2022

Essays in labor economics

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

ESSAYS IN LABOR ECONOMICS

by

VITTORIA DICANDIA

B.S., Bocconi University, 2012
M.S., Bocconi University, 2015
M.A., Boston University, 2019

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2022

Approved by

First Reader

Kevin Lang, PhD
Laurence A. Bloom Professor of Economics

Second Reader

M. Daniele Paserman, PhD
Professor of Economics

Third Reader

Joshua Goodman, PhD
Associate Professor of Education and Economics

Acknowledgments

I would like to thank my advisors, Kevin Lang, Daniele Paserman, and Josh Goodman. I am very grateful for your guidance during these years, I am a better researcher for it. Your insightful comments have not only improved my current work, but the way in which I think about research. Kevin, you have been an incredible main advisor. Thank you for being a role model as the economist and mentor I can only aspire to be. Daniele, your encouragement has been fundamental in getting me through the toughest times. Josh, your support has been extremely valuable.

I am indebted to Costas Cavounidis, Pascual Restrepo, and Chiara Margaria, who offered great advice on the first chapter of this dissertation.

I would like to thank Bart Lipman, Chiara Margaria, and Silvia Vannutelli, for their dedication and work in promoting together initiatives for women in economics at BU, and Stefania Garetto for her advice during the job market.

I also would like to thank Paola Profeta, my advisor at Bocconi University, who helped me start my PhD journey, and who continues to support me to this day.

I am very grateful for the friends who have accompanied me in these six years.

The members of my fantastic cohort: Enrico De Magistris, Pablo Ignacio Guzman Pinto, Stefano Pica, Alden Porter, Thomas Pearson, Hanna Schwank, and Lalit Sethia. Your friendship during the good and the bad has made this journey so much better. I have learned a lot from you, and I am sure I will keep on doing so in the future. Hanna, thank you so much for being there when I needed it the most.

Chiara Margaria, thank you for being a supportive friend and for exploring Boston with me; I will deeply miss being just a few doors away from you. Ilaria D'Angelis, thank you for our endless conversations, I am so glad that we met (again) in Boston and that we could share this path. Silvia Vannutelli, my partner in many adventures:

I am very lucky to have found a friend and colleague I can lean on and share so much with. Anastasiia Arbuzova, Pierre De Leo, Gemma Dipoppa, Leo Elias, Darin Kurti, Fernando Payro Chew, Gianluca Russo, Anna Weber, thank you all for making Boston home.

I also would like to thank the friends that are always with me, no matter where I go: Francesca Visinoni, Paola Azzarone, Lorenzo Barreca, Marco Forletta, Flora Gatti, Luca Grezio, Chiara Ragno, Simone Romagnoli, Matteo Polledri, Lucrezia Sanginesi, Giorgio Tripodi, Fabrizio Vasselli, Lorenzo Vasselli.

Last, but not least, I would like to thank my family. My brother Lorenzo, my mother Annamaria, my father Michelangelo, and my aunt Mattia. Your unconditional support makes me stronger. Your example makes me want to be better, always. A lot of who I am and of what I do, also as a researcher, is because of you.

All errors are my own.

Vittoria Dicandia

Department of Economics

ESSAYS IN LABOR ECONOMICS

VITTORIA DICANDIA

Boston University, Graduate School of Arts and Sciences, 2022

Major Professor: Kevin Lang, PhD
Laurence A. Bloom Professor of Economics

ABSTRACT

This dissertation consists of three chapters in labor economics. The first chapter explains why the wage gap between black and white Americans has stalled since 1980, after a period of significant narrowing during the 1960s and 1970s. I argue that routine-biased technological change (RBTC) dampened wage gap convergence in 1980-2000. It had a differential impact across races at different parts of the wage distribution. I present new evidence on occupational patterns by race and on determinants of wage disparities along the wage distribution and rationalize them with an RBTC model in which firms engage in statistical discrimination. I show that, surprisingly, the share of employment in routine-intensive occupations has increased for black workers, in contrast with a significant decrease for white workers. I decompose the wage gap changes using the Oaxaca-RIF method. I show that differences in occupational sorting increased wage disparities, thwarting wage convergence between races at the bottom of the wage distribution. Together, these new empirical findings and model allow us to better understand the mechanisms behind racial disparities at the end of the 20th century.

The second chapter (with Costas Cavounidis, Kevin Lang, and Raghav Malhotra) develops a tractable general equilibrium model to explain within- and between-occupation changes in skill use over time. We apply the model to skill-use measures

from the third, fourth, and revised fourth editions of the Dictionary of Occupational Titles and data from the 1960, 1970, and 1980 Censuses and March Current Population Surveys. We recover changes in skill productivity by exploiting between-occupation movements. We conclude that finger-dexterity productivity grew rapidly while abstract-skill productivity lagged, a form of 'skill bias.' Together with substitutability between abstract and routine inputs, these results explain changes in skill use within occupations.

In the third chapter (with Silvia Vannutelli), we exploit the enlargement of the European Union in 2007 to study the consequences for the Italian labor market of the permanent legalization of immigrants from Romania and Bulgaria. We use a unique administrative employer-employee dataset covering the universe of Italy's private-sector workers. We study firms' responses in terms of personnel choices. We find short-term effects on firm-level employment. Employment increased for EU07 migrants at the expense of natives, accompanied by a rise in hirings and separations for the former. We provide evidence that the findings are mainly driven by the migrants' change of legal status rather than by the arrival of new workers in the country. We also observe a reduction in per-capita revenues and operative added value, confirming that the legalization of previously undocumented workers likely drives the effects.

Contents

List of Tables

List of Figures

List of Abbreviations

ASEC	....	Annual Social and Economic Supplements
CERVED	....	Centri Elettronici Reteconnessi Valutazione Elaborazione Dati
CPS	....	Current Population Survey
CZ	....	Commuting Zone
DOT	....	Dictionary of Occupational Titles
EU	....	European Union
EU07	....	Romania and Bulgaria
FD	....	First Difference
FE	....	Fixed Effects
IBM	....	International Business Machines Corporation
INPS	....	Istituto Nazionale Previdenza Sociale
IPUMS	....	Integrated Public Use Microdata Series
ISTAT	....	Istituto Nazionale di Statistica
IV	....	Instrumental Variable
LLM	....	Local Labor Market
NBER	....	National Bureau of Economic Research
RBTC	....	Routine-Biased Technological Change
RIF	....	Recentered Influence Function
RTI	....	Routine Intensity Measure
SBTC	....	Skill-Biased Technical Change
UK	....	United Kingdom
US	....	United States of America
WP	....	Working Paper

Chapter 1

Technological Change and Racial Disparities

1.1 Introduction

After striking progress in the 1960s and 1970s, the hourly wage gap between black and white workers stalled at about 20% starting in the 1980s (see Figure 1·1), but there is no consensus in the literature as to why.[1]

Figure 1·1: Raw white-black hourly wage gap

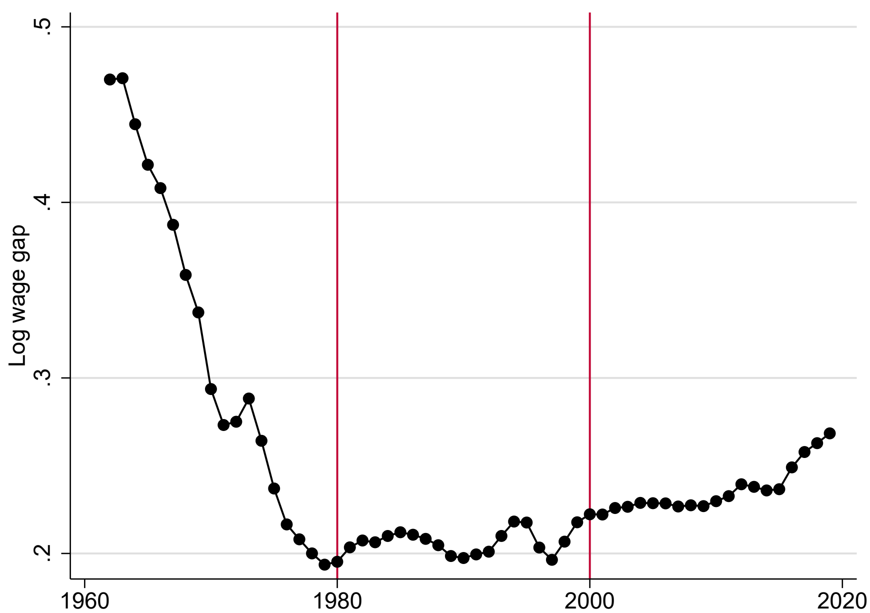

Source: CPS ASEC. Sample: All 20-64 black and white male workers with non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

[1] For trends in the gap of mean annual wage, see Figure A·1. Patterns are invariant in essence, but the gap from 1980 onward is slightly larger.

I argue that routine-biased technological change (RBTC) contributed significantly to the failure of the wage gap to decline because it affected black and white workers differently. After 1980, technological progress, particularly computerization, reduced employment and wages in occupations such as clerical jobs and jobs in operations and production, that used routine skills intensively (Autor et al. (2003), Autor and Dorn (2013)). The standard view is that routine-intensive jobs were mostly concentrated in the lower middle of the wage distribution and that RBTC polarized the US labor market by shifting employment to both lower and high-skill jobs.

Figure 1·2: Change in employment 1980-2000

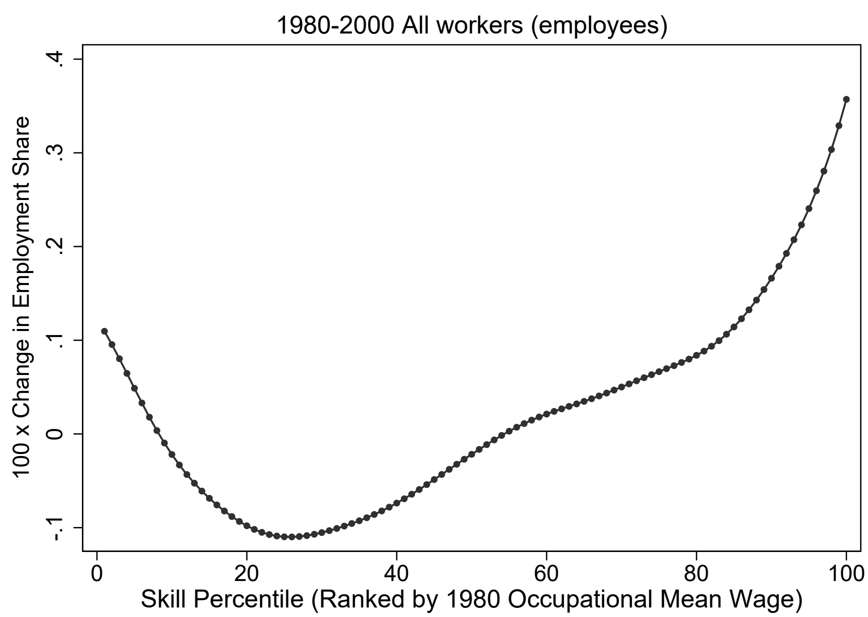

Source: IPUMS. Sample: All 20-64 black and white male workers with non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

I focus on male black and white workers between 1980 and 2000 and first show that, surprisingly, black workers increased their employment share in routine-intensive occupations. This shift is prevalent among those with a high school diploma or less, and younger cohorts. However, the diverging patterns between black and white workers is not explained by difference in their demographics. These findings are robust to various sample choices.

I then show how the wage gap changed at different points of the wage distribution. While the mean gap in hourly wage was stagnant, disparities decreased in the bottom half of the distribution and increased in the top half. Using a Oaxaca-RIF decomposition (Firpo et al. (2009)), I show that changes in occupational sorting associated with RBTC increased the wage gap across the entire wage distribution. In the bottom half of the wage distribution, it only partially counteracted the decreases in disparities due to increased educational achievement among black workers and reduced gaps in racially-specific returns to worker characteristics. Consistent with this pattern, the wage gap narrowed throughout the distribution in commuting zones with low proportions of routine-intensive workers.

Lastly, I show that when firms engage in statistical discrimination, RBTC can generate the racial resorting of workers that we observe. I consider a labor market with two types of workers, black and white, and three sectors: manual (lowest paying sector), routine (middle paying sector), and abstract (highest paying sector). I model technological change as a shock that reduces wage responsiveness to individual ability in occupations in the routine sector. When firms observe a signal of ability that is less precise for black than for white workers, this form of technological change can generate racial differences in net movements in and out of the routine sector. Thus, I rationalize my empirical findings as possible in this setting, something that is not possible in the model of RBTC developed in Acemoglu and Autor (2011).

This study is related to two main strands of literature in labor economics: racial disparities and routine biased technological change.

Racial disparities. This paper contributes to deepen our understanding of the evolution of the racial wage gap in the second half of the 20th century. It is therefore linked to an extensive body of literature studying the mechanisms behind racial disparities and discrimination (Lang and Lehmann (2012), Fang and Moro (2011), Lundberg and Startz (1983), Coate and Loury (1993)). Another consistent set of studies focuses on the evolution of wage inequality over time (Bound and Freeman (1992), Juhn et al. (1991), Altonji and Blank (1999), Western and Pettit (2005)), documenting a reduction in the gap between 1940s and 1980s, but a slow-down in the following period. Some studies instead assess the role of specific anti-discrimination policies, such as the Voting Rights Act (Aneja and Avenancio-Leon (2019)). Others, examine the effects of broad economic changes, such as the increase in education level of African-Americans (Smith and Welch (1986)). Derenoncourt and Montialoux (2021) provide evidence that the extension of the federal minimum wage (Fair Labor Standards Act of 1966) contributed to earning-gaps reductions in the 1960s and 1970s. Thus, like this paper, they use a racially-neutral economic shock to explain changes in racial disparities. Batistich and Bond (2019), instead, focus on the effects of the Japanese manufacturing trade shock of the mid-1970s, and show that black employment was negatively affected. As in this paper, they implement a spatial strategy for their analysis, by using differences across labor markets in initial exposure to identify the effects of the shock. Lastly, this paper is also related to Bayer and Charles (2018), in which they show how racial gap patterns differ along the wage distribution.

Routine biased technological change. In the last twenty years there has been a significant increase in contributions to the technological change literature, especially regarding the role played by technology in explaining the polarization of the US labor market, as well as of other markets in the developed world (Goos et al. (2014)). One strand has focused on understanding and correctly interpreting the mechanisms through which this technological change has affected the labor market (Autor et al. (2003), Autor and Dorn (2013), Cavounidis and Lang (2020), Cavounidis et al. (2021)). Other papers have explored how this shock has affected men and women differently (Autor and Price (2013), Black and Spitz-Oener (2010), Cavounidis et al. (2021)).

However, research on technological change in relation to racial disparities has been scant. The recent paper by Hurst et al. (2021) is a notable exception. Their work develops a task-based framework that shows how the effects of decreases in the racial-skill gap and taste-based discrimination were counteracted by the increase in returns to abstract skills, which favors white workers. My work can be viewed as complementary to theirs, in that it addresses other aspects of the impact of RBTC on racial disparities. It provides a more comprehensive picture of the changes that occurred along the entire wage distribution and focuses on racially different occupational sorting, rather than economy-wide skill use by race. Importantly, it presents new empirical findings on relevant different occupational patterns between black and white workers, and it rationalizes them by introducing a model of RBTC that, although different, is not incongruous with the existing literature.

1.2 Data and Sample

The focus of this paper is how, between 1980 and 2000, RBTC impacted disparities between male black and white workers. As stagnation of the racial wage gap continued after 2000 (see Figure A·2 in the appendix for Figure 1·1 for male workers only), I focus on the earlier period because it precedes the China trade shock. Such trade shocks are likely to differentially affect different racial groups (e.g., Batistich and Bond (2019) on the Japan trade shock). The later period is also greatly affected by the Great Recession and, more recently, the pandemic.

1.2.1 Data

I use three data sources for my analysis: the March Current Population Survey (CPS), from 1972 to 2000, the US Census decennial data, from 1980 to 2000, and the third edition of the Dictionary of Occupational Titles, published in 1977. I employ the CPS datasets to show time trends with high (annual) frequency. However, its sample size does not allow for a detailed examination of the evolution of racial disparities, especially changes along the wage distribution and geographic heterogeneity. I therefore use the Census decennial data for such analyses.

Lastly, I use the Dictionary of Occupational Titles to classify occupations in terms of skill-use intensity and whether they are more or less affected by routine-biased technological change.

March CPS

I use the March Annual Social and Economic Supplements (ASEC) of the Integrated Public Use Microdata Series, particularly. These files provide individual-level data on employment (weeks and hours worked, occupation, and industry) and income, and demographic information such as race, gender, education, age, marital status.

The CPS first used the 1970 Census occupation codes in 1972. Since this study requires observing occupations consistently across time, my sample begins in 1972. I adopt Dorn (2009) occupational crosswalk, which harmonizes Census decennial classifications into 330 consistent occupations.

Income is given by nominal annual wage and salary in the previous year, and is top coded to prevent identification of high earners. I substitute topcoded values using the dataset created by Larrimore et al. (2008), who built cell mean values by demographics using restricted use CPS data. Before 1976, I do not observe anyone with a top-coded income, and, therefore, make no adjustment. Lastly, I use the 2006 Personal Consumption Expenditures Price Index from Federal Reserve Economic Data to transform nominal into real values.

I calculate hourly wages by dividing annual income from wage and salaries and by the reported number of weeks multiplied by hours worked. For earlier years, weeks worked is reported in intervals. Following Derenoncourt and Montialoux (2021), I obtain weekly wages by dividing annual wages by the median number of weeks in the indicated interval, and then smooth it by adding a random number generated from a uniform distribution. I then compute hourly wages by dividing the smoothed weekly wages by the number of hours worked in the previous week, which is consistently reported throughout the period.

The sample used in the main analysis includes all black and white male employees of age 20-64 in the private and public sectors. I exclude self-employed, unpaid family workers, workers in group quarters, workers in agricultural occupations, and workers in the armed forces. I do not limit the sample to full-time full-year workers but show in the appendix that doing so does not change the result. I use ASEC individual weights for all computations.

Census

I use the 5% 1980 and 2000 samples of the US decennial Census. Like the CPS, the Census provides data on employment (weeks and hours worked, occupation, and industry), annual income, and demographic information such as race, gender, education, age, marital status. Because of its larger sample, the Census provides more detailed geographic information on individuals, including their county of residence. Using this information, I adopt Dorn (2009) geographic crosswalk and place all counties in 722 time-consistent commuting zones, geographic units characterized by a common labor market. Commuting zones are defined across the entire country, except for Alaska and Hawaii, which I, therefore, exclude from my sample.

Data on hours and weeks of work for these two samples is provided with non-intervalled measures, making it straightforward to calculate hourly wages by dividing annual wage income by the number of hours and weeks worked the previous year. Lacking a more accurate data source to deal with income topcoding issues, I make an *ad hoc* adjustment by multiplying top coded annual wages by 1.5.[2]

I use the same sample restriction that I used for the CPS data. All computations use individual weights provided by the Census multiplied by the product of the number of hours and weeks worked (Autor and Dorn (2013)).

Dictionary of Occupational Titles

I use the occupation-level measures of routine intensity computed by Autor and Dorn (2013). The data is originally from the Dictionary of Occupational Titles, which is the standard dataset for information on skill use. It provides numerical measures of attitudes, temperaments, and abilities needed to perform a job. The data is at the occupational-title level, but they aggregate it to Census-occupation level. Autor and Dorn (2013) define three measures, for manual, routine, and abstract skills, and use the following measure to define routine intensiveness of occupations:

$$RTI = ln(R) - ln(A) - ln(M) \tag{1.1}$$

where routine (R) is the mean of variables measuring *ability to work requiring set limits, tolerances, or standards* and *finger dexterity*; abstract (A) is the mean of variables measuring *quantitative reasoning requirements* and *direction, control, and planning of activities*; manual (M) is defined by the variable measuring *eye, hand, foot coordination*.

Like them, I use RTI to identify routine-intensive occupations, defined as the top third occupations in terms of RTI.[3] This measure, as well as its individual components, has been used in several studies of Routine-Biased Technological Change. Using it for

[2]Lemieux (2006) multiply top coded wages by 1.4, and Autor et al. (2008) multiply them by 1.5. There is a single top income value in 1980 but in 2000 it varies by state. Hence, I make this adjustment uniformly across states in 1980 and by state in 2000.

[3]The 10 occupations with the highest RTI scores are: Butchers and meat cutters; Secretaries and stenographers; Payroll and timekeeping clerks; Bank tellers; File clerks; Cashiers; Typists; Pharmacists; Bookkeepers, accounting clerks; Postal clerks, except mail carriers. (Autor and Dorn (2013))

my analysis allows for consistency with the existing literature. While there may be concerns about the validity of RTI, they should be less important when it is only used to identify a group of routine-intensive occupations.

1.3 Employment

Figure 1·3 shows race-specific employment shares in the top third RTI occupations over time. Black and white workers exhibit opposite patters: after 1980, we observe a sharp decrease for white workers, and a significant increase for black workers.

Figure 1·3: Race-specific employment share of top RTI occupations

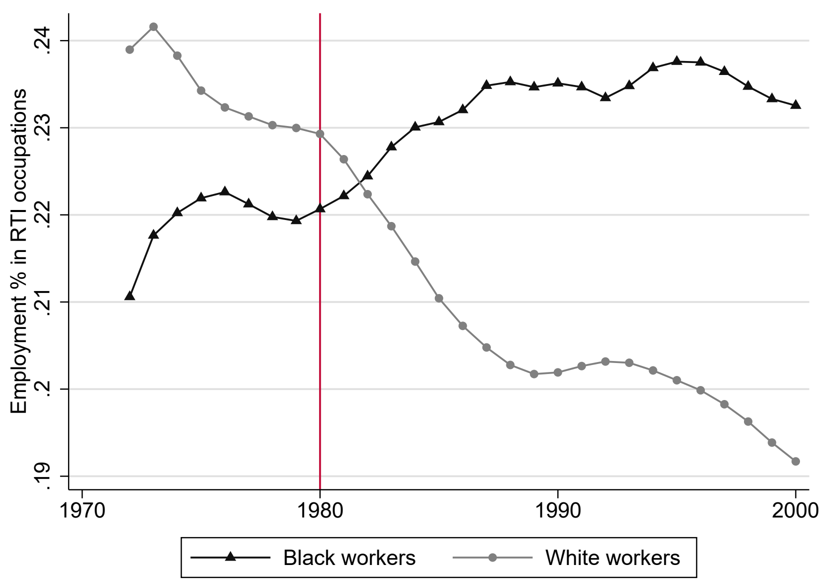

Source: CPS. Sample: All 20-64 black and white male workers with non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

Between 1980 and 2000, the share of white workers employed in the top third RTI occupations fell from 23% to 19%. This change, although striking, is expected and consistent with the standard RBTC framework, which predicts that workers will shift out of routine-intensive occupations in response to the computerization shock. The same cannot be said for black workers. Even before 1980, we observe a slight increase in their share of routine-intensive occupations, which, at the time, were viewed as good jobs, with salaries that would place workers broadly in the middle of the wage distribution. After 1980, notwithstanding the declining desirability of routine jobs, black workers' employment share in RTI occupations increased even more sharply.

These patterns are invariant to sample choice. Perhaps this anomaly reflects the over-representation of black workers in seasonal and/or part time jobs.[4] In Figure A·3 I show that these trends hold when I conservatively restrict the sample to individuals who work at least 30 hours per week for a minimum 40 weeks.

Perhaps differences in educational achievement can account for this unexpected finding. It is possible, for instance, that RBTC negatively affected only workers with at least some college education, causing their displacement from routine jobs and some level of replacement from workers with a high school degree or less. Given the higher college–non college ratio for white workers, this type of technological shock would be consistent with the patterns shown although it would still be inconsistent with the standard framework. Still, it would mean that the difference was due to education, not race. Figure 1·4 addresses this concerns by showing that the same trends apply to workers with a high school diploma or less. Thus, for workers with lower educational achievement there are diverging employment trends between black and white workers.

