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

Amanda Chuan and Weilong Zhang

October 13, 2021

Click here for latest version

Abstract

Women used to lag behind but now exceed men in college enrollment. We argue that changes in non-college job prospects contributed to these trends. We first document that routine-biased technical change disproportionately displaced non-college occupations held by women. We next employ a shift-share instrument for the impact of routinization to show that these lower non-college job prospects for women increase female enrollment. Results show that a one percentage point decline in the share of routine task intensive jobs leads to a 0.6 percentage point rise in female college enrollment, while the effect for male enrollment is directionally smaller and insignificant. We next embed this instrumental variation into a dynamic model that links education and occupation choices. The model finds that routinization decreased returns in non-college occupations for women, leading them to shift to cognitive work and increasing their college premiums. In contrast, non-college men's occupations were less susceptible to routinization. Altogether, our model estimates that workplace routinization accounted for 67% of the growth in female enrollment and 31% of the change in male enrollment between 1980 to 2000.

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

JEL Classifications: I23, I24, I26, J16, J24, I26

"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 over time, social scientists predicted that the college gender gap would eventually close, and that men and women would attend college at roughly equivalent rates afterwards. Women indeed closed the gap in 1970-1980, as shown in Figure 1. Contrary to expectations, they then reversed the gap by attending college at increasingly higher rates relative to men. It remains an open puzzle as to why women exceed men in enrollment, especially when men tend to work longer hours and earn higher median salaries than women. A large literature documents that men face greater struggles in formal human capital investment due to differences in the non-cognitive skills critical for paying attention, staying disciplined, and persevering through school (Becker et al., 2010; Bertrand and Pan, 2013; Goldin et al., 2006). It posits a greater *supply* of women prepared for college than men.

In contrast, this paper argues that women have greater *demand* for a college degree than men, given differences in job prospects with only a high school diploma ("non-college job prospects").[1] We observe that the non-college labor market exhibits severe gender polarization, in that almost all occupations are male- or female-dominated, and few are gender-equal. From this observation emerge two stylized facts. The first is that non-college occupations dominated by women tend to pay less than those dominated by men. The second is that many female-dominated occupations disappeared from the non-college labor market between

[1] To focus on the role of non-college job opportunities, this paper abstracts away from the myriad other explanations that could also contribute to the college gender gap, such as the marriage market premium from a college degree (see Ge, 2011 and Zhang, 2021) and the "motherhood wall" in more demanding occupations (for a recent review, see Juhn and McCue, 2017).

1970-2000. Together, these facts suggest that outside options to college-going were worse for women, but deteriorated even further over time. Based on this evidence, we posit that men have a comparative advantage in non-college work, since they dominated high-paying occupations. Women's comparative disadvantage in non-college work led them to sort into lower paying occupations, which were more vulnerable to displacement over time.

To evaluate this argument, we attempt to isolate plausibly exogenous shifters of non-college job prospects by investigating why women's non-college jobs disappeared. We produce descriptive evidence identifying one key cause to be *routinization*, the displacement of routine-intensive occupations by automation. 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 note that routinization had especially severe impacts for the job prospects of women (Autor and Wasserman, 2013; Beaudry and Lewis, 2014; Black and Spitz-Oener, 2010). We further highlight that *non-college* women were the most vulnerable to displacement. In 1970, over 70% of non-college working women between 18 to 30 held occupations that were highly susceptible to routinization. When exploring the change in labor share from 1970 to 2000, we find that routinization lowered labor share *only* for non-college women, but not for college men, non-college men, or college women. Motivated by these data patterns, we use routinization as a shifter of non-college job opportunities to examine the impact of the non-college labor market on the reversal of the college gender gap.

Following Autor and Dorn (2013), we measure local susceptibility to routinization using *routine task intensive (RTI) share*, the share of occupations that involve many routine tasks relative to other tasks. We use instrumental variation in routinization to overcome two challenges with causal inference. One is that RTI share in a local labor market could depend on the share of college and non-college workers, leading to simultaneity concerns. Another is that RTI share and college enrollment rates could both 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. Both sources of endogeneity would bias our estimates of how routine non-college work opportunities impact college enrollment decisions.

Our instrument predicts a local labor market's routinizability using job posting data on administrative activity. The intuition is that labor markets with industries intensive in administrative activity possess more routine-intensive work. Cross-sectional variation stems from 1950 industry composition, which pre-dates any labor market and educational changes that occur during our analysis period of 1960-2000. Time-series variation stems from within-occupation changes at the national level, which should not depend on changes in any particular commuting zone. 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 changes in RTI share. We test our identifying assumptions using additional specifications, which verify that our results are not driven by 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, which exploit different sources of identifying variation to predict vulnerability to routinization.

Our first set of results come from the two stage least squares (2SLS) regressions. The first stage regressions indicate that labor markets with higher levels of routine work in 1950 would experience greater declines in administrative activity in 1960-2000 as routine intensive industries underwent automation. The second stage results demonstrate that declines in RTI share led to increased college enrollment among young women. We find that a 1 percentage point decline in RTI share corresponds to a 0.58-0.61 percentage point rise in female enrollment. Equivalently, moving from a commuting zone in the 25th percentile of routinizability to one in the 75th percentile in 1970 (a decline in RTI share of 5.51 percentage points) leads to a 3.20-3.36 percentage point rise in the proportion of women attending college. For men, whose non-college job prospects experienced less displacement, coefficient estimates are directionally smaller and not systematically significant. We thus use routinization to establish that the deteriorating availability of non-college jobs increased college enrollment. Women's non-college jobs declined more over this period, leading female enrollment to grow at a faster rate than male enrollment.

To investigate the mechanism behind our 2SLS results, we develop a two-period Roy model with unobserved skill heterogeneity. In the model, forward-looking individuals choose their education level in the first period and their occupation in the second period. We allow men and women to have heterogeneous, three-dimensional skill endowments (cognitive, manual and administrative), which are measured by the Armed Services Vocational Aptitude Battery (ASVAB) from the National Longitudinal Survey of Youth 1979 (NLSY79). We estimate the model using maximum likelihood.

Our model allows skill prices to vary across genders and occupations. Gender differences in skill endowments and skill prices create different comparative advantages for men and women, leading to gender polarization among non-college occupations. In the presence of this polarization, we predict changes in skill price to have uneven impacts on the non-college occupational returns of men versus women. To test this prediction, we obtain plausibly exogenous changes in skill price by instrumenting for routinization. In particular, we assume skill prices to be functions of predicted RTI share generated from the first stage of our 2SLS approach. Rather than assuming an ad-hoc linear mapping between college enrollment and RTI share, this specification posits an explicit mechanism for the second stage relationship between predicted RTI share and enrollment. By allowing individuals to have heterogeneous responses to local RTI share, our model generates more precise quantitative predictions compared with a simple back-of-the-envelope calculation based on our 2SLS estimates.

Our model explains the tight connection between the gender polarization of the non-college labor market and the reversal of the college gender gap. Men are more likely to sort into manual occupations given their higher mechanical skill, and women are more likely to sort into administrative occupations given their higher administrative skill. Since manual occupations pay more relative to administrative occupations, men enjoy a comparative advantage in non-college work overall. Women's comparative disadvantage, on the other hand, led them into administrative occupations which were more susceptible to displacement over time. As the labor market routinized, the price of administrative skill declined, impacting occupations predominantly held by non-college women. Skill returns for non-college men experienced smaller changes, since the occupations they held were harder to routinize. Consequently, routinization increased female college enrollment but had little impact on male college enrollment. Simulations from our model estimate that, as routine tasks became automated, the change in occupational returns would increase female enrollment by 6.0 percentage points and male enrollment by 0.6 percentage points. This accounts for 66.7% of the change in college enrollment for women, but only 31.6% of the change in college enrollment for men.

Contributions to the literature. To our knowledge, this is the first paper that uses automation as a source of variation to explain how the non-college labor market shaped the college gender gap over time. We exploit the impact of automation to isolate plausibly exogenous changes in labor demand for non-college workers in routine-intensive jobs. Prior work on the impact of labor market returns on the college gender gap have 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; Rendall, 2017). In contrast to these approaches, our instrumental variable 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 technological change by quantifying automation's impact on the rise of female college-going. To our knowledge, this is the first paper to evaluate 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. Specifically, we show that routinization, a gender-neutral process, generates gender-specific changes in college enrollment by influencing the relative price of skills in which men and women are differentially endowed. Our findings illuminate the role of technological change in shaping gender disparities in human capital over time. Our model simulations estimate that routinization contributed to two-thirds of the change in female enrollment from 1980 to 2000.

Third, our paper uses a model-based approach to link gender-based occupation polarization with trends in the college gender gap. Most prior papers use job task requirements to indirectly infer gender differences in skill levels (Duran-Franch, 2020; Ngai and Petrongolo, 2017; Olivetti and Petrongolo, 2014; Rendall, 2017; Yamaguchi, 2018). As a result, their models are limited in disentangling the gender difference in skill endowments from the gender difference in skill returns. We overcome this limitation by separately measuring skill endowments using the ASVAB in the NLSY79 and task requirements using job posting data. The closest frameworks to ours are Prada and Urzúa (2017) and Roys and Taber (2019), but our model deviates from them in two ways. We study both male and female workers and focus on gender inequality as it pertains to college enrollment choices, while the other two papers only analyze male workers. Furthermore, inspired by Eisenhauer, Heckman, and Vytlacil (2015) and Heckman et al. (2018), we introduce instrumental variation from routinization to shift the college premium. This helps us separately identify the college premium and occupational choice, which are usually jointly determined in a classical Roy model.

The paper is organized as follows. Section 2 describes stylized facts and data. Sections 3 and 4 describe our methodology and results from the 2SLS approach. Sections 5 and 6 describe our methodology and results from the structural model approach. We conclude in Section 7.

2 Data and stylized facts

In this section, we begin with an overview of our data. We then discuss the descriptive evidence that motivates our analytical approach. First, we present two stylized facts regarding the gender disparity 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

We start our analysis with data from the 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. Following Autor and Dorn (2013), we specify 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 are defined for only certain areas and therefore may under-represent certain industries (e.g., metropolitan statistical areas may underrepresent industries in rural areas such as agriculture or mining). Labor force measures, such as labor force participation or manufacturing employment share, are constructed from workers 16-65 years old, excluding residents of institutional group quarters and unpaid family workers.[2]

The dependent variable is the college enrollment rate among 18-25 year olds. Individuals are considered college enrollees if they have ever enrolled in college. Since our paper investi-

[2]Following 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.

gates the decision to attend college among those prepared for college, 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 the impact of non-college job prospects on the *choice* to pursue further education. College completion is influenced by a number of factors other than non-college job prospects, such as academic ability or perseverance, which complicate the task of isolating how non-college job opportunities change demand for a college degree.

To measure the impact of routinization, we use data from Autor and Dorn (2013). We focus on measures of routine occupation share and routine task intensive (RTI) share at the commuting zone level, described further in subsection 2.3.1. Our main instrumental variable comes from data on job postings from Atalay et al. (2020). We use the share of occupations high in administrative activity, where administrative activity measures are constructed based on job postings from *The Boston Globe*, *The New York Times*, and *The Wall Street Journal* from 1950 to 2000.

Our structural model uses individual level data from the geocoded National Longitudinal Survey of Youth 1979 Cohort (NLSY79). The NLSY79 interviews the same 12,686 respondents annually from 1979-1994 and every two years from 1996 until present day. We focus on a binary college attendance decision which equals 1 if years of education exceed 12 and 0 otherwise. We designate the individual's occupation choice to be the modal occupation between ages 25 to 35, and the occupation's monetary return as the individual's average annual earnings when she worked in this occupation. The final sample contains 8,540 individuals, with 4,217 men and 4,323 women. We provide further details and summary statistics in Appendix A.1.

Two advantages of the NLSY79 make it a good complement to the census data. First, the NLSY79 contains information on the respondent's county of residence at age 14 and traces each individual up to age 35, allowing us to account for potential composition effects due to migration. Second, the NLSY79 enables us to capture individual skill heterogeneity, as measured by test scores. Our primary skill measures come from the Armed Services Vocational Aptitude Battery (ASVAB), a set of tests designed by the United States Department of Defense to measure a wide array of skills. These individual-level ability measures shed light on why men and women may have comparative advantages in different occupations, which is the key mechanism behind the gender gap in college enrollment rates.

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 by gender and education. "Male-dominated" occupations comprise of less than 30% women, "female-dominated" occupations comprise of more than 70% women, and "gender-equal" occupations comprise of 30-70% women. "Non-college occupations" are occupations with at least 50% high school graduates, and "college occupations" comprise 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. As shown in Figure 2a, there is a "missing quadrant" in the non-college labor market. Plenty of male-dominated occupations pay above the median income of all workers in our sample (including college graduates), while female-dominated occupations pay below the 20th percentile. Occupations such as miner, machinist, and truck driver are over 90% male and earn between the 40th to the 80th percentile of annual earnings. In contrast, occupations that are over 90% female, such as cashier, housekeeper, and cosmetologist, earn below the 10th percentile of annual earnings. Based on this descriptive evidence, a typical male high school graduate still has the potential for high earnings, whereas his female counterpart appears less likely to sort into occupations with high earnings potential.

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 occupations that are female-dominated. 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. On the other hand, it would be expected for many women to hold low-paying college occupations if their non-college job prospects were not particularly lucrative.[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 severe gender polarization. One third (34%) of non-college occupations were female-dominated, over half (53%) were male-dominated, and only 13% were 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, as shown in Figure 3b. The share of gender-equal occupations rose from 17% to 50%, while the share of male-dominated occupations dropped from 72% to 21%. The share of female-dominated occupations rose from only 12% to 29%. The descriptive evidence suggests that as female non-college job opportunities were declining, women were entering college occupations that were formerly male-dominated. Over time, men and women appeared more substitutable in college work, but non-college occupations remained polarized by gender. Guided by this evidence, we designate college occupations as "white-collar", female-dominated non-college occupations as "pink-collar", and male-dominated non-college occupations as "blue-collar".

To investigate the disappearance of female-dominated occupations from the non-college labor market, we note that female-dominated occupations tended to be intensive in routine tasks. Pink-collar occupations such as secretary, clerical worker, stenographer, or typist involved a great deal of repetitive tasks, such as filling out administrative forms or typing strings of letters, which were easy to codify using automated devices. The most routine-intensive occupations exhibited the largest decline in labor share. For example, from 1970

[3]In Section 5, we will show this sorting mechanism can arise naturally given gender difference in comparative advantages of their non-college work.

to 2000, there was a 66% decline in the share of secretaries and a 95% decline in the share of typists in the census data. We next examine the role of routinization – the displacement of routine-intensive occupations by automation – on non-college job prospects over time.

2.3 Routinization and occupational composition

2.3.1 Measuring routinization

To show how the non-college labor market contributes to the college gender gap, we use plausibly exogenous changes in job opportunities that stem from automation. The main challenge is measurement, because automation is a gradual, continuous process that has existed since the dawn of the modern labor market, taking on forms that are often unknown, unrecorded, or unquantifiable.[4] Instead of investigating all forms of automation and their impacts on education, we narrowly focus on the role of automation in the *routinization of occupations* during 1960-2000.

We follow Autor and Dorn (2013) in using an occupation's "routine task intensity" (RTI) to measure its vulnerability to routinization. The RTI of an occupation k is calculated using the logged index of its routine, manual, and abstract tasks:

$$RTI_k = \ln(routine_k) - \ln(manual_k) - \ln(abstract_k)$$

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. "Routine" tasks are defined as codifiable tasks that can be executed following an explicit set of rules. As technology progressed, automat-

[4]Indeed, automation could affect college-going through other channels that are beyond the scope of this paper. Work on the emergence of robots since the 1990s shows that roboticization substituted for manually intensive work, disproportionately affecting the job prospects of men (Acemoglu et al., 2020; Acemoglu and Restrepo, 2019, 2020). Greenwood et al. (2005) posit that the impact of technological progress on household productivity is an important factor behind the increase in female labor force participation during the 20th century. Automation may also have direct effects on the education decision (e.g., computers facilitating learning and increasing college preparedness among students).

ing devices replaced human labor in executing these tasks, decreasing employer demand for workers that specialized in these tasks. Examples include electric typewriters and carbon paper obviating 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 tasks. Routinizability declines with manual job content, which involves the handling of objects across space, such as lifting materials or moving one's body around. It was challenging to program the devices available in 1960-2000 to perform such tasks, since navigating space in environments with other moving objects makes each task unique and hard to codify. For example, it was difficult to program a machine to wait tables at a restaurant, a highly manual task which required navigating around furniture and other moving bodies in unpredictable situations. Such technology only emerged after the 1990s (Acemoglu and Restrepo, 2020). Lastly, "abstract" tasks involve complex mental processes that are not easily programmable, such as problem solving, management, and complex communication. If two occupations have the same routine and manual job content, the one with greater abstract content would have lower predicted routinizability, since the execution of routine tasks would occur in conjunction with cognitively demanding tasks that could not be completed using automated devices. Prior work has also found that automation directly substituted for routine tasks while complementing abstract and manual tasks.[5]

To measure the impact of routinization at the commuting zone level, we construct the labor share of high RTI occupations ("RTI share"):

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

The term L_{ckt} is the total number of workers 16-64 years of age in commuting zone c, occupation k, and year t. An occupation k is designated high routine task intensity ("high

[5]Brynjolfsson 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).

