating out the effect of the other two dimensions of skills, while the middle panel and the lower panel present analogous results for mechanical skill and administrative skill, respectively.

Figure 9: College Attendance Rates by Skill Quintiles



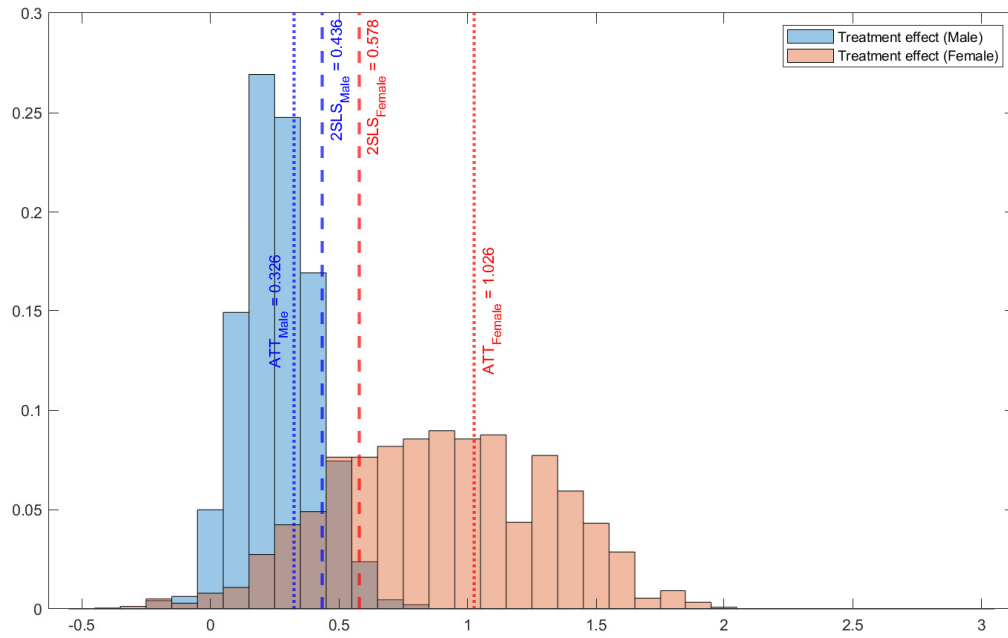
We simulate each individual 200 times based on the estimates of the model to calculate the college attendance rate by skill quintile and gender. The vertical axis is the fraction of workers in each skill group. The upper panels present the effect of cognitive skill, integrating out the effect of the other two dimensions of skills. The middle panel and lower panel present analogous results for mechanical skill and administrative skill, respectively.

Figure 10: Distribution of Occupations by Gender and Education



We simulate each individual 200 times based on the estimates of the model to calculate the occupation distribution by gender and education levels. The vertical axis is the fraction of workers in each occupation group. The upper panels present occupation distribution for college-goer, with blue bars for men and red bars for women. The lower panel present occupation distribution for high school graduates, with blue bars for men and red bars for women.

Figure 11: College Attendance Responses to Routinization by Gender



The figure illustrates the distribution of college attendance responses to routinization, using blue bars for men and red bars for women. Long dashed vertical lines represent the ATT, while short dashed vertical lines denote the 2SLS estimates.

Online Appendix - For Online Publication

A Data Appendix

A.1 Census microdata

Our first data sets come from the decennial census microdata from 1950 to 2000, which are conducted by the U.S. Census Bureau and made publicly available through the Integrated Public Use Microdata Series (IPUMS, Ruggles et al., 2021). For enrollment, we only examine 18-25 year olds to ensure that we only detect changes in education among those closest to college enrollment age. Following Acemoglu and Autor (2011), we restrict the sample to full-time (at least 35 hours worked per week), full-year (at least 40 weeks worked per year) workers.

The college enrollment variable is constructed using the harmonized EDUCD variable. Individuals are coded as college enrollees if they report having at least some college education. They are coded as never having enrolled in college if their highest reported level of educational attainment was a high school diploma or equivalent. Those who did not report an education level were excluded from the analysis.

Annual earnings data is obtained from the variable INCWAGE, the pre-tax individual income from wages and salary. Annual earnings are only computed for workers who report working for wages or salary. We exclude individuals who report being self-employed or unpaid family workers and individuals who report working no weeks in the previous year. Annual earnings are top-coded at the pre-determined Census top-code levels, which vary from year to year. They are bottom-coded as the 1st percentile of reported earnings for each year. All earnings are inflated to 2008 dollars.

All regressions are conducted at the commuting zone-year level. We merge the census data to corresponding commuting zones using the crosswalks provided by Autor and Dorn (2013). Demographic characteristics, occupations, education, earnings, and work variables are collapsed to the commuting zone level using labor supply weights calculated following the method of Acemoglu and Autor (2011).

Appendix Table A.1 presents summary statistics by decade from 1960 to 2000. Each variable represents the average across commuting zones. Female enrollment increases steadily over the decades, while male enrollment quickly rises from 1960-1970, then declines in 1980 before rising again. The proportion of women in each commuting zone stays constant at 50-51%, and the proportion of blacks also hold constant at 8% over our analysis period. The share of Hispanics grows steadily over time, from 3% in 1960 to 8% in 2000.

A.2 Data from Autor and Dorn (2013)

To obtain information on work content, we merge the census data to the occupational task intensity data compiled by Autor and Dorn (2013) using the OCC1990 variable, which is harmonized across all years. Autor and Dorn (2013)'s Routine Task Intensity (RTI) measure is the primary measure we use to determine how routine-intensive an occupation is. Following Autor and Dorn (2013), we classify an occupation as highly routine-intensive

if its RTI measure falls in the top third of all RTI in 1980. Out of 330 total occupations, 113 occupations fit this criterion.