[4]Even if that were the case, it would still be a result incompatible with the current consensus on routine occupations employment patterns.

Figure 1·4: Race-education employment share of top RTI occupations

High school diploma or less

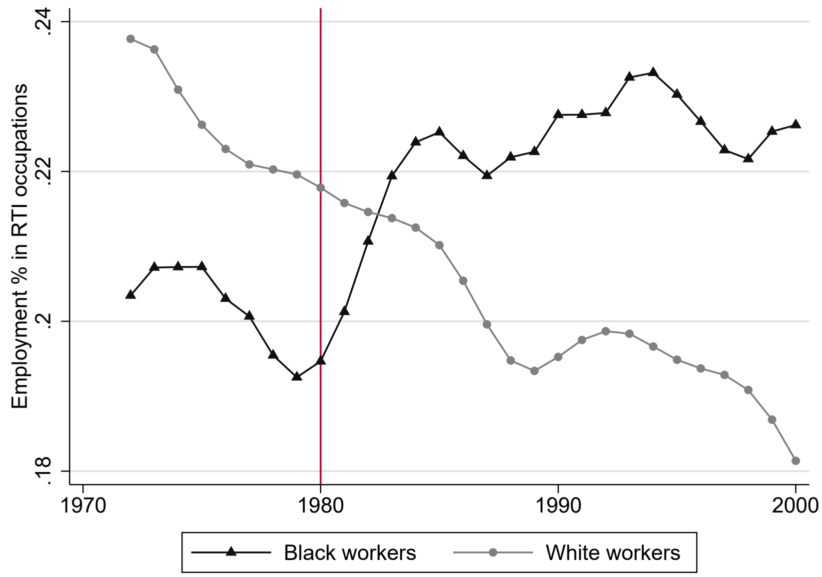

Source: CPS. Sample: All 20-64 black and white workers with high school diploma or less, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

For workers with at least some college education, we only observe differences in the extent to which black and white workers exit routine intensive occupations (see Figure A·4). Thus, workers with a high school diploma or less drive the diverging patterns. The higher prevalence of these workers among the black workforce exacerbates observed differences in the overall working population.

Lastly, I look at whether there are differences in the patterns by worker age (and consequently, experience). I partition workers into three groups: new entrants (age 20-30), prime-age workers (31-50), and those closer to retirement (51-64). Figure 1·5 shows that workers belonging to different age groups exhibit diverse magnitudes in employment share shifts, but that the racial discrepancies are consistent. For white workers, there is a fall for all age groups, with the extent of the drop being larger the older the age group. For all black workers there is a rise in their employment share for routine occupations, but younger workers are those with the bigger magnitude.

Figure 1·5: Race-age group employment share of top RTI occupations

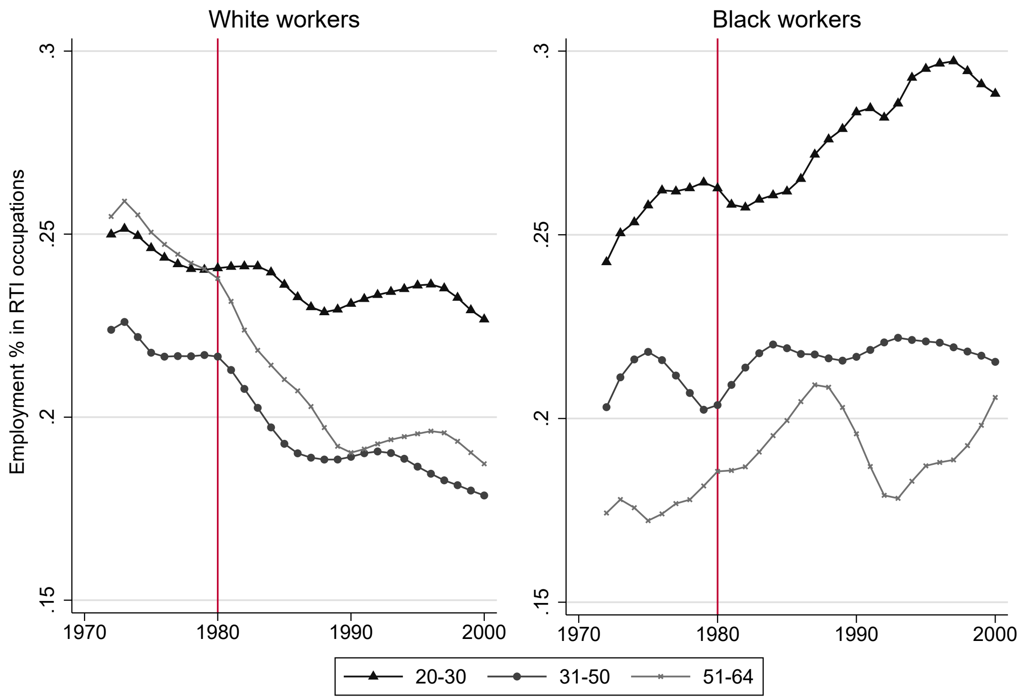

Source: CPS. Sample: All 20-64 black and white workers, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

This finding helps to address a widespread concern about missing black men in survey data. In fact, a well-know phenomenon that impacts differently the composition of the black and white workforce, especially among males, is mass incarceration. The period under analysis in this paper has been characterized by a significant rise in black males incarceration rates.[5]. The consequent higher rate of missing men in survey data (Sabety and Spitzer (2021)) among the youngest group could have potentially be a reason for concern, had we observed a higher decrease in routine jobs employment for this group. Consider an alternative scenario in which the pattern

[5]From about 1% of the black male population in 1980 to roughly 4% in 2000 (National Research Council (2014))

for young black workers was actually decreasing. The overall increase in routine employment for black workers could have resulted from survey data sample selection, with a relative over-representation of older workers. While available data does not allow to determine the types of jobs in which missing men are employed, observing an increase in routine employment for black workers of all age groups provides suggestive evidence that the reliability of these findings is not threatened by sample selection.

1.4 Wage disparities

We have seen that between 1980 and 2000 black and white workers exhibited different employment patterns. In what follows, as advocated by Bayer and Charles (2018), I document how the wage gap has changed along the distribution and relate the changes to the effect of RBTC.

Figure 1·6 illustrates how black and white employment in RTI jobs changed at different points along the wage distribution. Employment of white workers in RTI jobs fell at all points in the distribution except jobs in the bottom 3% and top 10% of the wage distribution[6]. Black workers, instead, increased RTI employment in the bottom part of the distribution, and decreased such employment in the top half. This finding reflects the different patterns observed earlier between high and low-education workers.

[6] The increase at the top of the distribution reflects employment in law and other high-pay RTI occupations.

Figure 1·6: 1980-2000 top RTI change over wage distribution

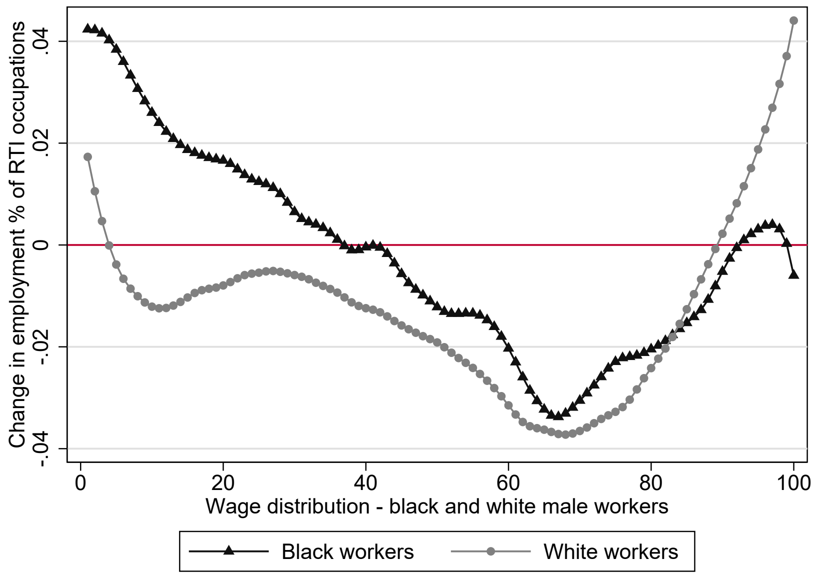

Source: IPUMS. Sample: All 20-64 black and white male workers, with hourly wage $\geq$ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Wage distribution defined on the entire described sample. Smoothing bandwidth: 0.2.

1.4.1 Oaxaca-RIF decomposition

I use a Oaxaca-RIF decomposition to better explain how the racial wage gap in 1980 changed along the entire distribution between 1980 and 2000. This approach extends the Oaxaca decomposition to address statistics other than the mean. The standard decomposition is:

$$\Delta \text{Hourly wage}_{W-B} = X_W\beta_W - X_B\beta_B = \underbrace{(X_W - X_B)\beta_W}_{\text{Composition gap}} + \underbrace{X_B(\beta_W - \beta_B)}_{\text{Differential returns gap}} \tag{1.2}$$

where X is a set of worker's characteristics, the subscript W and B indicate white and black workers.

Equation (2) divides the total change in wage gap into a *composition* element, which captures the difference in characteristics, such as educational achievement or the occupations in which workers are employed, and a *differential returns* element, which captures the racial disparities in the effects of characteristics on earnings. We can further decompose each of the two elements into specific characteristics, in order to, for instance, show the role played by education and occupations separately.

The Oaxaca-RIF method (Firpo et al. (2009)) relies on Recentered Influence Functions. It involves choosing a baseline group on which the effect of the economic shocks of interest is minimized. The relevant statistics (percentiles of the wage distribution, in this paper) are expressed as the average of the conditional expectation of the RIF given the covariates (Firpo et al. (2018)). A major advantage of this method is that it allows for non-sequential decomposition, as in the standard Oaxaca decomposition.

I control for the following characteristics: demographic - 9 5-year age groups and a dummy for married individuals; education - 4 schooling levels; occupation - 12 categories (Autor and Dorn (2013)); industry - 14 categories (Firpo et al. (2018)) and a dummy indicating employment in the private sector; geography - 4 regions and 3 categories for rural areas, metropolitan areas, and mixed ones. The baseline group consists of white non-married workers, age 40-45, with a high school diploma, employed in mechanical/mining/construction occupations in the construction industry in the private sector, living in the northeastern region of the country and in a non-metropolitan area.

Figure 1·7 shows that in 1980 the black-white wage gap was greater at lower percentiles of the wage distribution. This heterogeneity was mainly driven by the differential returns component: in 1980, racial differences in pay were higher for those who were paid less, and these workers tended to have less education. Therefore, it is not surprising that these black workers would be more highly penalized in the labor market. Of course, these disparities need not reflect only race discrimination.

Figure 1·7: 1980 white-black wage gap along wage distribution

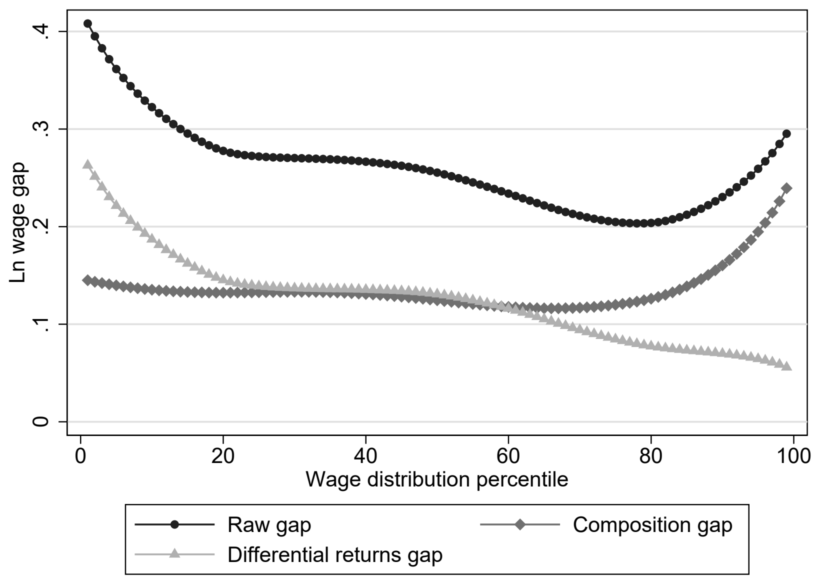

Source: IPUMS. Sample: All 20-64 black and white male workers, with hourly wage ≥ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Race-specific wage distribution. Smoothing bandwidth: 0.4.

Next, I obtain Oaxaca-RIF decompositions for 1980 and 2000 and show how the gaps have changed. This allows me to calculate the changes in the composition and differential returns gaps:

$$\begin{aligned}
&\Delta\text{Hourly wage}^{2000}_{W-B} - \Delta\text{Hourly wage}^{1980}_{W-B} = \\
&\underbrace{(X^{2000}_W - X^{2000}_B)\beta^{2000}_W - (X^{1980}_W - X^{1980}_B)\beta^{1980}_W}_{\text{Composition gap}} \\
&+ \underbrace{X^{2000}_B(\beta^{2000}_W - \beta^{2000}{}_B) - X^{1980}_B(\beta^{1980}_W - \beta^{1980}_B)}_{\text{Differential returns gap}}
\end{aligned} \tag{1.3}$$

Figure 1·8 reports the outcome of this exercise. Recall that because it reports a white-black wage gap change, a *negative* coefficient implies a *reduction* in the gap, and vice versa.

Figure 1·8: 1980-2000 white-black wage gap change along wage distribution

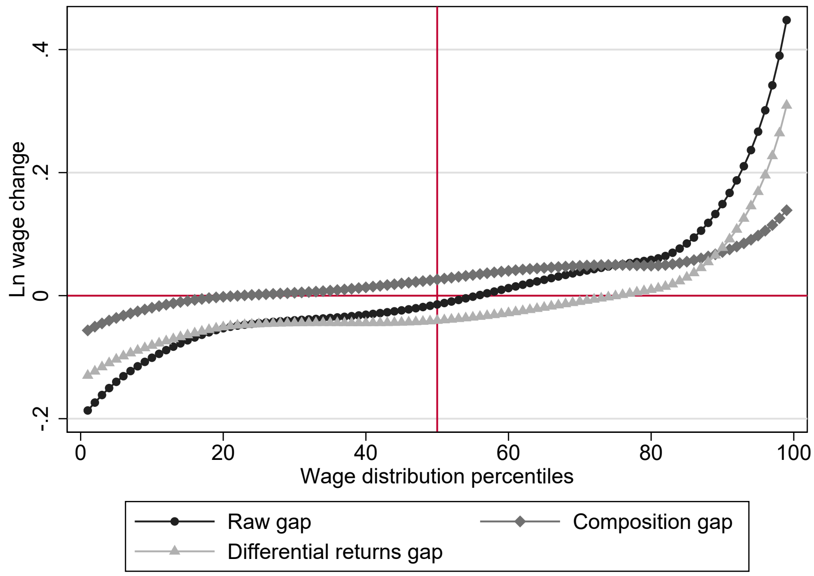

Source: IPUMS. Sample: All 20-64 black and white male workers, with hourly wage $\geq$ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Race-specific wage distribution. Smoothing bandwidth: 0.4.

Between 1980 and 2000, the racial wage gap decreased in the bottom half of the distribution and increased in the top half. Consistent with the prior literature, the change in the gap at the median is approximately 0. Strikingly, Figure 1·8 shows the composition element does reduce the gap. While, throughout most of the wage distribution, differential returns declined[7], the different composition of the black and white workforce increased the earnings gap.

1.4.2 Detailed Oaxaca-RIF decomposition

In Figure 1·9 I further decompose the composition change. Here, I show only the education and occupation components, the two major drivers of the change (see Appendix Figure A·6 for all components).

[7]Given the focus of this paper, I will not devote space to the changes in the top 20% of the distribution. The increased gap in this range is consistent with Hurst et al. (2021) and Bayer and Charles (2018) who ascribe it to the increased gap in college graduation and the increased return to abstract skill.

Figure 1·9: 1980-2000 white-black detailed *composition* change

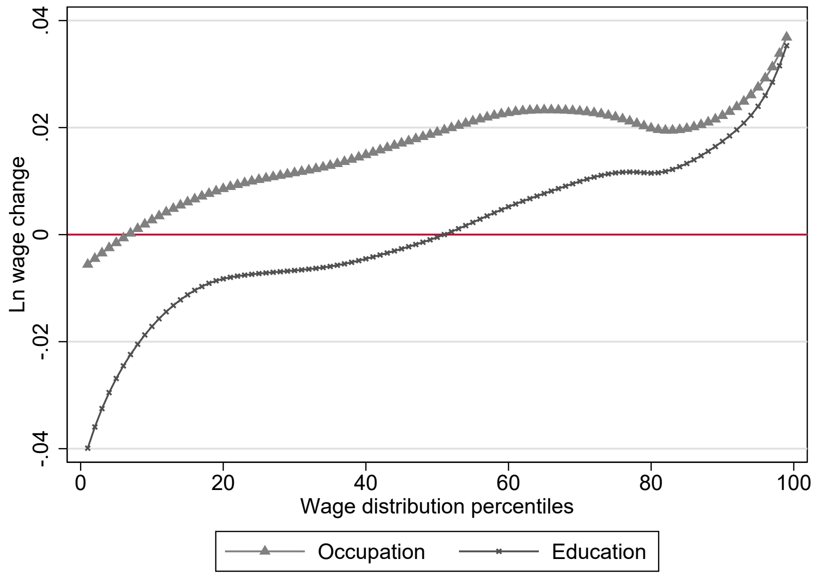

Source: IPUMS. Sample: All 20-64 black and white male workers, with hourly wage $\geq$ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Race-specific wage distribution. Smoothing bandwidth: 0.4.

The most striking finding is that shifts in the occupational component drive the increase in the wage gap explained by composition change. The education component alone, instead, would decrease earnings disparities in the bottom half of the distribution, consistent with the increase in black educational attainment. However, in the top half of the distribution, the education component increases gap (although less so that the occupational component), consistent with the increase in the college graduation gap between black and white men.

From equation (3), we know that the change in the occupational component is itself a combination of two changes: the change in the return to each occupation group and the change in the difference in occupations held by white and black workers.

$$\underbrace{(X_W^{2000} - X_B^{2000})\beta_W^{2000} - (X_W^{1980} - X_B^{1980})\beta_W^{1980}}_{\text{Composition gap}} =$$

$$\underbrace{[(X_W^{2000} - X_B^{2000}) - (X_W^{1980} - X_B^{1980})]\beta_W^{2000}}_{\text{Change in sorting differences}} - \underbrace{(\beta_W^{2000} - \beta_W^{1980})(X_W^{1980} - X_B^{1980})}_{\text{Change in occupational returns}}$$

Figure 1·10 shows the each component's effect. I show findings separately for each occupational group included in the Oaxaca decomposition specifications. Groups labeled "A" are those in which abstract is the most important skill, those labeled "R" involve the use routine skills mostly, and those labeled "M" are jobs that require the use of manual skills most of all.

Changes in the occupations held by black and white workers drive the 1980-2000 difference in occupational returns. Interestingly, the only group for which changes in occupational sorting lead to notable increase in the wage gap is clerical jobs, which are the most routine intensive ones. Overall, Figure 1·10 shows important shifts in returns to occupations over the observed period. Black workers were penalized by these changes, as they were employed in occupations that were more likely to be negatively affected, and did not leave this occupations in sufficient numbers to take advantage of the aftermath of RBTC.

Figure 1·10: 1980-2000 detailed occupation components

Source: IPUMS. Sample: All 20-64 black and white male workers, with hourly wage ≥ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Race-specific wage distribution. Positive coefficients are equal to an increase in the wage gap

Figure A·7 in the appendix shows the detailed decomposition of the differential returns component. Most of the fall in the wage gap explained by this component is driven by the race fixed effect, which can be seen as a proxy (although imperfect) of a decrease in racial discrimination over this time period. The only element that has a contrasting pattern is geography, which by itself would lead to an increase in the differential returns component of the wage gap.

1.4.3 Spatial analysis

Next, I show whether we observe different convergence patterns for geographical areas that were more or less exposed to RBTC. In fact, given that this shock had effects on employment sorting and occupational returns, we would expect to see a narrower convergence in racial wage gaps in labor markets that were affected more by RBTC. As mentioned in the Data section, I will use Commuting Zones (CZ) as the geographic unit of interest, as they identify local labor markets. In the spirit of Autor and Dorn (2013), I split the sample in two: top RTI CZs, which are the top third commuting zones for routine employment share in 1980, and bottom RTI CZs, which are the bottom two thirds in terms of the same measure.

Figure 1·11: 1980-2000 white-black wage gap total change by CZ groups

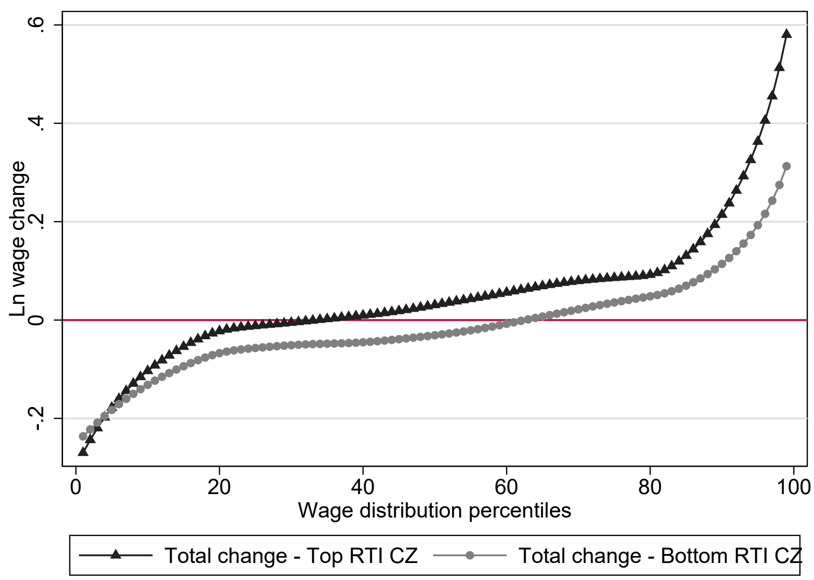

Source: IPUMS. Sample: All 20-64 black and white male workers, with hourly wage ≥ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Race-specific wage distribution

Figure 1·11 shows that, indeed, the wage gap narrowed only for the bottom 40% of the wage distribution in top routine CZs, while the reduction in bottom routine CZs affected the bottom 60% of the distribution. Figure 1·12 below, shows the distribution of the composition and of the differential returns component for the two groups of CZs.

Figure 1·12: 1980-2000 white-black wage gap change decomposition by CZ groups

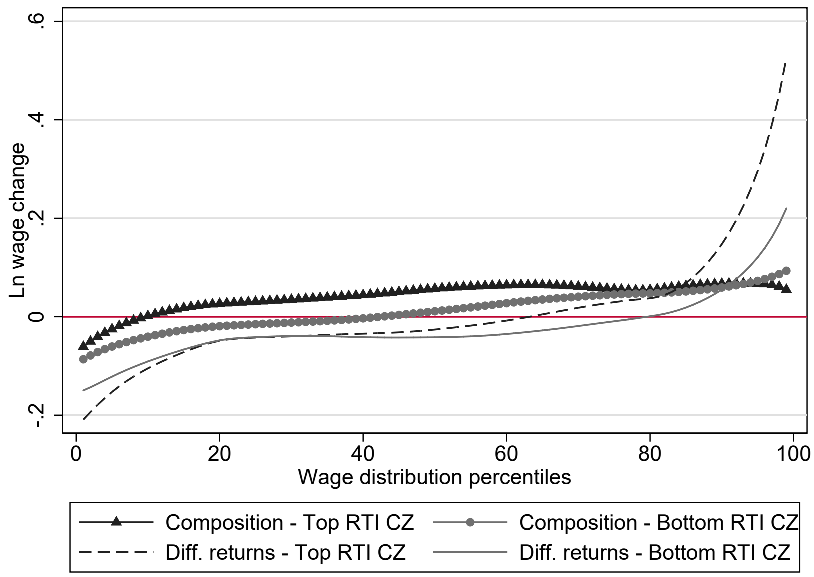

Source: IPUMS. Sample: All 20-64 black and white male workers, with hourly wage $\geq$ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Race-specific wage distribution. Smoothing bandwidth: 0.4.