RTI") if it exceeds the 66th percentile of routine task intensity for all occupations in 1980: $RTI_k > RTI^{P66}$.

2.3.2 The link between routinization, job polarization, and college enrollment

Prior work shows 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 young workers, whose college-going would be most directly impacted by job prospects, these changes are borne by women. Figure 4 panel A graphs standardized routine task intensity (RTI) across all jobs held by 18-30 year old men and women. While the RTI of women's jobs was consistently higher than men's, it declined substantially from 0.4 standard deviations above the average in 1960 to 0.2 standard deviations by 2000. In contrast, the RTI of men's jobs held steady at -0.2 standard deviations over the same time period.

Since RTI is comprised of three dimensions of task content (routine, abstract, and manual), we further investigate which one(s) drove the gender differences in trends over time. Figure 4 panels B and C shows that average routine content is the only task content measure to exhibit stark gender differences. Manual task content for men and women stayed constant during this period, while abstract task content grew by similar rates for both men and women's jobs. Similarly, Appendix Table A.4 shows that the raw correlation between female non-college work and routine task content was 18 times higher in 1970 than in 2000.

We then zero in on the types of jobs that are affected by exploring the labor share of occupations based on their routinizability. Panel A of Figure 5 graphs occupations by high and low RTI (top and bottom third, respectively). Consistent with prior literature documenting the vulnerability of routine occupations to automation, we find a decline in the labor share of high-RTI occupations but not low-RTI occupations. Among youth 18-30 years of age, the share of high-RTI occupations peaked at 40.0% in 1970 before declining to 33.9% by 2000. The share of low-RTI occupations, in contrast, rose from 28.9% to 36.5%

over this time period.

Panel B breaks this down by gender. Among young women, the share of high-RTI occupations peaked at 55.8% in 1970 and then plummeted by over 10 percentage points to 44.1% by 2000. The share of low-RTI occupations, on the other hand, grew from 22.5% in 1970 to 32.3% in 2000. The differential time trends for high-RTI (routinizable) compared to low-RTI (non-routinizable) occupations suggest that automation had a displacing impact on certain jobs held by women. Remarkably, these divergent trajectories are not observed for men. Among young men, the labor shares of high- and low-RTI jobs follow parallel trajectories: both grew about 3-5 percentage points from 1980 to 2000. Automation's displacement of high-RTI jobs appears to have largely affected the jobs held by young women, without noticeably affecting the aggregate labor share of young men.

The natural next question is whether this affected the college-going margin for women. Panel C of Figure 5 depicts labor share by RTI and educational status for women. For non-college women, there are stark differences in how labor share changed over time in high- versus low-RTI jobs. First, the share of young women who work in non-college low-RTI jobs is tiny and constant at 5.4-7.3% over the entire time period. The share in non-college high-RTI occupations in 1970 is far higher at 31.8% of all working women 18-30 years old. From 1970 on, however, the labor share plummeted from 31.8% to 14.1%. The decline by over half mirrors the decline in high-RTI labor share among all women in panel B, suggesting that automation's impact on women's jobs was concentrated in non-college jobs. Indeed, high-RTI college jobs did not appear to experience this same displacement. Panel C shows that for college women, the labor share of both high- and low-RTI jobs followed parallel trajectories, increasing by 10-12 percentage points from 1970 to 2000.

Together, Figures 4-5 indicate that the displacing impact of automation coincided with a decline in the routinizable jobs held by non-college women, but not college women, non-college men, or college men. The evidence suggests that the non-college occupations traditionally held by women were most vulnerable to displacement, lowering women's outside options to college-going. On the other hand, men's outside options were less vulnerable, making their college enrollment decisions less dependent on the labor market impacts of routinization. The descriptive evidence suggests that automation encouraged the growth of female enrollment through displacing women's non-college jobs. In the next section, we discuss our methodology to assess the causal impact of routinization on college enrollment.

3 Two Stage Least Squares Approach

Before presenting the two stage least squares (2SLS) results, we first report the OLS results of college enrollment regressed on the RTI share of non-college workers (Table 2). Across multiple specifications, we find that the relationship between RTI share and college enrollment is positive for men and insignificant for women. However, the OLS estimates do not necessarily isolate the causal impact of RTI share on college enrollment, since other explanations could contribute to our estimates. For example, if declines in parental income led fewer men to attend college, they may choose to work in manually-intensive jobs instead, leading both RTI share and college enrollment to fall. We therefore turn to a 2SLS approach to isolate plausibly exogenous variation in RTI share arising from routinization.

Our primary instrument predicts the share of highly administrative occupations in an industry. We calculate this "administrative share" using data from newspaper job postings. Atalay et al. (2020) extract information about occupational characteristics from job postings that appeared in *the Boston Globe*, *the New York Times*, and *the Wall Street Journal* in 1940-2000. Using their data, we construct an instrument based on the frequency with which job postings mentioned *administrative activity*, as measured by the occurrence of keywords such as "administrative", "paperwork", "filing", and "typing" for each decade from 1950 to 2000. We predict administrative share at the commuting zone level by interacting the frequency measure with the commuting zone's 1950 industry share:

$$\text{admin}_{ct} = \sum_{i=1}^{I} E_{i,c,1950} \frac{\sum_{k \in i} L_{kt} \mathbf{1}(\text{admin}_k > \text{admin}^{P66})}{\sum_{k \in i} L_{kt}}$$

where i indexes industry, k indexes occupation, t indexes year, and c indexes commuting zone. The variable L_{kt} represents the number of workers in occupation k in year t, while $E_{i,c,1950}$ represents the share of industry i in commuting zone c in 1950. The indicator $\mathbf{1}(\text{admin}_k > \text{admin}^{P66})$ equals 1 if occupation k is in the top third of administrative activity in 1980.[6]

Time-series variation is derived from within-occupation changes in administrative activity. Figure 6 shows a large dropoff in the share of highly administrative occupations from 1950 to 2000 for all major occupation categories, which is to be expected if automation decreased the need for employers to hire workers for routine tasks. The fall was especially severe for office-related occupations, as activities like typing another's handwritten notes declined with the advent of word processing software which allowed cognitive workers to type as they thought (Atalay et al., 2020). Among the remaining occupation groups, the decline in administrative activity over time was also pronounced among management-related occupations.

To predict routinizability at the commuting zone level, we fix industry shares in 1950 and interact them with the administrative share within each industry. The intuition behind the instrument is that commuting zones with high shares of routinizable industries should have experienced greater declines in administrative activity as these industries automated over time. In Appendix Figure A.1, we correlate our instrument with personal computer adoption to confirm that commuting zones with steeper declines in administrative activity did in fact experience more intensive automation in 1980-1990.

With this administrative share IV, we then perform the following two stage least squares regression. The first stage regression captures the relationship between RTI share and the

[6]Following Autor and Dorn (2013), we peg our measure of "high administrative activity" to the 1980 distribution to ensure that our definition of routinizable is constant over time.

administrative share within commuting zone c, year t:

$$\text{RTI share}_{ct} = \alpha_0 + \alpha_1 \text{admin}_{ct} + \alpha_2 W_{ct} + \theta_c + \phi_t + u_{ct} \tag{1}$$

We control for commuting zone-year level controls W_{ct}, commuting zone dummies, and year dummies. The matrix of control variables W_{ct} includes proportion female, black, and Hispanic. It also includes the proportion of people by 10-year age bin. We control for census division and year. For reasons described below, we control for labor force participation, manual share, and the 10-year lagged share of the service sector and major industries: manufacturing, retail, and mining. In some specifications, we control for the 10-year lag of RTI share and for the median log earnings of abstract-intensive work.

Our premise is that 18-25 year olds make their college-going decisions based on the job prospects of *others*. If younger women observe deteriorating job prospects for older women without a college degree, their beliefs about the returns to non-college occupations should correspondingly worsen. We therefore exclude 18-25 year olds in our measure of RTI share. The variable *RTI share* represents the share of workers in routine task intensive occupations among all workers between the ages of 25 to 65. Furthermore, if enrollment among 18-25 year olds rose for other reasons during this time (e.g., the rise in social norms favoring college attendance), fewer workers would take routine task intensive jobs, and RTI share would mechanically decline. This would overestate the true impact of RTI share on college enrollment. By excluding 18-25 year olds in our measure of RTI share, we address this potential source of endogeneity. Our preferred specification focuses on RTI share among only *non-college* occupations, but in robustness checks we use RTI share among both college and non-college workers.[7]

The second stage regression then uses the first stage linear prediction $\widehat{\text{RTI share}}_{ct}$ to

[7] To directly measure outside options to college-going, our preferred specification focuses on high-RTI non-college jobs as a share of all non-college jobs. However, this measure depends on the number of high school graduates, which is endogenous to supply-side considerations such as social norms regarding education or the ease of progressing through high school. In section 4.1, we apply our 2SLS specification to the RTI share among both college and non-college workers, which is less dependent on such concerns.

isolate the impact on college enrollment in commuting zone c, year t for gender g:

$$\text{college enrollment}^{g}_{ct} = \beta_0 + \beta_1 \widehat{\text{RTI share}}_{ct} + \beta_2 W_{ct} + \theta_c + \phi_t + \epsilon^{g}_{ct} \tag{2}$$

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. Note 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 how the instruments impact the RTI share for men and women.

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 the RTI share at the commuting zone level. This restriction is met if no commuting zone plays a large role in determining administrative share in an 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.[8]

The general threat to the exclusion restriction is that industry-level changes in routine activity, measured by administrative activity, 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. Below, we discuss plausible time-varying confounders that could generate the gender differences in college-going response we report in Section 4.

[8] We also use an alternate instrument with a leave-one-out specification, which nets out local labor market shocks that may be correlated with contemporaneous employment and education.

These confounders motivate the inclusion of certain controls into the W_{ct} matrix.

One possibility is that non-automation factors could drive industry level changes correlated with both enrollment and routine share in a commuting zone. For instance, the decline in manufacturing over this period could change both college enrollment and the proportion of routine 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.[9] We also control for lagged service sector shares, given Autor and Dorn (2013)'s finding that automation raised service sector employment. Using the lagged shares is preferable to current shares, since it nets out the effect of contemporaneous omitted variables related to both industry share and college enrollment.[10]

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. Non-college jobs for women may be especially plentiful, which would then increase 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 typically beyond college age, their labor force participation should not directly depend on the college enrollment of 18-25 year olds.[11]

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 high RTI jobs in 1950 may have more favorable social norms regarding women's schooling in 1960-2000. We control for lagged RTI share to capture the effects of

[9] A 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, which comprise a large share of the overall labor force.

[10] Excluding lagged industry and service sector shares does not change our point estimates (results available upon request). This is consistent with the rationale behind our two stage least squares approach, which is designed to isolate the variation in RTI share arising from time-series changes in administrative share that do not depend on industry composition in any commuting zone.

[11] As with lagged industry and service sector shares, specifications that do not control for labor force participation do not appreciably change our estimates.

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.

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).[12] In addition, 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 RTI share, 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 on 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.

We find that commuting zones with high historic shares of administrative industries experienced greater declines in RTI share over time. On average, a commuting zone with a

[12] The Kleibergen Paap F-statistic cannot formally test for weak instruments when errors are heterskedastic, serially correlated, or clustered (Pflueger and Wang, 2015). Another limitation is that the Kleibergen-Paap and Cragg-Donaldson F-statistics may be high even under weak instruments (Lee and Wolpin, 2006; Olea and Pflueger, 2013).

1 percentage point higher share of administrative industries in 1950 experienced a 0.38-0.39 percentage point greater decline in RTI share 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 are unlikely to explain these relationships. Across all specifications, Montiel Olea-Pflueger F-statistics hover at 41.44-49.93, indicating that the Nagar bias from the first stage regression is below 5% of the worst case benchmark.[13]

To assess fit, we plot the first stage linear prediction against the raw data in Figure 7. The raw data exhibit a clear negative relationship between RTI share and administrative share, corroborating the first stage regression results in Table 3. Both the raw data and the first stage predictions suggest that local labor markets with high levels of administrative activity in 1950 experienced greater routinization in later years, leading to sharper declines in RTI share.

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, which is consistent with the premise that labor markets more vulnerable to routinization experienced greater displacement of women's non-college job opportunities. Commuting zones with a 1 percentage point higher share of administrative activity in 1950 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. Panel C demonstrates that commuting zones which undergo more routinization experience higher

[13]The Nagar bias is the approximate asymptotic bias under weak instruments. The Montiel-Pflueger F-statistics enable us to test whether this bias exceeds a certain fraction of the "worst case" benchmark, where the instruments are uninformative and when the first- and second-stage errors are perfectly correlated (Olea and Pflueger, 2013; Pflueger and Wang, 2015).

female enrollment. Our estimates indicate that a 1 percentage point decline in RTI share leads 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, but are only marginally significant ($p < 0.10$). Overall, panels C and D indicate that moving from the 75th percentile to the 25th percentile of RTI share, about a 5.51 percentage point decline, increases the female enrollment rate by 3.18-3.33 percentage points and the male enrollment rate by 2.40-2.44 percentage points. We also estimate Anderson-Rubin weak instrument-robust 90% confidence intervals around our coefficient estimate, which 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.

In comparing across specifications, we find that including median earnings for cognitive occupations does not change our estimates. This is consistent with the evidence in Figure 4 that abstract task content changed at similar rates for both men and women, and therefore cannot explain the gender differential in college enrollment trends. Adding lagged routine share also does not change point estimates across specifications, indicating that persistence in occupational composition across time within commuting zones is unlikely to drive our results.

Overall, we find consistently negative impacts of RTI share on female enrollment, while any impacts on male enrollment are marginal at best. It is possible that the erosion of routine jobs also impacted male college-going. After all, some men worked in occupations that were vulnerable to automation. In addition to the few men who worked in secretarial and clerical occupations, high-RTI occupations that were dominated by men include shipping clerks, meter readers, security guards, machinists, and machinery repairers. Yet, even if men and women had equivalent responses to changes in RTI share, far more women worked in high-RTI jobs than men (around 70% of non-college women compared to 40% of non-college men during 1960-2000), so 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 6.

4.1 Additional specifications

We next address potential concerns regarding our main regression specification from Table 4 panels C and D.

Local shocks in Boston and New York. The content of job postings may be endogenous to the supply of skills in the local labor market. For example, if a commuting zone has a large share of college workers skilled in abstract tasks, employers may specify more abstract tasks and fewer routine or manual tasks in their job postings. The advantage of our approach is that we exploit trends in administrative activity over time in Boston and New York City, so local shocks from other commuting zones should not directly affect our job posting data. To ensure that local shocks in Boston and New York City are not driving our results, we exclude the commuting zones containing these two cities. The results are shown in column (1), Table 5. Our point estimates of -0.608 ($p < 0.01$) for female enrollment and -0.503 ($p < 0.10$) for male enrollment are similar to our main estimates.

Changes in abstract occupation share. In our main specifications, we instrument for the share of routine task intensive (RTI) occupations, which measure the routine content of an occupation relative to its manual and abstract content. We control for manual content, but allow routine and abstract content to vary freely since prior work has found that routinization coupled the decline in routine content with a rise in abstract content over time (see Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002). However, this raises the question of whether our results are driven by deteriorating job prospects in routine-intensive occupations or by improving job prospects in abstract-intensive occupations. While we already control for abstract median earnings in Table 4, we go further by controlling for abstract occupation share in column (2) of Table 5. 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. Despite the stringency of this assumption, estimates are similar to the main results, leading us to conclude that the response of female enrollment to changes in RTI share are not driven by improving returns to abstract-intensive occupations alone. We find point estimates of -0.628 ($p < 0.01$) for women in panel A and -0.441 ($p > 0.10$) for men in panel B.

Changes in the composition of non-college workers. Our sample period witnessed substantial growth in college enrollment due to many supply-side factors, such as greater high school completion rates, social norms encouraging college graduation, and more generous financing options for education. Since our main specification uses RTI share among non-college workers, the rise in college enrollment over this period could lower the denominator of our RTI share variable. This will lead to an underestimate of the true effect, since the covariance between RTI share and the instrument is inversely proportional to the second stage estimate. To address this concern, we instead use RTI share among both non-college and college workers, which is less sensitive to college enrollment changes. Column (3) reports our results, which do not quantitatively change from our main estimates for men or women (although the marginally significant estimate for male enrollment becomes insignificant). Comparing the two sets of results suggests that despite the decline in non-college worker share during our sample period, our 2SLS approach appears effective in netting out the role of supply-side changes on non-college RTI share.