Our main analysis concerns RTI share, the proportion of jobs in a commuting zone that are highly routine intensive. Routinization, the main endogenous regressor in the two stage least squares approach, is the reduction in RTI share from 1950 base levels. We restrict the RTI share measure to only 25-65 year olds. If youth choose to enroll in college for reasons not captured by our data, that would mechanically lower labor share and bias the estimated causal relationship between labor share and college enrollment. We therefore exclude 18-25 year olds to avoid these simultaneity concerns. In our main specifications, we focus on the RTI share among non-college workers, since we aim to isolate the impact of routinization on non-college employment opportunities. Appendix Table A.1 summarizes this RTI share measure, averaged over all commuting zones. The RTI share among non-college workers rises from 1960 to 1980, from an average of 15.4% across all commuting zones to 21.5%. It falls from 1980 on, reaching 13.6% in 2000. These trends are roughly consistent with the change in RTI share depicted in Figure 4.

We also use Autor and Dorn (2013)’s measures on the manual and abstract task content of occupations as control variables in our two stage least squares (2SLS) approach. Specifically, we control for *predicted manual and abstract occupation share*, which are constructed in parallel ways. For both measures, we interact a commuting zone’s 1950 industry composition with the share of occupations in the top third of manual or abstract content.

$$\text{manual occupation share}_{ct} = \sum_i E_{i,c,1950} \frac{\sum_k L_{i,k,1950} \mathbf{1}[\text{manual}_{k,1980} > \text{manual}_{1980}^{P66}]}{\sum_k L_{i,k,1950}}$$

$$\text{abstract occupation share}_{ct} = \sum_i E_{i,c,1950} \frac{\sum_k L_{i,k,1950} \mathbf{1}[\text{abstract}_{k,1980} > \text{abstract}_{1980}^{P66}]}{\sum_k L_{i,k,1950}}$$

where i indexes industry, c indexes commuting zone, and k indexes occupation. $E_{i,c,1950}$ is the share of industry i in commuting zone c in 1950. $L_{i,k,1950}$ is the number of workers in industry i , occupation k in 1950. We follow Autor and Dorn (2013) and define highly manual and highly abstract occupations based on the 1980 distribution, which was when RTI peaked in the census data. $\mathbf{1}[\text{manual}_{k,1980} > \text{manual}_{1980}^{P66}]$ equals 1 for occupations in the top third of manual content in 1980 and 0 otherwise. $\mathbf{1}[\text{abstract}_{k,1980} > \text{abstract}_{1980}^{P66}]$ equals 1 for occupations in the top third of abstract content in 1980 and 0 otherwise.

One of our alternate instruments, the routine share instrument, is adapted from the instrument in Autor and Dorn (2013). It is constructed in a parallel way to the manual and abstract occupation share measures above. Section B.1 elaborates.

A.3 Data from Atalay et al. (2020)

Three of our four instrumental variables come from Atalay et al. (2020). To extract occupational characteristics, Atalay et al. (2020) perform textual analysis on advertisements for job vacancies from *The Boston Globe*, *The New York Times*, and *The Wall Street Journal* from 1940 to 2000. For each occupation in each year, they characterize the work styles, knowl-

edge requirements, and task content desired by employers based on measures used in the literature. They compile one set of measures to match the information in the Occupational Information Network (O*NET), which describes the activities, tasks, and skills associated with thousands of jobs throughout the U.S. economy (see Occupational Information Network, 2022 and Hershbein and Kahn, 2018).

Using this set of measures, we construct our main instrumental variable, which predicts the administrative share in a commuting zone. We define administrative share as the proportion of jobs that are in the top third of administrative activity based on the 1950 occupational distribution, when administrative activity was at its highest. According to O*NET, the administrative activity consists of “performing day-to-day administrative tasks such as maintaining information files and processing paperwork” (O*NET Work Activity 4Ac1; Occupational Information Network, 2022). Occupations that involve high amounts of administrative activity include receptionists, information clerks, secretaries, and administrative assistants. Atalay et al. (2020) compile an occupation-level measure of administrative activity based on mentions per job posting, using keywords such as “filing,” “paperwork,” “administrative,” and “typing”. Summary statistics in Appendix Table A.1 show that the administrative share instrument exhibits a sizable decline over time, from 0.298 in 1960 to 0.0775 in 2000. This is consistent with the decline in RTI share due to routinization during this time period.

We also use the predicted administrative activity as a separate instrument. Rather than as a share, this instrument is measured as the frequency of keyword mentions per job posting. We construct the administrative activity instrument as follows:

$$\text{administrative activity}_{ct} = \sum_i E_{i,c,1950} \sum_k \text{admin}_{i,k,t}$$

where $\text{admin}_{i,k,t}$ represents the average number of keywords for administrative activity per job posting associated with occupation k in industry i at year t . $E_{i,c,1950}$ is the share of industry i in commuting zone c in 1950, based on Census data.

Our last instrument is constructed from Atalay et al. (2020)’s data on *clerical requirements*, which corresponds to whether an occupation requires “knowledge of administrative and clerical procedures and systems such as word processing, managing files and records, stenography and transcription, designing forms, and other office procedures and terminology” (O*NET Knowledge Requirement 2C1b; Occupational Information Network, 2022). Examples of occupations high in clerical requirements are word processors, typists, secretaries, administrative assistants, and office clerks. Atalay et al. (2020) classify a job ad as specifying clerical requirements if it includes words such as “clerical,” “secretarial,” “stenography,” or “typing”.²⁷ It is constructed in a parallel form to the administrative activity instrument:

²⁷The data set has a few other variables related to routine work, but they do not isolate routine tasks as cleanly as the administrative activity or clerical requirements variables. For example, Atalay et al. (2020) characterize occupations based on the task content classification of Spitz-Oener (2006). Specifically, Spitz-Oener (2006) found that *routine cognitive* tasks made an occupation more susceptible to automation, ceteris paribus. However, in the Atalay et al. (2020) data, an occupation’s routine cognitive task content depends on ad words such as “correcting,” “calculating,” “measuring,” “fixing,” and “rectifying,” which encompass a greater variety of tasks than those that were routinized.

$$\text{clerical requirements}_{ct} = \sum_i E_{i,c,1950} \sum_k \text{clerreq}_{i,k,t}$$

where $\text{clerreq}_{i,k,t}$ represents the average number of keywords for clerical requirements per job posting associated with occupation k in industry i at year t . $E_{i,c,1950}$ is the share of industry i in commuting zone c in 1950.