We observe that at the bottom of the distribution, the differences in changes for the differential returns component are very slim, and for both the top and the bottom routine CZs this component alone would lead to a reduction of the wage gap for the bottom 60%. For the composition element, instead, the difference between the two groups in more significant, and in line with our priors: in CZs that were more affected by RBTC, the change in the composition element thwarts wage gap convergence, while in less affected CZs this change contributes to the narrowing of the wage gap observed in the bottom half of the distribution.

1.5 Model

I now develop a 2-period 3-sector statistical discrimination model to rationalize the empirical findings. The standard model (Acemoglu and Autor (2011)) of routine-biased technological change (RBTC) predicts that no group of workers increases its routine-intensive employment although some may be affected less than others (e.g., men vs. women). However, I show that the interaction of statistical discrimination with RBTC can increase the share of black workers in routine-intensive occupations.

1.5.1 Setting - Labor Demand

There is a continuum of perfectly competitive employers, each operating in one of three sectors. The matching of employers to sectors is exogenous and time-invariant. Hence, any technological change taking place between the two analyzed periods affects employment through overall allocation of workers across sectors, ignoring possible firms entry and exit, as well as changes of sector in which they operate. This simplifies the problem at hand, implying that firms are affected by technological shocks only insofar as the production function is affected, and consequently the type and number of workers they employ.

As is common in the RBTC literature, employment is partitioned into three sectors: manual, routine, and abstract. Manual occupations are the least skill sensitive and abstract the most, with routine occupations in between. Routine jobs involve completing standardized tasks, and, hence, higher ability raises productivity less than in abstract jobs. In equilibrium, the highest-skill workers will be matched with abstract jobs and have the highest wages.

The sector-specific production functions are given by:

$$v_{ij}(\theta_i) = \alpha_j + \beta_j\theta_i \tag{1.4}$$

where v_{ij} is the value of the output produced by a worker of ability θ_i, and α_j and β_j are sector-specific parameters where β_j captures the sensitivity of output value to individual ability. Denoting the manual, routine, and abstract sectors as M, R, and A, I assume

$$0 = \beta_M < \beta_R < \beta_A \tag{1.5}$$

$$\alpha_M > \alpha_R > \alpha_A \tag{1.6}$$

Condition (2) ensures the routine-sector productivity is less responsive than abstract-sector productivity to workers' ability. For simplicity, I assume that manual-sector productivity is independent of worker's ability.

Condition (3) ensures that for some workers it is optimal to be employed in a sector with lower productivity and wage. Consider for instance the case in which $\alpha_M \leq \alpha_R$. Employment in the manual sector will be equal to 0, because for all individuals, regardless of their signal, it will be optimal to work in the routine sector, given the linearity of the wage schedule and the higher return of $w_R(\theta)$ to signals.

1.5.2 Setting - Labor Supply

Workers are endowed with ability $\theta_i \sim \mathcal{N}(\mu, \sigma)$. Note that ability is unidimensional. If there are multiple skills, they can be combined into a single interval scale.

Employers observe a noisy signal of ability given by:

$$s_{ir} = \theta_i + \varepsilon_{ir} \tag{1.7}$$

where ε_{ir} is a race-specific (black or white) error and is independently and normally distributed: $\varepsilon_{ir} \sim \mathcal{N}(0, \sigma_r)$. I make the standard assumption that the ability signal is less informative is less informative for black than for white workers, i.e. $\sigma_B > \sigma_W$. The distribution of s_{ir} is therefore:

$$s_{ir} \sim \mathcal{N}(\mu, \sigma + \sigma_r) \tag{1.8}$$

1.5.3 Equilibrium

Workers choose employment in the sector that maximizes their wage. This implies that the wage in each sector, the wage is given by:

$$w_{ij}(s) = E[v_{ij}(\theta|s)] = E[\alpha_j + \beta_j\theta_i|s] = \alpha_j + \beta_j E[\theta_i|s]]$$

and by the standard properties of the bivariate normal distribution:

$$w_i(s) = \alpha_j + \beta_j \left[\left(1 - \frac{\sigma^2}{\sigma^2 + \sigma_r^2} \right) \mu + \frac{\sigma^2}{\sigma^2 + \sigma_r^2} s \right] \tag{1.9}$$

The wage is a weighted average of mean ability, which is not race specific, and of the signal. Because the signal is less precise for black workers compared to white workers, their wage puts more weight on the mean. Thus black and white workers with same signal will (almost always) receive different wages.

The wage schedule ultimately depends on signal s, and the cutoff values that determine sorting into the three sector are:

$$\begin{aligned}
&w_M(s^*_{rM}) : \alpha_M \\
&w_R(s^*_{rR}) : \alpha_M = \alpha_R + \beta_R \left[\left(1 - \frac{\sigma^2}{\sigma^2 + \sigma_r^2} \right) \mu + \frac{\sigma^2}{\sigma^2 + \sigma_r^2} s^*_{rR} \right] \\
&w_A(s^*_{rA}) : \alpha_R + \beta_R \left[\left(1 - \frac{\sigma^2}{\sigma^2 + \sigma_r^2} \right) \mu + \frac{\sigma^2}{\sigma^2 + \sigma_r^2} s^*_{rR} \right] \\
&= \alpha_A + \beta_A \left[\left(1 - \frac{\sigma^2}{\sigma^2 + \sigma_r^2} \right) \mu + \frac{\sigma^2}{\sigma^2 + \sigma_r^2} s^*_{rA} \right]
\end{aligned} \tag{1.10}$$

The equilibrium is depicted in the figure below, where the thick line indicates the wage schedule in equilibrium:

Figure 1·13: Labor market equilibrium in period 1

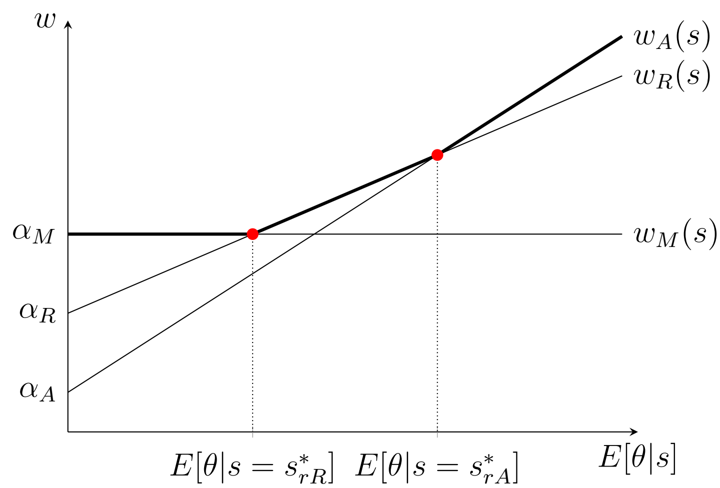

All workers are employed in equilibrium and they are sorted into sectors on the basis of their signal. $E[\theta|s = s^*_{rR}]$, $E[\theta|s = s^*_{rA}]$ are the conditional expectations of ability corresponding to the signal cutoff values indicated in (10).

What are the implications of this equilibrium for racial disparities in terms of employment sorting and wages? Given the wage schedules derived in (10), we can focus on the distribution of signals, rather then the distribution of $E[\theta|s]$. For any given cutoff wage, the racial difference in signal variance implies that the underlining ability signal is also different. Using (10) the relation between the signals for black and white workers at a given cutoff is:

$$s_W = \frac{\sigma_B^2 - \sigma_W^2}{\sigma^2 + \sigma_B^2}\mu + \frac{\sigma^2 + \sigma_W^2}{\sigma^2 + \sigma_B^2} s_B \tag{1.11}$$

Therefore, for any expected productivity cutoff above the mean, $s^*_B > s^*_W$, while $s^*_B < s^*_W$ for cutoffs below the mean. This is a straightforward consequence of the racial difference in signal noise; because the signal is less precise for black workers, a given signal implies expected productivity closer to the mean than it would for white workers.

I assume that no sector employs more than half the workforce. Therefore, $E[\theta|s = s^*_{rR}] < \mu$ and $E[\theta|s = s^*_{rA}] > \mu$, and the distance between the two is such that the routine sector share of employment is below 50% .[8] The equilibrium is shown in the figure below.

Figure 1·14: Equilibrium in period 1 - workers' signals

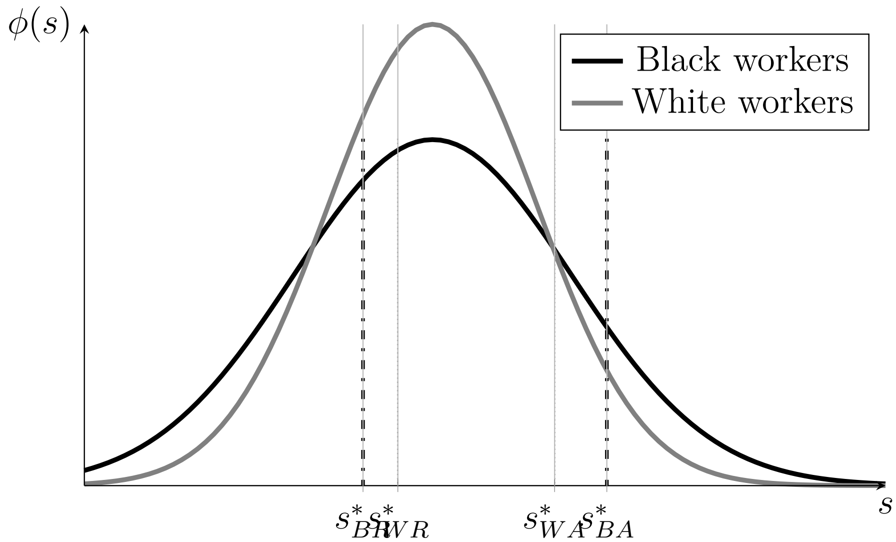

Notice that for each cutoff ($E[\theta|s]$) and, therefore, implicitly each wage cutoff, there are different cutoff signals for each group. Cutoffs above the mean require higher signals for black workers than for white workers, while the opposite is true for cutoffs below the mean. On either side of the mean, the closer the original cutoff to the center of the distributions, the narrower the gap between the required racial signals.

1.5.4 Technological Change and New Equilibrium

How does routine-biased technological change affect this labor market framework? From the empirical analysis, the model should be consistent with 1) an overall con-

[8]In this setting, counterfactually, relative to black workers, white workers always have a higher proportion of workers in each of abstract and manual employment. As discussed in the *Discussion* subsection, this concern is easily addressed by allowing μ to be race specific.

traction of the routine sector, and 2) an increase in the routine-employment share among black workers.

I model RBTC as reducing the responsiveness of output to individual skill in routine jobs. Personal computers allowed less skilled workers to do jobs for which they were previously unqualified. Typesetting use to be done by highly skilled typographers. Desktop publishing programs allowed other, less skilled, workers to do the same job. While technology also increased the productivity of skilled typographers, it lowered the slope of the output/skill relation, and by increasing output, lowered the price. Therefore, I model technology as rotating the wage/skill line segment in routine jobs. Such a technological shock corresponds to a change in period 2 of the value production function for the routine sector, with an increase in α_R (worker-independent productivity) and a decrease in β_R (sensitivity to worker-specific ability).

The new equilibrium is shown below. The shift and pivot of $v_{R,t1}$, associated with the changes in α_R and β_R are represented by the dashed line $v_{R,t2}$:

Figure 1·15: Labor market equilibrium in period 2

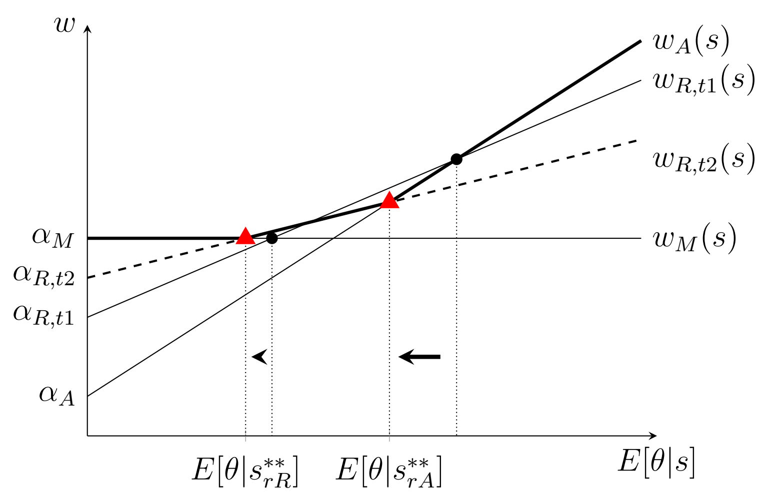

The new wage schedule changes the cutoffs that determine employment sorting to $E[\theta|s = s_{rR}^{**}]$, $E[\theta|s = s_{rA}^{**}]$, shifting both to the left. In the example, the parameter changes lead to a shrinkage of the routine sector. This ultimately depends on the densities of the signal distributions, hence on whether the density change caused by the shift for the cutoff between manual and routine sectors is smaller than the one brought about by the cutoff between routine and abstract sectors.

These parameter assumptions imply that employment share of the manual sector falls. This appears to contradict evidence that low-skill service employment increased (Autor and Dorn (2013)). However, this growth is primarily after 2000[9] and, therefore, does not apply to the period I study. In addition, I abstract from the distinction between service occupations and occupations in transportation and blue collar trades; during the 1980s and 1990s the decline in the latter outweighed the increase in the former occupation.

These changes in the threshold values for signals are not sufficient to generate the racial differences changes in employment-shares changes observed in the data. The following conditions are necessary and sufficient for an increase in the routine_employment share for black workers and a decrease for white workers:

$$|\Phi(s^{**}_{WA}) - \Phi(s^{*}_{WA})| > |\Phi(s^{**}_{WR}) - \Phi(s^{*}_{WR})| \tag{1.12}$$

$$|\Phi(s^{**}_{BA}) - \Phi(s^{*}_{BA})| < |\Phi(s^{**}_{BR}) - \Phi(s^{*}_{BR})| \tag{1.13}$$

The first condition states that the changes in $E[\theta|s = s^{*}_{rR}]$ and $E[\theta|s = s^{*}_{rA}]$ have to be such that for white workers the shift out of the routine sector (LHS) is larger than the shift into the routine sector (RHS) while the second condition requires the opposite for black workers.

In the example below, I illustrate how these conditions are portrayed in terms of signal distributions.

[9]See Figure 10 in Acemoglu and Autor (2011)

Figure 1·16: Equilibrium in period 2 - Black workers

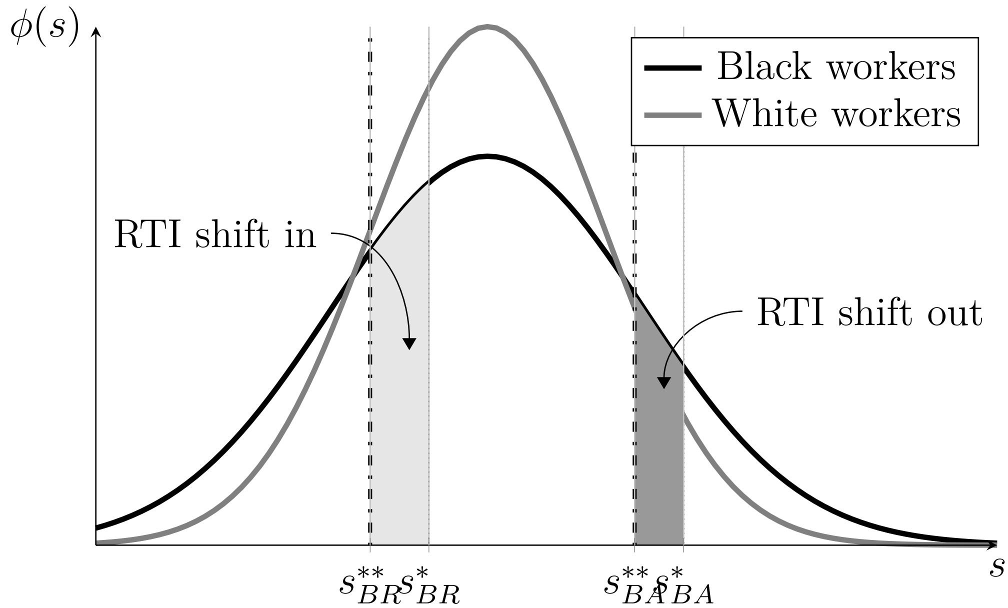

Figure 1·17: Equilibrium in period 2 - White workers

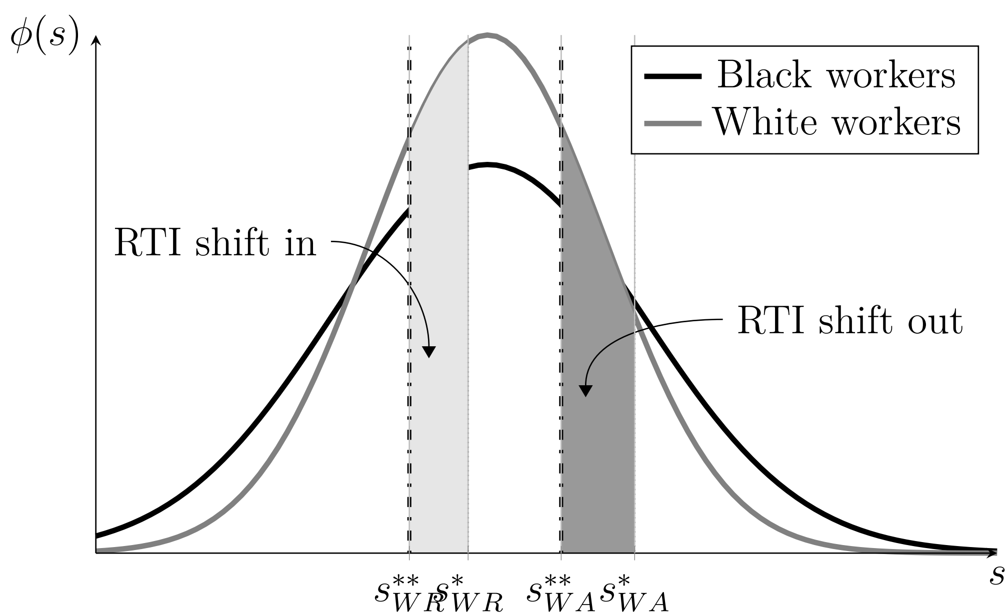

In both Figures 1·16 and 1·17, dashed lines denote cutoffs in period 2 (labelled with **), while lighter dotted lines represent cutoffs from period 1. As previously stated, a labor market shift in wage cutoffs corresponds to different changes for signals of black and white workers. In Figure 1·16, it is possible to observe workforce shifts for black workers, with the darker shaded area representing movement out of the routine sector, and the lighter shaded area movement into it. Similarly, Figure 1·17 represents employment shifts for white workers. Overall, the two pictures show a case in which a change in the wage cutoffs of routine and abstract sectors can cause a net shift into the routine sector for black workers, and out of it for white workers.

1.5.5 Discussion

The model is designed to rationalize the racially different patterns with respect to routine-intensive occupations by enriching a framework of routine-biased technological change with statistical discrimination.

I have modeled RBTC differently from the approach in the task literature. This literature (Autor et al. (2003), Acemoglu and Autor (2011), Autor and Dorn (2013)) models RBTC as a fall in computer prices that leads to substitution of workers by machines in the routine sector. Consequently, this sector contracts and wages fall. The standard model predicts that employment shifts from low-wage routine to manual jobs, which does not allow for different racial patterns.

The narrative I introduce is more closely related to Cavounidis et al. (2021). There, we show how increases in productivity of a given skill can lead to a decrease in employment in occupations that use it intensively.

The model is highly stylized, but provides a simple framework that explains the otherwise counter-intuitive empirical findings I presented earlier. At the same time, without modification, it is inconsistent with some other regularities. In what follows, I acknowledge and briefly address these concerns.

1. *Increase in returns to abstract occupations*

I have not addressed the increased return to abstract skills, which exacerbated racial wage disparity (Hurst et al. (2021)). My model can easily be extended to allow for an increase in β_A, which would further reduce routine employment, especially among white workers. Provided the density of black workers around the cutoff is sufficiently small, nothing else would change.

2. *Racial-specific share of employment in manual occupations*

The model predicts counterfactually that higher share of white than of black workers will be in manual employment. This reflects the assumption that the two ability distributions have the same mean. The model could easily be extended to allow $\mu_B < \mu_W$. This would make it easier to have a much higher density of black workers around the manual-routine cutoff and a lower density around the routine-abstract cutoff, reinforcing the main result. It would not, however, alter the fundamental features of the model. Note that $\mu_B < \mu_W$ can be endogenous to statistical discrimination as in Lundberg and Startz (1983), where statistical discrimination leads to less unobserved investment among black workers.

3. *Static model*

Finally, the model is static and thus abstracts from changes over time in the composition of labor force.[10] The focus on male workers in the empirical work makes this issue less problematic, since, over time, they did not change their labor force participation as much as women did.

1.6 Conclusion

This paper studies the role of Routine biased technological change (RBTC) in hindering the black-white male wage gap convergence in 1980-2000. I show that this economic shock has a differential impact across races and along the wage distribution. This results in unfavorable effects on racial disparities despite the advancements

[10] This implies that the framework doesn't contemplate training on the job, nor updating beliefs for observed ability by either the employer or the worker after job-matching.

achieved through increased levels of education for black workers. I present three new major empirical facts. First, I observe an increase in the share of employment in routine intensive occupations for black workers, while there is a significant decrease for white workers. This is a surprising pattern in light of the current literature, which predicts instead that employment in these jobs should decline for all workers. Second, I show that this racially-different trend holds when conditioning for levels of education and age brackets, with the increase in black routine intensive employment being driven workers with a high school diploma or less, and stronger for younger cohorts. Lastly, using the Oaxaca-RIF decomposition methodology, I show that changes in the composition of the workforce increase wage disparities, opposing the observed wage convergence between races at the bottom of the wage distribution. This finding is mainly explained by the occupational sorting differences among the two groups, and counters the concurrent decrease in racially-differential returns. I rationalize these empirical findings with a statistical discrimination model, characterized by three sectors (manual, routine, and abstract) with different productivity sensitivity to individual skills, and two races of workers whose ability is noisily observed. I demonstrate that a technological shock decreasing skill-responsiveness for the middle sector has predictions in line with empirical patterns: it shrinks the overall share of employment in routine occupations, while causing a net shift into (out of) routine jobs for workers with higher (lower) ability noise.

This paper contributes to our understanding of the determinants of trends in racial disparities. In future research, I plan to expand the time horizon of this analysis and study how RBTC has interacted with other major economic events, such as the fall of unionization and the Chinese trade shock, and how their concurrence has influenced the racial wage gap. Another interesting area for future work involves extending the analysis on RBTC impacts to black and white women.

Chapter 2

Estimating the Nature of Technological Change: Exploiting Shifts in Skill Use Within and Between Occupations

2.1 Introduction

Consider the IBM Selectric, an electronic typewriter introduced in 1961. It replaced the traditional strikebars with a golf-ball-like element and, in later versions, was even 'self-correcting.' The Selectric made typing much more productive. Secretaries and typists could produce many more and more attractive typewritten pages. Typists who quickly caught a mistake could correct it invisibly, where previously, they retyped the page entirely.