The routine share instrument. Our administrative share instrument primarily exploits time-series variation in the fall in administrative activities. We next use an alternative instrument, modified from Autor and Dorn (2013), which exploits cross-sectional variation across commuting zones using a leave-one-out construction. Unlike the administrative share instrument, which predicts the fall in RTI share over time, this "routine share instrument" directly predicts future RTI share. It isolates the "long-run, quasi-fixed component of industrial structure that determines [a] commuting zone's...routine occupation share" (Autor and Dorn, 2013; see Appendix B.1 for details on instrument construction and identification assumptions). Despite leveraging a different source of variation from the administrative share instrument, point estimates are close to the main results for women (panel A column 4). A 1 percentage point decline in RTI share corresponds to a 0.57 percentage point decline in female enrollment ($p < 0.01$). In contrast, we find a small and insignificant estimate of -0.267 for men ($p > 0.10$, panel B column 4), which is slightly lower than the estimates of -0.436 to -0.444 in our main specification. Overall, our findings remain unchanged. Since the variation in the routine share instrument uses a leave-one-out specification to net out local labor market shocks and does not depend on job characteristics in any particular city, it provides an additional check that our main results are not driven by characteristics local to Boston or New York City.

The administrative activities instrument and clerical requirements instrument. Lastly, we look to alternate methods of using the job posting data to predict routinization. In column (5), we construct the "administrative activity instrument" using the predicted frequency of administrative activities, rather than administrative share. The units are the number of mentions of an administrative activity per job posting, rather than the share of occupations that require intensive amounts of administrative activity. In column (6), we use the "clerical requirements instrument", constructed from the number of times a clerical knowledge requirement is specified per job posting for an occupation. With both instruments, mentions were most frequent in office-related occupations, and declines were steepest in these occupations as they automated over time.[14] Estimated effects in specification (5) and (6) are comparable with the effects in our baseline results, although the results in column (6) are slightly higher than the main estimates. These comparisons indicate that our results are

[14]However, Appendix Figure A.2 shows that the two instruments follow different trajectories over time, since they draw on different keywords in the text of job postings. Keywords for administrative activities include "paperwork", "typing", and "filing". Keywords for clerical requirements include "clerical", "secretarial", "typing", "stenography", "word processing", or "dictaphone" (Atalay et al., 2020).

not contingent on the particular structure of our administrative share instrument. Rather, we arrive at the same results using multiple measures of routine work from the job posting data.

Overall, the results indicate that the negative relationship between RTI share and female enrollment is consistently significant across different forms of instrumental variation and different model specifications. In contrast, the impact on male enrollment appears weaker. Across all specifications in Table 5, Anderson-Rubin weak instrument robust confidence intervals are squarely negative for women but include 0 for men. We rule out a null effect in characterizing the relationship between routine intensive work and female college-going, but fail to reject the null hypothesis of no relationship between routine intensive work and male college-going. Taken together, the contrasting results for men and women align with the fact that female enrollment grew at a faster rate than male enrollment during 1950-2000, a period which experienced steady decline in the routine content of labor.

5 Structural Model Approach

Our 2SLS approach uses a local labor market approach to argue that declines in RTI share raised female enrollment. However, it is limited in evaluating the mechanism by which this occurs. We next use an augmented Roy model with latent skills to delve into how individual choices can change based on non-college job prospects. Following the dynamic discrete choice literature (Eisenhauer, Heckman, and Mosso, 2015; Keane and Wolpin, 1997; Roys and Taber, 2019; Todd and Zhang, 2020), we explicitly model sequential education and occupation decisions. Our innovation is that we incorporate instrumental variation in routinization into occupation-specific skill prices.[15] By doing so, we leverage instrumental variation to exogenously shift skill prices and identify the causal effects of routinization at well-defined margins of the education and occupation choices. These estimates is then used

[15] Eisenhauer, Heckman, and Vytlacil (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.

to simulate how male and female enrollment would change based only on changes in RTI share from 1980 to 2000, enabling us to quantify the importance of routinization in explaining why the reverse college gender gap grew so large.

The model has two periods with transitions and nodes shown in Figure 8. Individuals are forward looking and sequentially choose their education D_i in period 1 and their occupation O_i in period 2. The second period starts immediately after individual i finishes their education. 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), our model is then augmented with a set of test scores, which comprise the measurement system to identify workers' unobserved skills. We use $\theta_i = [\theta_{ci}, \theta_{mi}, \theta_{ai}]$ to represent a vector of three-dimensional skill sets for individual i, where subscripts c, m, and a are used to denote cognitive, mechanical, and administrative skills, respectively. We allow men and women to differ 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 the 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 the contrast between the college and non-college labor markets shown in Figure 3, where gender polarization is severe in non-college occupations but not in college occupations. Men and women appear to sort into similar jobs if they have a college degree, but different jobs if they only have high school diplomas. This classification enables our model to capture, for instance, the notion that blue collar jobs tend to be more brawn-intensive, leading to a comparative advantage for men due to their higher mechanical skill endowments. Lastly, we allow for home-staying as an outside option to working ($O_i = 4$).

Our specification is intentionally more parsimonious than typical life-cycle dynamic discrete choice models (Keane and Wolpin, 1997, 2001; Roys and Taber, 2019; Todd and Zhang, 2020). It assumes that attending college is the only binary education choice, that occupation choices are made once and permanent, and that individuals cannot return to school after entering the labor market. Our model is intentionally simple so as to hone in on the connection between college attendance decisions and the heterogeneous college wage premium across different occupations. This simplicity enables us to specify an explicit mechanism by which instrumental variation in RTI share shifts skill prices. As a result, the model deepens our understanding as to the mechanisms behind how routinization changes college-going decisions at the individual level.

5.1 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 $U(O_i|D_i)$:

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

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., job amenities, job flexibility, potential discrimination costs). The term $\epsilon_{O,D,i}$ is an idiosyncratic preference shock that follows the extreme value type I distribution.[16] Earnings in occupation O_i are expressed as

$$\log Y(O_i|D_i) = X_i^Y \beta^g_{O,X} + D_i \beta^g_{O,D} + \theta_i \beta^g_{O,\theta} + \theta_i D_i \beta^g_{O,D,\theta} + u^g_{O,i} \tag{4}$$

[16] Notice 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.

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 $D_i \beta^g_{O,D}$ and $\theta_i D_i \beta^g_{O,D,\theta}$, in which $\beta^g_{O,D}$ captures the common return to education while $\beta^g_{O,D,\theta}$ captures the component varying by the skill level θ_i. Lastly, $u^g_{O,i}$ 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

$$\log P(O_i|D_i) = X_i^Y \alpha^g_{O,X} + D_i \alpha^g_{O,D} + \theta_i \alpha^g_{O,\theta} + \theta_i D_i \alpha^g_{O,D,\theta} \tag{5}$$

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

In the first period, individual i decide college attendance depending on the perceived value of the flow utility and expected value from the second period.

$$\begin{aligned} D_i &= \mathbf{1}[V_i^1 + \xi^g_{D,i} > V_i^0] \\ V_i^0 &= E_\epsilon[U(O_i|D_i = 0)] \\ V_i^1 &= X_i^D \lambda^g_X + \theta_i \lambda^g_\theta + \rho E_\epsilon[U(O_i|D_i = 1)] \end{aligned} \tag{6}$$

where D_i denotes a binary variable equal to 1 if the individual chooses to attend college and 0 otherwise. X_i^D captures a vector of characteristics that are commonly believed as relevant factors for education choice.[17] The term $\theta_i \lambda^g_\theta$ captures the heterogeneous cost of attendance for individual i with skill θ_i and gender g.[18] The preference shock on education $\xi^g_{D,i}$ is assumed to be orthogonal to X_i^D and θ_i.

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

[18]For identification purposes, we normalize the flow utility of not attending college to 0.

5.2 Incorporating routinization

One of the biggest challenges in the generalized Roy model is the identification of skill prices, as they are endogenous outcomes jointly determined by supply and demand. Existing literature addresses this challenge by either using general equilibrium models (Lee and Wolpin, 2006) or assuming exogenous skill demand functions (Roys and Taber, 2019). Inspired by Heckman et al. (2018), we instead use the instrument for routinization defined in Section 3 to shift job prospects, specifically occupation-specific skill prices. In particular, changes in routinization impose different changes in skill returns based on pre-existing skill endowments, yielding different incentives to attend college. Therefore, rather than relying on variation across local labor markets alone, our model identifies the heterogeneous causal impact of routinization at the individual level.

We incorporate automation by specifying the skill prices as functions of the first stage prediction of local RTI share estimated from Equation 1. In particular, we assume that the vector of pecuniary and non-pecuniary returns to skills to be functions of the local RTI labor share $\widehat{\text{RTI share}}^{g}_{c(i),t}$:

$$\begin{aligned}
\beta^{g}_{O,\theta}(c,t) &= \beta^{g,0}_{O\theta} + \beta^{g,1}_{O,\theta}\widehat{\text{RTI share}}^{g}_{c(i),t} \\
\beta^{g}_{O,D,\theta}(c,t) &= \beta^{g,0}_{O,D,\theta} + \beta^{g,1}_{O,D,\theta}\widehat{\text{RTI share}}^{g}_{c(i),t} \\
\alpha^{g}_{O,\theta}(c,t) &= \alpha^{g,0}_{O,\theta} + \alpha^{g1}_{O,\theta}\widehat{\text{RTI share}}^{g}_{c(i),t} \\
\alpha^{g}_{O,D,\theta}(c,t) &= \alpha^{g,0}_{O,D,\theta} + \alpha^{g,1}_{O,D,\theta}\widehat{\text{RTI share}}^{g}_{c(i),t}
\end{aligned} \tag{7}$$

where $\widehat{\text{RTI share}}^{g}_{c(i),t}$ is the first stage predicted labor share for commuting zone $c(i)$, year t, gender g, which is defined in equation 1. Therefore, $\{\beta^{g}_{O,\theta}(c,t), \beta^{g}_{O,D,\theta}(c,t), \alpha^{g}_{O,D}(c,t), \alpha^{g}_{O,D,\theta}(c,t)\}$ is the skill price vector that individuals in commuting zone c would adopt when making their education choices at period t. Substituting Equation (7) into Equations (4) and (5) demonstrates that:

$$\log Y(O_i|D_i) = X_i^Y \beta^g_{O,X} + D_i \beta^g_{O,\theta} + \theta_i \beta^{g,0}_{O,\theta} + \theta_i D_i \beta^{g,0}_{O,D,\theta}$$
$$+\theta_i \beta^{g,1}_{O,D,\theta} \widehat{\text{RTI share}}^g_{c(i),t} + \theta_i D_i \beta^{g,1}_{O,D,\theta} \widehat{\text{RTI share}}^g_{c(i),t} + u^g_{O,i}$$

$$\log P(O_i|D_i) = X_i^Y \alpha^g_{O,X} + D_i \alpha^g_{O,\theta} + \theta_i \alpha^{g,0}_{O,\theta} + \theta_i D_i \alpha^{g,0}_{O,D,\theta} +$$
$$+\theta_i \alpha^{g,1}_{O,D,\theta} \widehat{\text{RTI share}}^g_{c(i),t} + \theta_i D_i \alpha^{g,1}_{O,D,\theta} \widehat{\text{RTI share}}^g_{c(i),t}$$

Based on the above equation, returns to different occupations depend on both individual characteristics (e.g., gender, education and skill levels) as well as predicted RTI share in the resident commuting zone. Therefore, even identical workers working in the same occupation may have different returns if they live in areas with differing vulnerability to routinization.

It is worth noting that we assume that routinization must only impact college-going in ways reflected by changes in skill prices. Our model effectively uses the skill price as a summary statistic to capture the returns to different skills within different occupations. This assumption is appropriate because our focus is on the impact of workplace routinization, as opposed to all forms of automation. It allows our model to abstract away from other forms of technological progress, which have been shown to influence work outcomes via channels other than occupational returns.[19]

5.3 Structural model estimation strategy

5.3.1 Latent abilities

We use the NLSY79's ASVAB tests to construct multi-dimensional skill profiles at the individual level. The ASVAB is comprised of nine subtests: arithmetic reasoning, word knowledge, paragraph comprehension, mathematics knowledge, numerical operations, coding speed, automotive and shop information, electronics information, and mechanical comprehension. Following Prada and Urzúa (2017), we perform Exploratory Factor Analysis (EFA)

[19]For example, Greenwood et al., 2005 demonstrate the role of household technologies in increasing female labor force participation.

analysis on the NLSY79's ASVAB tests to construct multi-dimensional skill profiles at the individual level. The analysis suggests that two separate skills ("factors") are necessary to 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, there are gender differences in factor loadings for the second factor. For men, the loadings are statistically significant only for the three mechanical subtests: automotive and shop information, electronics information, and mechanical comprehension. For women, loadings for the second factor are statistically significant only for the two administrative subtests: coding speed and numerical operations. Figure 9 displays the estimated factor loadings.

Based on our results, we characterize each individual's skill set θ_i by three dimensions: the common first factor as cognitive ability $\theta_{c,i}$, men's second factor as mechanical skill $\theta_{m,i}$, and women's second factor as administrative skill $\theta_{a,i}$.[20] This particular skill structure sheds light on how men and women can have different comparative advantages in different occupations, leading to the occupational sorting shown in Figure 2. Men tend to have higher mechanical skill, which would give them a comparative advantage in manually intensive tasks. Women tend to have higher administrative skills, which provide a comparative advantage in routine office work. Appendix A.4.1 provides more information on the EFA implementation.

Guided by the exploratory factor analysis, we specify the measurement equations for an individual i with latent skill vector $\theta_i = [\theta_{c,i}, \theta_{m,i}, \theta_{a,i}]$ as follows:

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

where $C_{j,i}$ denotes the four subtests exclusive for the cognitive ability measure, M_{ji} denotes the three mechanical subtests, and $A_{j,i}$ denotes the two administrative subtests.[21] We restrict

[20]Our EFA results match perfectly with Prada and Urzúa (2017) regarding the definition of mechanical skills for men. However, the results on administrative skill for women are novel.

[21]In particular, $C_{j,i}$ $\in${arithmetic reasoning, word knowledge, paragraph comprehension, mathematics

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 all error terms $\{e_{1,i}^c, ..., e_4^c, e_{5,i}^m, e_{6,i}^m, e_{7,i}^m, e_{8,i}^a, e_{9,i}^a\}$ to be 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.[22] For the purpose of identification, we follow Carneiro et al., 2003; Eisenhauer, Heckman, and Mosso, 2015; Heckman et al., 2006; Prada and Urzúa, 2017 to assume that at least one measure in $M_{j,i}$ is exclusively driven by mechanical skill, and one measure in $A_{j,i}$ is exclusively driven by administrative skill, and a set of standard normalizations.[23] We refer the interested readers to the aforementioned papers or Appendix B.2 for further details on identification.

5.3.2 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, 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 $A_{j,i}$. The likelihood

knowledge}, $M_{j,t}$ ∈{automotive and shop information, electronics information, and mechanical comprehension} and $A_{j,i}$ ∈{coding speed and numerical operation}.

[22] This assumption is consistent with recent papers using similar empirical findings (e.g. Eisenhauer, Heckman, and Mosso, 2015; Prada and Urzúa, 2017).

[23] In 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 operation (λ_9^a) are standardized to be 1.

function for individual i is given by

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

where $\Pr(.)$ represents the probability of occupation choice O_i or education choice D_i defined in equation 3 and equation 6, $f_j(.)$ is the probability density function for test j defined by equations 8, $f_Y(.)$ is the probability density function of earnings Y_i in equation 4, and $F_\theta()$ is the joint cumulative distribution of the latent skill vector $\theta \in \Theta$. After taking the logarithm of Equation (9) and summing across all individuals, we obtain the sample log likelihood $\log L$:

$$\log L = \sum_{i=1}^N \log \ell_i(\Omega_i|\psi)$$

Lastly, we need to impose some distributional assumptions to complete our likelihood function. 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 provides minimal restrictions on the underlying distributions of $[\theta_c, \theta_m, \theta_a]$.[24] Following Prada and Urzúa, 2017, we use mixtures of two normal distributions and assume $E[\theta_c] = E[\theta_m] = E[s] = 0$.[25] After plugging the distribution assumptions into Equation (9), $\Pr(O_i)$ will be a multinomial logit function and $\Pr(D_i)$ will be a probit function. We can then obtain the estimates $\hat{\psi}$ by maximizing the total likelihood function

[24] Ferguson (1983) argues 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.