A.4 National Longitudinal Survey of Youth 1979

The National Longitudinal Survey of Youth 1979 (NLSY79) has tracked the same 12,686 participants since 1979 to the present. Initially, surveys were conducted annually up until 1994, after which they were conducted biennially. Our study focuses on the 11,155 participants who completed at least 12th grade or obtained a GED. We further refine our sample by excluding those who were employed but lacked wage information between the ages of 25 and 35, resulting in a sample size of 8,540. Lastly, we omit individuals missing ASVAB test scores or pertinent family background information. Our final sample comprises 2,505 men and 2,490 women. Appendix Table A.2 presents summary statistics for key variables in the model.

A.4.1 Measuring skill heterogeneity

We utilize the NLSY79’s Armed Services Vocational Aptitude Battery (ASVAB) test scores to create multi-dimensional skill profiles for individuals. In 1981, over 90% of NLSY79 respondents completed the ASVAB. The test consists of nine subtests: arithmetic reasoning, word knowledge, paragraph comprehension, mathematics knowledge, numerical operations, coding speed, automotive and shop information, electronics information, and mechanical comprehension. Some of these subtests contribute to the Armed Forces Qualification Test (AFQT) score, a widely used measure of cognitive ability.²⁸

Instead of using the AFQT scores directly, we apply exploratory factor analysis (EFA) to all nine subtests to construct multiple skill dimensions. EFA is often used to eliminate ambiguity in the number of latent factors and the underlying factor structure of a set of variables (Diekhoff, 1992). EFA allows us to leverage the correlation structure in scores across the nine ASVAB subtests when developing our skill measures. The analysis indicates that two distinct skills (“factors”) are necessary to explain the variation in ASVAB scores. Our EFA approach aligns with that of Prada and Urzúa (2017), who also found a two-factor structure to be most suitable for explaining the variance in ASVAB test scores for men.

Figure 6 displays the estimated factor loadings. For both men and women, the first factor has significant loadings on all subtests, with the highest values for arithmetic reasoning, word knowledge, mathematics knowledge, and paragraph comprehension. These subtests are designed to assess cognitive ability and form the primary components of the AFQT.

²⁸Different studies employ varying subtests to calculate AFQT scores, with arithmetic reasoning, paragraph comprehension, and word knowledge being the most common. However, mathematics knowledge, numerical operations, and coding speed have also been used (see, among many others, Neal and Johnson, 1996; Cameron and Heckman, 1998; Heckman and Cameron, 2001; Heckman et al., 2006).

Factor loadings for the second factor show gender differences. For men, loadings are statistically significant only for automotive and shop information, electronics information, and mechanical comprehension.²⁹ The United States Department of Defense designed these subtests to measure mechanical skill, as they evaluate the ability to solve simple mechanics problems and understand basic mechanics principles (Welsh et al., 1990). For women, loadings for the second factor are statistically significant only for coding speed and numerical operations. The Department of Defense classifies these subtests into the administrative qualification area, as they measure the ability to memorize letter strings or perform rapid arithmetic operations (ASVAB Prep Tests, 2022).

B Robustness Appendix

B.1 Two Stage Least Squares Approach: Additional Instruments

B.1.1 Routine share instrument

Our main instrument uses keywords from job postings. Because job postings are occupation-specific, the instrument may better detect routinization in certain occupations over others. To ensure our results are not driven by specific occupations, we construct the routine share instrument, which captures the extent of routinization across *all* occupations.

$$Z_{ct} = \left[\sum_{i=1}^I \gamma_{i,c,1950} \widehat{\text{routine}}_{i,-c,1950} \right] \times \text{year}$$

Here, i indexes industry, t indexes year, and c indexes commuting zone. The variable $\gamma_{i,c,1950}$ is the share of industry i in commuting zone c in 1950. The variable $\widehat{\text{routine}}_{i,-c,1950}$ is the share of high routine occupations in industry i in 1950 for all states except the state that holds commuting zone c . It is constructed based on

$$\widehat{\text{routine}}_{i,-c,1950} = \frac{\sum_{k \in i} L_{-c,k,1950} \mathbf{1}(\text{routine}_k > \text{routine}^{P66})}{\sum_{k \in i} L_{-c,k,1950}}$$

The indicator $\mathbf{1}(\text{routine}_k > \text{routine}^{P66})$ equals 1 if occupation k is in the top third of routine task content in 1980.³⁰

Since it is constructed from only 1950 characteristics, the routine share instrument pre-dates potential omitted variables that would influence both enrollment and employment decisions. Furthermore, the leave-one-out construction nets out contemporaneous local labor market shocks. We follow Autor and Dorn (2013) and interact it with a matrix of year dummies to nonparametrically estimate how the labor share of routine occupations in 1950 impacts RTI share in each year from 1960 to 2000. The intuition is that commuting zones with historically high routine labor shares will continue to have high RTI shares in 1960-2000, making them the areas where routinization actively occurred.

²⁹Factor loadings exceeding 0.3 are considered statistically significant (see Diekhoff, 1992; Sheskin, 2004).

³⁰Using the occupation distribution in other years does not alter the main results.

B.1.2 Administrative activity and clerical requirement instruments

We describe the construction of the administrative activity and clerical requirement instruments in Appendix A.3. They use similar identification assumptions. The intuition is that commuting zones with historically high shares of industries that are intensive in administrative activity or clerical requirements would experience greater routinization over time. The identifying assumption for these instruments is similar to the identifying assumption for the administrative share instrument. The administrative activity or clerical requirements in an occupation at the national level should only influence enrollment in ways captured by routinization at the commuting zone level. In other words, local omitted variables that influence both routinization and college enrollment should have negligible influence on the administrative activity or clerical requirements of an occupation at the national level.