One could say this change was biased towards specific tasks in the economy. In particular, it made typists far more productive than they were earlier. However, the interaction with task demand seems ambiguous. The Selectric did not replace typists or substitute for typing outputs, the approach used by the task model (Autor et al. (2003), Acemoglu and Autor (2011), Acemoglu and Restrepo (2018)). If anything, the demand for typed pages increased after the Selectric's introduction.

However, the Selectric did not make high-skilled labor more productive or more in demand, as in the canonical SBTC models (Katz and Murphy (1992), Berman et al. (1994), Berman et al. (1998), Juhn (1999)). Instead, it simply increased the speed at which everyone could type. Thus, the individuals whose productivity increased dramatically were primarily middle-skilled workers in typing-intensive jobs. Further, this technological innovation might have changed the skills used in specific jobs, e.g., by expanding the set of occupations involving some typing.

Thus, the Selectric had two effects within occupations: a primitive effect, which made typing significantly more productive, and an adaptation effect, in which workers responded to the primitive effect by altering their use of typing skills. At first glance, one may expect that the second effect acts as a confounding factor, making it hard to recover the primitive technological change. However, we demonstrate that, in a relatively general model, a first-order approach allows pure employment changes to act as a "sufficient statistic" to identify relative skill productivity changes.

We integrate insights from both task-based and skill-biased approaches in a tractable general equilibrium framework. We model occupations as combining skills, akin to tasks in the tasks model, to produce intermediate goods. Occupations utilize workers' skills in heterogeneous ways, and we impose nearly no structure on their production technologies. Our model allows workers to choose first the skills they develop and then their occupations. The market prices intermediate goods that aggregate to a single final good.

As in the SBTC model, we allow for skill-enhancing technological change. The Selectric made typing - or 'finger dexterity' - more productive. The number of typed pages produced must increase, but whether workers deepen their typing skills or not depends on the elasticity of substitution between skills in each specific occupation. Finger-dexterity use (typing) can decline in one occupation (secretaries) but increase in another (economics professors). Employment in typing-intensive occupations can increase or decrease. If the elasticity of demand for such an occupation's output is less than one, as we believe plausible, demand for that occupation falls.

We also allow for shifts in product demand (possibly due to trade shocks) or outside competition (possibly due to robots or offshoring) that alter the demand for workers with different skills. The model thus clarifies the distinction between technological changes to the productivity of individual skills and changes to demand for particular kinds of workers.

The model provides us with a simple approach to measuring the relative increase in the technological productivity of different skills while taking demand shifts into account. In effect, we develop a transparent structural model for interpreting within and between-occupation changes in skill use that, for local estimation, relies only on ordinary least squares and weighted means, using readily observable variables.

Note this is orthogonal to the exercises conducted by Autor et al. (2003) and Goos et al. (2014). They use routine intensity as a measure of vulnerability to technological change, and study the degree to which it drives employment changes. Our approach instead leverages employment changes to identify what is the relevant technological change. Autor et al. (2003) observe the correlation between computerization and performance of routine tasks, and show that this type of technological change can provide relevant insights in understanding changes in employment in the United States. Goos et al. (2014) study the role of routine-biased technological change and offshoring in explaining changes in employment in 16 European countries, providing evidence of a much bigger role played by the former.

To estimate the model, we use a subset of the skills studied by Autor et al. (2003) and measured in the Dictionary of Occupational Titles, using the third edition for skill use in 1960, the original fourth edition for 1971, and the revised fourth edition for 1983. We combine these measures with data from the Current Population Surveys and Censuses to measure between and within-occupation changes in skill use from 1960 to 1983.

We find that workers moved into abstract-intensive occupations in both periods, with the shift for women being slower in the earlier period but much faster in the later period. Our estimation exercise shows that relatively rapid growth of finger-dexterity productivity and slow growth of abstract-skill productivity explain between-occupation shifts, especially among women.

We also show that within-occupation shifts can dwarf those due to movement across occupations. In the earlier period, abstract-skill use among men grew within occupations while routine-skill use fell. This pattern slowed for men but accelerated for women in the later period. We show how complementarity and substitutability of skills with respect to their own and other skills' growth in productivity explain these patterns.

We are not the first to look at within-occupation changes in skill use. Black and Spitz-Oener (2010), using German data, and Deming and Noray (2020), using Burning Glass data, track significant within-occupation shifts in skill use, but for a later period. Atalay et al. (2020), using keyword frequencies from three newspapers' job ads over an impressively long period, show that within-occupation changes account for most task variation over time.[1] However, we develop a model to help us interpret the results. Moreover, Atalay et al. (2020) are unable to examine gender differences. Autor and Price (2013) also study a very long period and decompose changes by gender but do not allow for within-occupation changes in skill use.

This paper can be read in two ways. Those interested solely in a better accounting of the changes in the 1960s and 1970s can jump to the data section and then examine Tables 2.1 and 2.2 and the accompanying text in the results section. We think this analysis is a contribution in its own right. However, we are hopeful that readers will find that the model presents a simple, versatile framework allowing for different kinds of technological shocks, and therefore assists in thinking about our results and the

[1] Autor et al. (2003) examine the relation between computer use and within-occupation change in task use between the 1977 and 1991 revisions of the *DOT*, but do not discuss the magnitudes of these changes.

large literature in this area.

2.2 A model of skill and job choice in general equilibrium

2.2.1 Skill acquisition and intermediate good production

Before employment, each worker chooses a vector of skills $S \in \mathbb{R}^n_+$, where each component S_i reflects ability at task i. Once workers have acquired skills, each chooses a job $J \in \mathcal{J}$, where $\mathcal{J}$ is the set of all jobs. If a worker with skills S is employed at job J, she produces a quantity $y((A_iS_i)_{i\leq n}, J)$ of intermediate good J, where each $A_i > 0$ is common to all jobs and is a measure of the general productivity of skill i. Thus, each A_iS_i is the 'effective' amount of input i.[2]

We place as little structure on $\mathcal{J}$ and y as possible. We assume only that $\mathcal{J}$ is a compact subset of a Euclidian space, that $y(\cdot, J)$ is a constant-returns standard neoclassical production function,[3] and that y is continuous.

For simplicity, we assume that workers have a fixed budget for skills, which we normalize to 1, so that for any individual $\Sigma_i S_i = 1$. This assumption captures the idea that a worker can study plumbing or philosophy, but if she chooses to spend more time on philosophy, she must spend less time learning plumbing. We do not allow her to choose to spend more time on learning.[4]

[2]Thus output y depends on the vector of effective inputs $(A_iS_i)_{i\leq n}$.

[3]$y(\cdot, J)$ is strictly increasing in each A_iS_i on $\mathbb{R}^n_{++}$, is twice continuously differentiable, features a bordered Hessian with non-vanishing determinant on $\mathbb{R}^n_{++}$, is strictly quasiconcave, and $y((A_iS_i)_{i\leq n}, J) = 0$ iff $A_iS_i = 0$ for some i. This will imply that optimal skills are continuously differentiable in A and, more importantly, interior. If skills are quite occupation-specific, e.g., plumbing or surgery skills, this may be a bad assumption; however, the skills used in our empirical section are relatively general. We thus think that excluding corner solutions is unproblematic for our application.

[4]This is without loss of generality since we can always normalize the time she chooses to spend on learning to 1. This could affect comparative statics on total production through a labor/leisure/learning trade-off. That said, since this only adjusts the effective number of labor units each worker provides, with a constant returns to scale aggregate production function, it will not affect the objects of interest to us.

A worker who anticipates holding job J will therefore

$$\max_{S\geq 0} y((A_i S_i)_{i\leq n}, J) \tag{2.1}$$

$$\text{subject to} \sum_i S_i = 1. \tag{2.2}$$

The optimal $S^*(J)$ and $y^*(J) := y((A_i S_i^*(J))_{i\leq n}, J)$ are given by solving the Lagrangian. The Lagrangian's first order condition at the optimum with respect to any S_i is

$$A_i y_i'((A_i S_i^*(J))_{i\leq n}, J) = \lambda = y^*(J) \tag{2.3}$$

where the second equality follows straightforwardly from constant returns to scale. We assume that workers always have skills that are optimal for the job they perform. Although this assumption is strong, we maintain that in the sort of timescales our empirics cover, workers will at the least endeavor to develop the right skills for the careers they select. Allowing for investment while employed, as in Cavounidis and Lang (2020), would make this a sensible assumption for workers not too far advanced in their work lives.

How do optimal output and skills change with A? From the Envelope Theorem,

$$\frac{\partial y^*(J)}{\partial A_i} = S_i^*(J) y_i'((A_i S_i^*(J))_{i\leq n}, J) \tag{2.4}$$

so that substituting for y_i' using (2.3), we get

$$\frac{\partial \ln y^*(J)}{\partial \ln A_i} = S_i^*(J). \tag{2.5}$$

This is effectively an application of Roy's Identity, with our skill constraint playing the role of the budget constraint in standard utility maximization.

To speak sensibly about the effect of changes in A on $S^*(J)$, we proceed by inspecting $y(\cdot, J)$'s i-j elasticity of substitution for any two inputs at the optimum

$$\sigma_{i,j}((A_i S_i^*(J))_{i\leq n}, J) = \frac{\partial ln\left(\frac{A_i S_i^*(J)}{A_j S_j^*(J)}\right)}{\partial ln \frac{A_i}{A_j}} = 1 + \frac{\partial \ln(S_i^*(J)/S_j^*(J))}{\partial \ln(A_i/A_j)} \tag{2.6}$$

which we can rearrange as

$$\frac{\partial \ln(S_i^*(J)/S_j^*(J))}{\partial \ln(A_i/A_j)} = \sigma_{i,j}((A_i S_i^*(J))_{i\leq n}, J) - 1. \tag{2.7}$$

Thus, if inputs i and j are gross substitutes (complements) in job J at the optimal skill bundle, a relative increase in the productivity of skill i will cause workers to acquire relatively more (less) of it. If all inputs are gross substitutes (complements) in job J at the optimal skill bundle, the constraint that $\sum_i S_i^*(J) = 1$ further implies that $\frac{\partial S_i^*(J)}{\partial A_i} > 0$ (< 0).

2.2.2 Final good production and worker allocation

So far, the model somewhat resembles Cavounidis and Lang (2020) in the sense that workers are aligning their skill choices and occupation choices. We extend it by assuming that instead of goods of intrinsic value, workers produce inputs in a CES final good production function

$$Y(q) = \left[\int_{\mathcal{J}} h(J) q(J)^{\varepsilon} dJ\right]^{\frac{1}{\varepsilon}}. \tag{2.8}$$

Here, $h(J)$ is the relative importance of input J for final production, and $q(J)$ is the total quantity of intermediate good J used as an input. We assume h is continuous. The economy has workers of total measure 1, and each worker acquires skills, subject to the constraint, and may choose any job in $\mathcal{J}$.

The model satisfies conditions under which the decentralized equilibrium is Pareto efficient. Therefore, we solve for the equilibrium by solving the planner's problem subject to the skill acquisition and worker measure constraints. Efficiency implies that workers producing good J will all be identical and acquire skills $S^*(J)$; therefore, $q(J) = y^*(J)f(J)$, where $f(J)$ is the density of workers assigned to producing intermediate good J.

Therefore, we can write the planner's problem as

$$\max_f \left[\int_{\mathcal{J}} h(J) \left[y^*(J) f(J) \right]^{\varepsilon} \right]^{\frac{1}{\varepsilon}} \tag{2.9}$$

$$\text{subject to} \quad \int_{\mathcal{J}} f(J) = 1. \tag{2.10}$$

We can then pointwise differentiate the Lagrangian and obtain

$$h(J) y^*(J)^{\varepsilon} f(J)^{\varepsilon - 1} = h(J') y^*(J')^{\varepsilon} f(J')^{\varepsilon - 1}, \tag{2.11}$$

which we can write as

$$f(J) h(J')^{\frac{1}{1-\varepsilon}} y^*(J')^{\frac{\varepsilon}{1-\varepsilon}} = f(J') h(J)^{\frac{1}{1-\varepsilon}} y^*(J)^{\frac{\varepsilon}{1-\varepsilon}} \tag{2.12}$$

so that we can now integrate out J' and using constraint (2.10) get

$$f(J) = \frac{h(J)^{\frac{1}{1-\varepsilon}} y^*(J)^{\frac{\varepsilon}{1-\varepsilon}}}{\int_{\mathcal{J}} h(J')^{\frac{1}{1-\varepsilon}} y^*(J')^{\frac{\varepsilon}{1-\varepsilon}}}. \tag{2.13}$$

2.2.3 Comparative statics

We consider the effect of technological progress that is broadly skill enhancing, as measured by A, and changes in the demand for intermediate goods, as measured by h. The distinction is imperfect. For example, the reduction in transportation costs, at least partly due to technological change, reduced demand for some locally produced intermediate goods that had hitherto been too expensive to import. Still, we think of changes in A as capturing broad-based technological progress such as electronic calculators rather than adding machines for routine-cognitive skills and electric rather than manual drills for manual skills, and h as capturing the effects of trade and, more recently, robots.

The effect of skill-augmenting technological change

What happens if skill i becomes more productive? Taking the derivative of (2.13) with respect to A_i gives

$$\frac{\partial f(J)}{\partial A_i} = \frac{\varepsilon}{1-\varepsilon} f(J) \left[\frac{\partial \ln y^*(J)}{\partial A_i} - \int_{\mathcal{J}} \frac{\partial \ln y^*(J')}{\partial A_i} f(J') \right] \tag{2.14}$$

or simply, using (2.5),

$$\frac{\partial \ln f(J)}{\partial \ln A_i} = \frac{\varepsilon}{1-\varepsilon} \left[S_i^*(J) - \int_{\mathcal{J}} S_i^*(J') f(J') \right]. \tag{2.15}$$

In other words, if and only if the elasticity of substitution among intermediate goods $1/(1-\varepsilon)$ is less than 1, will an increase in the productivity of skill i move workers away from jobs where it is used more than average, and towards jobs where it is used less than average. So, for example, if routine skill is a complement to other skills in intermediate good production, and intermediate good demand is inelastic, an increase in A_R (a technological change that makes routine skill more productive) will (a) reduce routine use in all jobs (within) and (b) shift workers to less routine-intensive jobs (across).

The idea that sectors experiencing slower productivity growth also experience faster employment growth is old (Baumol (1967), see also Ngai and Pissarides (2007) and Acemoglu and Guerrieri (2008)). We build on that idea In our case, jobs making more use of skills whose productivity grows slowly will experience more employment growth.

The effect of changes in demand for intermediate goods

What about changes in h? In our setup, these will move workers around but have no effect on skill use within a job. A decrease in horseshoe demand merely alters how many people shoe horses, not how they shoe them.

To see the effect of changes in h on employment, we take the log of each side in (2.13) and totally differentiate to get

$$d\ln f(J)=\frac{1}{1-\varepsilon}d\ln h\left(J\right)+\frac{\varepsilon}{1-\varepsilon}d\ln y^{*}(J)-d\ln\left(\int\limits_{\mathcal{J}}h(J')^{\frac{1}{1-\varepsilon}}y^{*}(J')^{\frac{\varepsilon}{1-\varepsilon}}\right). \quad (2.16)$$

For a change in h, the second term in (2.16) is 0 and the third term does not depend on J. A few manipulations yield

$$d\ln f(J)=\frac{1}{1-\varepsilon}\left[d\ln h\left(J\right)-\int\limits_{\mathcal{J}}d\ln h(J')f(J')\right]. \quad (2.17)$$

Thus, the percentage employment growth in job J is proportional to the deviation of the percentage change in $h\left(J\right)$ from the employment-weighted average.

Putting it all together

Combining (2.15) and (2.17), we have

$$\begin{aligned} d\ln f(J) \quad &= \quad \frac{\varepsilon}{1-\varepsilon}\sum_{i}\left[S_i^{*}(J)-\int\limits_{\mathcal{J}}S_i^{*}(J')f(J')\right]d\ln A_i \\ &\quad +\frac{1}{1-\varepsilon}\left[d\ln h\left(J\right)-\int\limits_{\mathcal{J}}d\ln h(J')f(J')\right]. \quad (2.18)\end{aligned}$$

The model distinguishes between changes that replace (or reduce demand for) occupations by automating or offshoring them (a decline in h) as when data input is imported from abroad, and those in which technology makes relevant skills more productive as when keypunch machines are replaced by input at computer terminals.

When h declines, the number of workers employed in data entry in the home country falls, but any workers engaged in data input continue to input data using the same skill set. Suppose the productivity A_i of a skill i important to data entry increases. If skill inputs are complements at data entry and intermediate-good demand is inelastic, workers in data entry jobs end up with less of skill i, and fewer workers are hired to input data.

Interpreted within our model, Autor et al. (2003) found that, in a later period, technological innovation increased the productivity of routine skills. Since the demand for these skills was inelastic, the amount of time individual workers spent on them declined as did total employment in routine-intensive occupations. Our interpretation of the period that we study will be that the productivity of abstract skill use did not increase as rapidly as the productivity of other skills, most notably finger dexterity. This caused a shift towards abstract-skill use because the elasticity of substitution between intermediate goods is less than one, thereby shifting employment to abstract-intensive occupations. Within occupations, declining relative abstract-skill productivity shifted skill use toward greater abstract and less routine-skill use. Strikingly, within occupations increased productivity of finger dexterity, reduced the use of both abstract and finger-dexterity skills, and increased the use of routine skills.

We note that our model assumes *ex-ante* identical workers. In a richer model with *ex-ante* heterogeneous workers, demand changes might alter how jobs are done. Intuition suggests that workers "better at routine tasks" do jobs more routinely than other workers. In such a world, a reduction in demand for routine-intensive outputs *would* shift such workers to less-routine jobs who would then perform them *more* routinely than before, which is the reverse of what we observe.

2.2.4 Implications for empirical work

For empirical analysis, we rewrite (2.18) as

$$\Delta \ln(emp_{I,J}) = \frac{\varepsilon}{1-\varepsilon}\Sigma_i \left(d \ln A_i \left(S_{i,J} - \overline{S}_i\right)\right) + \gamma_I + \mu_{I,J} \tag{2.19}$$

where $\Delta \ln(emp_{I,J})$ is the change in the employment level in industry I in occupation J, the empirical counterpart of $f(J)$ and γ_I is the coefficient on an industry that captures demand changes due to shifts in industry demand. We note that this is an imperfect proxy for changes in h. It will capture changes in demand for an occupation resulting from, for example, import competition **and product demand** but not will capture changes due to occupation-specific factors such as robots. If we performed our analysis in a later period, we would want to include measures of robot adoption or potential for robot adoption. We measure $S_{i,J}$ by its average in two proximate editions of the *DOT*. μ is a mean-zero error term. We estimate (2.19) separately for each gender/time-period pair.

Since each worker's skills sum to 1; skill use on a job sums to 1, as does mean skill use. Therefore, (2.19) still applies if we add a constant term to each $d \ln A$, and we can rewrite the equation as

$$\begin{aligned}
\Delta \ln(emp_{I,J}) &= \frac{\varepsilon}{1-\varepsilon}\Sigma_i \left(\left(d \ln A_i - d \ln \overline{A}\right)\left(S_{i,J} - \overline{S}_i\right)\right) + \gamma_I + \mu_{I,J} && (2.20)\\
&= \frac{\varepsilon}{1-\varepsilon}\Sigma_i \left(\left(d \ln A_i - d \ln \overline{A}\right) S_{i,J}\right) + \gamma_I + \mu_{I,J} && (2.21)\\
&: \;= \Sigma_i S_{i,J}\beta_i + \gamma_I + \mu_{I,J}. && (2.22)
\end{aligned}$$

Equation (2.22) describes a regression of the (approximate) percentage change of employment in an occupation/industry cell on the skills used in that occupation and industry dummies. The coefficients show the change in each skill's productivity relative to the average up to a factor of proportionality. This factor is negative if the elasticity of substitution between intermediate goods is less than 1, which we assume. Thus, a negative coefficient means that the productivity of that skill grew faster than the average of the skills.

Assuming an elasticity less than 1 seems natural. As Jones (2011) notes in a somewhat different context, intermediate goods are unlikely to be substitutes. As he puts it, computers are close to essential for producing some goods. Consistent with this argument, Goos et al. (2014) estimate that the elasticity of substitution across industry outputs is 0.42. Our case is even stronger; the outputs of secretaries, sales workers, plumbers, and truck drivers cannot easily substitute for each other. Note that this is different from the statement that someone who works as a secretary might be almost as productive if he worked in sales. In our model, this is entirely plausible if the underlying skills required are close.

Note that we must drop a skill because the skills sum to 1. Therefore, we can interpret the coefficients as the rate of growth of productivity of each skill relative to the excluded skill, again up to a multiplicative factor. Together with the requirement that the sum of the deviations from average productivity growth equals 0, this fully identifies the relative productivity of all the skills.

We estimate (2.22) by ordinary least squares. Consistent with Solon et al. (2015) and Dickens (1990), we experimented with feasible weighted least squares and found no evidence of important heteroskedasticity with respect to occupation size.

Although derived quite differently, our equation (2.22) is similar to the one in Goos et al. (2014). Their theoretical model includes wages in the equivalent of (2.22), which they proxy by industry-year and occupation dummies. [5] Since we first-difference the data and estimate the model separately for each pair of years, we implicitly control for occupation and year, while explicitly controlling for industry. They also use an alternative specification in which they explicitly control for wages, but do not include

[5] We ignore the country component since we study only one country.

it in the main text as there are concerns about endogeneity. While we agree with such concerns, we perform the same exercise, and observe that the inclusion of wages does not alter the outcome of our analysis.[6] The major difference in our specifications is that they include only routine-task intensity and not the other skills but also include offshorability, something of limited importance in our period.

2.3 Data

Following Autor et al. (2003), our skill-use measures come from the *Dictionary of Occupational Titles (DOT).* We use the third edition, issued in 1965 but compiled starting sometime after the release of the second edition in 1949, as our measure of skill use in an occupation in 1960 although it may be centered more on the late 1950s. The 1965 *DOT* has not, to the best of our knowledge, been used previously for this type of analysis. We use the fourth edition, published in 1977 and based on data starting in 1965 for job use in 1970-72 ('1971'). Finally, we use the last revision of the fourth edition, based on revisions from 1977 to 1991 for skill use in 1982-84 ('1983'). The files for the 4th and revised 4th versions of the *DOT* come from Autor et al. (2003). As others have noted, the revised fourth edition is not a 'fifth' edition; many occupations were not revisited between the fourth edition and the revised 1991 edition because the revision addressed only occupations believed to have changed the skills they used. Therefore, we probably underestimate the extent of within-occupation changes in skill use between 1971 and 1983. However, we observe differences between the 4th and revised 4th editions for all but 41 of the 528 gender/occupations cells present in both 1971 and in 1983.[7]

[6] We do not model wages in our current framework. A case in favor of their inclusion could be made if we assumed the labor supply to an occupation to be less than infinitely elastic. However, this is unnecessary given that, empirically, including wages does not affect our results.

[7] We do not observe a change in the use of abstract skills for 94 cells, of routine skills for 112 cells, of manual skills for 118 cells, and of finger dexterity skills for 119 cells. It should be noted that these are changes for Census occupations, while skill use is reported for *DOT* occupations. Therefore, we observe a change for a Census occupation even if only one of the *DOT* occupations it comprises has been updated.

The *DOT* identifies aptitudes, temperaments, and abilities used in a job and measures them numerically. Observations are at the occupation-title level. Therefore, at a point in time, differences in skill use by sex reflect only differences in employment shares across occupation titles.