[25] However, we allow mean values for men and women to differ and not equal 0.

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

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

6 Structural model results

6.1 Goodness of model fit

To assess model fit, we compare simulated occupation and education choices with those from the real data in Table 6. The upper panel shows that moments from the model simulation are close to real data on occupational choice. The simulation replicates that the two most common choices for men are white and blue collar occupations, while the two most common choices for women are white and pink collar occupations. The middle panel shows that for average log wages, simulated wages are reasonably close to actual wages. The average wage is highest in white collar occupations and lowest in pink collar occupations, both for men and for women. The lower panel summarizes education choices. Although our model slightly overpredicts the overall college attendance rate, it captures the pattern that women attend the college at a much greater rate than men. The fraction of women enrolled in college is around 60%, while the fraction of men is around a half.

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

Our model estimates reveal notable gender differences in skill profiles, depicted in Figure 10. First, Figure 10a demonstrates similar distributions of cognitive skill for men and women, although the variance is lower for women than men.[26] The similar distribution of cognitive

[26]This 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 end up 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

skills provides further evidence that men and women are substitutable in white collar work, and can explain why the majority of college occupations had similar proportions of men and women in 2000 (see Figure 3).

Gender differences in mechanical and administrative skill are substantial. Figure 10b shows that the mechanical skill distribution for men is higher in mean and variance than women, and that mechanical skills for women appear to max out near the male mean. In contrast, Figure 10c shows that women on average have higher administrative skills than men, given that the distribution for women falls to the right of the distribution for men. These differences in mechanical and administrative skill 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 men's higher mechanical skill or women's higher administrative skill.[27]

Aside from the difference in skill profiles, women and men may also receive different returns for the same skill in the same occupation. Figure 11 plots the returns to different occupations by skill quintiles for men and women (left and right panels, respectively). Blue bars denote returns from blue-collar occupations, pink bars denote returns from pink-collar occupations, and white bars denote returns from white-collar occupations. Comparing the blue bars between the left and right panels reveals that men receive higher returns from blue-collar jobs than women do, even among those with the same level of mechanical skill. On the other hand, the pink bars show that women receive much higher rewards from pink collar jobs than men do among those with the same level of administrative skill. Lastly, the white bars show that average returns for white-collar jobs are similar between men and women.[28]

also higher among women than men.

[27]Some studies in the literature also discuss the gender difference in social skill as one reason why men and women sort into different occupations (Black and Spitz-Oener, 2010; Borghans et al., 2014; Ngai and Petrongolo, 2017). Our measure of administrative skill overlaps with social skill, since administrative tasks often involve social elements. For example, secretaries and clerical workers often closely interacted with, planned for, and coordinated with coworkers and supervisors. Since our focus is on the routinization of occupations, we chose to focus on administrative skill rather than social skill.

[28]It is unclear why returns differ between men and women who possess the same skill in the same occupa-

Second, Figure 11 shows that different occupations reward different skills. Returns to blue-collar occupations tend to increase with mechanical skill for men, possibly because jobs such as HVAC engineer, material mover, or equipment repairer tend to be manually intense and therefore require a great degree of mechanical skill. Compensating wage differentials contribute to the high pay of these occupations, since they are manually challenging even if not cognitively intense. Returns to pink-collar occupations increase with administrative skill for women, possibly because office roles such as secretary or clerical worker reward the ability to file paperwork, coordinate others' schedules, and quickly enter strings of letters repeatedly into administrative forms. Returns to white-collar occupations increase with cognitive skill for both men and women, given that they tend to be intense in abstract tasks such as problem-solving, computation, and critical thinking.

Together, Figures 10 and 11 suggest that gender differences in skill endowments lead to comparative advantages at different occupations. This then creates gender differences in occupational sorting, as shown in Figure 12. 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. For women, high mechanical skill is positively associated with home-staying and negatively associated with white-collar occupations. Lastly, administrative skill is more relevant for women's occupation choices than men's. As administrative skill increases, the share of women entering pink-collar occupations grows while the share entering white-collar occupations declines. For men, on the other hand, administrative skill has little impact on

tion. We speculate that differences in skill returns could be due to occupational sorting. That is, controlling for mechanical skill, returns will be greater for men than women in blue-collar jobs since more men tend to sort into these jobs, making the non-pecuniary amenities of the job higher for men than women. For example, blue-collar jobs such as HVAC engineer, material mover, or equipment repairer have adapted to a majority male workforce, which may affect how comfortable women feel in these occupations regardless of mechanical ability.

the likelihood of sorting into any of the four occupational choices.

We then examine the correlation between skill endowment and college attendance in Figure 13. While cognitive skill predicts college-going for both men and women, it explains more of the variation in men's college-going. Women with low cognitive skill still attend college at high rates, while comparable men exhibit low attendance rates. The disparity is highest among individuals in the first and second quintiles of cognitive skill, where the proportion of women who attend college far exceeds that of men. This gap declines as cognitive skill increases, eventually flipping in favor of men for the highest quintile. The patterns are consistent with the idea that women have worse outside options to attending college than men. Men with low cognitive skills still have the option of entering blue-collar work, which can pay well, especially for men with high mechanical skills. Therefore, the compensation from attending college must be sufficiently high to warrant giving up the high pay from a blue-collar job. In other words, college is worthwhile only for 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 13 shows that enrollment declines with mechanical skill for both men and women, but that this decline is steeper for men. The evidence suggests that mechanical skill presents a sharper trade-off between college and non-college work for men than women. This interpretation is consistent with the result in Figure 11 that higher mechanical skill plays a larger role in whether men enter blue collar work, which presents especially lucrative outside options to attending college. Lastly, as administrative skill increases, female enrollment slightly declines but male enrollment does not change. It appears that high administrative skill presents some trade-off between college and non-college work for women, in that returns to pink collar work rise for women with high administrative skill. However, this trade-off is not nearly as stark as the trade-off that mechanical skill presents for men.

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 14. The non-college labor market exhibits severe gender polarization. Few non-college workers hold white collar occupations, given the complementarity between white collar occupations and college degrees. Instead, non-college men specialize in blue collar jobs given their higher mechanical skills, whereas non-college women specialize in pink collar jobs since they tend to have higher administrative skills. In contrast, the college labor market exhibits less gender polarization. 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 in the raw data which motivated our study to begin with, shown in Figure 3.

6.3 The effect of automation on occupation choice and college enrollment

In this subsection, we use our estimated model to quantitatively assess how much of the gender gap is attributable to changes in automation between 1980 to 2000. We incorporate local variation in predicted RTI share to assess the impact of routinization on individuals' college-going based on the commuting zone of residence. We then simulate the counterfactual trajectory of occupation choices for the 1979 cohort assuming that automation was the *only* change from 1980 to 2000 that impacted skill prices. All other primitive parameters, including the utility value for home-staying, are kept constant.[29]

Table 7 reports our simulated college enrollment rates between 1980 to 2000. Although automation increased the college attendance rate for both men and women, the growth rate for women is more substantial. Female enrollment grew from 63.2% to 69.2%, accounting for 66.7% of the observed 9 percentage point change between 1980 to 2000. In contrast, male enrollment grew only slightly from 49.6% to 50.2%, accounting for 31.6% of the observed

[29] However, labor force participation may still evolve over time if predicted RTI share changes the difference in utility from working versus not working.

2 percentage point change. At the same time, women shifted from pink- to white-collar occupations. The share of women working in white-collar jobs increased from 40.73% to 59.2%, while the share of women working in pink-collar jobs decreased from 32.8% to 15.1%. The simulated change in occupation shares is consistent with the empirical fact that many female-dominated occupations disappeared from the non-college labor market over time, highlighted in Figure 3. Lastly, our simulation predicts a slight rise in the proportion of people who stay at home. This finding suggests that automation decreased demand for certain workers, leading them to exit the labor force, and is consistent with Grigoli et al., 2020, which documents the negative effects of automation on labor force participation.

7 Conclusion

The college gender gap reversed in 1970-1980 when women exceeded men in college enrollment. This came as a surprise to social scientists, who anticipated that male and female enrollment would eventually converge. We argue that women have greater demand for a college degree due 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, with most occupations being male- or female-dominated and few occupations being gender-equal. 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 working young non-college women.

Informed by these stylized facts, we instrument for the share of jobs vulnerable to routinization in a local labor market. 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 declines in routine-intensive labor as workplaces automated. This decline led to significant enrollment growth among 18-25 year old women, but effects for men are directionally smaller and not systematically significant. We estimate that moving from a commuting zone in the 75th percentile of RTI share to one in the 25th percentile corresponds to a 3.34 rise in female enrollment. The corresponding estimate for male enrollment, 2.45 percentage points, is only marginally significant or insignificant, depending on the specification.

To investigate the mechanisms that explain these results at the individual level, we develop a two-period discrete choice model. Individuals choose whether to obtain a college degree in the first period based on occupational returns in the second period. 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 men to sort into manual-intensive work and women into routine-intensive work. The resulting gender polarization among non-college occupations translates to a comparative advantage for men in non-college work in general, 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. The model argues 1) that women's college premium increased relative to men's over time and 2) that the *efficient* college enrollment rate for women is higher than for men given men's comparative advantage in non-college work. Counterfactual simulations from the model reveal that routine-biased technical change accounts for 66.7% of the change in female college enrollment but only 31.6% of the change in male enrollment between 1980 to 2000.

Our paper has two implications. Given prior research showing that boys face greater struggles in school (Becker et al., 2010; Bertrand and Pan, 2013; Cappelen et al., 2019;

Goldin et al., 2006), popular media has framed the college gender gap as a problem rooted in men's "under-investment" in college. Our results indicate that men's relative "under-investment" is natural given that their job options are plentiful and lucrative even with only a high school diploma. Similarly, women's relative "over-investment" is a rational response to their bleak non-college job options. Given the gender-based sorting we document in the non-college labor market, we argue that it is efficient for a gender gap to exist.

Second, this paper speaks broadly to the role of technological transformation on the skill composition of the future workforce. We show that broad-based automation in 1960-2000 led to gender-asymmetric impacts on college enrollment decisions. Our argument can be applied to assess other forms of skill-biased technological change, such as roboticization during the 1990s and artificial intelligence during the 2000s (Acemoglu et al., 2020; Acemoglu and Restrepo, 2019, 2020; Hershbein and Kahn, 2018). Further work that examines these various forms of technological transformation will be valuable in explaining the changing demand for human capital investments for men versus women.

References

Acemoglu, D., and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *4*, 1043–1171.

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 us labor markets. *Journal of Political Economy*, *128*(6), 2188–2244.

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.

Asvab fact sheet. (n.d.). Retrieved February 17, 2021, from https://www.officialasvab.com/wp-content/uploads/2020/07/ASVAB-Fact_Sheet.pdf

Atalay, E., Phongthiengtham, 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., and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, *103*(5), 1553–97.

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., 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.

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.

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.

Cappelen, A. W., Falch, R., and Tungodden, B. (2019). The boy crisis: Experimental evidence on the acceptance of males falling behind. *NHH Dept. of Economics Discussion Paper*, (06).

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.

Charles, K. K., and Luoh, M.-C. (2003). Gender differences in completed schooling. *Review of Economics and statistics*, *85*(3), 559–577.

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: Univariante, bivariate, multivariante*.

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.

Ellwood, D., Kane, T. J. et al. (2000). Who is getting a college education? family background and the growing gaps in enrollment. *Securing the future: Investing in children from birth to college*, 283–324.

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.

Grigoli, F., Koczan, Z., and Topalova, P. (2020). Automation and labor force participation in advanced economies: Macro and micro evidence. *European Economic Review*, *126*, 103443.

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., and McCue, K. (2017). Specialization then and now: Marriage, children, and the gender earnings gap across cohorts. *Journal of Economic Perspectives*, *31*(1), 183–204.

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.

Kautz, T., and Heckman, J. J. (2014). *9. fostering and mea sur ing skills: Interventions that improve character and cognition*. University of Chicago Press.

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*.

Lee, D., and Wolpin, K. I. (2006). Intersectoral labor mobility and the growth of the service sector. *Econometrica*, *74*(1), 1–46.

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.

Network, O. I. (n.d.-a). *Knowledge: Clerical*. Retrieved September 14, 2021, from https://www.onetonline.org/find/descriptor/result/2.C.1.b

Network, O. I. (n.d.-b). *O*net resource center*. Retrieved September 14, 2021, from https://www.onetcenter.org/overview.html

Network, O. I. (n.d.-c). *Work activities: Performing administrative activities*. Retrieved September 14, 2021, from https://www.onetonline.org/find/descriptor/result/4.A.4.c.1?d=1

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.

Pflueger, C. E., and Wang, S. (2015). A robust test for weak instruments in stata. *The Stata Journal*, *15*(1), 216–225.

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.

Prep, A. (n.d.). *Air force and asvab*. Retrieved February 17, 2021, from http://asvabpreptests.com/index.php/asvab/airforce.html

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).

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

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

Welsh Jr, 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: Examples of Non-College Occupations, 2010

Occupation (1990 BLS classification)	% female workers	Earnings percentile
Cashiers	75%	4%
Housekeepers, maids, and cleaners	87%	7%
Hairdressers and cosmetologists	91%	11%
Miners	3%	82%
Machinists	4%	60%
Truck, delivery, and tractor drivers	7%	43%

Notes: Examples of "non-college occupations", where less than half of workers hold a college degree. Column (2) displays the proportion of workers that are women, and column (3) displays the earnings percentile among all workers with at least a high school degree.

Table 2: OLS Regression of College Enrollment on RTI Labor Share

	College Enrollment			
	(1)	(2)	(3)	(4)
	A: Women			
RTI Share	-0.416	-0.448	-0.431	-0.467
	(0.096)***	(0.098)***	(0.094)***	(0.095)***
Observations	3610	3610	3610	3610
	B: Men			
RTI Share	-0.215	-0.221	-0.235	-0.243
	(0.120)*	(0.122)*	(0.128)*	(0.130)*
Observations	3610	3610	3610	3610
Commuting zone FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Median cognitive earnings		Yes		Yes
Lagged RTI share				Yes

Notes: OLS regressions of enrollment on instruments at the commuting zone-year level. All regressions include demographic controls: proportion female, black, hispanic, and by 10-year age bin. All regressions also control for 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), (4), (6) and (8) add median annual log earnings for occupations in the top third of abstract-intensive tasks. Columns (3), (4), (7), and (8) additionally control for the 10-year lag of the share of high-RTI occupations. Standard errors clustered at commuting zone level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: First Stage Regression of RTI Share on Instruments

	RTI Share			
	(1)	(2)	(3)	(4)
Administrative Share IV	-0.387	-0.383	-0.388	-0.383
	(0.060)***	(0.056)***	(0.059)***	(0.054)***
F-statistic	41.441	47.387	43.035	49.933
Observations	3610	3610	3610	3610
Commuting zone FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Median cognitive earnings		Yes		Yes
Lagged RTI share			Yes	Yes

Notes: First stage regression of RTI share on instruments. All regressions include demographic controls: proportion female, black, hispanic, and by 10-year age bin. All regressions also control for 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)
	A: Reduced Form Regression, Women			
Administrative Share IV	0.224	0.232	0.224	0.232
	(0.068)***	(0.065)***	(0.068)***	(0.065)***
Observations	3610	3610	3610	3610
	B: Reduced Form Regression, Men			
Administrative Share IV	0.169	0.170	0.169	0.170
	(0.100)*	(0.099)*	(0.100)*	(0.100)*
Observations	3610	3610	3610	3610
	C: Second Stage Regression, Women			
RTI Share	-0.578	-0.606	-0.578	-0.606
	(0.202)***	(0.186)***	(0.201)***	(0.184)***
	[-0.974,-0.183]	[-0.970,-0.242]	[-0.971,-0.185]	[-0.967,-0.246]
F-statistic	41.441	47.387	43.035	49.933
Observations	3610	3610	3610	3610
	D: Second Stage Regression, Men			
RTI Share	-0.436	-0.444	-0.436	-0.444
	(0.263)*	(0.262)*	(0.265)	(0.264)*
	[-0.952,0.080]	[-0.958,0.070]	[-0.955,0.084]	[-0.962,0.073]
F-statistic	41.441	47.387	43.035	49.933
Observations	3610	3610	3610	3610
Commuting zone FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Median cognitive earnings		Yes		Yes
Lagged RTI share			Yes	Yes

Notes: 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: proportion female, black, hispanic, and by 10-year age bin. All regressions also control for 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Second Stage Regressions, Additional Specifications

	(1)	(2)	(3)	(4)	(5)	
	A: Second Stage Regression of Female Enrollment on RTI Share					
RTI Share	-0.607	-0.628	-0.736	-0.574	-0.548	-0.784
	(0.185)***	(0.205)***	(0.307)**	(0.184)***	(0.184)***	(0.184)***
	[-0.970,-0.244]	[-1.030,-0.227]	[-1.339,-0.134]	[-0.934,-0.214]	[-0.955,-0.141]	[-1.310,-0.259]
	B: Second Stage Regression of Male Enrollment on RTI Share					
RTI Share	-0.503	-0.441	-0.540	-0.267	-0.432	-0.616
	(0.285)*	(0.281)	(0.417)	(0.263)	(0.264)	(0.264)*
	[-1.060,0.055]	[-0.991,0.109]	[-1.356,0.277]	[-0.783,0.249]	[-1.046,0.183]	[-1.344,0.111]
Observations	3600	3610	3610	3610	3610	3610
First Stage F-statistic	256.985	40.740	45.288	42.788	34.246	25.202
Excluding Boston and NYC	Yes					
Control for abstract occupation share		Yes				
RTI share: non-college workers	Yes	Yes		Yes	Yes	Yes
RTI share: college and non-college workers			Yes			
IV: Administrative Share	Yes	Yes	Yes			
IV: Routine Share				Yes		
IV: Administrative Activities					Yes	
IV: Clerical Requirements						Yes

Notes: Two stage least squares regressions, additional specifications. All regressions also control for census division, year, commuting zone, labor force participation, manual occupation share, median cognitive earnings, lagged RTI share, and lagged major industry shares: services, manufacturing, retail, and mining. Column (1) excludes commuting zones that contain Boston and New York City. Column (2) controls for abstract occupation share. Column (3) uses the RTI share of all workers, instead of the RTI share of only non-college workers used in the main specification. Column (4) uses the gender-specific non-college RTI share, rather than pooling men and women. Columns (1)-(4) use the administrative share IV, while column (5) uses the routine share IV, column (6) 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-Pflueger 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: Goodness of model fit

	Women		Men	
	Data	Sim	Data	Sim
Occupation choices				
White collar	0.409	0.407	0.369	0.384
Blue collar	0.055	0.056	0.509	0.497
Pink collar	0.337	0.328	0.059	0.050
Not working	0.199	0.209	0.064	0.069
Average log wages by occupation				
White collar	1.907	1.956	2.069	2.110
Blue collar	1.631	1.622	1.779	1.801
Pink collar	1.416	1.444	1.570	1.571
Education choices				
High school	0.395	0.368	0.517	0.504
College	0.605	0.632	0.483	0.496

Notes: This table compares conditional moments from the model simulation with those from the real data. The first two columns compare moments for female workers and the last two 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.