B.2 Structural Model Identification

Our model identification strategy is based on the formal frameworks established in Carneiro et al. (2003) and Prada and Urzúa (2017). Here, we outline the main components of the strategy.

We start by identifying the loading factors exclusive to cognitive skill measures:

$$C_{j,i} = \lambda_j^c \theta_{c,i} + e_{j,i}^c, j = 1, 2, 3, 4$$

We normalize the loading associated with mathematics knowledge to 1 ($\lambda_2^c = 1$) to nonparametrically identify the other three loading factors $\{\lambda_1^c, \lambda_3^c, \lambda_4^c\}$. For example, $\lambda_1^c = \frac{Cov(C_j, C_1)}{Cov(C_j, C_2)} = \frac{\lambda_j^c \lambda_1^c var(\theta_c)}{\lambda_j^c \lambda_2^c var(\theta_c)} = \frac{\lambda_1^c}{\lambda_2^c}$ because λ_2^c has been normalized to be 1. We can then apply Klotarski's theorem to secure nonparametric identification of the distributions of θ_c and $e_{j,i}^c$, with $j = 1, 2, 3, 4$ (Carneiro et al., 2003).

Next, we identify the loading factors in the mechanical skill measures:

$$M_{j,i} = \lambda_j^c \theta_{c,i} + \lambda_j^m \theta_{m,i} + e_{j,i}^m, j = 5, 6, 7$$

We specify a linear correlation between $\theta_{c,i}$ and $\theta_{m,i}$:

$$\theta_{m,i} = \alpha_1 \theta_{c,i} + \theta_{1,i}$$

where θ_1 is an additional factor, assumed to be independent of θ_c . The above mechanical skill measure equation can be written as

$$\begin{aligned} M_{j,i} &= \lambda_j^c \theta_{c,i} + \lambda_j^m \theta_{m,i} + e_{j,i}^m \\ &= \lambda_j^c \theta_{c,i} + \lambda_j^m (\alpha_1 \theta_{c,i} + \theta_{1,i}) + e_{j,i}^m, j = 5, 6, 7 \\ &= \beta_j \theta_{c,i} + \lambda_j^m \theta_{1,i} + e_{j,i}^m \end{aligned}$$

where $\beta_j = \lambda_j^c + \lambda_j^m \alpha_1, j = 5, 6, 7$. Under this setup, we can decompose the identification strategy into three steps.

1. Once we identify the variance of cognitive skill $var(\theta_c)$ and the loading factors associated with the cognitive measures, we can recover β_j from $Cov(M_j, C_{j'}) = \lambda_j^c \beta_j var(\theta_c)$.

2. We normalize mathematics knowledge: $\lambda_7^m = 1$. This secures the identification of the other factor loadings λ_5^m and λ_6^m in the mechanical test score system: $\lambda_5^m = \frac{\text{cov}(M_5, M_6)}{\text{cov}(M_6, M_7)}$ and $\lambda_6^m = \frac{\text{cov}(M_5, M_6)}{\text{cov}(M_5, M_7)}$. We can then apply Klotarski's theorem to nonparametrically identify the distributions of θ_1 and $e_{j,i}^m$, with $j = 5, 6, 7$.
3. To identify α_1 , we assume the factor loading of cognitive skill on automotive shop information test is 0 ($\lambda_5^c = 0$). This implies that the cognitive factor θ_c affects the first mechanical test score M_5 only indirectly, through its correlation with the mechanical factor θ_m . We can then recover α_1 from the equation $\beta_5 = \lambda_5^m \alpha_1$.

Identification for the loading factors in the administrative skill equations follows a similar process. We first impose

$$\theta_{a,i} = \alpha_2 \theta_{c,i} + \theta_{2,i}$$

where θ_2 is an additional factor, assumed to be independent of θ_c . The administrative measure equations can be rewritten as follows:

$$\begin{aligned} A_{j,i} &= \lambda_j^c \theta_{c,i} + \lambda_j^a \theta_{a,i} + e_{j,i}^a \\ &= \lambda_j^c \theta_{c,i} + \lambda_j^a (\alpha_2 \theta_{c,i} + \theta_{2,i}) + e_{j,i}^a \quad j = 8, 9 \\ &= \gamma_j \theta_{c,i} + \lambda_j^a \theta_{2,i} + e_{j,i}^a \end{aligned}$$

where $\gamma_j = \lambda_j^c + \lambda_j^a \alpha_2$, $j = 8, 9$. Finally, we impose the normalization assumptions $\lambda_9^c = 0$, $\lambda_9^a = 1$, where $j = 9$ denotes the numerical operations subtest.