The 1965 *DOT* includes all of the skill-use (task) measures used in Autor et al. (2003). With some small caveats discussed below, it recorded them on the same scales as the later edition, allowing us to have consistent skill measures. Of course, we cannot be sure that individuals interpreted the measures in the same way in the 1950s, 60s, and 70s, but we see no reason that this concern should be greater than for many measures used to compare time periods or geographies.

The one small change is that the earlier edition provides a single measure of "General Education Development" while the later releases measure reasoning, mathematical, and language development separately. We experimented with using the average or the maximum of these three to generate a single measure comparable to the 1965 measure and checked whether this affected the correlation between the third and fourth edition measures. The correlations were similar. Looking across groups did not create a strong case for either. We present results using the average of the reasoning, mathematical, and language development measures for General Education Development in the 1977 and 1991 *DOT*s. In addition, the 1965 *DOT* sometimes provides more than one value of an aptitude, temperament, or ability for a single job title. In such cases, we use a simple average of the values reported.

Like Autor et al. (2003), we measure routine-cognitive skill using the variable "adaptability to situations requiring the precise attainment of set limit, tolerances, or standards," routine-manual skill (hereafter, finger-dexterity skill) by "finger dexterity," manual skill by "eye-hand-foot coordination." For our measure of abstract skill, we use "General Education Development." rather than only its mathematical component, to allow for consistency across all *DOTs*. We drop interactive skills from the analysis, partly for simplicity and partly because the explosion in the demand for social skills (Deming (2017)) appears to date from a later period. In addition, given data limitations, adding a fifth skill would prevent us from measuring patterns of substitution among skills without additional very strong restrictions on parameters. For each census occupation, we use a weighted average (by employment share) of the skill use in the *DOT* occupations comprising that census occupation.

For consistency with our theoretical model, we depart from Autor et al. (2003) and Autor and Dorn (2013) in how we use these measures. Autor et al. (2003) use the absolute value of each skill, while Autor and Dorn (2013) focus on routine intensity defined as $(RTI = ln(R) - ln(M) - ln(A))$.[8] Instead, we first scale the absolute level of skill use by where it lies between the maximum and minimum of that skill's use in any occupation over our sample period. Thus, use of skill i in occupation J at time t is:

$$\widetilde{skill}_{i,J,t} = \frac{skill_{i,J,t} - skill_i^{min}}{skill_i^{max} - skill_i^{min}} \tag{2.23}$$

where $skill_{i,J,t}$ is the value obtained directly from the *DOT* measures aggregated at the occupation level, $skill_i^{min}$ and $skill_i^{max}$ are the minimum and maximum absolute values (at the occupation level) for skill i in any version of the *DOT*. Finally, we compute the share of each skill in the overall sum

$$S_{i,J,t} = \frac{\widetilde{skill}_{i,J,t}}{\Sigma_k \widetilde{skill}_{k,J,t}} \tag{2.24}$$

so that our four skill measures sum to 1.

Census occupations are more highly aggregated than the *DOT's* job titles. Fol-

[8] We, like everyone else in this literature, have to treat the ordinal measures in the DOT as measured on an interval scale. We do so with an unusual level of chagrin given that one of us has pointed out (Bond and Lang (2013), Bond and Lang (2019)) that findings can be sensitive to how an ordinal scale is converted to an interval scale. Unfortunately, the approaches in Bond and Lang (2013) are not available to us in this setting.

lowing Autor et al. (2003)'s treatment of the 1977 *DOT* and the 1991 revision, we construct gender-specific skill measures for the 1965 *DOT* by aggregating the *DOT* titles to the census occupations separately for men and women. This accounts for the different distribution of workers by gender across job titles within each census occupation. Following Autor et al. (2003), we use the *DOT*-augmented version of the April 1971 Current Population Survey for this aggregation since this is the only dataset with both *DOT* and census codes.

We use the consistent occupation system created by Dorn (2009) and the crosswalk files provided by Autor and Dorn (2013), linking these occupations to previous census classifications. This gives us 212 occupations in the initial period, 265 in the intermediate period, and 329 in the later period. We create the occupation skill measures using occupation weights from all full-time workers not living in group quarters between age 18 and 64 in the IPUMS 1960 5% sample, in the IPUMS 1970 1% State sample, and the IPUMS 1980 5% sample.

Despite the tremendous insights measures of these skills have provided, about six and seven percent of workers work in jobs that purportedly make no use of manual and routine skill. We leave it to the reader to assess whether this is plausible.

Our data on the occupation distribution by sex come from the Census (IPUMS) and from March (Annual Social and Economic Supplement) Current Population Surveys (CPS) and are limited to workers age 25-64, but otherwise, our sample restrictions are the same as for the calculation of the skill weights. Since economists know these data well, we do not describe them here. Our choice of which sources to use for different purposes reflects an admittedly arbitrary trade-off between sample size and proximity of the employment data to the timing of the *DOTs*. Before 1968, the CPS coded occupations in fewer than forty categories and did not use the Census classification. Therefore, we use the 1960 1% Census sample for our initial period.

We rely on the 1970 and 1980 Census samples for the two later periods when we believe greater accuracy in estimating the employment cells is critical. Thus, we use the censuses to aggregate from *DOT* to census occupations and when using occupation/industry cells as observations in our regressions. Our decomposition of skill use into within and between-occupation changes relies on occupation, not industry, and therefore, uses larger cells. Consequently, we use the current occupation in the 1970-72 and 1982-84 March CPS for this purpose.

2.4 Results

Table 2.1 shows the evolution of average skill use over our period. There are four panels, one for each skill. Within each panel, we show the mean and standard deviation of skill use for all workers, for men, and for women.

In contrast with Autor et al. (2003) and Autor and Price (2013), we find that the decline in routine skill use started in the earlier period. The difference is that we use the *DOT 3rd edition* to measure skill use in the earlier period, and therefore account for within-occupation shifts. This decrease is much less pronounced among women than among men, consistent with the relative direction of changes in Autor and Price (2013). Consistent with earlier work, the use of abstract skills increased in the earlier period. Our results suggest that this change was solely among men. In contrast with earlier work, we find a decrease in finger dexterity (routine manual) and an increase in (nonroutine) manual, but with noticeable differences in the patterns between men and women.

The later period corresponds most closely to the 1970-80 change in Autor et al. (2003) and Autor and Price (2013). Like these papers, we find a decline in routine (cognitive) skill use and increased abstract-skill use, but these changes are much more pronounced among women. Finally, overall the changes in manual and finger dexterity reverse the signs of the changes in the earlier period although again, the pattern is

Table 2.1: Skills use levels by year

	Routine skills			Abstract skills		
	1960	**1971**	**1983**	**1960**	**1971**	**1983**
All						
Mean	0.314	0.276	0.242	0.298	0.342	0.378
Std. Dev.	(0.164)	(0.198)	(0.185)	(0.142)	(0.179)	(0.182)
Women						
Mean	0.299	0.288	0.233	0.312	0.317	0.375
Std. Dev.	(0.179)	(0.209)	(0.190)	(0.123)	(0.171)	(0.175)
Men						
Mean	0.319	0.271	0.248	0.293	0.353	0.379
Std. Dev.	(0.157)	(0.192)	(0.181)	(0.148)	(0.181)	(0.185)
	Manual skills			**Finger dexterity skills**		
	1960	**1971**	**1983**	**1960**	**1971**	**1983**
All						
Mean	0.084	0.097	0.083	0.305	0.285	0.298
Std. Dev.	(0.062)	(0.110)	(0.105)	(0.067)	(0.080)	(0.088)
Women						
Mean	0.058	0.070	0.049	0.331	0.325	0.343
Std. Dev.	(0.059)	(0.092)	(0.079)	(0.072)	(0.096)	(0.101)
Men						
Mean	0.093	0.109	0.103	0.296	0.267	0.270
Std. Dev.	(0.060)	(0.115)	(0.113)	(0.062)	(0.064)	(0.064)

Notes: Estimates use the occupation distributions from the 1960 Census, the March 1970-72, and 1982-84 Current Population Surveys. The skills used in each occupation are taken from the third, fourth, and revised fourth editions of the Dictionary of Occupational Titles. DOT occupations are aggregated to census occupations using the April 1971 Current Population Survey.

somewhat different between men and women.

We treat the results for manual and finger dexterity with some caution. The correlation between the measures in the 3rd and 4th editions of the *DOT* are somewhat low, only .46 for finger dexterity and .49 for nonroutine manual compared with .68 for abstract and .63 for routine. While it is certainly possible that the 1960s saw dramatic change in the importance of the two manual skills in a way that changed their ranking of importance across occupations, it is also possible that, despite defining the skills similarly, the two editions measured them differently.

2.4.1 Within-occupation changes are important (sometimes)

Table 2.2 decomposes skill-use changes into within and across-occupation changes using the following decomposition:

$$Skill_{e+1,t+1} - Skill_{e,t} = \underbrace{(Skill_{e+1,t+1} - Skill_{e+1,t})}_{\Delta \text{ across}} + \underbrace{(Skill_{e+1,t} - Skill_{e,t})}_{\Delta \text{ within}} \quad (2.25)$$

where e indicates the *DOT* edition, and t indicates the period considered. Thus, Δ within shows how skill use would have changed had the occupations in which, for example, males worked been the same in 1960 and 1971. In parallel, Δ across shows how much skill use would have changed had skill use in each occupation remained constant between 1960 and 1971 and only the occupations where workers were employed shifted. This latter measure corresponds to that typically presented in the literature, primarily because of the limitations of the *DOT*. Black and Spitz-Oener (2007) , which uses German data on a later period is an exception.

We begin by looking at across-occupation changes since these are akin to what the literature most frequently measures. We remind the reader that any differences from the prior literature may reflect our use of different editions of the *DOT* and/or our somewhat different use of the skill measures. All across-occupation changes seem quite modest in the early period, with the largest change for abstract-skill use. Still, this change amounts to only 0.06 standard deviations. In contrast, across-occupation changes are much more important in the later period. The .022 increase in abstract-skill use corresponds to roughly one-eighth of a standard deviation and the corresponding declines in manual and routine-skill use to declines of .10 and .05 standard deviations.[9]

Perhaps the most important message of Table 2.2 is that between-occupation shifts miss a great deal of the action. In the earlier period, we observe, at most, very modest shifts in skill use across occupations, but there are large within-occupation changes; within occupation, routine-skill and finger-dexterity use decline by more

[9] We use the standard deviation in the base year, 1960 or 1971, in all cases.

Table 2.2: Within- and across-occupation components

	Routine skills		Abstract skills		Manual skills		Fingdex skills	
	60-71	**71-83**	**60-71**	**71-83**	**60-71**	**71-83**	**60-71**	**71-83**
All	-0.037	-0.034	0.044	0.035	0.013	-0.014	-0.020	0.013
Δ within	-0.035	-0.024	0.035	0.013	0.018	-0.003	-0.018	0.014
Δ across	-0.003	-0.010	0.009	0.022	-0.004	-0.011	-0.002	-0.001
Std. Dev.	(0.164)	(0.198)	(0.142)	(0.179)	(0.062)	(0.110)	(0.067)	(0.080)
Women	-0.011	-0.056	0.005	0.058	0.013	-0.021	-0.007	0.019
Δ within	-0.009	-0.047	-0.003	0.026	0.014	-0.010	-0.002	0.031
Δ across	-0.001	-0.008	0.007	0.032	-0.001	-0.011	-0.005	-0.012
Std. Dev.	(0.179)	(0.209)	(0.123)	(0.171)	(0.059)	(0.092)	(0.072)	(0.096)
Men	-0.048	-0.023	0.061	0.026	0.016	-0.006	-0.029	0.003
Δ within	-0.044	-0.014	0.049	0.007	0.019	0.000	-0.024	0.006
Δ across	-0.004	-0.009	0.012	0.019	-0.003	-0.006	-0.004	-0.004
Std. Dev.	(0.157)	(0.192)	(0.148)	(0.181)	(0.060)	(0.115)	(0.062)	(0.064)

Notes: This table decomposes the change in the use of each of four skills into the change that would have been observed if the occupation distribution had been the same at the end of the period as at the beginning of the period (Δ within) and what would have been observed if the skill use were always the skill use at the end of the period but the occupation distribution had changed. Fingdex refers to finger dexterity. Estimates use the occupation distributions from the 1960 Census, the March 1970-72, and 1982-84 Current Population Surveys. The skills used in each occupation come from the decennial censuses. DOT occupations are aggregated to census occupations using the April 1971 Current Population Survey. Standard deviation in the base years in parenthesis.

than one-fifth of a standard deviation, offset by similar increases in abstract-skill and manual-skill use.

Thus, between 1960 and 1971, men experience a very substantial reduction in routine-skill use, with the overall decline (-.048) due almost entirely (-.044) to within-occupation changes. Notably, the modest decline in routine-skill use among women in this period was *not* the result of within and between changes offsetting each other. Instead, we observe that each was largely unchanged.

There are also notable differences between men and women in the skill shifts, which, in the early period, are much larger for men, particularly when we focus on within-occupation shifts. Except for a .2 standard deviation increase in manual-skill use within occupations, all of the shifts experienced by women are small. In contrast, during this period, men increased their abstract-skill use by almost .4 standard deviations, of which over 80% was within occupation. Similarly, their routine-skill-use decline of about .3 standard deviations, occurred almost entirely within occupation. Their manual-skill use increased within occupation by about .3 standard deviations, more than offsetting a small between-occupation decrease. Finally, within-occupation changes account for more than 80% of their almost .5 standard deviation decrease in the use of finger dexterity.

The table tells a notably different story about the later period. When we do not separate the results by gender, changes in skill use remain large (between .1 and .2 standard deviations) but are smaller than in the early period by this metric. Most of the change in abstract and manual-skill use is between occupations. However, within-occupation shifts are the more important source of changes in routine and finger-dexterity use.

Nevertheless, as in the earlier period, there are notable differences in the changes we observe among men and women. The overall changes are consistently much larger for women than for men. Most importantly, women see an increase in abstract-skill use both within and between occupations (.15 and .19 standard deviations), roughly on par with the increase for men in the early period. This is largely offset by a reduction in routine-skill use of .27 standard deviations, almost entirely within-occupations. At the same time, women use more finger-dexterity within occupations, but move to occupations that make less use of it; while not ruled out by our model, this is somewhat surprising.

Our analysis would be misleading if within-occupation changes reflected shifts in the distribution of more disaggregated occupations within an occupation. The problem does not arise for aggregating *DOT* occupations to census occupations. We have only a single crosswalk for this aggregation so that the relative weight of legal and medical secretaries in the census occupation does not change over time. The problem arises if, for example, secretaries who work for litigators and those who work for bond lawyers use different skills, if one grows faster than the other, and if the shift in the relative importance of the more disaggregated occupations affects the skills the various *DOT* editions report for legal secretaries.

2.4.2 Relative skill-productivity growth matters (sometimes)

Recall that estimating (2.22) and imposing that the coefficients sum to 0 allows us to identify the relative growth of skill productivity.[10] Table 2.3 shows the results of this exercise.[11]

Perhaps the most striking result is the rapid relative growth of the productivity of finger dexterity among women, as reflected in its negative coefficient. This is consistent with the importance of the IBM Selectric typewriter discussed in the introduction, the early versions of word processors that appeared towards the end of this period, and electronic calculators, which became widely available in the 1970s.

Recall that the coefficients in the table measure the relative growth rate of the productivity of the skills multiplied by $\varepsilon/(1-\varepsilon)$. Assuming that the elasticity of substitution is less than one, then $0 > \varepsilon/(1-\varepsilon) > -1$, and we can bound the difference relative to the average in the annualized rate of growth over the twelve years by the coefficient divided by twelve. The implied growth rate of the relative productivity of finger dexterity is large, at least about 8% per year among women in both periods, although the 95% confidence intervals include differences of less than

[10] To reduce measurement error, we restrict the sample to occupation/industry combinations comprising at least .0001% of employment in each year included in the pair and at least an average of .0002% over the two years. We impose this requirement separately for men and women so that an occupation might, for example, be included in the regression for men but not for women. The second requirement ensures that we do not create this bias by dropping observations near the threshold that saw a modest change in employment that caused it to cross the .01% threshold but keep similarly small occupation/industry observations that happen not to cross the threshold. Nevertheless, many of the employment changes we observe remain implausible. Since occupations are coded consistently across periods, we are not concerned that changes in occupation drive these changes. We winsorize the data fairly severely at the 20th and 80th percentiles. Winsorizing at the 10th and 90th percentiles gives results with a similar interpretation but that are generally larger in absolute value and much more imprecise. Finally, we average our skill-use measures from the two editions (or the revision) corresponding to the pair of years in our analysis.

[11] See Table in the Appendix for non-transformed coefficients

Table 2.3: Skill Productivity Growth Relative to Average

	(1) women 60-70	(2) women 70-80	(3) men 60-70	(4) men 70-80
Routine	-0.169	0.027	-0.078	-0.045
	(0.149)	(0.162)	(0.157)	(0.103)
Abstract	0.246	0.923	0.281	0.417
	(0.157)	(0.185)	(0.189)	(0.123)
Manual	0.916	0.022	0.065	0.150
	(0.349)	(0.361)	(0.301)	(0.150)
Fingdex	-0.993	-0.971	-0.268	-0.523
	(0.284)	(0.259)	(0.314)	(0.207)
r2	0.16	0.16	0.15	0.12
proportion due to skills	0.16	0.47	0.07	0.18
N	3089	4628	4853	7013
p(all skill coefs=0)	0.006	0.000	0.428	0.005
p(rout=man=findex)	0.004	0.005	0.824	0.101

Notes: Standard errors in parentheses, clustered at the occupation level. Estimates are transformed from regression of change in log employment in an occupation/industry cell on average skill (routine, manual, finger dexterity) use in that cell over the period (equation (2.22) in the text) and imposing that the mean deviation from mean skill growth for all four skills is 0. Proportion due to skills is the proportion of the R-squared attributable to the three skills in the regression using the Shapley-Owen decomposition.

4% per year. As noted earlier, Goos et al. (2014) estimate an elasticity of substitution across industry outputs of .42. This would entail multiplying the differences by 1.7. Since we believe there should be less substitutability across occupations, we find a somewhat lower multiplier more plausible, but the reader is not bound by our intuition.

The second striking result is the difference between our early and later periods. In the early period, differences in the growth of skill productivity play little role in explaining employment changes. We cannot reject that all skills grew at the same rate for men. While we can reject this hypothesis for women, the differences explain little of the between-occupation differences in employment growth. Using the Shapley-Owen decomposition, we find that the skill composition of occupations accounts for only about 16% of the explained sum of squares or about 2% of the total variance.

The later period is very different. The coefficients on skills are highly significant.

Moreover, they account for a notable proportion of the explained sum of squares, 46% among women, although less so (18%) among men. When we recognize that we have many more industry dummies than skills, it is apparent that we probably noticeably underestimate the relative importance of the skills measure.[12] For both men and women we cannot reject that routine and manual skill productivity grew at the same rate as the average of the skills. However, in both cases, we see evidence of faster growth of the productivity of finger dexterity and slower growth of abstract skills.

As previously mentioned, we perform this exercise also including the % change in average wages. Table 2A in the appendix shows that this has no meaningful impact on our estimates.[13]

2.4.3 Slow growth of abstract productivity and faster growth of other skills (mostly) explains the within shifts

To understand what our model says about within-occupation skill shifts, we take a linear expansion of $S_i\left(J\right)$ with respect to relative changes in skill productivities:

$$dS_i(J) = \Sigma_k \frac{\partial S_i(J)}{\partial \ln A_k} d\ln A_k. \tag{2.26}$$

Now, we multiply by $f(J)$ and integrate over all jobs

$$\int\limits_{\mathcal{J}} dS_i\left(J\right) f(J)dJ = \Sigma_k \left(d\ln A_k \int\limits_{\mathcal{J}} \frac{\partial S_i\left(J\right)}{\partial \ln A_k} f(J)dJ \right). \tag{2.27}$$

[12] Intuitively, while asymptotically a coefficient on a variable with no effect on the dependent variable has an expected value of 0, in finite samples it has a non-zero value with probability 1 and therefore contributes to explaining the variance of the dependent variable.

[13] Although not reported, we also performed this analysis using the change in the average log wage, and reached the same conclusion, with point estimates for skills productivity that are even closer to those in Table 2.3. Table available upon request.

Now, we use the fact that $\Sigma_k S_k = 1$ to get

$$\Sigma_k \frac{\partial S_k(J)}{\partial \ln A_i} = 0. \tag{2.28}$$

In addition, we know from (2.5) that

$$\frac{\partial S_i(J)}{\partial \ln A_k} = \frac{\partial^2 \ln y^*(J)}{\partial \ln A_i \partial \ln A_k} = \frac{\partial S_k(J)}{\partial \ln A_i} \tag{2.29}$$

so that we can rewrite (2.28) as

$$\Sigma_k \frac{\partial S_i(J)}{\partial \ln A_k} = 0. \tag{2.30}$$

Thus, we can normalize (2.27) with respect to an arbitrary $d \ln A_n$:

$$\int_{\mathcal{J}} dS_i(J) f(J) dJ = \Sigma_{k \neq n} (d \ln A_k - d \ln A_n) \int_{\mathcal{J}} \frac{\partial S_i(J)}{\partial \ln A_k} f(J) dJ. \tag{2.31}$$

Denoting the integral on the right by $\partial \overline{S_i} / \partial \ln A_k$, and replacing the left-hand-side with the within estimates in Table 2.2 and the $d \ln A_k$ terms with the estimates in Table 2.3, we arrive at

$$\widehat{\text{within}}_i = \Sigma_{k \neq n} \left(\widehat{d \ln A_k} - \widehat{d \ln A_n} \right) \frac{\partial \overline{S_i}}{\partial \ln A_k}. \tag{2.32}$$

These $\partial \overline{S_i} / \partial \ln A_k$ terms represent the average changes in workers' skills brought on by isolated productivity changes, and we are most interested in extracting them. As Section 2.4.2 suggests, however, more than one A_k changed in each of our periods, making this exercise nontrivial.

Assuming that these derivative terms do not change over time, after imposing symmetry per (2.29), we have six equations and six unknowns for men and similarly for women. Unfortunately, one of the six equations is redundant. This is not a generic problem. If we had three skills rather than four, we would have three derivatives and four equations, of which one would be redundant, giving us a unique solution. If we had three sets of changes and four skills, the problem would be overidentified.

For illustrative purposes, we impose that there is no substitutability between manual and abstract skill.[14] With this restriction, in theory, the derivatives are just identified. However, the system has no solution since the within-occupation changes are estimated with error and the equations are only a first-order approximation. We choose the parameter estimates that minimize the sum of the squared differences between the calculated within change and the predicted within change.

The derivatives, $\partial\overline{S_i}/\partial \ln A_k$, capture a concept analogous to p and q complementarity and substitutability. If the derivative is positive, an increase in the productivity of skill k increases the amount of skill i acquired by workers. We refer to this case as A-complementarity. Note that, unlike p-complementarity, a skill may be A-complementary or A-substitutable with itself.[15]

Recall that in Table 2.3, we estimate $\varepsilon/\left(1-\varepsilon\right) * d\ln A_i$. So, as ε is unknown, with a change of sign, the coefficients represent lower bounds on the absolute values of the skill-productivity changes. Therefore, using these coefficients yields upper bounds on the derivatives. Consequently, we focus on the signs of the estimated derivatives rather than their precise magnitude and ignore the $\varepsilon/\left(1-\varepsilon\right)$ term other than to assume that it is negative. Thus in reading Table 2.4, which displays the results of this exercise, readers can rely on their intuition to divide the estimated derivative by something in the range 1.3 to 1.7.

Although the precise values of the estimated derivatives in Table 2.4 differ between men and women, their interpretation is broadly similar. All skills are, on average, A-substitutes for themselves. However, the derivative is about an order of magnitude

[14]In the appendix, we report the outcomes of this exercise with the alternate assumption that there is no substitutability between manual and routine skill. The results change only in the details. See Tables and .