Table 7: Simulated changes in occupation choices and education choices due to the routinization trend

	Women			Men		
Year	1980	1990	2000	1980	1990	2000
Occupation choices						
White collar	0.407	0.542	0.592	0.384	0.396	0.402
Blue collar	0.056	0.063	0.064	0.497	0.489	0.485
Pink collar	0.328	0.194	0.151	0.050	0.050	0.050
Not working	0.209	0.201	0.193	0.069	0.065	0.063
Education choices						
High school	0.368	0.322	0.308	0.504	0.500	0.498
College	0.632	0.678	0.692	0.496	0.500	0.502

Notes: This table simulates education and occupation choices for the NLSY79 cohort using predicted RTI share. The first three columns provide simulated choices for women, while the last three columns provide simulated choices for men. The "1980" columns report the baseline choices for both genders. The "1990" columns report counterfactual choices based on predicted RTI share in 1990. The "2000" columns report counterfactual choices based predicted RTI in 2000.

Figure 1: College Enrollment by Gender, 1950-2000

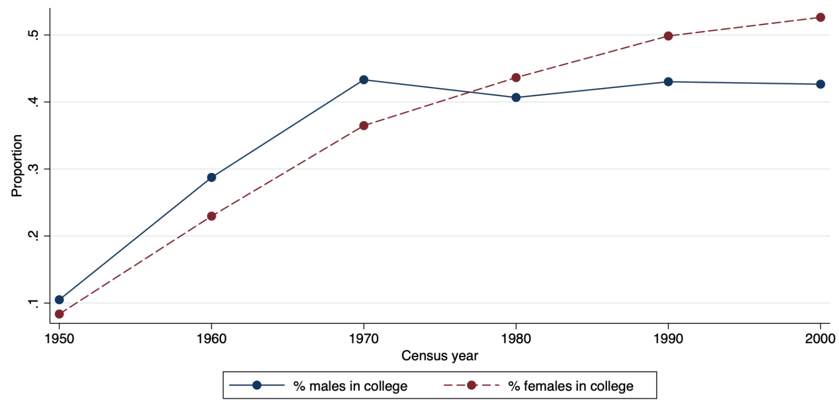

Notes: Proportion of 18-25 year olds ever enrolled in college. Solid lines represent male enrollment, while dashed lines represent female enrollment. Data from Census (1950-2000).

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

(a) Non-College Labor Market

(b) College Labor Market

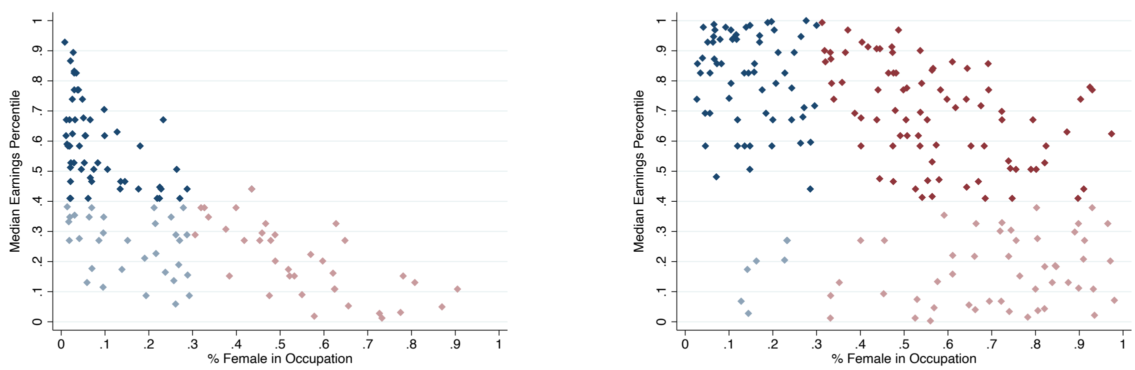

Notes: Occupations by proportion female and median annual earnings percentile in 2000. Panel A depicts occupations with 50% or fewer college graduates. Panel B depicts occupations with 50% or more college graduates. Data from Census.

Figure 3: Occupational Dispersion by Gender Composition

(a) Non-College Labor Market

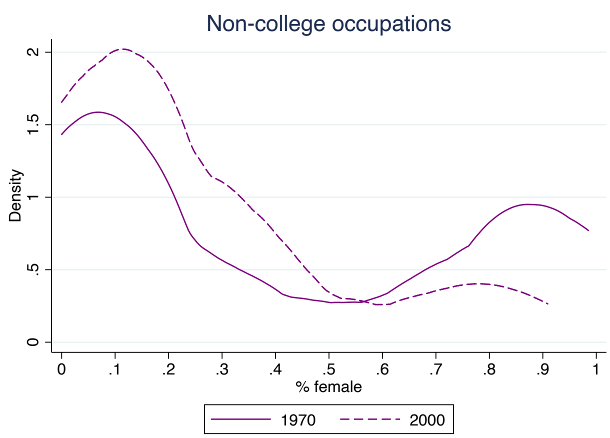

(b) College Labor Market

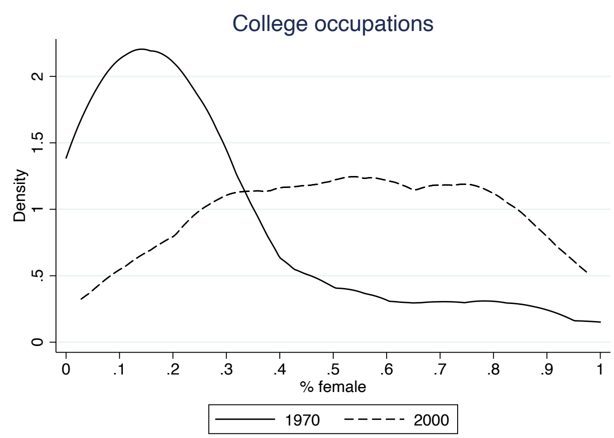

Notes: Distribution of occupations by proportion female in 1970 and 2000 for "non-college" occupations (A) and "college" occupations (B). Individuals aged 18-30. Data from Census.

Figure 4: Routine Task Intensity (RTI)

(a) RTI by Gender

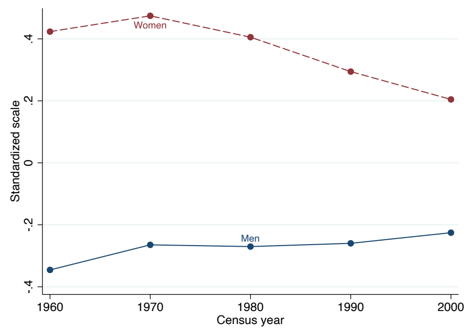

(b) RTI and Task Content, Women

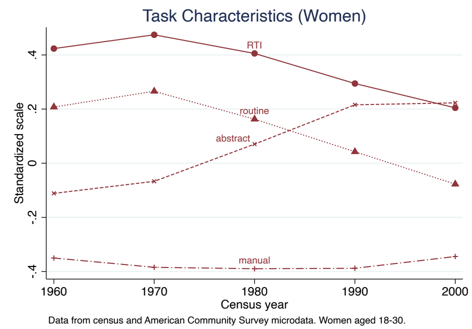

(c) RTI and Task Content, Men

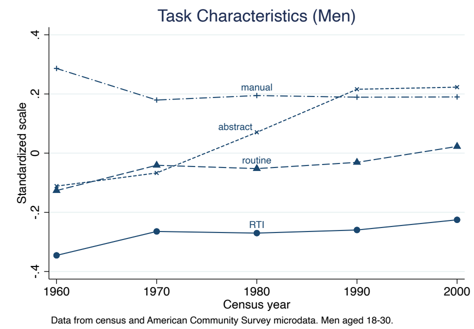

Notes: Panel A plots routine task intensity (RTI) of men's and women's work over time among 18-30 year olds. RTI is then broken up into its component task content measures in panels B (women) and C (men). All variables are standardized. Data from Census and Autor and Dorn (2013).

Figure 5: Labor Share by High vs. Low Routine Task Intensity (RTI)

(a) All

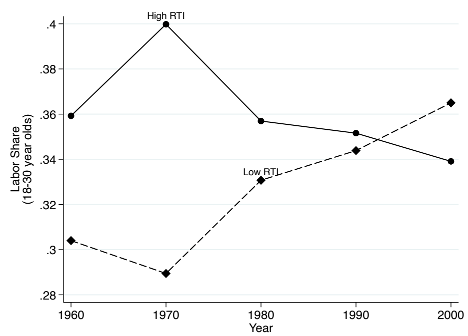

(b) By Gender

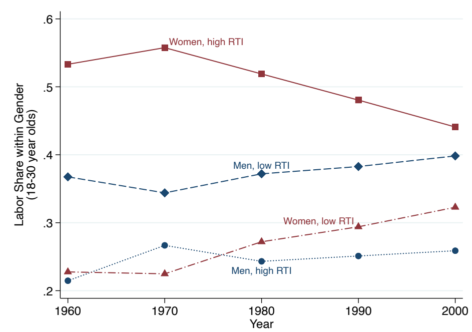

(c) By Education, Women

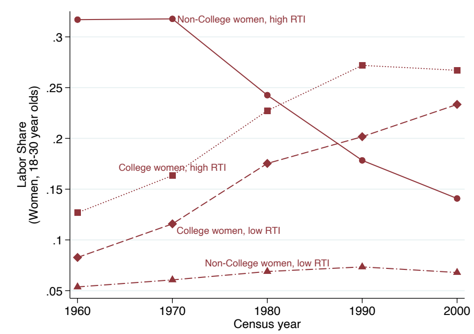

Notes: Labor share by high vs low susceptibility to automation, as measured by RTI. Panel A depicts labor share by RTI among 18-30 year olds. Panel B plots the labor share by RTI among young women (red) and the labor share by RTI among young men (blue). Panel C plots the labor share among young women by RTI and education. Data from Census and Autor and Dorn (2013).

Figure 6: Share of high administrative occupations by major occupation group

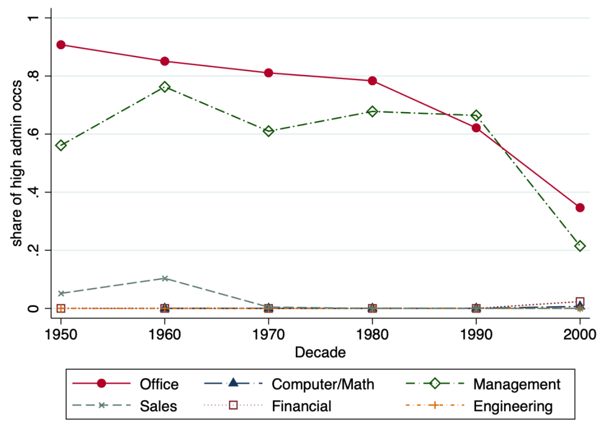

Notes: Share of occupations in top third of administrative activity for the six major occupation groups. Data from Atalay et al. (2020).

Figure 7: First Stage Prediction between RTI Share and Administrative Share Instruments

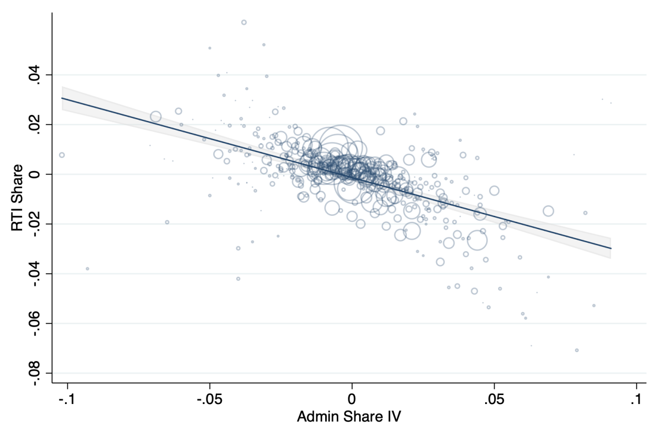

Notes: First stage prediction. The figure depicts the residual plot of high RTI labor and predicted automation susceptibility after partialling out the controls in Table 3 column (4). In predicting automation susceptibility, panel A uses the predicted share of occupations with high administrative activity as the instrument. The solid line shows the correlation estimated from an OLS regression using labor supply weights. The shaded gray area depicts 95% confidence intervals. Data from Census, Autor and Dorn (2013), and Atalay et al. (2020).

Figure 8: Two period dynamic discrete choice model

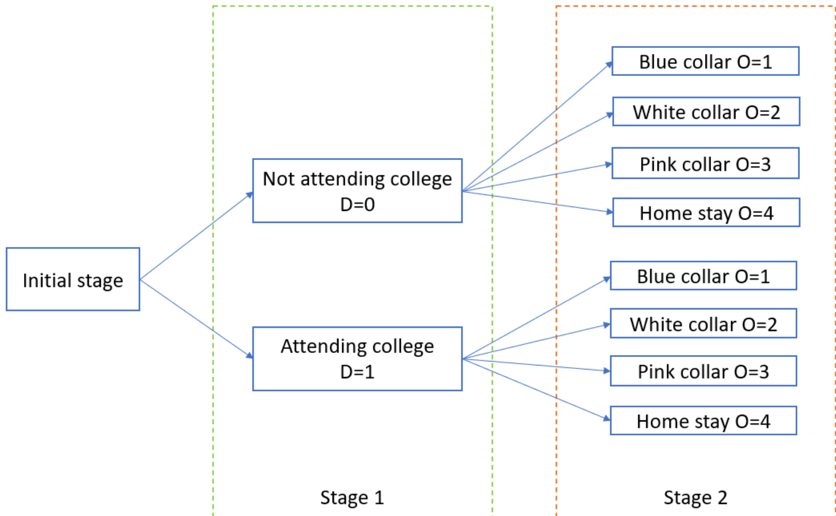

Notes: Description of structural discrete choice model. In Stage 1, individuals choose their college attendance decision. In Stage 2, they choose their occupation from four choices: blue collar, white collar, pink collar, or home staying. The model is solved via backwards induction.