C Additional Tables and Figures

Table A.1: Summary Statistics, U.S. Census Data

	1960	1970	1980	1990	2000	All years
Female enrollment	0.217 (0.00306)	0.348 (0.00361)	0.407 (0.00348)	0.502 (0.00363)	0.529 (0.00341)	0.376 (0.00252)
Male enrollment	0.228 (0.00377)	0.381 (0.00404)	0.313 (0.00337)	0.388 (0.00390)	0.397 (0.00363)	0.305 (0.00214)
RTI share	0.154 (0.00136)	0.215 (0.00170)	0.179 (0.00153)	0.152 (0.00125)	0.136 (0.00117)	0.161 (0.000741)
Admin share IV	0.298 (0.00180)	0.189 (0.00133)	0.175 (0.00132)	0.180 (0.00105)	0.0775 (0.000452)	0.228 (0.00189)
Population	565149.2 (82541.9)	555278.7 (59610.9)	310933.0 (31270.0)	340498.1 (34956.9)	386447.3 (39302.9)	394666.5 (20208.0)
% female	0.502 (0.000450)	0.510 (0.000360)	0.511 (0.000382)	0.511 (0.000366)	0.506 (0.000387)	0.505 (0.000193)
% black	0.0842 (0.00497)	0.0801 (0.00425)	0.0760 (0.00431)	0.0769 (0.00430)	0.0815 (0.00445)	0.0808 (0.00187)
% Hispanic	0.0317 (0.00339)	0.0326 (0.00310)	0.0487 (0.00400)	0.0575 (0.00437)	0.0800 (0.00492)	0.0460 (0.00159)
% ages 18-25	0.0858 (0.000656)	0.114 (0.000679)	0.129 (0.000788)	0.0988 (0.000834)	0.0969 (0.000820)	0.105 (0.000359)
% ages 25-35	0.117 (0.000528)	0.113 (0.000392)	0.152 (0.000628)	0.156 (0.000593)	0.123 (0.000609)	0.135 (0.000349)
% ages 35-45	0.123 (0.000386)	0.107 (0.000288)	0.106 (0.000324)	0.143 (0.000469)	0.154 (0.000368)	0.128 (0.000312)
% ages 45-55	0.111 (0.000362)	0.108 (0.000286)	0.0961 (0.000265)	0.0997 (0.000305)	0.134 (0.000417)	0.109 (0.000237)
% ages 55-65	0.0864 (0.000539)	0.0942 (0.000395)	0.0958 (0.000454)	0.0892 (0.000382)	0.0925 (0.000421)	0.0903 (0.000204)
% ages 65 or older	0.0969 (0.000943)	0.111 (0.000934)	0.126 (0.00110)	0.143 (0.00110)	0.143 (0.00106)	0.117 (0.000530)

Summary statistics for U.S. census sample, 1960-2000. The sample is restricted to individuals who have finished high school or hold a GED. All summary statistics represent the average across commuting zones. Standard errors in parentheses.

Table A.2: Summary Statistics, NLSY79 Data

	Men		Women		Difference	
	Mean	Std. Dev.	Mean	Std. Dev.	Diff	P-value
College by age 25	0.485	0.500	0.609	0.488	-0.123	0.000
Cohort 1 (born 1957-1958)	0.267	0.442	0.254	0.435	0.013	0.300
Cohort 2 (born 1959-1960)	0.225	0.418	0.244	0.430	-0.019	0.105
Cohort 3 (born 1961-1962)	0.253	0.434	0.268	0.443	-0.015	0.222
Cohort 4 (born 1963-1964)	0.247	0.432	0.231	0.422	0.017	0.170
Father completed high school	0.269	0.443	0.269	0.444	-0.001	0.974
Mother completed high school	0.208	0.406	0.21	0.407	-0.003	0.831
Living in urban area at age 14	0.780	0.414	0.779	0.415	0.001	0.938
Living in the South at age 14	0.330	0.470	0.356	0.479	-0.027	0.045
Family income in 1979	11.31	0.935	11.31	0.895	-0.001	0.971
Number of siblings in 1979	3.40	2.394	3.51	2.442	-0.104	0.129
<i>Occupation choices between 25 to 35</i>						
White collar	0.074	0.262	0.441	0.497	-0.366	0.000
Blue collar	0.542	0.498	0.093	0.290	0.450	0.000
Pink collar	0.384	0.486	0.467	0.499	-0.083	0.000
Home staying	0.066	0.248	0.200	0.400	-0.134	0.000
<i>Average annual earnings between 25 to 35</i>						
White collar	23,579	15,904	15,233	8,969	8346	0.000
Blue collar	14,461	9,075	11,201	6,278	3260	0.000
Pink collar	11,138	7,694	8,119	5,319	3019	0.000

Summary statistics for the NLSY79 sample. The sample is restricted to individuals who have finished high school (12th grade) or hold a GED degree. Their occupation choice is defined as the modal occupation between ages 25 to 35. College by age 25 is a dummy variable that equals 1 if the individual's years of education exceeds 12 by age 25. The sample only includes individuals with complete family background information and test score information.

Table A.3: Estimates of Wage Coefficients by Occupation and Gender

	Men			Women		
	White	Blue	Pink	White	Blue	Pink
College	0.159 (0.001)	0.011 (0.001)	0.011 (0.002)	0.095 (0.001)	0.134 (0.003)	0.003 (0.002)
$\widehat{Routinization}$	0.207 (0.007)	-0.034 (0.006)	0.201 (0.018)	0.599 (0.007)	0.459 (0.018)	-1.005 (0.011)
Cognitive	0.086 (0.001)	-0.042 (0.001)	0.092 (0.002)	0.121 (0.001)	0.226 (0.003)	-0.047 (0.001)
Cognitive*college	0.061 (0.001)	-0.165 (0.001)	-0.037 (0.002)	0.009 (0.001)	-0.154 (0.003)	0.083 (0.001)
Cognitive* $\widehat{Routinization}$	-0.094 (0.019)	-0.187 (0.007)	-0.033 (0.021)	-0.722 (0.013)	0.393 (0.028)	0.011 (0.009)
Cognitive*college* $\widehat{Routinization}$	0.006 (0.021)	0.035 (0.014)	0.006 (0.038)	0.028 (0.014)	-0.002 (0.034)	0.009 (0.012)
Manual	-0.042 (0.001)	0.108 (0.001)	-0.069 (0.003)	-0.123 (0.002)	-0.109 (0.005)	-0.086 (0.002)
Manual*college	-0.001 (0.002)	0.123 (0.002)	0.010 (0.004)	0.055 (0.002)	-0.110 (0.006)	-0.289 (0.002)
Manual* $\widehat{Routinization}$	0.072 (0.025)	0.228 (0.008)	0.253 (0.030)	0.727 (0.028)	-0.308 (0.066)	-0.341 (0.018)
Manual*college* $\widehat{Routinization}$	0.015 (0.028)	0.004 (0.016)	-0.018 (0.056)	0.010 (0.029)	0.029 (0.083)	0.089 (0.025)
Admin	0.232 (0.005)	0.130 (0.002)	0.181 (0.007)	-0.401 (0.006)	-0.080 (0.015)	0.122 (0.005)
Admin*college	-0.085 (0.005)	0.063 (0.005)	-0.192 (0.010)	-0.197 (0.007)	0.075 (0.018)	-0.154 (0.008)
Admin* $\widehat{Routinization}$	-0.166 (0.071)	-0.088 (0.026)	-0.026 (0.105)	0.135 (0.085)	-0.261 (0.145)	-0.208 (0.053)
Admin*college* $\widehat{Routinization}$	-0.033 (0.071)	0.019 (0.043)	-0.001 (0.104)	-0.143 (0.087)	-0.065 (0.164)	-0.019 (0.055)
Constant	1.881 (0.001)	1.693 (0.001)	1.572 (0.002)	1.795 (0.001)	1.523 (0.003)	1.245 (0.001)
Standard deviation	0.455 (0.000)	0.446 (0.000)	0.402 (0.000)	0.398 (0.000)	0.419 (0.001)	0.437 (0.000)

Parameter estimates for the wage coefficients in Equation 4, reported by occupation and gender. Standard errors in parentheses.