[15]In contrast A_iS_i, the 'effective' amount of skill i supplied by the worker, must increase with A_i.

Table 2.4: Derivatives of Skill Use with Respect to Skill Productivity

	Skill Used			
$\Delta\ lnA_i$	Routine	Abstract	Manual	Finger Dexterity
		Women		
Routine	-0.263			
Abstract	0.151	-0.095		
Manual	0.021	0	-0.015	
Finger Dexterity	0.910	-0.056	-0.006	-0.039
		Men		
Routine	-0.576			
Abstract	0.246	-0.143		
Manual	0.081	0	-0.066	
Finger Dexterity	0.249	-0.103	-0.015	-0.131

Notes: Each cell shows the derivative of the average use of the column skill with respect to a change in the relative productivity of the row skill. Estimates are up to a factor of proportionality of $\frac{-\varepsilon}{1-\varepsilon}$ (which is strictly between 0 and 1). The estimates are derived from combining changes in skill use across time with estimates of relative productivity growth from Table 3. See equation (2.32) in the text for the precise formulation. The cross-derivative between abstract and manual is set to 0. See the text for more detail.

greater for routine skill than for finger dexterity or manual skill and noticeably larger for routine than for abstract skill. Routine and all other skills are A-complements, again averaged across occupations.

Table 2.5 leverages these results to show how the change in the productivity of each skill accounts for the overall within-occupation shift in skill use. It also compares the predictions of the model with the data. Not surprisingly, given the imprecision of the skill-growth estimates for men in the earlier period, the model does much better for women than for men. For women, the largest gaps are for the shifts in the use of finger dexterity, which we over-predict in the earlier period and under-predict in the later period. For men, we under-predict the growth of abstract-skill use in the early period and over-predict it in the later period.

The large shift from routine to abstract-skill use among men in the early period is accounted for by the slow growth of abstract-skill productivity and the somewhat above-average growth of routine-skill productivity, which the effect of the very rapid growth in the productivity of finger dexterity partially offsets.

Table 2.5: Decomposition of Within-Occupation Changes in Skill Use

	Women 1960-1971 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	-0.044	0.026	0.004	0.015
Abstract	-0.037	0.023	0	0.014
Manual	-0.019	0	0.014	0.006
Finger Dexterity	0.090	-0.056	-0.006	-0.029
Total Predicted	-0.010	-0.007	0.011	0.006
Data	-0.009	-0.003	0.014	-0.002
	Women 1971-1983 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	0.007	-0.004	-0.001	-0.002
Abstract	-0.139	0.088	0	0.052
Manual	0.000	0	0.000	0.000
Finger Dexterity	0.088	-0.054	-0.006	-0.028
Total Predicted	-0.044	0.029	-0.006	0.021
Data	-0.047	0.026	-0.010	0.031
	Men 1960-1971 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	-0.045	0.019	0.006	0.019
Abstract	-0.069	0.040	0	0.029
Manual	-0.005	0	0.004	0.001
Finger Dexterity	0.067	-0.028	-0.004	-0.035
Total Predicted	-0.053	0.032	0.007	0.014
Data	-0.044	0.049	0.019	-0.024
	Men 1971-1983 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	-0.026	0.011	0.004	0.011
Abstract	-0.103	0.060	0	0.043
Manual	-0.012	0	0.010	0.002
Finger Dexterity	0.130	-0.054	-0.008	-0.069
Total Predicted	-0.010	0.017	0.006	-0.012
Data	-0.014	0.007	0.000	0.006

Notes: Each entry is the predicted change in the within-occupation use of the column skill due to changes in the productivity of the row skill according to equation (2.32) in the text and using the values from Tables 2.3 and 2.4. Total predicted is the sum of the four values above. The predictions can be compared with the within changes reported in Table 2.2 and repeated in the line labelled Data.

Similarly, among women in the later period, the large decline in routine-skill use and the offsetting increases in abstract-skill and finger-dexterity use are driven by the slow growth of abstract-skill productivity that is not fully offset by the rapid growth of the productivity of finger dexterity.

2.5 Summary and conclusion

We make two contributions. First, at a purely empirical level, we provide new evidence on changes in skill use in the 1960s and 1970s. We show that in the 1960s, such changes were important but were particularly important for men and much more pronounced within than between occupations. In contrast, in the 1970s, skill use shifted both within and between occupations, and changes were particularly pronounced among women.

Second, we develop a simple model that reconciles or combines two approaches to technological change, the SBTC and task-based literatures, by modeling technological change as increasing the productivity of individual skills such as finger dexterity rather than, for example, college-educated workers. While our model also allows us to account for technological change that replaces occupations, we focus on detecting changes in skill productivity; we capture changing demand for occupations only through changes in industry demand.

We use the insights from the model to measure the pattern of skill-productivity growth needed to explain the employment shifts that we observe. For women in the 1960s, we find that differences in the productivity growth of skills account for very little of the employment changes that we observe. In contrast, in the 1970s, they account for almost half the explained difference among women and a fifth among men.

Our empirical results suggest that if a skill's productivity increases, use of that skill within an occupation generally decreases. Thus, skills generally are A-substitutes for themselves. Abstract and routine skills are A-complements, as are finger dexterity and routine skills. Among women in the later period, the very slow growth of abstract-skill productivity shifted skill use within occupations away from routine-skill use and towards abstract-skill use. The rapid growth of the productivity of finger dexterity, which shifted skill use towards routine and away from finger dexterity, only partially offset the decline in routine-skill use.

We hope and believe that we have demonstrated that our simple model provides a useful framework for understanding changes in skill use both between and within occupations. Obviously, readers must make that judgment for themselves.

Chapter 3

Immigrants' Legalization and Firms: Evidence from the 2007 EU Enlargement

[1]

3.1 Introduction

Most countries restrict migrants' labor market access, usually by requiring employer-sponsored visas. These policies have distort the labor market, by pushing migrants in the informal sector, and put downward pressure on formal labor-market compensation. Hence, the combination of strong legal requirements and a sizeable informal sector allows firms to exert labor market power over migrant labor, and negatively affects native workers. It is thus unsurprising that there is a large policy debate on whether countries should relax work-visa requirements in order to downsize the informal sector and promote fairer competition in the labor market. Yet, the evidence on the effects of such policy changes remain scant.

We study the 2007 European Union enlargement that admitted Romania and Bulgaria (which we refer to as EU07), one of the most significant episodes relaxing migrants' labor market restrictions. This event was particularly relevant for Italy because Romanians are its largest immigrant group due to the strong language similarity between the two countries. Currently, there are over 1 million Romanians residing in Italy, representing a fifth of the total foreign-born population.

[1]This paper is part of the VisitInps initiative. We would like to thank all the members of INPS for their support. The findings do not represent the views of INPS

Our empirical approach combines the natural experiment induced by the EU enlargement in 2007, with a shift-share spatial approach. This method builds upon a large body of literature (see Dustmann et al. (2017) for a comprehensive review) using the historical settlement pattern of migrants from a given country to predict current settlements. This approach requires that ethnic enclaves predict current migrants' location decisions without being correlated with current economic shocks.

We focus on how firms adjust their employment and compensation decisions in response to the change of regulatory status for EU07 migrants. Consistent with the expansion of formal employment induced by the change in migrants' legal status, we observe a small increase in overall employment at the firm level. We also find evidence that the relative employment of EU07 migrants increased at the expense of natives. This pattern is mainly driven by a change in firms' hiring decisions: compared to the pre-reform period, they hire more EU07 workers and fewer native workers. We also observe a short-term increase in the rate of job separations for EU07 migrants. This is consistent with Romanian and Bulgarian citizens no longer needing employer-sponsored visas that limited job-to-job mobility. These results are consistent with the findings of Naidu et al. (2016), who study a reform in the United Arab Emirates that relaxed restrictions on job-to-job transitions for migrant workers.[2]

The short-term nature of our findings on employment suggests that the effects are mainly driven by movement to the formal sector of migrants previously employed in the informal sector.

We observe a short-run decrease in average (log) firm-level wages. EU07 workers experience a relatively small but persistent decrease in average earnings, while we do not observe an effect discernible from 0 for native workers.

[2]Naidu et al. (2016) document changes in a quite different setting from ours, both in terms of the context and of the type of reform. Italy, unlike the UAE, has a sizeable informal labor market. Secondly, the reform we analyze has a much wider scope since it eliminates visa requirements for migrants from Romania and Bulgaria, rather than just affecting job-to-job transitions.

We also perform a series of heterogeneity analyses aimed at gaining a deeper understanding of the mechanisms behind the observed patterns, particularly by determining which types of firms respond most to the EU enlargement.

Our results are largely driven by firms that participated in a previous amnesty legalizing migrants, or that employed EU07 workers before 2007. This reflects the tendency of firms employing migrants in the informal sector to also employ migrants in the formal one (Porto et al. (2018)). Moreover, this validates the hypothesis that the change in legal status for migrants increases their bargaining power in the labor market: while it is more convenient for employers to employ workers in the informal sector because it decreases their costs and the risk of being caught is quite low, workers prefer formal employment, which provides a range of social benefits.

Lastly, we analyze the impact of the enlargement on firms' business performance. We find a reduction in per-capita revenues and operative added value, providing evidence that the employment expansion is not associated with overall firm growth. This provides further suggestive evidence that firms formalized existing informal employment relationships.

This paper contributes to three main strands of the literature. We speak to the study of how labor restrictions affect labor market outcomes, particularly for migrant workers (e.g., Naidu et al. (2016)).

We also contribute to a recent strand of papers studying the impacts of migration in the presence of imperfect competition (Amior and Manning (2020)). In our setting, the interaction of labor restrictions on migrants and the presence of a sizeable informal market give firms labor market power. In this sense, our work is closely related to a recent working paper of Elias et al. (2019), studying the labor market implications of an amnesty program for undocumented migrant workers in Spain in 2004. Our work is distinct in that we focus on firm-level responses, and a broader legalization that applies to all immigrant workers of a given nationality, regardless of whether they were working as undocumented workers.

Lastly, we contribute to the vast body of research on migration and the labor market. By adding to an exiguous literature on its effects on firms, papers on this topic, discussed below, focus on reforms or contexts involving changes in the supply of highly skilled migrants. Dustmann and Glitz (2015), Ottaviano and Peri (2012), and Ottaviano et al. (2018) each study the impact of highly educated migrants in a single European country (respectively West Germany, France, and the UK). These papers all use the standard shift-share approach and do not exploit any policy changes. Ottaviano et al. (2018) focus on the effects on the imports, exports, and productivity of firms operating in the service sector, finding that immigration has a positive impact on all these dimensions. Dustmann and Glitz (2015) focus on wage responses and on the underlying mechanisms explaining them, such as the variation in the skill mix in the labor force, the change in returns to different skill types, and the creation and destruction of firms. Mitaritonna et al. (2017) is more closely related to our study, since it investigates migration's effect on business outcomes, as well as impacts on wages and employment.

Another set of papers, focuses on labor supply shocks involving highly educated migrants, but exploits one-time policy shocks changing the pool of migrants legally allowed to work. In particular, Signorelli (2020) studies the effect on firm-level wages and employment in France of a reform that ease the hiring of foreign workers in certain technical occupations . Beerli et al. (2021), by contrast, analyze the removal of restrictions on hiring EU workers in Switzerland, and find that it had positive effects on measures of firms' innovation initiatives and investments. Ours is the first paper focusing on a labor supply shock of low-education migrants and its effects on firms' outcomes. Thanks to the richness of our data, we can study the resulting changes in firm-level wages and employment, as well as in business performance measures, providing a wide ranging - although not comprehensive - first collection of evidence. We study Italy for the following three reasons: 1) we can leverage a particularly rich dataset that allows us to perform accurate analyses at the firm level, 2) we can exploit a convenient institutional setting involving an exogenous policy change that caused a large labor supply shock, and 3) despite an extensive literature on migration in Italy, it is not yet comprehensive. Partly, this is due to immigration being a relatively new phenomenon for the country, and partly to the only recent availability of appropriate data to examine its impacts. The existing research mostly focuses on the effects of smaller legalization reforms on individual economic outcomes (Porto et al. (2018), Elias et al. (2019)) or on outcomes that are not strictly economic (Mastrobuoni and Pinotti (2015), Pinotti (2017)). To the best of our knowledge, this is the first paper to perform a detailed firm-level analysis of the impact of immigration on Italy.

3.2 Institutional context

Long a country of emigrants, Italy recently experienced a role reversal, becoming a destination for many foreign workers. The share of foreign-born in Italy rose up from 1% to 9% from the early 1990s to 2017.

Immigration could represent an economic and demographic opportunity for Italy, since it is experiencing low and decreasing fertility rates and population aging. There are, however, concerns driven by the conjunction of increasing migration and persistent economic downturns. Natives worry about downward pressure on their wages and reduced economic opportunities. In addition, Alesina et al. (2018) find that Italians are among those with the strongest misperceptions about migration. They both overestimate the number of migrants in the country, and underestimate their education level and employment rates, and overestimate their dependence on welfare.

For these historic and economic reasons, Italian legislation has been quite conservative. Migration policy is regulated mainly by Laws *40/1998* and *189/2002*. Economic migrants may enter Italy lawfully with a work permit. Each year, the number of available work permits by province and type of contract is established nationally. Firms have to sponsor foreign workers and apply for the permit. Successful migrants may then be employed by their sponsors. They can change jobs or start their own businesses, but can only be unemployed for six months, after which they must leave Italy (for further details, see Pinotti (2017)).

These restrictions have not stopped economic migrants from coming to the country. The absence of a process for legalizing such workers has led many migrants to work in the black market, receiving very low salaries and lacking decent standards of job, health, and social security protection.

To address this problem, like other European countries, in the past 20 or so years Italy has implemented one-time general amnesties, allowing illegal workers already living and working in the country to work. Applications for amnesty must be submitted by employers: a successful outcome grants the same permit status as the quota system. The number of individuals affected by these legalizations has substantially increased over the years. The latest amnesty before 2007, which took place in 2002, was Italy's largest, causing more than 600,000 foreign workers to emerge from the black labor market, as shown in Figure 3·1. Many of these migrants were Romanians: their language similarities with Italian make the country a good haven for job opportunities. EU enlargements have also changed the pool of potential legal migrants. Whenever a new country earns EU membership, its citizens have the right to reside and work in any of the countries within the Union [3]. The two largest enlargements to date took place in 2004 and in 2007, in both cases extending EU membership to Eastern

[3]There are a few exceptions to this statement. For instance, some countries like the UK or Germany, implemented temporary restrictions when the EU extended the membership to Eastern European countries. Such restrictions usually implied that free movement of workers was granted only to those working in certain economic sectors in the first years of EU membership.

Figure 3·1: Evolution of foreign-born population over time

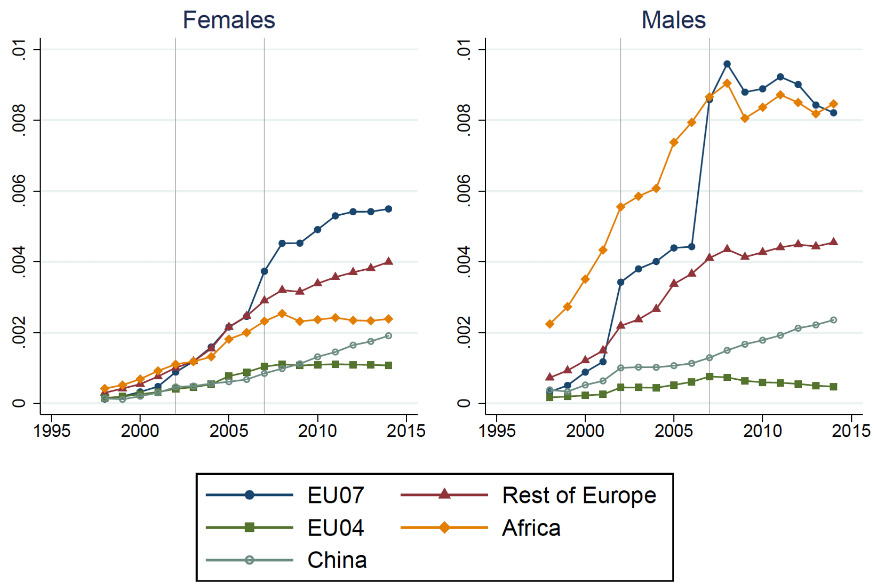

Notes: The graphs show the evolution of the share of foreign-born resident population for different countries of origin, divided by the total population in 1991. "EU07" refers to foreign-born from Romania and Bulgaria, "EU04" refers to foreign-born from the countries that joined the European Union in 2004 (Cyprus, Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Slovakia, and Slovenia), "Rest of Europe" refers to foreign-born residents from non-EU European countries (Albania, Belarus, Bosnia, Croatia, Macedonia, Moldova, Montenegro, Russia, Serbia, Ukraine).

European countries. The 2004 expansion, which admitted Cyprus, the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Slovakia, and Slovenia, had little impact on the Italian economy, (see Figure 3·1) since the great majority of workers from these newly added countries migrated to other countries, such as Germany and France. The 2007 expansion only affected Bulgaria and Romania, a nation with migration ties and interest in Italy.

While the new EU membership became effective on January 1^{st} 2007, the upcoming expansion was announced in December 2005. Figure 3·1 shows that, as expected, there was a sudden increase in the share of EU07 migrants in 2007. Perhaps more importantly, we do not observe an anomalous rise in EU07 workers in 2006, suggesting that there were no anticipation effects, at least in the legal labor market. This is not surprising: work permits are usually approved in the Fall, thus the effects of the announcement of the EU07 expansion would not be evident until the last couple of months of 2006.

While Romanian and Bulgarian migrants were immediately granted freedom of access and circulation in Italy, their access to the labor market was delayed in some sectors. However, the sectors in which Romanians workers were mostly employed, domestic work, seasonal jobs, agriculture, construction, tourism and services, were immediately accessible. In other sectors, the barriers to access were much lower than before; the employer had to notify the prefecture of its intention to hire a migrant but did not have to wait for authorization[4].

We expect legalization following EU enlargement to differ from one resulting from a general amnesty.

The right to work in the country is not employer-bound. General amnesties grant foreign workers a work permit linked to the employer, and in case of unemployment for more than 6 months the migrant is expected to leave the country. Migrants with EU citizenship, by contrast, can look for the most suitable job match and most favourable employment conditions. Thus, EU citizenship allows for greater job, sector, and geographic mobility. Moreover, it also grants stronger status security for those willing to embark on self-employment.

Secondly, the policy change is permanent. All Romanian and Bulgarian citizens were affected, regardless of whether they were already working in Italy on January 1st 2007. The effect on the economy will therefore not only stem from the emergence of irregular workers in the legal market, but also from the arrival of new migrants who are allowed to work in the formal market.

Lastly, the migrants' characteristics differ under the two policies. Amnestied workerstend to be more risk-loving than the average, since they accept living and working

[4] `https://www1.interno.gov.it/mininterno/export/sites/default/it/sezioni/sala_stampa/notizie/europa/app_notizia_23478.html`

in circumstances that place them outside the law. For the same reason, they are also more likely to come from a place of economic distress, which is correlated with a lower educational level and set of qualifications.

These considerations point to EU enlargements having potentially different economic impacts from the more commonly studied one-time amnesties.

3.3 Data and descriptive statistics

We use high-quality and restricted-access administrative data available at the Italian Social Security Institute (INPS). Matched employer-employee records are available over the period 1983-2015 for the universe of non-agricultural firms with at least one employee. The data covers 74% of private employment in Italy and 93% of private-sector employees.

For each worker-firm record, the following information is available: beginning and end date of the contract, type of contract (permanent or fixed-term, full-time or part-time), broad occupation group (blue-collar, white-collar, or manager), wages, days and hours of work, sector in which the firm operates, and location of the firm. Moreover, we also have access to demographic information, such as gender, age, nationality, municipality of residence.

In the second part of our paper, we also use CERVED data, which provides detailed firm-level balance sheet information for the universe of registered firms in Italy. About one third of the firms present in the INPS dataset have also records in CERVED, as not every company is legally bound to be registered. Even though this causes a significant reduction in the sample size for our analysis of business outcomes, it still gives us the unique opportunity to study key firm-level outcomes that are usually not available. For our analysis, we restrict our sample to firms that have a record in every year included in our period of interest (2004-2009). We include only workers between the ages of 18 and 64, who worked for at least 5 weeks in a given year. For each

Table 3.1: Summary statistics for the outcomes in 2006

	Mean (1)	SD (2)	P50 (3)	P10 (4)	P90 (5)
Total Employment	11.34	156.3	3	1	16
Native Employment	10.37	146.5	3	1	15
EU07 Employment	.130	2.201	0	0	0
Native Employment (Rel)	.903	.219	1	.667	1
EU07 Employment (Rel)	.015	.087	0	0	0
Native Hires (Rel)	.201	.296	0	0	.667
Native Separations (Rel)	.192	.291	0	0	.643
EU07 Hires (Rel)	.007	.05393	0	0	0
EU07 Separations (Rel)	.006	.052	0	0	0
Average Wage	418.8	136.1	395.7	299	559.2
Eu07 Wage	383	108.1	370.1	287.2	490.1
Native Wage	423.2	138.7	399.5	299.6	568.5
Number of firms	779552	779552	779552	779552	779552

worker, we only keep the record of that person's main job in any given year, defined as the most remunerative one; if there exist two jobs with the same salary, we select the one with the longest spell. Given that firms are our unit of interest, we collapse all information at the firm level.

In our final sample, we have a total of 779,552 firms, roughly corresponding to **47%** of the firms in the full dataset. Table 3.1 and 3.2 provide information on the firms in our sample. On average, a firm has 11 employees, which is to be expected in an economy characterized by a strong presence of small and medium enterprises (the median size is in fact much smaller, equal to 3). In 2006, the year prior to the EU expansion, 95%

Table 3.2: Summary statistics for the outcomes in 2006, by firm type

	Firms w. EU07 (1)	Firms w. only natives (2)
Total Employment	45.086	8.246
EU07 Employment	1.502	.
Native Employment	39.217	7.762
Average Wage	413.289	420.301
EU07 Wage	382.978	.
Native Wage	427.763	422.777

of our firms had at least a native worker among their employees, while 9% employed at least one EU07 migrant. As can be seen from Table 3.2, these latter firms tended to be larger (they have a mean of 45 employees) and to have a higher average wage, while wages for EU07 workers are on average 10% lower than natives' ones.

3.4 Methodology

Our empirical framework analyzes firms' response to the shock in local supply of (formal) EU07 workers induced by the enlargement of the European Union. In order to trace the dynamic evolution of the response, we adopt a flexible dynamic difference-in-differences regression. We regress the change in firms' outcomes in area lmm between year t and $t-1$ on the change in the EU07 migrants' share of the workforce in the same geographical unit.

Our baseline model is:

$$\Delta Y_{ft} = \alpha_{ft} + \sum_{t=2005}^{t=2009} \beta_t \times D_t \times \Delta Share E\hat{U}07_{07-06,llm} + \delta_{r,t} + \gamma_{s,t} + \theta_{d,t} + \phi_{f,t} * A_{f,t} + \mu_{llm} + \varepsilon_{ft} \tag{3.1}$$

where $\Delta Share E\hat{U}07_{07-06,llm}$ is the exogenous shock in migrants' labor supply and our main variable of interest.

$\delta_{r,t}$ controls for regional trends, $\gamma_{s,t}$ for broad sector trends (9 categories) , $\theta_{d,t}$ is a time-varying effect of the firm's size[5] fixed at base year 2006, $\phi_{f,t}$ is the effect of firm's age, μ_{llm} is a local labor market fixed effect, and ε_{ft} is the error term.

The period included in our analysis is 2004-2009. Although our data allows us to extend the time horizon of the study, we choose to restrict it for practical reasons: on the one hand, the years prior to 2004 are too proximate to the aforementioned 2002 amnesty for illegal migrants and are affected by it (Di Porto et al., 2018); on the other, Italy faced a deep economic downturn after 2009 during the Great Recession.