Figure 9: Factor Analysis Loadings

(a) Male (b) Female

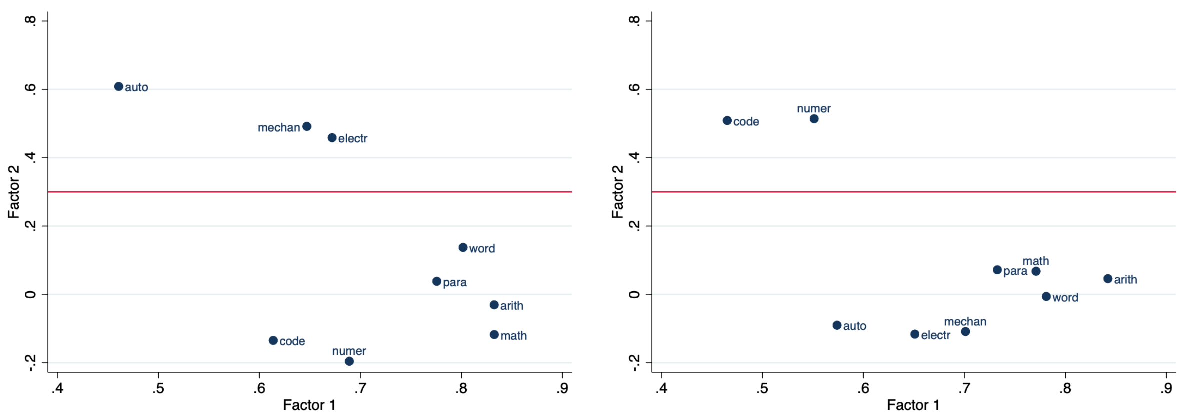

Note: Loadings calculated from exploratory factor analysis (quartimax rotation). The red horizontal line marks the statistically significant threshold (see Diekhoff, 1992 and Sheskin, 2004). arith = arithmetic reasoning; auto= automotive information and shop information; code = coding speed; electr =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

Figure 10: Distribution of skills by gender

(a) Cognitive Skills

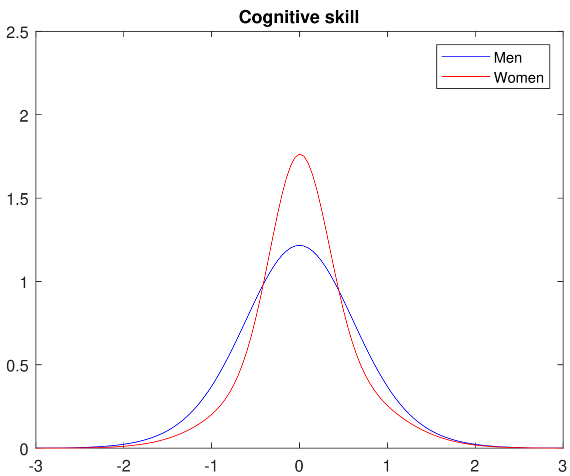

(b) Mechanical Skills

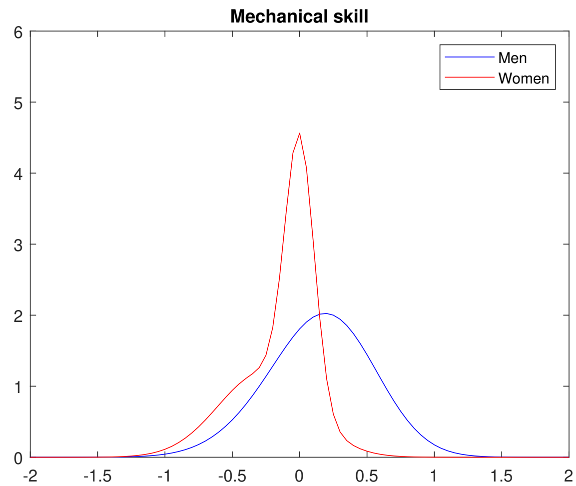

(c) Administrative Skills

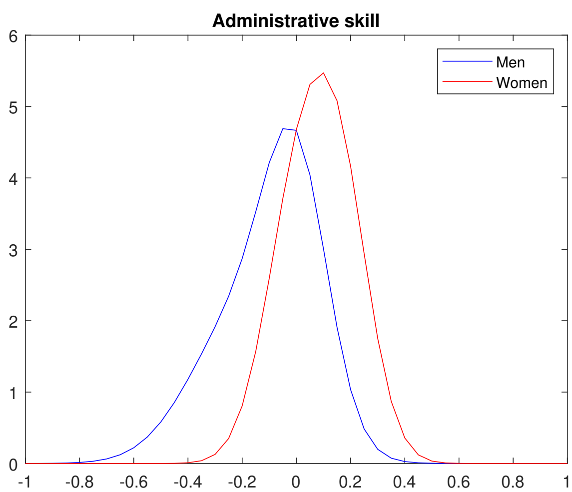

Notes: Distribution of skills by gender. The blue distribution is for men while the red distribution for women. Panel (a) presents the estimated distribution of cognitive skill, while panels (b) and (c) present analogous results for mechanical skill and administrative skill, respectively.

Figure 11: Returns to each occupation by skill quintiles and gender

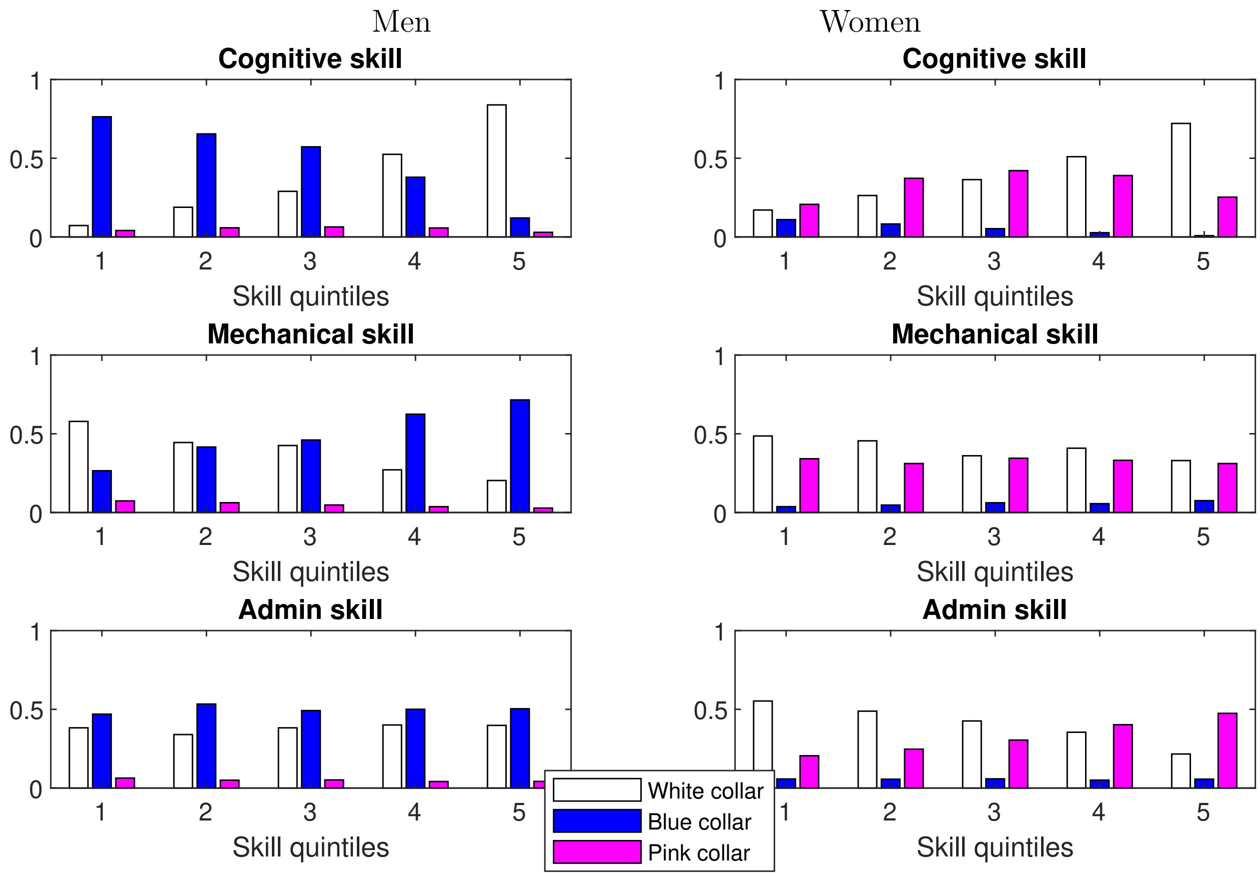

Notes: We simulate each individual 200 times based on the estimates of the model to calculate average returns to each occupations by skill quintiles and gender. Returns include both the wage return as well as the non-pecuniary returns. The upper panels present the effect of cognitive skill by gender, 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.

Figure 12: Occupation choice distribution by skill quintiles and gender

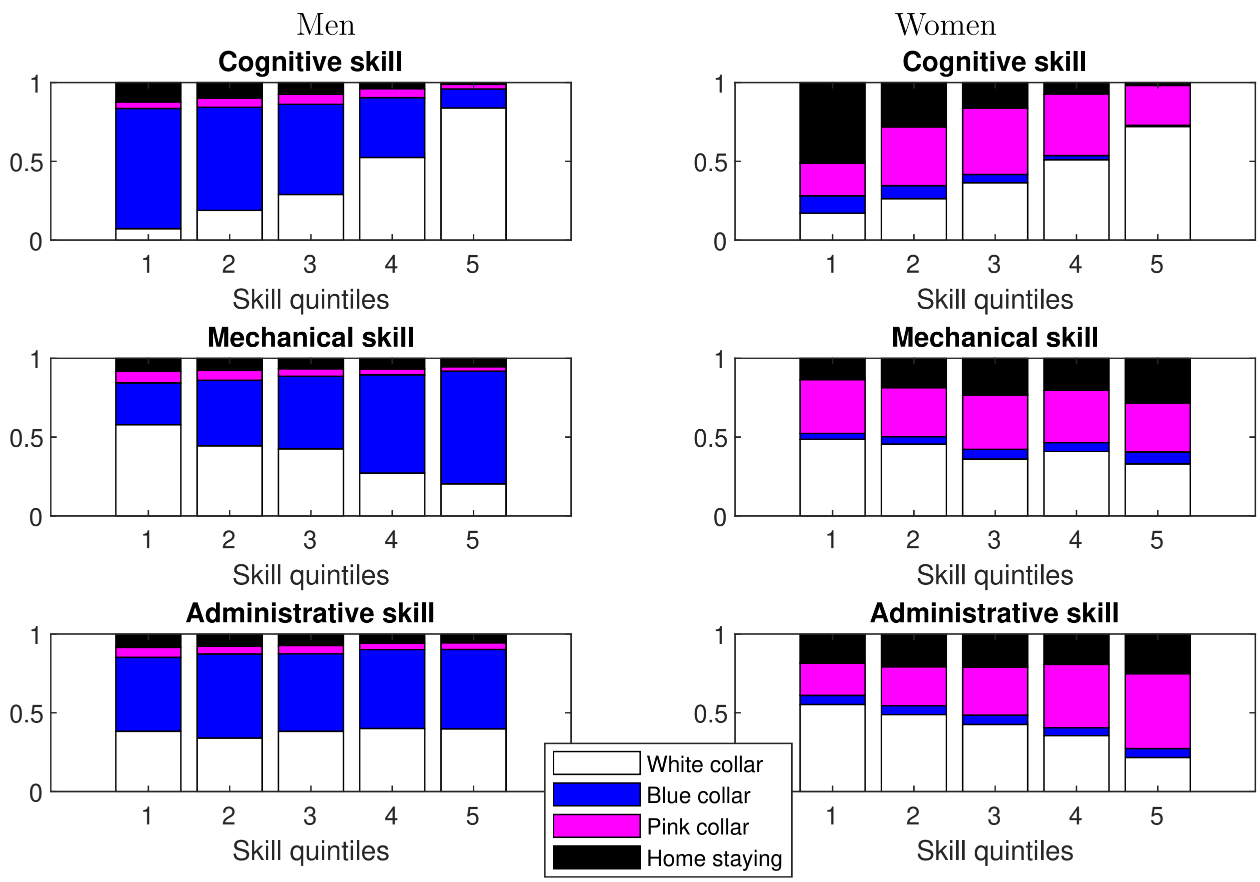

Notes: We simulate each individual 200 times based on the estimates of the model to calculate the distribution of occupation choices by skill quintiles and gender. The upper panels present the effect of cognitive skill by gender, integrating out the effect of the other two dimensions of skills, while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

Figure 13: The college attendance rate by skill quintiles

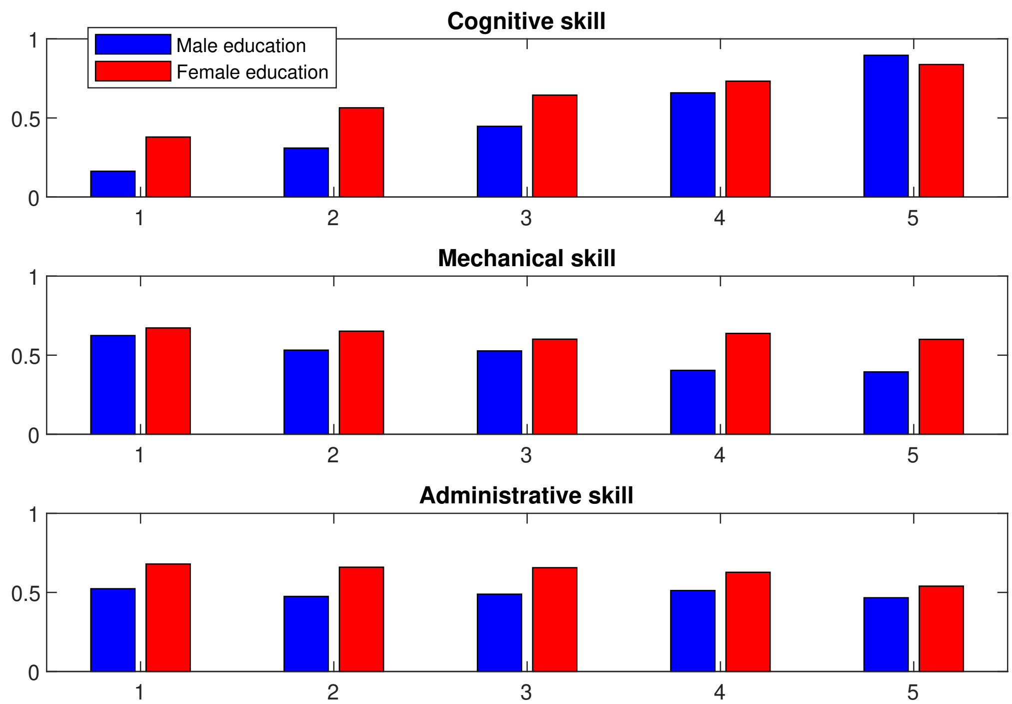

Notes: We simulate each individual 200 times based on the estimates of the model to calculate the college attendance rate by skill quintile and gender. The vertical axis is the fraction of workers in each skill group. The upper panels present the effect of cognitive skill, integrating out the effect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

Figure 14: The distribution of occupations by gender and education degree

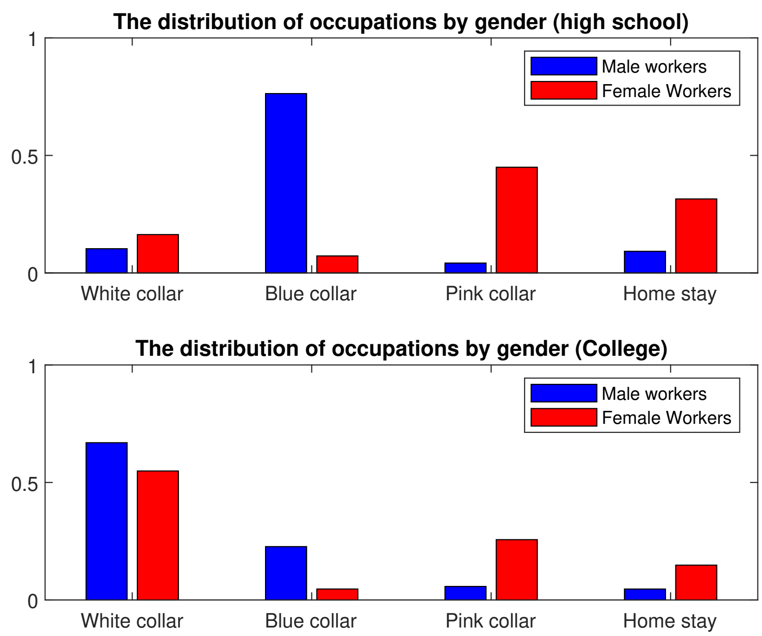

Notes: We simulate each individual 200 times based on the estimates of the model to calculate the occupation distribution by gender and education levels. The vertical axis is the fraction of workers in each occupation group. The upper panels present occupation distribution for college-goer, with blue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

Online Appendix (Not for Publication)

A Appendix: data

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 an unpaid family worker and individuals who report working no weeks in the previous year. Annual earnings are topcoded at the pre-determined Census topcode 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 and ACS 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 occupation 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 endogenous regressor is RTI share, the proportion of jobs in a commuting zone that are highly routine-intensive. 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 lead to downward bias in the estimated 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 5.