Table A.4: Estimates for Utility Parameters by Occupation and Gender

	Men			Women		
	White	Blue	Pink	White	Blue	Pink
College	0.415 (0.009)	-0.878 (0.009)	0.354 (0.016)	0.832 (0.009)	-1.228 (0.020)	-0.475 (0.009)
$\widehat{Routinization}$	-0.139 (0.071)	-0.459 (0.053)	-0.173 (0.134)	13.473 (0.060)	6.359 (0.138)	-10.198 (0.058)
Cognitive	0.469 (0.010)	-0.038 (0.005)	0.261 (0.010)	0.927 (0.007)	0.464 (0.017)	0.392 (0.005)
Cognitive*college	0.453 (0.012)	0.755 (0.009)	0.447 (0.014)	0.303 (0.008)	-0.291 (0.022)	-0.319 (0.006)
Cognitive* $\widehat{Routinization}$	-4.243 (0.142)	-3.927 (0.049)	1.253 (0.135)	7.827 (0.092)	2.678 (0.164)	-9.299 (0.045)
Cognitive*college* $\widehat{Routinization}$	-0.030 (0.166)	0.002 (0.099)	0.013 (0.252)	-4.117 (0.102)	0.007 (0.212)	4.145 (0.060)
Manual	-0.001 (0.012)	0.339 (0.006)	-0.154 (0.015)	-0.827 (0.016)	-0.180 (0.028)	-0.632 (0.010)
Manual*college	-1.118 (0.014)	-0.435 (0.012)	-0.219 (0.021)	0.150 (0.018)	1.455 (0.036)	0.532 (0.013)
Manual* $\widehat{Routinization}$	13.205 (0.187)	-4.154 (0.064)	2.848 (0.204)	-1.334 (0.200)	-0.550 (0.360)	0.162 (0.100)
Manual*college* $\widehat{Routinization}$	-3.868 (0.222)	4.207 (0.119)	-0.380 (0.332)	0.011 (0.213)	0.026 (0.505)	0.000 (0.127)
Admin	0.327 (0.035)	-0.049 (0.020)	0.017 (0.037)	0.840 (0.042)	-0.169 (0.097)	-0.088 (0.030)
Admin*college	-0.576 (0.043)	0.131 (0.040)	-0.379 (0.063)	1.038 (0.050)	-0.792 (0.111)	0.223 (0.047)
Admin* $\widehat{Routinization}$	0.946 (0.516)	0.182 (0.208)	-0.979 (0.650)	4.441 (0.523)	0.831 (0.998)	-6.406 (0.301)
Admin*college* $\widehat{Routinization}$	-0.009 (0.509)	0.009 (0.279)	0.046 (0.633)	0.577 (0.516)	0.079 (0.925)	-0.899 (0.281)
Constant	-6.484 (0.006)	-4.313 (0.005)	-5.480 (0.010)	-6.733 (0.007)	-5.997 (0.017)	-4.150 (0.006)

Parameter estimates for the non-pecuniary utility coefficients in Equation 5, reported by occupation and gender. Standard errors are in parentheses.

Table A.5: Estimates for the Education Equation by Gender

	Men		Women	
	Estimate	Std. Error	Estimate	Std. Error
Cognitive	1.033	0.006	1.123	0.004
Manual	-0.541	0.006	-0.267	0.006
Admin	0.182	0.016	0.170	0.020
Cohort 2	-0.153	0.002	-0.154	0.002
Cohort 3	-0.003	0.002	0.111	0.002
Cohort 4	-0.020	0.002	0.260	0.002
Father's education	0.935	0.003	0.308	0.002
Mother's education	0.207	0.002	0.866	0.003
Urban	0.366	0.002	0.262	0.002
South	0.304	0.002	0.212	0.002
Intact family	0.507	0.001	0.156	0.001
Number of siblings	-0.033	0.000	0.003	0.000
Constant	-5.406	0.014	-1.652	0.012
Standard deviation	0.949	0.003	0.978	0.003

Parameter estimates for the education decision in Equation 6 are reported in columns (1) and (3) for men and women, respectively. Columns (2) and (4) report the associated standard errors.

Table A.6: Parameters for Skill Distributions and Measurement Equations

	Skill distribution		Measurement equation			
	Men	Women	Loadings		Std. Dev.	
	(1)	(2)	(3)		(4)	
μ_{cog}	-0.002 (0.0005)	0.076 (0.0005)	λ_1^m	1.516 (0.0005)	$\sigma_{c,1}$	0.458 (0.0002)
μ_{manual}	0.283 (0.0002)	-0.265 (0.0002)	λ_2^c	0.546 (0.0002)	$\sigma_{c,2}$	0.530 (0.0001)
μ_{admin}	-0.169 (0.0002)	0.165 (0.0002)	λ_2^m	0.919 (0.0004)	$\sigma_{c,3}$	0.532 (0.0001)
$\sigma_{cog}^{(1)}$	0.844 (0.0006)	0.836 (0.0006)	λ_3^c	0.472 (0.0002)	$\sigma_{c,4}$	0.446 (0.0001)
$\sigma_{manual}^{(1)}$	0.320 (0.0005)	0.365 (0.0002)	λ_4^c	1.037 (0.0002)	$\sigma_{m,5}$	0.480 (0.0001)
$\sigma_{admin}^{(1)}$	0.229 (0.0006)	0.123 (0.0010)	λ_6^c	0.914 (0.0002)	$\sigma_{m,6}$	0.587 (0.0001)
$\sigma_{cog}^{(2)}$	0.618 (0.0006)	0.286 (0.0012)	λ_7^c	0.895 (0.0002)	$\sigma_{m,7}$	0.637 (0.0001)
$\sigma_{manual}^{(2)}$	0.426 (0.0004)	0.066 (0.0007)	λ_8^c	0.738 (0.0003)	$\sigma_{a,8}$	0.699 (0.0002)
$\sigma_{admin}^{(2)}$	0.124 (0.0008)	0.129 (0.0009)	λ_9^a	1.122 (0.0021)	$\sigma_{a,9}$	0.947 (0.0003)