[5] The firm's size variable is a dummy of value 1 if there are 15 or more employees in 2006, and 0 otherwise. Italy has a dual labor market, in which firms in these two categories are subject to different regulations concerning hirings and separations.

These events could have an impact on our findings and consequentially affect the validity of our analysis. The baseline period for the event-study is 2006-2005, i.e. the one prior to the EU enlargement of 2007.

Our geographic unit of interest is the local labor market (LLM), a non-administrative area defined by the Italian National Institute of Statistics (ISTAT) on the basis of daily commuting flows of the Italian population. It is, therefore, an area that transcends municipality borders and is characterized by a common labor market, making it the perfect candidate to study the effect of a sudden rise in labor supply. In total, there are 611 LLMs in Italy, that constitute the level at which we observe the variation in our treatment.

Because we are dealing with a firm-level panel, we can make the most of it by implementing strategies that allow us to get rid of unobserved time-invariant components, which could potentially create omitted variable bias. One way to address this is to use firm fixed effects, another is to use a first-difference specification. We prefer the latter for the conceptual and econometric reasons outlined below.

Conceptually, we believe that the most effective way to study how the EU07 enlargement has affected local labor markets is by looking at how the presence of migrants has changed, rather than by observing the EU07 share of workers in a particular year. It is this variation over time, i.e. the comparison between the new share of migrants as a result of the EU enlargement and the old share, that best captures the shock that firms in a specific local labor market have to face. This claim is supported by the data, as we observe that places that had a relatively lower share of EU07 workers are those that experienced the largest relative influx of migrants. The explanation behind this phenomenon is not as mysterious as one might think: the sizeable increase in the number of EU07 migrants between 2006 and 2007 led to a diffusion of their presence throughout the country.

From an econometric standpoint, we perform the standard test to verify that a FD estimation is the most appropriate one for our data. In Figure C·1 we show that the $\Delta_{f,t}$ values, i.e. the first differences in the error term that result from our main specification, are uncorrelated, making FD our preferred approach at the expense of a FE estimation.

As already mentioned, we have chosen to use the difference in shares of EU07 migrants as our main explanatory variable. In addition to this option, we have considered two alternatives, but we deemed them to be more problematic. The first possible choice was to use the absolute number of EU07 migrants in a local labor market, which might be a misleading measure because it intrinsically contains geographic characteristics - such as population size and labor market sector composition - that are correlated with the outcome of interest and might ultimately confound our findings. The second option was to use the percentage growth in the share of EU07 workers. This measure falls short in terms of capturing the change in migrants' relevance with respect to the overall population. In fact, a growth of 100% or more is likely to be related to a negligible number of migrants in a local labor market. The absolute difference in shares, instead, allows us to fully account for the change in migrants' labor market penetration in the aftermath of the EU enlargement.

Clearly, the usual concerns regarding endogeneity apply: areas in which we observe a larger increase in the share of EU07 migrants might be ones that experience some macro-economic shock that could affect and confound our findings. We therefore instrument the change observed in 2006-2007 with the change observed in 2001-2002, when there was a horizontal amnesty legalization of migrant workers who were employed in the informal labor market. We argue that places that experienced a large change in migrants' share as a result of the past legalization shock would be more prone to experience a large change with the EU07 enlargement. Moreover, the variation taking place between 2001 and 2002 is the most relevant for our period of interest, because it too was characterized by an expansion of migrants throughout the country.

3.4.1 Instrument validity

We perform a series of tests to assess the validity of our chosen instrument. In Figure 3·2, we show that there is a strong correlation between the change in the share of EU07 migrants in the workforce in 2001-2002 with the change in 2006-2007. In appendix Table C.1, we also present the results of our first stage regression, showing that our chosen IV is a strong predictor for the exogenous shock of interest. We then provide

Figure 3·2: Evolution of EU07 migration over time: first stage

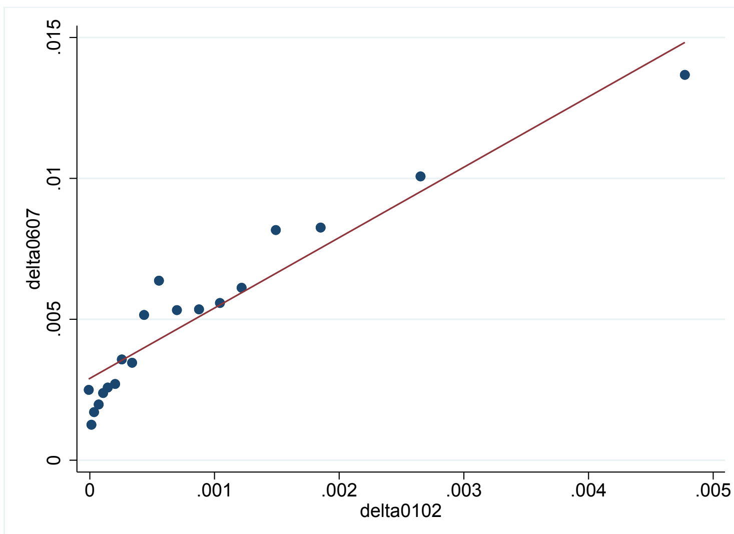

Notes: The graph is a binned scatterplot showing the evolution of the share of foreign-born resident population from EU07 countries over time. On the y-axis, we show the change in the share of EU07 over total population (measured in 1991) between 2006 and 2007. On the x-axis, we show the same change, but for the period between 2001 and 2002.

evidence suggesting that the exclusion restriction holds. We want to verify that the change in 2001-2002 is not correlated with current economic unobserved shocks. In order to do so, we regress our instrument on the change in total employment and wage (separately) at the local labor market level in our baseline period, 2005-2006, while controlling for region fixed effects. The lack of significance of our estimates in

Panel A of Table C.2 suggests that current economic shocks are not correlated with our instrument.

We might also be concerned that persistent economic shocks taking place prior to 2001-2002 can both predict the change observed in our instrument as well as firm outcomes, which would cause a spurious correlation between the two. In Panel B of Table C.2, we show that changes in economic outcomes at the local labor market level in the period 1992-1997 also do not predict our instrument. Finally, in Panel C, we also show that our instrument seems to be uncorrelated with changes in labor market outcomes between 2002 and 2006.

3.5 Results

3.5.1 Employment

We start by investigating the impact on overall firm-level employment, measured in logs, estimating our dynamic Difference-in-Differences specification and reporting the coefficients on the interaction term between year and the main treatment variable in Figure C·2. We report both the OLS and the IV coefficients. We find a small positive impact on overall employment in the short-run, that fades away in 2009. We do not detect any significant pre-trend.

We then move to evaluate the effects of relative employment for our two main groups of interest, namely EU07 migrants and natives, in Figure 3·3. Here, two opposite patterns emerge. On the one hand, we find an immediate, short-run increase in EU07 relative employment in 2007. The effect is sizeable, but it is very short-lived: the estimated coefficient for 2008 is very small in magnitude and statistically indistinguishable from 0, and in 2009 we even find a negative effect of negligible magnitude. Looking at the relative employment for natives, we find the opposite trend: on impact, relative employment significantly decreases, but then the effect disappears in 2008.

Figure 3·3: Firm-level Relative Employment, Natives and EU07 Workers

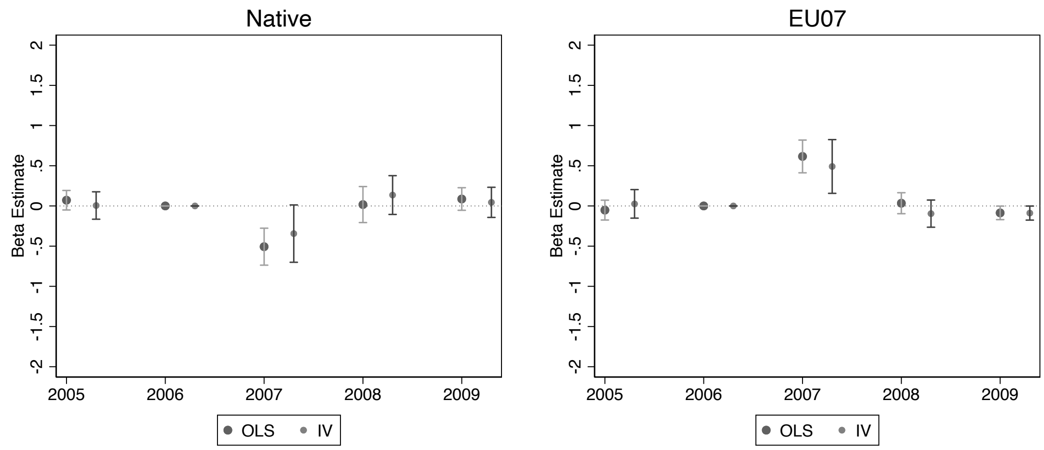

To better understand the drivers behind overall employment effects, we look at the

Figure 3·4: Firm-Level Hires and Separations, EU07 Workers

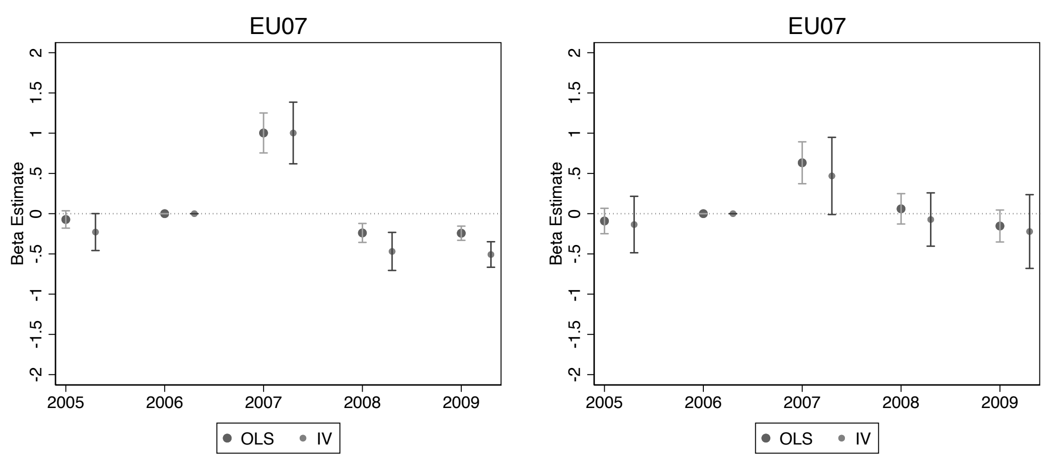

patterns for hires and separations. In Figure 3·4, we look at the dynamics for hires and separations of EU07 workers, expressed as a share of total employment. Looking at the left panel of the figure, we find a sizeable increase in EU07 hires in 2007 that is roughly twice as large as the effect that we find for the effect on the relative share of EU07 workers on total employment. The difference between the two effects is explained by looking at the right panel of Figure 3·4, where we find a sizeable increase in separations of EU07 migrants in 2007. Interestingly, the effect on hires is reversed in the medium run, as we find a small and significant reduction of EU07 hires in 2008 and 2009.

Overall, the two panels suggest a significantly higher dynamism of EU07 workers, who likely move across firms in response to the EU enlargement. This is perfectly consistent with the fact that it removes a crucial barrier to the job mobility of workers, by de-facto removing the relevance of employer-sponsored work permits.

Figure 3·5: Firm-Level Hires and Separations, Native Workers

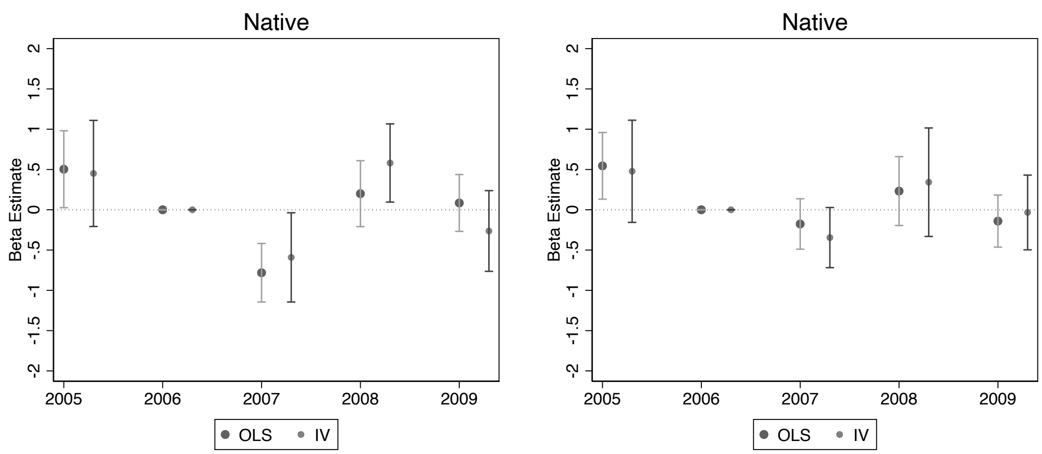

In Figure 3·5, we look at hirings and separations of natives. Firms are significantly less likely to hire natives in 2007, but the effect is reversed in 2008 and 2009. On the right, we do not observe significant changes in separations. Hence, firms are more likely to hire EU07 workers in the formal sector after 2007. This is likely to be the results of the combination of two factors. On one hand, a change in the legal status of EU07 workers makes them more similar to natives because the former have gained bargaining power. On the other hand, EU07 workers were likely to be already employed by firms, albeit informally, so starting in 2007 the legalization and emergence of EU07 workers in the formal sector prevailed in terms of hiring choices.

To gain a deeper understanding of the mechanisms through which these effects took place, we perform a series of heterogeneity analyses that show which types of firms are driving the results. In particular, we are going to look at two crucial margins. First, we investigate whether firms with previous experience of legalization of migrants through an amnesty are more likely to respond to the change in legal status. Second, we study whether the effects observed are different for firms that employed at least one EU07 worker prior to 2007.

Having participated in a previous legalization episode might be a proxy for the propensity of a given firm to hire workers illegally and then formalize them only when faced with lower administrative costs. We thus exploit the existence of a very large amnesty that took place in September 2002 and legalized over 700,000 undocumented migrants (see Porto et al. (2018) for a detailed description of the policy). We are able to identify the firms that participated in the amnesty, namely those who applied for the legalization of at least one migrant worker at that time.

In Table 3.3 we show the heterogeneous dynamics of employment for natives and EU07 workers.

Consistent with our hypothesis, across all outcomes, we find effects that are on average twice as large in magnitude for firms that had participated in the 2002 amnesty program.

The second types of firms we identify are those that already had at least one formally employed EU07 worker prior to 2007. Here, the direction of the expected effect is less obvious. On the one hand, these firms may be less responsive to the EU07 enlargement, as they were already willing to undertake the cost of formalizing workers before the change in their legal status. On the other hand, firms employing migrants in the

Table 3.3: Heterogeneity by Past experience with Migrants, Employment

	Native Relative Employment		
	(1) All	(2) W. Eu07 pre-2006	(3) Regularizing in 2002
2005×Δ ShareEU07_llm,07-06	0.0054	-0.0877	-0.3732
	(0.0870)	(0.2268)	(0.5814)
2007×Δ ShareEU07_llm,07-06	-0.3441*	-0.4194	-1.3847**
	(0.1819)	(0.2937)	(0.6298)
2008×Δ ShareEU07_llm,07-06	0.1355	0.2775	-0.1429
	(0.1230)	(0.2759)	(0.5924)
2009×Δ ShareEU07_llm,07-06	0.0448	-0.2406	0.3090
	(0.0956)	(0.1977)	(0.5303)
	EU07 Relative Employment		
	(1) All	(2) W. Eu07 pre-2006	(3) Regularizing in 2002
2005×Δ ShareEU07_llm,07-06	0.0257	0.1382	-0.6387
	(0.0906)	(0.2686)	(0.4337)
2007×Δ ShareEU07_llm,07-06	0.4909**	0.6786*	1.6301**
	(0.1705)	(0.3826)	(0.6008)
2008×Δ ShareEU07_llm,07-06	-0.0956	-0.3302	0.2142
	(0.0862)	(0.2046)	(0.5620)
2009×Δ ShareEU07_llm,07-06	-0.0887**	0.1096	0.1260
	(0.0446)	(0.1731)	(0.4383)
N	3,827,560	336,675	164,230

Note: Robust standard errors clustered at the Local Labor Market level reported in parentheses. $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

formal sector may be more likely to employ migrants also in the informal one, as well as have a higher access to these workers' network. We find that the second direction prevails, with patterns being stronger for firms employing EU07 workers prior to 2007 than the average.

3.5.2 Other outcomes

Using once again our preferred specification, we show how the EU07 enlargement has affected wages for natives and EU07 workers. In 3·6 we observe that for both types of workers there is no effect. Moreover, unlike the patterns reported for employment, we do not observe significant differences across years in the post period. Therefore, henceforth we will just show pre- and post-2007 coefficients for the analyses of wages.

Similarly to what we have done for the analysis on employment, we look separately at firms that were already employing at least one EU07 migrant before 2007 or that participated to the 2002 amnesty. We observe a decrease in the wage for both migrants and EU07 workers for the first type of firm, indicating that higher labor supply and

Figure 3·6: Firm-Level Wage, Native and EU07 Workers

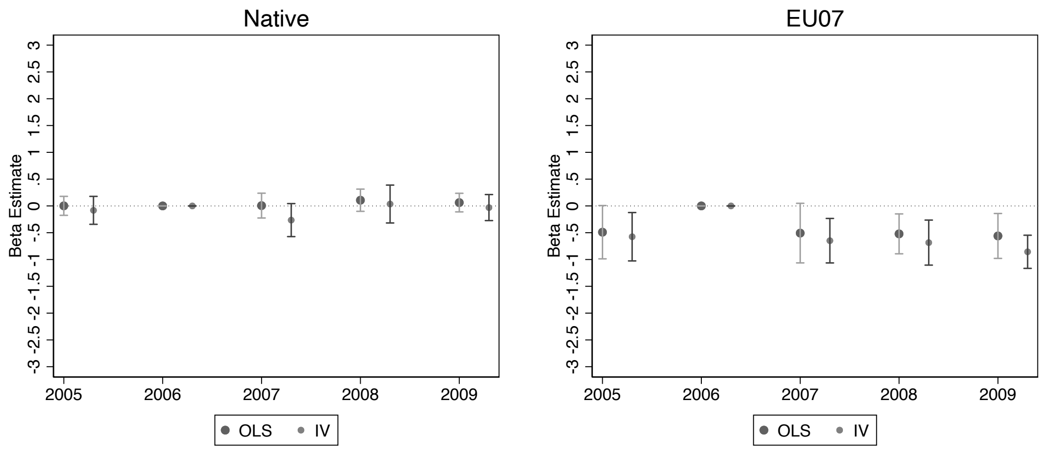

Table 3.4: Heterogeneity by Past experience with Migrants, Wage

		Native Wage (logs)	
	(1) All	(2) W. Eu07 pre-2006	(3) Regularizing in 2002
Post×Δ ShareEU07_llm,07-06	-0.0458	-0.2007**	0.1063
	(0.0838)	(0.0797)	(0.2980)
Post	-0.0096**	-0.0218**	-0.0108**
	(0.0022)	(0.0025)	(0.0047)
N	3,686,101	305,851	138,199
		EU07 Wage (logs)	
	(1) All	(2) W. Eu07 pre-2006	(3) Regularizing in 2002
Post×Δ ShareEU07_llm,07-06	-0.4591**	-0.2903**	-1.0870**
	(0.1041)	(0.1177)	(0.5349)
Post	-0.0228**	-0.0335**	-0.0009
	(0.0052)	(0.0061)	(0.0073)
N	243,010	201,035	48,591
		Average Wage (logs)	
	(1) All	(2) W. Eu07 pre-2006	(3) Regularizing in 2002
Post×Δ ShareEU07_llm,07-06	-0.0899	-0.3044**	-0.1414
	(0.0755)	(0.0709)	(0.2139)
Post	-0.0090**	-0.0170**	-0.0074*
	(0.0021)	(0.0033)	(0.0038)
N	3,827,437	336,672	164,229

Note: Robust standard errors clustered at the Local Labor Market level reported in parentheses. $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

competition had a small negative effect on all workers. For the second type of firm, we do not observe any effect on native workers, while there is an effect much stronger than in the average firm for EU07 workers. This difference can result from the fact that migrant workers operating with certainty (at least in the past) in the informal sector have even less bargaining power than their fellow citizens and are more likely to have lower qualifications.

It may come as a surprise that we observe a decrease in the wage for migrants affected by the EU expansion, given the stronger bargaining power acquired as a result of the changes in legal status. However, wages are not the only margin that might be relevant for workers. In fact, the Italian market is characterized by an important dualism in terms of type of contract, permanent or temporary. Given the strong preference on the employees' side for permanent contracts, it is important to investigate whether they made progress in their labor market performance along this dimension. Table C·3 in the Appendix shows that it is indeed the case that post 2007 a higher share of EU07 migrants is employed through a permanent contract (also referred to as open-eneded contract).

3.5.3 Business performance

Lastly, we study the effects of the EU07 expansion on measures related to business performance: revenues and worker value added. As stated previously, we do not have this type of information for all firms in our original dataset, therefore our analysis is limited to those that can be linked to the CERVED data, which provides business-related information on firms.

Figure 3·7: Firm-Level Business Performance Outcomes

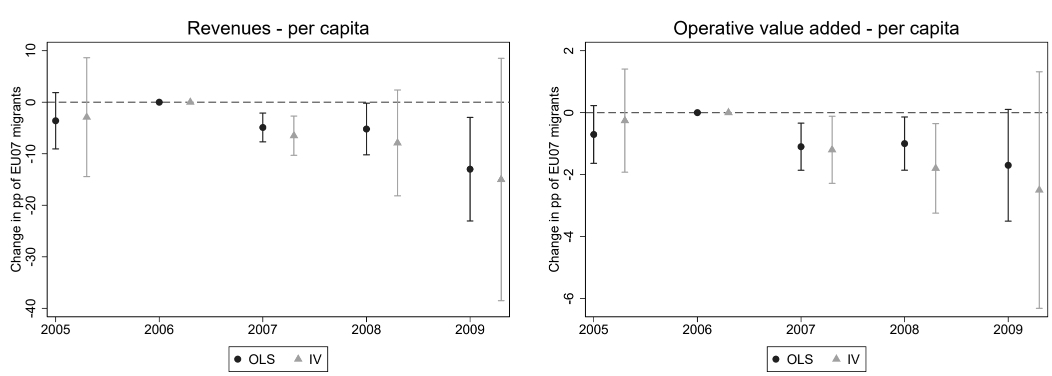

Figure 3·7 shows the result of this analysis: both revenues and worker value added per capita decrease after 2007. Both of these results indicate that firms hiring more workers were not increasing the volume of their operations, and therefore that the choice of expanding in terms of employment was not business driven. Thus, this provides further suggestive evidence that the increase in EU07 employment was mainly driven by higher bargaining power of EU07 workers who were previously employed informally, and after 2007 were able to achieve formal employment status in the labor market.

3.6 Conclusion

We study the consequences on the Italian labor market of the inclusion of Romania and Bulgaria in the European Union. The enlargement of the EU, which took place in 2007, resulted in a change in the legal status for a group of migrants (from Romania in particular) that had already a strong presence in Italy, due to the language similarities between Italian and Romanian. We study firms responses in terms of personnel choices using a unique administrative employer-employee dataset that covers the universe of private sector workers in Italy.

We observe short-term effects on firm-level employment: an increase for EU07 migrants at the expenses of natives. This is explained by a rise in hirings and separations for migrant workers, and by a decrease in hirings of native workers. Simultaneously, employment per-firm increases, while per-capita revenues and operative added value contract.