We also use Autor and Dorn (2013)'s measures on the routine, manual, and abstract task content of occupations as instruments or control variables in our two stage least squares (2SLS) approach. We use data on routine task content to construct the routine share instrument, which is designed to predict RTI share in a commuting zone in a future year based on the commuting zone's industry composition in 1950 (see Section B.1 for details on instrument construction). We construct predicted manual and abstract shares in the same fashion. As discussed in Section 3, manual and abstract shares are used as control variables in the first and second stage regressions.

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

Three out 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 1940-2000. For each occupation in each year, they characterize the work styles, knowledge requirements, and task content desired by employers based on measures used in the literature. They compile one set of measures to match 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 Hershbein and Kahn (2018) and Network (n.d.-b)).

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 1980 distribution). According to O*NET, administrative activity consists of "performing day-to-day administrative tasks such as maintaining information files and processing paperwork" (O*NET Work Activity 4Ac1, Network, n.d.-c). 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. The trends are consistent with the substantial dropoff in administrative share depicted in Figure 6, which is driven by office and management occupations. Based on this descriptive evidence, we anticipate that commuting zones with high 1950 shares of industries that involved many office and management-related occupations would experience the largest declines in RTI share. Our first stage regression results in Table 3 are consistent with these expectations.

We also use predicted administrative activity as a separate instrument. Rather than as a share, this instrument is measured as the frequency of mentions per job posting. Appendix Figure A.2(a) shows that time trends differ between administrative activity and administrative share. Administrative activity starts out highest in office occupations and second highest in management occupations. However, its decline is pronounced only for office occupations; for other occupations, administrative activity shows only mild declines from 1950 to 2000. This stands in contrast to administrative share, which experienced pronounced declines in office and management occupations but hovered at or near 0 for all other occupation groups.

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, Network, n.d.-a). 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". Appendix Figure A.2(b) shows that variation in clerical requirements follows similar trends to the administrative activity variable, although there are some differences. Mentions per ad are most frequent in the office, sales, and management occupations in 1950. While office occupations experienced the most pronounced decline in clerical requirements over time, there are clear declines in clerical requirements among the other major occupation groups, in contrast to the more muted decline in administrative activity shown in Appendix Figure A.2(a). Both the administrative activity and the clerical requirements instruments provide alternate ways to predict the negative impact of routinization on labor demand for routine-intensive work.[30]

A.4 National Longitudinal Survey of Youth 1979

The National Longitudinal Survey of Youth 1979 (NLSY79) surveys the same 12,686 from 1979 until present day. Surveys were conducted annually until 1994, and then once every two years. We restrict our sample to the 11,155 individuals who finished at least 12th grade or hold a GED degree. We then further drop individuals who were employed but did not have wage information between 25 to 35, leaving a sample size of 8,540. Finally, we exclude individuals who were missing ASVAB test scores or relevant family background information, making a final sample of 2,505 men and 2,490 women. Appendix Table A.3 presents summary statistics for key variables in the model.

[30]The 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. O*NET includes descriptions of whether an occupation requires *knowledge of administration and management* (O*NET Knowledge Requirement 2C1a). It involves overseeing, managing, and coordinating with others, which are considered abstract tasks that would make an occupation harder to automate. Atalay et al. (2020) also 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 are quite vague and encompass a greater variety of tasks than those that were routinized.

A.4.1 Measuring skill heterogeneity

We use the NLSY79's ASVAB tests to construct multi-dimensional skill profiles at the individual level. In 1981, over 90% of NLSY79 respondents completed the ASVAB. The ASVAB is comprised 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 are used to construct the Armed Forces Qualification Test (AFQT) score, a common measure of cognitive ability in the literature on skill returns.[31] Rather than use AFQT directly, we take a different approach by using exploratory factor analysis (EFA) on all nine subtests to construct multiple dimensions of skill.[32] This technique is commonly used in the literature to avoid the ambiguity of the number of latent factors and the underlying factor structure of a set of variables (Diekhoff (1992)).

Exploratory Factor Analysis (EFA) enables us to make use of the correlation structure among the nine ASVAB subtests when constructing our skill measures. The analysis suggests that two separate skills ("factors") are necessary to explain the variation in ASVAB scores.[33] Figure 9 displays the estimated factor loadings. For both men and women, the first factor has significant loadings on all subtests. It is highest for arithmetic reasoning, word knowledge, mathematics knowledge, and paragraph comprehension, which are designed to measure cognitive ability and comprise the main components of the AFQT.

There are gender differences in factor loadings for the second factor. For men, loadings are statistically significant (>0.3) only for the three technical subtests: automotive and shop information, electronics information, and mechanical comprehension.[34] The United States Department of Defense designed these subtests to measure mechanical skill, in that they evaluate the ability to solve simple mechanics problems and understand basic mechan-

[31] Different studies use slightly different subtests to construct AFQT scores. While arithmetic reasoning, paragraph comprehension and word knowledge are commonly used, mathematics knowledge, numerical operations and coding speed have also been adopted to construct the AFQT (see, among many others, (Cameron and Heckman, 1998; Ellwood, Kane, et al., 2000; Heckman and Cameron, 2001; Heckman et al., 2006; Kautz and Heckman, 2014; Neal and Johnson, 1996)).

[32] Besides the five components used to construct AFQT scores, automotive and shop information, electronics information, and mechanical comprehension are commonly referred as the technical composites, since they measure mechanical skill rather than pure IQ (see (Welsh Jr et al., 1990)). Coding speed and numerical operations are separately considered the "administrative qualification area", since they measure the ability to memorize and type strings of letters or numbers or perform quick arithmetic operations on the fly (Prep, n.d.).

[33] Our EFA approach follows that of Prada and Urzúa (2017), who also find that a two-factor structure was most appropriate for explaining the variance in ASVAB test scores for men.

[34] Factor loadings exceeding 0.3 can be considered statistically significant (see Diekhoff, 1992 and Sheskin, 2004).

ics principles ("ASVAB Fact Sheet", n.d.). For women, loadings for the second factor are statistically significant (>0.3) only for the administrative subtests: coding speed and numerical operations. The Department of Defense classifies these subtests into the administrative qualification area, since they are designed to capture efficiency in completing routine administrative tasks (Prep, n.d.).

B Robustness Appendix

B.1 Two Stage Least Squares Approach: Additional Specifications

Routine share instrument. In this subsection, we discuss the construction and identification assumptions of our instruments from Section 4.1. We first use the routine share instrument, which is a modified version of the instrument used in Autor and Dorn (2013). The instrument is constructed as follows:

$$\text{routine share IV}_c = \sum_{i=1}^{I} E_{i,c,1950} \frac{\sum_{k=1}^{K} L_{i,-c,t} \mathbf{1}[\text{routine}_k > \text{routine}^{P66}]}{\sum_{k=1}^{K} L_{k,-c,t}} \tag{10}$$

where i denotes industry, c commuting zone, k occupation, and t year. $E_{i,c,1950}$ is the share of industry i in commuting zone c in 1950. The expression $\frac{\sum_{k=1}^{K} L_{i,-c,t} \mathbf{1}[\text{routine}_k > \text{routine}^{P66}]}{\sum_{k=1}^{K} L_{k,-c,t}}$ calculates the labor share of all occupations in the top third of routine intensity based on the 1980 distribution. We construct predicted manual and abstract occupation shares in parallel ways.

The routine share instrument predicts the RTI share of commuting zone c by isolating the "long-run, quasi-fixed component" of a commuting zone's industrial structure that affects routine share in future decades Autor and Dorn (2013). The intuition is that commuting zones with high shares of routine industries in 1950 will continue to have high RTI share in 1960-2000, despite the displacing impact of automation.[35] Since the instrument is constructed from only 1950 characteristics, it pre-dates omitted variables that could influence employment and education in 1960-2000. Moreover, the instrument leaves out the commuting zone of interest in calculating the routine intensity of an industry, netting out contemporaneous local labor market shocks.

[35] All four of our instruments are premised on the notion that automation decreases routine activity over time, but not to the point where high-routine commuting zones become low-routine commuting zones over time.

In the first stage regression, we interact the routine share instrument with a matrix of year dummies to non-parametrically estimate the impact of the instrument on future years. Since the instrument predicts RTI in a future year, we expect first stage coefficients to be positive, in contrast to the administrative share instrument, which predicts how RTI will decline. Appendix Table A.2 column (4) displays the coefficient estimates using the full set of controls. For a 1 percentage point rise in predicted routine share, RTI share rises by 0.382-0.525 percentage points in 1970-2000, with 1960 as the omitted year ($p < 0.01$). Effect sizes are largest at 0.525 percentage points in 1980 and second largest at 0.474 in 1970, before slightly diminishing to 0.382-0.393 in 1990-2000. These trends roughly match the raw data in Figure 5, where RTI share was highest for women in 1970-1980 but steadily declined by 1990-2000 as industries automated. First stage F-statistics are 42.79, indicating that the Nagar bias is less than 5% of the worst case benchmark. Given recent work arguing that first stage relationships may be weak even if F-statistics exceed 10(Lee et al. (2020)), we also estimate Anderson-Rubin weak instrument robust confidence intervals. They provide further support that the second stage effect for female enrollment is significantly negative, but cannot rule out the possibility that the second stage effect for male enrollment is 0.

Administrative activity and clerical requirement instruments . The administrative activity and clerical requirement instruments are constructed similarly and use similar identification assumptions, so we discuss them together. In both cases, we obtain variation at the commuting zone level by interacting the frequency of mentions with the industry share in 1950:

$$\mathrm{IV}_{ct} = \sum_{i=1}^{I} E_{i,c,1950} \sum_{k \in i} Z_{kt} \tag{11}$$

where Z_{kt} represents the number of mentions of administrative activity or clerical requirements per job posting for occupation k in year t. All other indices are defined as above.

The intuition is that commuting zones with high historic shares of industries intensive in administrative activity or clerical requirements would experience greater routinization over time. Appendix Figure A.2 shows the time-series variation exploited by both instruments, with panel a displaying the change in administrative activity and panel b displaying the change in clerical requirements. As industries automated over time, mentions of administrative activity or clerical requirements in job postings grew less frequent. The sharpest decline occurred for office occupations, but notable declines also happened for clerical requirements among the sales and management occupations. Our instrument predicts greater declines in RTI share among commuting zones with high historic shares of industries where office, management, and sales occupations were prevalent.

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 RTI share at the commuting zone level. That is, local omitted variables that influence both RTI share and college enrollment should have negligible influence on the administrative activity or clerical requirements of an occupation at the national level.

Appendix Table A.2 shows the first stage regression estimates in columns (5) and (6). Point estimates are -3.217 for the administrative activities instrument and -1.460 for the clerical requirements instrument ($p < 0.01$). They are larger than the -0.315 to -0.389 estimated using the administrative share instrument, since the units are in terms of mentions per job posting rather than shares. Our estimates indicate that commuting zones predicted to have 1 more mention of administrative activity per 100 job postings in 1950 will experience a 3.22 percentage point greater reduction in RTI share in future years. Commuting zones predicted to have 1 more mention of clerical requirements per 100 job postings in 1950 will experience a 1.46 percentage point greater reduction in RTI share in future years. Montiel Olea-Pflueger F-statistics are 34.25 for the administrative activities instrument and 25.20 for the clerical requirements instrument, indicating that the Nagar bias is less than 10% of the worst case benchmark. We also estimate Anderson-Rubin weak instrument robust confidence intervals in Table 5, which are squarely negative for female enrollment but include 0 for male enrollment.

B.2 Structural Model identification

Carneiro et al., 2003 and Prada and Urzúa, 2017 provide the formal the formal nonparametric identification arguments of models very close to ours. Therefore, we only sketch the main components that secure the model's identification.

We first identify the loading factors in the equations exclusive for the cognitive skill measures

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

We need to normalize the loading associated with mathematics knowledge to one. ($\lambda_2^c = 1$) Then the other three loading factors $\{\lambda_1^c, \lambda_3^c, \lambda_4^c\}$ can be nonparametric identified. 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 the nonparametric identification of the distributions of θ_c and $e^c_{j,i}$, with $j = 1, 2, ...4$. (Carneiro et al., 2003)

We then argue how to identify the loading factors in the mechanical skill measures

$$M_{j,i} = \lambda^c_j \theta_{c,i} + \lambda^m_j \theta_{m,i} + e^m_{j,i}, j = 5, 6, 7$$

we have assumed $\theta_{c,i}$ and $\theta_{m,i}$ to be correlated. Therefore, we can rewrite 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, which is assumed to be independent of θ_c. Thus, the above mechanical skill measure equation has the following expression

$$\begin{aligned} M_{j,i} &= \lambda^c_j \theta_{c,i} + \lambda^m_j \theta_{m,i} + e^m_{j,i} \\ &= \lambda^c_j \theta_{c,i} + \lambda^m_j (\alpha_1 \theta_{c,i} + \theta_{1,i}) + e^m_{j,i} \;\; j = 5, 6, 7 \\ &= \beta_j \theta_{c,i} + \lambda^m_j \theta_{1,i} + e^m_{j,i} \end{aligned}$$

where $\beta_j = \lambda^c_j + \lambda^m_j \alpha_1, j = 5, 6, 7$. We then decompose the identification into three steps.

1. given the variation of the cognitive skill $var(\theta_c)$ and the loading factors associated with cognitive measures have been identified, we can recover β_j from the covariance $Cov(M_j, C_{j'}) = \lambda^c_{j'} \beta_j var(\theta_c)$.

2. We then apply the normalization for the mathematics knowledge $\lambda^m_7 = 1$. This secures the identification of the other factor loadings λ^m_5 and λ^m_6 in the mechanical test score system. For example, $\lambda^m_5 = \frac{cov(M_5,M_6)}{cov(M_6,M_7)}$ and $\lambda^m_6 = \frac{cov(M_5,M_6)}{cov(M_5,M_7)}$. We can then apply Klotarski's theorem to secure the nonparametric identification of the distributions of θ_1 and $e^m_{j,i}$, with $j = 5, 6, 7$.

3. Finally, we would like to identify α_1, in which we need to impose one additional normalization constraint. Specifically, we assume the factor loadings of cognitive skill on automotive shop information test ($\lambda^c_5 = 0$). 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. Therefore, α_1 is recovered from the equation $\beta_5 = \lambda^m_5 \alpha_1$.

Lastly, our identification for the loading factors in administrative equations is very similar to the identification for the loading factors in the mechanical equations. We first impose is

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

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

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

where $\gamma_j = \lambda_j^c + \lambda_j^m \alpha_2, j = 8, 9$ as we consider two administrative measures. Finally, we impose the normalization assumptions $\lambda_9^c = 0, \lambda_9^a = 1$, in which $j = 9$ denotes administrative measures on numerical operation.

C Additional tables and figures

Table A.1: Summary Statistics, Census Data (Average across Commuting Zones)

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

Notes: Summary statistics for 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: First Stage Regressions, Additional Specifications

	RTI Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Admin Share IV	-0.389	-0.382	-0.315			
	(0.024)***	(0.060)***	(0.047)***			
Routine Share IV∗2000				0.393		
				(0.065)***		
Routine Share IV∗1990				0.382		
				(0.062)***		
Routine Share IV∗1980				0.525		
				(0.070)***		
Routine Share IV∗1970				0.474		
				(0.058)***		
Admin Activities IV					-3.217	
					(0.550)***	
Clerical Requirements IV						-1.460
						(0.291)***
Observations	3600	3610	3610	3610	3610	3610
First Stage F-statistic	256.985	40.740	45.288	42.788	34.246	25.202
Excluding Boston and NYC	Yes					
Control for abstract occupation share		Yes				
RTI share: non-college workers	Yes	Yes		Yes	Yes	Yes
RTI share: college and non-college workers			Yes			
IV: Administrative Share	Yes	Yes	Yes			
IV: Routine Share				Yes		
IV: Administrative Activities					Yes	
IV: Clerical Requirements						Yes

Notes: First stage regressions, additional specifications. All regressions also control for census division, year, commuting zone, labor force participation, manual occupation share, median cognitive earnings, lagged RTI share, and lagged major industry shares: services, manufacturing, retail, and mining. Column (1) excludes commuting zones that contain Boston and New York City. Column (2) controls for abstract occupation share. Column (3) uses the RTI share of all workers, instead of the RTI share of only non-college workers used in the main specification. Column (4) uses the gender-specific non-college RTI share, rather than pooling men and women. Columns (1)-(4) use the administrative share IV, while column (5) uses the routine share IV, column (6) 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-Pflueger 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 A.3: Summary statistics, NLSY79 Data

	Men		Women		Difference	
	Mean	Std	Mean	Std	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 completes high school	0.269	0.443	0.269	0.444	-0.001	0.974
Mother completes 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
Occupation choices between 25 to 35						
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
Average annual earnings between 25 to 35						
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

Notes: Summary statistics for 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 skill test score information.