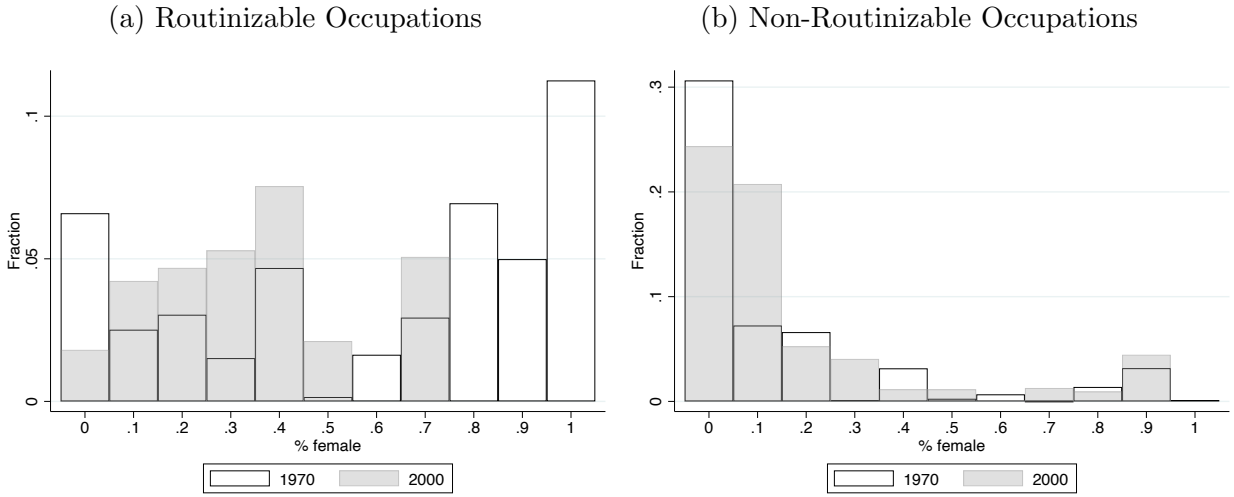
The left panel, “Skill distribution”, reports the distribution of skills by gender. Each skill is a mixture of two normal distributions. μ_{cog} denotes the mean of the first normal distribution for cognitive skill. The mean of the second normal distribution is pre-determined to be 0. $\sigma_{cog}^{(1)}$ reports the standard deviation of the first normal distribution for cognitive skill and $\sigma_{cog}^{(2)}$ reports the standard deviation of the second normal distribution for cognitive skill. μ_{manual} denotes the mean of the first normal distribution for manual skill. The mean of the second normal distribution is pre-determined to be 0. $\sigma_{manual}^{(1)}$ reports the standard deviation of the first normal distribution for manual skill and $\sigma_{manual}^{(2)}$ reports the standard deviation of the second normal distribution for manual skill. μ_{admin} denotes the mean of the first normal distribution for administrative skill. The mean of the second normal distribution is pre-determined to be 0. $\sigma_{admin}^{(1)}$ reports the standard deviation of the first normal distribution for administrative skill and $\sigma_{admin}^{(2)}$ reports the standard deviation of the second normal distribution for administrative skill. The right panel, “Measurement Equation” reports the estimates of the loading factors associated with Equation 9 in column (3). It reports the standard deviation of the residual term in each test score measurement equation in the column (4). Standard errors are in parentheses.

Table A.7: Goodness of Model Fit

	Women		Men	
	NLSY79 (1)	Sim (2)	NLSY79 (3)	Sim (4)
<i>Occupation choices</i>				
White collar	0.409	0.404	0.369	0.392
Blue collar	0.055	0.038	0.509	0.485
Pink collar	0.337	0.348	0.059	0.048
Not working	0.199	0.209	0.064	0.075
<i>Average log wages by occupation</i>				
White collar	1.907	1.890	2.069	2.092
Blue collar	1.631	1.592	1.779	1.779
Pink collar	1.416	1.422	1.570	1.570
<i>Education choices</i>				
High school	0.395	0.389	0.517	0.481
College	0.605	0.611	0.483	0.519

This table compares conditional moments from the model simulation with those from the NLSY79 data. Columns (1)-(2) compare moments for female workers and Columns (3)-(4) compare moments for male workers. The top panel displays occupation choices, the middle panel displays log average wages, and the bottom panel displays education choices.