Table A.4: Correlations between female and work type

Year	Corr(non-college women, routine)	Corr(college women, abstract)
1950	0.142***	0.0433***
1970	0.129***	0.107***
1980	0.0775***	0.117***
1990	0.0341***	0.147***
2000	0.00420***	0.156***

Notes: The table presents pairwise correlations of worker type and task-intensity. Non-college women are significantly more likely to work in routine-intensive occupations, but this likelihood declines over time. College women are significantly more likely to work in abstract-intensive occupations, and this likelihood increases over time. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Estimates for wage coefficients by occupation and gender

	Men			Women		
	White	Blue	Pink	White	Blue	Pink
College	0.141	0.011	0.013	0.102	0.152	0.003
	(0.001)	(0.001)	(0.003)	(0.002)	(0.002)	(0.002)
$\widehat{RTI}$	-0.200	0.035	-0.212	-0.661	-0.455	1.381
	(0.007)	(0.008)	(0.018)	(0.009)	(0.017)	(0.017)
Cognitive	0.119	-0.037	0.105	0.119	0.259	-0.045
	(0.004)	(0.001)	(0.006)	(0.004)	(0.003)	(0.003)
Cognitive*college	0.062	-0.173	-0.032	0.007	-0.149	0.076
	(0.004)	(0.002)	(0.006)	(0.004)	(0.005)	(0.004)
Cognitive*$\widehat{RTI}$	0.105	0.242	0.043	0.657	-0.414	-0.012
	(0.022)	(0.007)	(0.030)	(0.019)	(0.019)	(0.016)
Cognitive*college*$\widehat{RTI}$	-0.006	-0.033	-0.008	-0.026	0.002	-0.008
	(0.022)	(0.010)	(0.032)	(0.020)	(0.029)	(0.018)
Manual	-0.037	0.088	-0.058	-0.123	-0.138	-0.103
	(0.003)	(0.003)	(0.009)	(0.007)	(0.019)	(0.009)
Manual*college	-0.001	0.126	0.008	0.083	-0.112	-0.286
	(0.003)	(0.002)	(0.007)	(0.007)	(0.029)	(0.011)
Manual*$\widehat{RTI}$	-0.071	-0.191	-0.215	-0.561	0.320	0.375
	(0.019)	(0.013)	(0.048)	(0.035)	(0.106)	(0.046)
Manual*college*$\widehat{RTI}$	-0.014	-0.004	0.022	-0.012	-0.041	-0.092
	(0.016)	(0.013)	(0.044)	(0.035)	(0.158)	(0.054)
Admin	0.246	0.128	0.216	-0.311	-0.094	0.131
	(0.014)	(0.008)	(0.025)	(0.023)	(0.043)	(0.015)
Admin*college	-0.103	0.065	-0.190	-0.182	0.072	-0.197
	(0.011)	(0.009)	(0.023)	(0.023)	(0.051)	(0.018)
Admin*$\widehat{RTI}$	0.133	0.098	0.030	-0.144	0.273	0.251
	(0.078)	(0.041)	(0.142)	(0.134)	(0.238)	(0.081)
Admin*college*$\widehat{RTI}$	0.031	-0.020	0.001	0.127	0.061	0.018
	(0.068)	(0.050)	(0.121)	(0.130)	(0.309)	(0.098)
Constant	1.940	1.776	1.655	1.892	1.657	1.057
	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.003)
Standard deviation	0.457	0.411	0.475	0.409	0.452	0.444
	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)

Notes: The parameter estimates for wage coefficients in Equation 4 are reported by occupation and gender. Standard errors in parentheses.

Table A.6: Estimates for utility parameters by occupation and gender

	Men			Women		
	White	Blue	Pink	White	Blue	Pink
College	0.478	-0.904	0.419	0.890	-1.001	-0.493
	(0.011)	(0.009)	(0.020)	(0.012)	(0.016)	(0.011)
$\widehat{RTI}$	0.148	0.535	0.167	-6.689	-3.194	5.267
	(0.204)	(0.164)	(0.276)	(0.198)	(0.355)	(0.188)
Cognitive	0.479	-0.053	0.335	1.162	0.573	0.398
	(0.052)	(0.024)	(0.063)	(0.041)	(0.057)	(0.022)
Cognitive*college	0.531	0.987	0.564	0.560	-0.350	-0.288
	(0.057)	(0.031)	(0.089)	(0.042)	(0.080)	(0.031)
Cognitive*$\widehat{RTI}$	-0.839	0.838	-0.449	-8.097	-2.628	9.567
	(0.289)	(0.127)	(0.347)	(0.222)	(0.308)	(0.111)
Cognitive*college*$\widehat{RTI}$	0.030	-0.003	-0.012	3.976	-0.006	-4.103
	(0.312)	(0.164)	(0.484)	(0.229)	(0.428)	(0.159)
Manual	-0.001	0.497	-0.175	-0.726	-0.186	-0.629
	(0.059)	(0.035)	(0.095)	(0.081)	(0.155)	(0.063)
Manual*college	-1.094	-0.413	-0.217	0.154	1.596	0.537
	(0.062)	(0.039)	(0.123)	(0.076)	(0.219)	(0.080)
Manual*$\widehat{RTI}$	-1.094	0.298	-0.608	1.193	0.559	-0.174
	(0.330)	(0.186)	(0.518)	(0.427)	(0.858)	(0.323)
Manual*college*$\widehat{RTI}$	-0.005	0.013	0.016	-0.009	-0.023	0.000
	(0.345)	(0.202)	(0.664)	(0.407)	(1.179)	(0.413)
Admin	0.358	-0.073	0.015	0.735	-0.208	-0.087
	(0.220)	(0.134)	(0.306)	(0.262)	(0.450)	(0.190)
Admin*college	-0.508	0.109	-0.320	1.107	-0.674	0.212
	(0.176)	(0.125)	(0.242)	(0.211)	(0.405)	(0.128)
Admin*$\widehat{RTI}$	0.041	0.117	-0.442	-4.632	-0.801	6.643
	(1.223)	(0.702)	(1.706)	(1.430)	(2.503)	(1.012)
Admin*college*$\widehat{RTI}$	0.010	-0.010	-0.042	-0.612	-0.100	0.830
	(0.967)	(0.654)	(1.342)	(1.128)	(2.351)	(0.678)
Constant	-6.304	-4.497	-5.647	-5.377	-5.572	-4.852
	(0.037)	(0.031)	(0.047)	(0.039)	(0.066)	(0.036)

Notes: The parameter estimates for non-pecuniary utility coefficients in Equation 5 are reported by occupation and gender. Standard errors in parentheses.

Table A.7: Estimates for education equation by gender

	Men		Women	
	Estimate	Std. Error	Estimate	Std. Error
Cognitive	1.18	0.37	1.11	0.38
Manual	-0.39	0.39	-0.22	0.39
Admin	0.16	1.49	0.17	1.64
Cohort 2	-0.15	0.13	-0.20	0.12
Cohort 3	0.00	0.13	0.15	0.11
Cohort 4	-0.02	0.14	0.25	0.14
Father's Education	0.85	0.17	0.34	0.15
Mother's Education	0.25	0.13	0.80	0.29
Urban	0.39	0.13	0.19	0.11
South	0.38	0.14	0.21	0.13
Family intactness	0.48	0.04	0.15	0.03
Number of siblings	-0.03	0.01	0.00	0.01
Constant	-5.28	0.67	-1.51	0.62
Standard deviation	0.85	0.19	0.71	0.41

Notes: The parameter estimates for coefficients associated with 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.8: Parameters in skill distributions and measure equations

	Skill distribution		Measurement equation			
	Men	Women	Loadings		STD	
	(1)	(2)		(3)		(4)
μ_{cog}	-0.003	0.068	λ_1^m	1.552	$\sigma_{c,1}$	0.465
	(0.051)	(0.088)		(0.046)		(0.023)
μ_{manual}	0.296	-0.271	λ_2^c	0.565	$\sigma_{c,2}$	0.527
	(0.038)	(0.023)		(0.015)		(0.016)
μ_{admin}	-0.194	0.160	λ_2^m	0.929	$\sigma_{c,3}$	0.540
	(0.022)	(0.024)		(0.027)		(0.016)
$\sigma_{cog}^{(1)}$	0.800	0.745	λ_3^c	0.505	$\sigma_{c,4}$	0.479
	(0.090)	(0.126)		(0.016)		(0.016)
$\sigma_{manual}^{(1)}$	0.336	0.335	λ_4^c	1.064	$\sigma_{m,5}$	0.502
	(0.091)	(0.101)		(0.021)		(0.017)
$\sigma_{admin}^{(1)}$	0.191	0.111	λ_6^c	0.998	$\sigma_{m,6}$	0.557
	(0.084)	(0.143)		(0.020)		(0.016)
$\sigma_{cog}^{(2)}$	0.556	0.324	λ_7^c	0.936	$\sigma_{m,7}$	0.619
	(0.082)	(0.135)		(0.019)		(0.017)
$\sigma_{manual}^{(2)}$	0.398	0.108	λ_8^c	0.815	$\sigma_{a,8}$	0.699
	(0.091)	(0.121)		(0.024)		(0.028)
$\sigma_{admin}^{(2)}$	0.117	0.117	λ_9^a	0.945	$\sigma_{a,9}$	0.953
	(0.090)	(0.132)		(0.136)		(0.027)

Notes: 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 7 in column (3). It reports the standard deviation of the residual term in each test score measurement equation in column (4). Standard errors in parentheses.

Figure A.1: Instruments predict automation intensity

(a) Administrative Share IV

(b) Routine Share IV

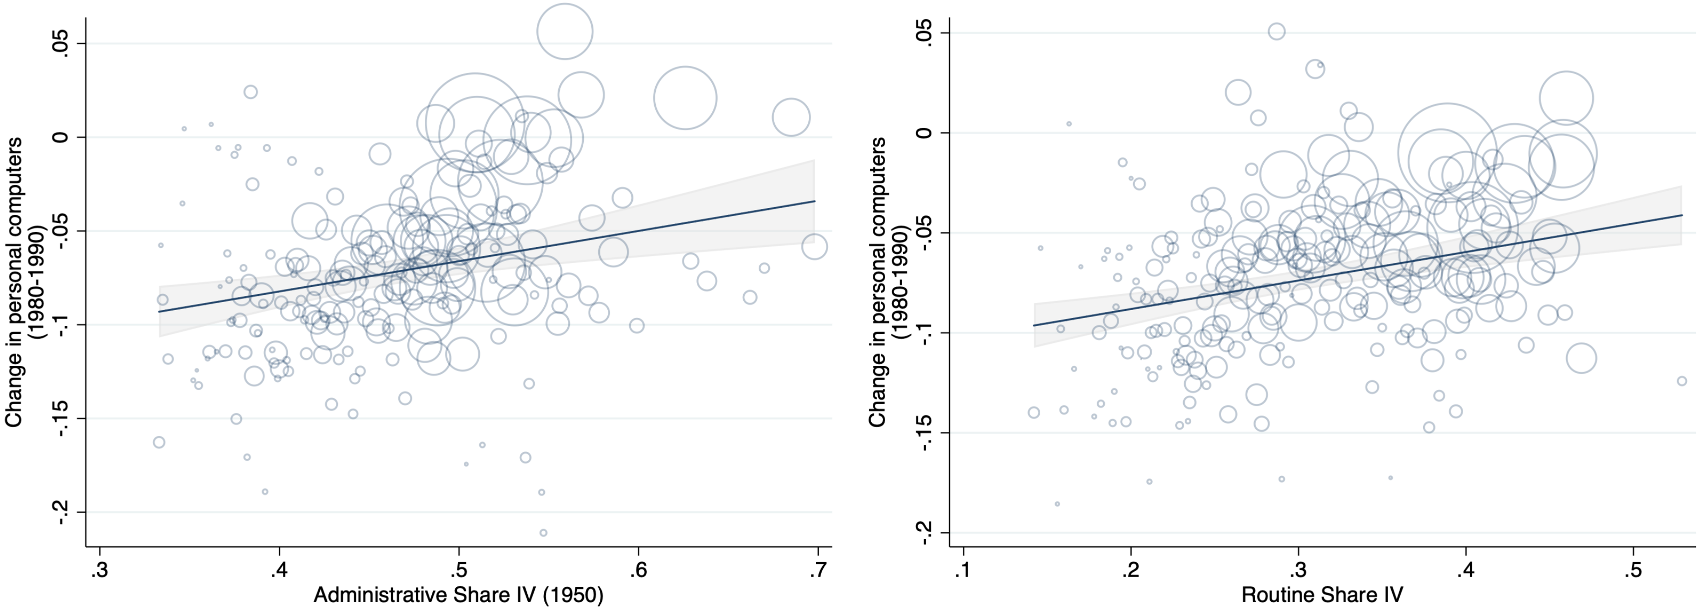

Notes: Raw correlation between instruments and future automation intensity, as measured by change in personal computers in 1980-1990. In panel (a), the instrument is predicted administrative share. In panel (b), the instrument is predicted routine share. The solid line shows the correlation estimated from an OLS regression using labor supply weights. The shaded gray area depicts 95% confidence intervals. Data from Census, Autor and Dorn (2013), and Atalay et al. (2020).

Figure A.2: Administrative activities and clerical requirements by major occupation group

(a) Administrative activities

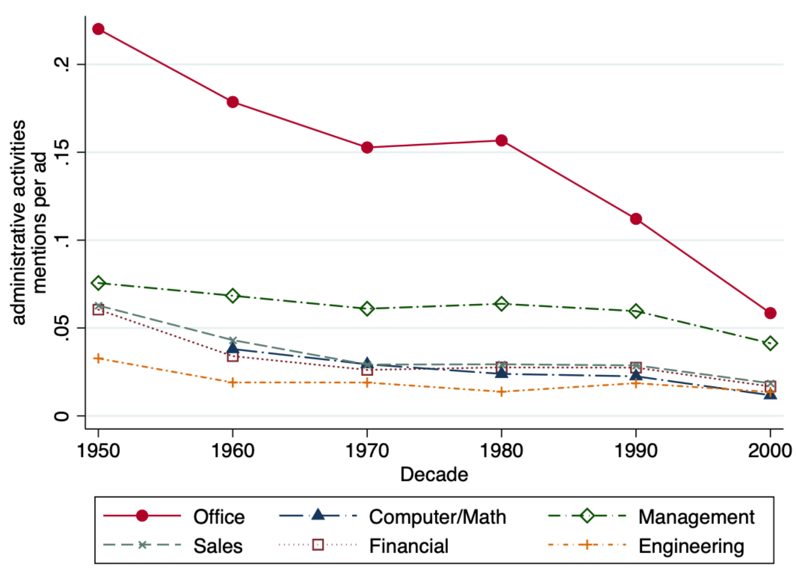

(b) Clerical requirements

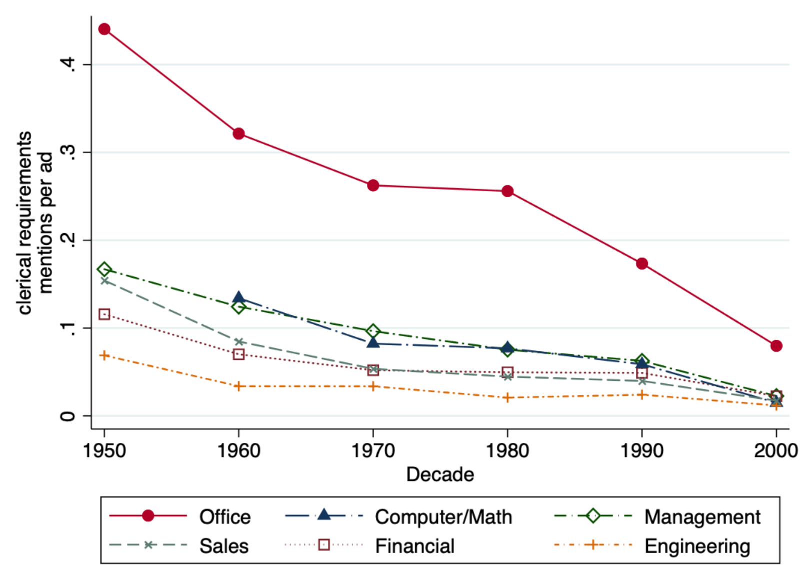

Notes: Frequency with which administrative activities (a) and clerical requirements (b) are mentioned per job posting. Data from Census (1950-2000) and Atalay et al. (2020).