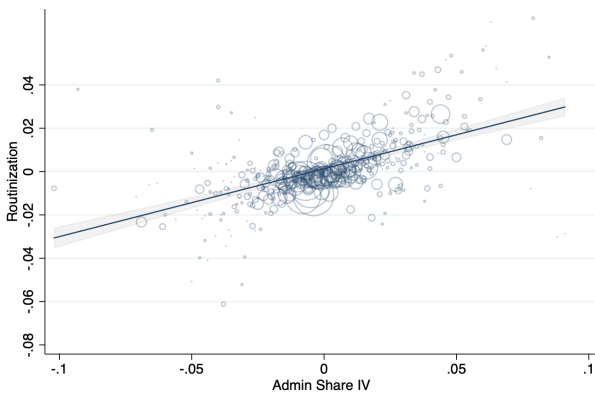
Figure A.1: Occupational Dispersion by Gender Composition, Non-College Occupations



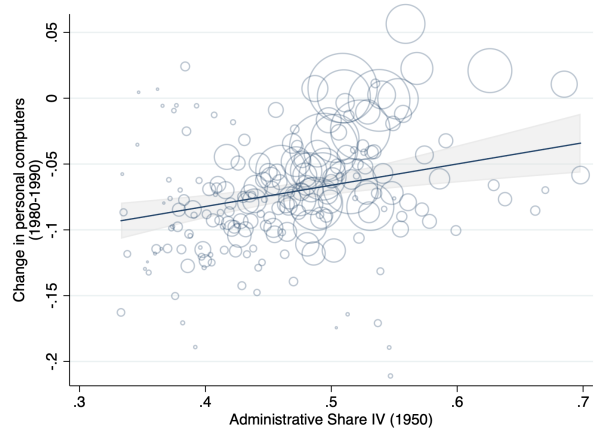
Distribution of occupations by proportion female in 1970 and 2000 for non-college occupations. Panel a shows routinizable occupations (top third of RTI), while panel b shows non-routinizable occupations (below the top 3rd of RTI). Individuals aged 18-30 years old. Data from the U.S. census and Autor and Dorn (2013).

Figure A.2: Assessing Administrative Share Instrument

(a) Instrument Predicts Routinization

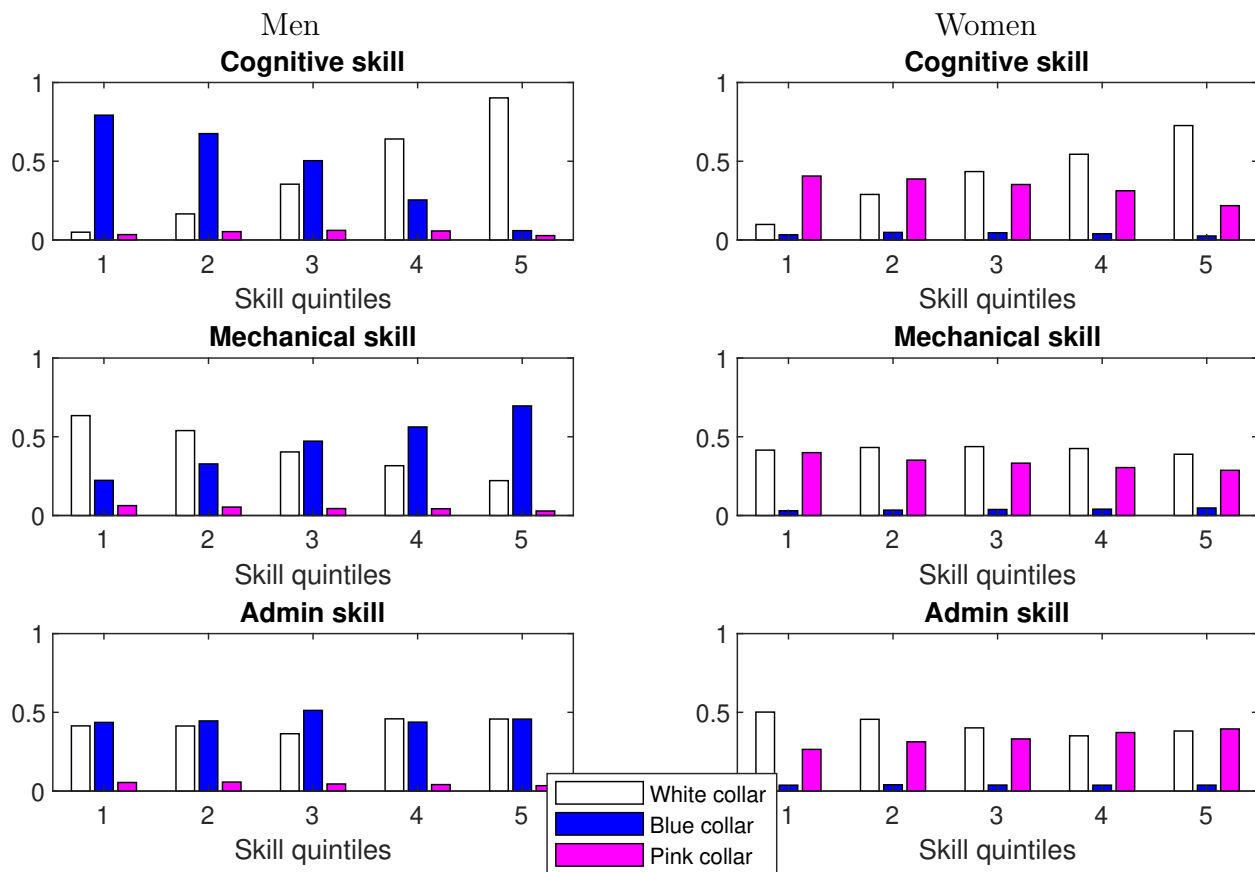


(b) Instrument Predicts Computer Adoption



Both panels assess the predictive power of the administrative share instrument. Panel (a) depicts the first stage prediction. It depicts the residual plot of routinization and the administrative share instrument after partialling out the controls in Table 3, column (4). Panel (b) plots the instrument in 1950 against the change in personal computers in 1980-1990. In both panels, the solid line represents the correlation estimated from an OLS regression using labor supply weights. The shaded gray area depicts 95% confidence intervals. Data from the U.S. census, Autor and Dorn (2013), and Atalay et al. (2020).

Figure A.3: Occupational Returns by Skill Quintile and Gender



We simulate each individual 200 times based on the estimates of the model to calculate average returns to each occupations by skill quintiles and gender. Returns include both the wage return and non-pecuniary returns. The upper panels present the effect of cognitive skill by gender, integrating out the effect of the other two dimensions of ability. The middle panel and the lower panel present analogous results for mechanical skill and administrative skill, respectively.