Our evidence suggest that the observed findings are mainly driven by the emergence into the formal labor market of unregistered migrant workers already present in the informal market. On one hand, the effects are short-lived and mainly present in 2007, the year in which the EU enlargement took place. On the other, increase in overall firm employment is not accompanied by evidence of firms' business growth.

Our results also suggest that there has been a change in market power in favor of EU07 migrants. With the EU enlargement, in fact, their ability of residing in Italy is no longer linked to a specific employer, which is reflected in the increase in separations of this type of workers. Interestingly, the change in legal status did not bring about a rise in wages for EU07 migrants, but it increased the share of workers among them with a more stable type of employment (permanent contracts), which is a very desirable contract feature in the Italian labor market.

Appendix A

Supplementary Material for Chapter 1

Figure A·1: Raw white-black annual wage gap

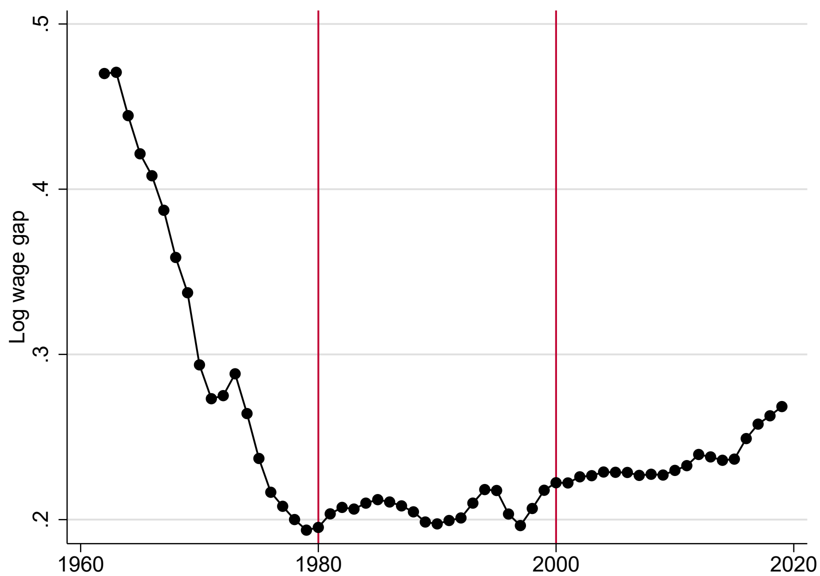

Source: CPS ASEC. Sample: All 20-64 black and white workers with non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

Figure A·2: Raw white-black hourly wage gap – Males

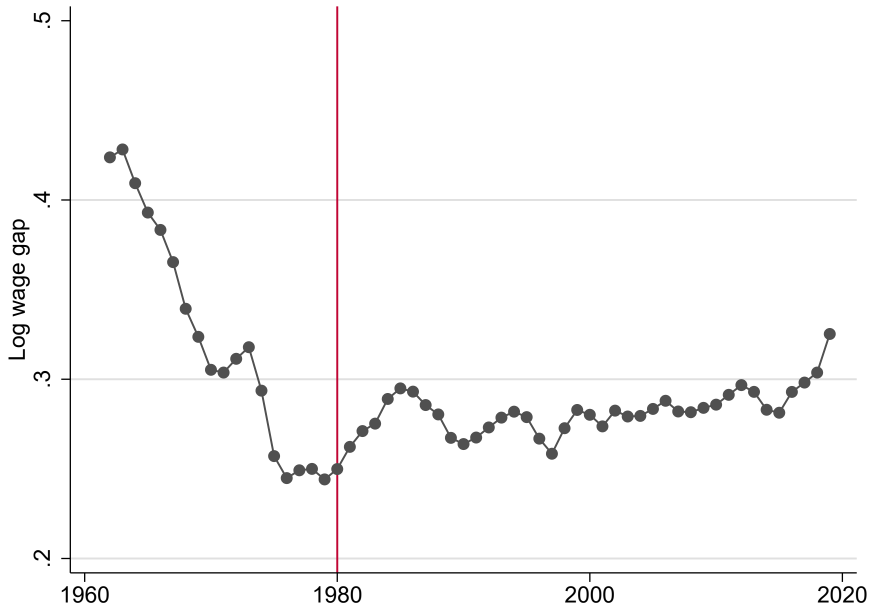

Source: CPS ASEC. Sample: Male 20-64 black and white workers with non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

Figure A·3: Race-specific employment share in top RTI occupations

(a) Full-time full-year workers

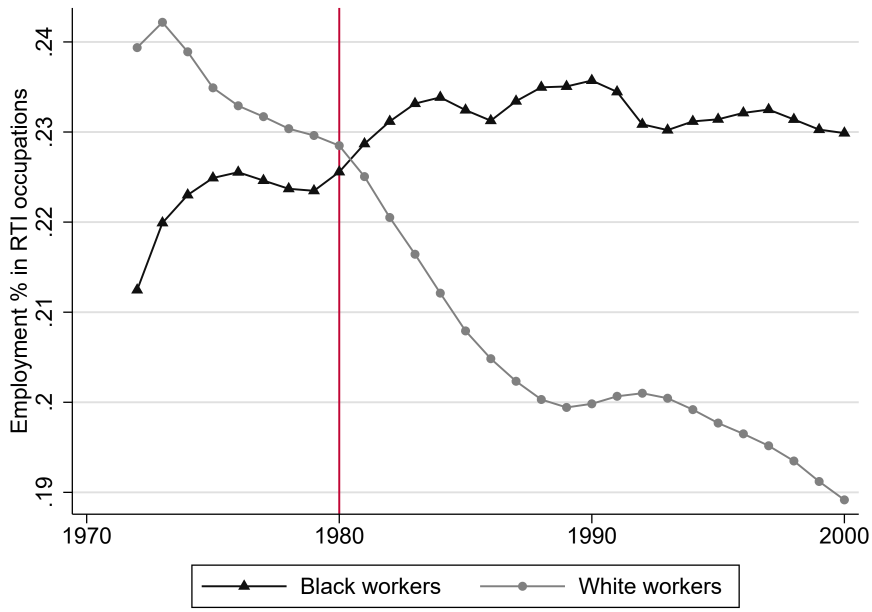

Source: CPS. Sample: All 20-64 black and white workers working at least 30 hours per week, for at least 40 weeks with non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

Figure A·4: Race-education employment share in top RTI occupations

(a) Some college or more

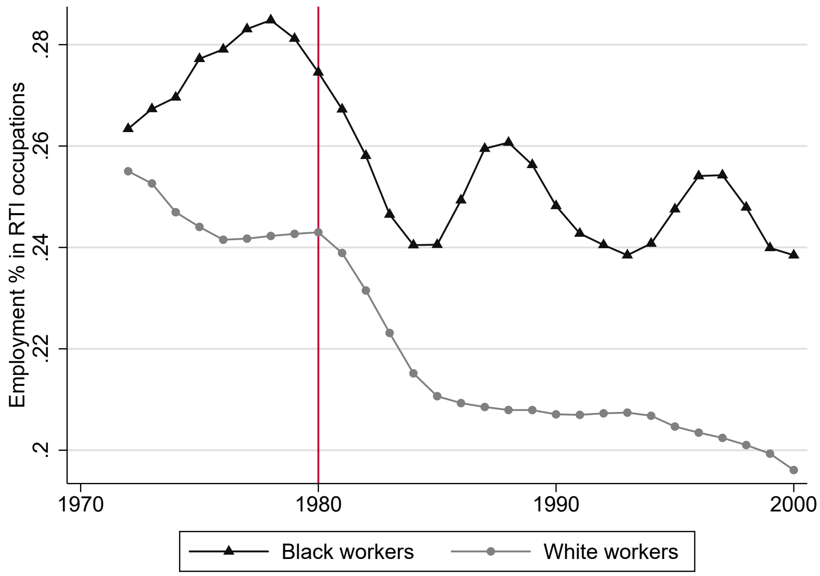

Source: CPS. Sample: All 20-64 black and white workers with some college or more, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Smoothing bandwidth: 0.2.

Figure A·5: 1980-2000 top RTI change along wage distribution

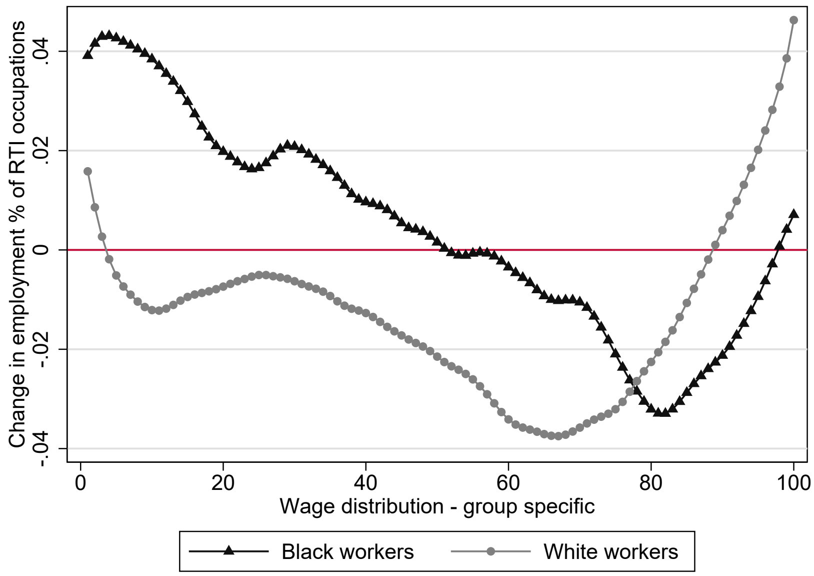

Source: IPUMS. Sample: All 20-64 black and white workers, with hourly wage ≥ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Wage distribution defined separately by race-gender groups. Smoothing bandwidth: 0.2.

Figure A·6: 1980-2000 white-black detailed *composition* change

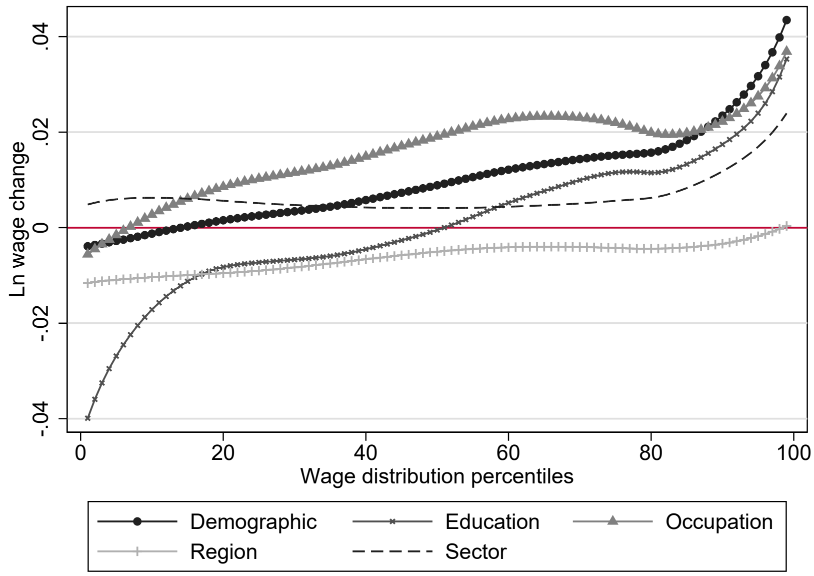

Source: IPUMS. Sample: All 20-64 black and white workers, with hourly wage ≥ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Wage distribution defined separately by race-gender groups. Smoothing bandwidth: 0.2.

Figure A·7: 1980-2000 white-black detailed *diff returns* change

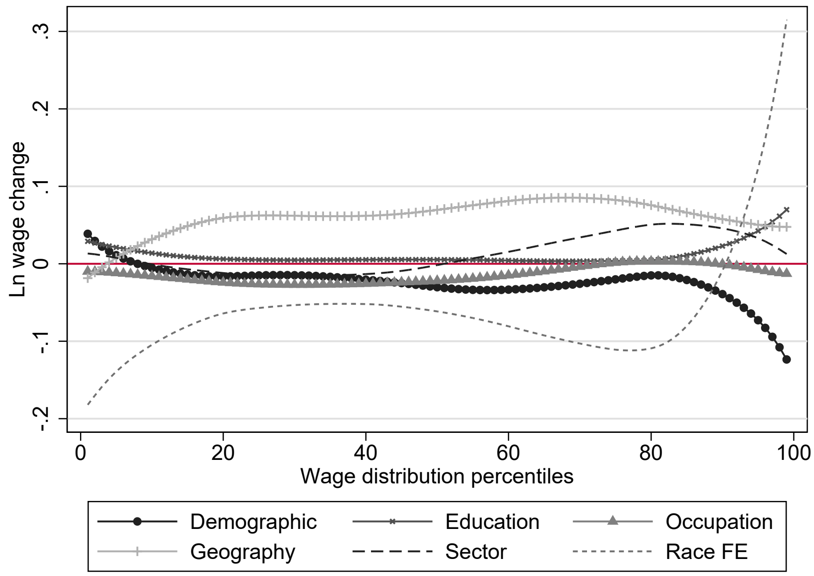

Source: IPUMS. Sample: All 20-64 black and white workers, with hourly wage ≥ 1, non-negative personal weights, not in group quarters, not self-employed, not unpaid family workers, not in agricultural occupations, no missing occupation or region. Wage distribution defined separately by race-gender groups. Smoothing bandwidth: 0.2.

Appendix B

Supplementary Material for Chapter 2

Table B.1: Skill Productivity Growth Relative to Average - non transformed coefficients

	(1) women 60-70	(2) women 70-80	(3) men 60-70	(4) men 70-80
Routine	-0.415	-0.896	-0.360	-0.462
	(0.183)	(0.186)	(0.226)	(0.144)
Manual	0.671	-0.901	-0.216	-0.268
	(0.421)	(0.484)	(0.366)	(0.196)
Finger dexterity	-1.238	-1.894	-0.549	-0.940
	(0.372)	(0.356)	(0.455)	(0.305)

Notes: Standard errors in parentheses, clustered at the occupation level. Estimates from regression of change in log employment in an occupation/industry cell on average skill (routine, manual, finger dexterity) use in that cell over the period (equation (2.22) in the text) before the transformation shown in Table 2.3.

Table B.2: Skill Productivity Growth Relative to Average - including wages

	(1) women 60-70	(2) women 70-80	(3) men 60-70	(4) men 70-80
Routine	-0.157	0.033	-0.076	-0.044
	0.151	0.163	0.156	0.103
Abstract	0.257	0.937	0.284	0.418
	0.160	0.186	0.188	0.123
Manual	0.903	0.002	0.073	0.149
	0.356	0.366	0.300	0.151
Finger dexterity	-1.002	-0.972	-0.280	-0.524
	0.289	0.260	0.313	0.207
% Change mean wage 60	0.000		0.012	
	(0.000)		(0.006)	
% Change mean wage 70		0.000		0.007
		(0.002)		(0.003)
F(3, N3)	21.97	10.16	22.42	14.63
p	0.000	0.000	0.000	0.000
r2	0.17	0.16	0.15	0.12
N	3046	4588	4848	6997

Notes: Standard errors in parentheses, clustered at the occupation level. Estimates are transformed from regression of change in log employment in an occupation/industry cell on average skill (routine, manual, finger dexterity) use in that cell over the period (equation (2.22) in the text) and imposing that the mean deviation from mean skill growth for all four skills is 0. Proportion due to skills is the proportion of the R-squared attributable to the three skills in the regression using the Shapley-Owen decomposition.

Table B.3: Derivatives of Skill Use with Respect to Skill Productivity

		Skill Used		
$\Delta \ lnA_i$	Routine	Abstract	Manual	Finger Dexterity
			Women	
Routine	-0.108			
Abstract	0.081	-0.063		
Manual	0	0.009	-0.012	
Finger Dexterity	0.027	0.027	-0.003	-0.051
			Men	
Routine	-0.522			
Abstract	0.277	-0.127		
Manual	0	-0.045	0.054	
Finger Dexterity	0.245	-0.105	-0.009	-0.131

Notes: Each cell shows the derivative of the average use of the column skill with respect to a change in the relative productivity of the row skill. Estimates are up to a factor of proportionality of $\frac{-\varepsilon}{1-\varepsilon}$ (which is strictly between 0 and 1). The estimates are derived from combining changes in skill use across time with estimates of relative productivity growth from Table 2.3. See equation (2.32) in the text for the precise formulation. The cross-derivative between routine and manual is set to 0. See the text for more detail.

Table B.4: Decomposition of Within-Occupation Changes in Skill Use

	Women 1960-1971 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	-0.018	0.014	0	0.004
Abstract	-0.020	0.016	-0.002	0.006
Manual	0	-0.008	0.011	-0.003
Finger Dexterity	0.027	-0.027	0.003	-0.003
Total Predicted	-0.011	-0.005	0.012	0.004
Data	-0.009	-0.003	0.014	-0.002

	Women 1971-1983 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	0.003	-0.002	0	-0.001
Abstract	-0.075	0.058	-0.008	0.025
Manual	0	0.000	0.000	0.000
Finger Dexterity	0.026	-0.026	0.003	-0.003
Total Predicted	-0.046	0.030	-0.005	0.021
Data	-0.047	0.026	-0.010	0.031

	Men 1960-1971 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	-0.041	0.022	0	0.019
Abstract	-0.078	0.036	0.013	0.029
Manual	0	0.003	-0.004	0.001
Finger Dexterity	0.066	-0.028	-0.002	-0.036
Total Predicted	-0.053	0.033	0.007	0.014
Data	-0.044	0.049	0.019	-0.024

	Men 1971-1983 Predicted Skill-Use Change			
Source of Change	Routine	Abstract	Manual	Finger Dexterity
Routine	-0.023	0.013	0	0.010
Abstract	-0.116	0.053	0.019	0.044
Manual	0	0.007	-0.008	0.001
Finger Dexterity	0.128	-0.055	-0.008	-0.065
Total Predicted	-0.011	0.018	0.003	-0.010
Data	-0.014	0.007	0.000	0.006

Notes: Each entry is the predicted change in the within-occupation use of the column skill due to changes in the productivity of the row skill according to equation (2.32) in the text and using the values from Tables 2.3 and 2.4. Total predicted is the sum of the four values above. The predictions can be compared with the within changes reported in Table 2.2 and repeated in the line labelled Data.

Appendix C

Supplementary Material for Chapter 3

Figure C·1: Specification validity

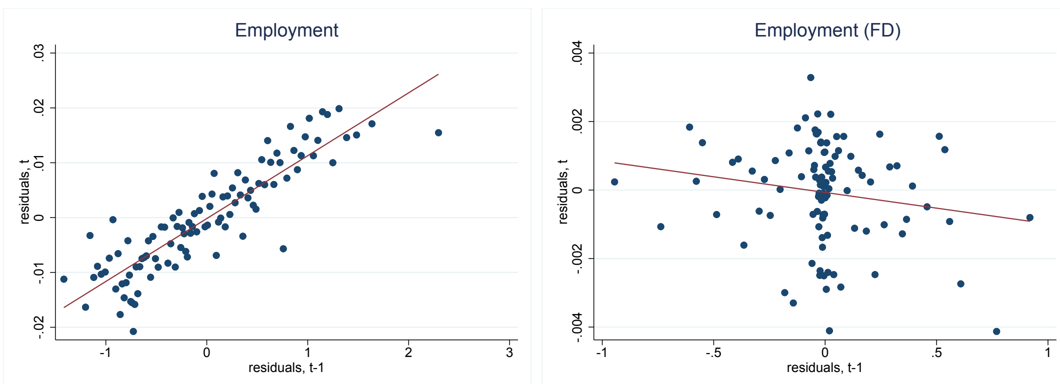

Figure C·2: Firm-level total employment (in logs)

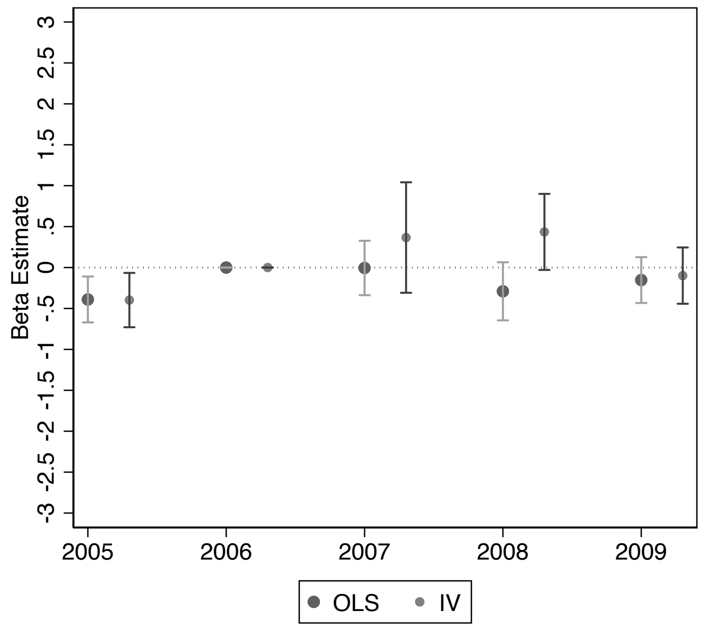

Figure C·3: Share of open-ended (permanent) contracts, EU07 workers

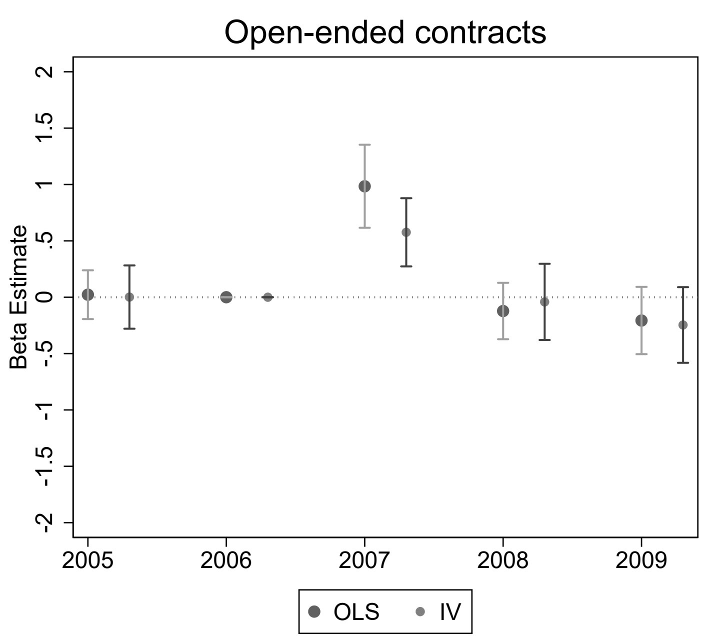

Table C.1: First-stage regression results

	(1) $\Delta ShareEU07_{llm,07-06}$
$\Delta ShareEU07_{llm,02-01}$	2.7472***
Dep. Var. Mean	0.0056
Observations	611
R-sq	0.455

Table C.2: Instrument validity

PANEL A:	$\Delta ShareEU07_{llm,02-01}$ (1)	S.E. (2)
Native Empl	-2.3626	(2.5682)
Native Wage	-0.9826	(1.7760)
EU07 Empl	0.3375	(0.2347)
EU07Wage	-0.2511	(0.2907)
PANEL B:	$\Delta ShareEU07_{llm,02-01}$ (1)	S.E. (2)
Native Empl (1992-1997)	-0.0825	(0.3121)
Native Wage (1992-1997)	0.0717	(0.1120)
EU07 Empl (1992-1997)	-0.4498	(0.2905)
EU07Wage (1992-1997)	-1.9322	(1.7914)
PANEL C:	$\Delta ShareEU07_{llm,02-01}$ (1)	S.E. (2)
Native Empl (2002-2006)	-2.3626	(2.5682)
Native Wage (2002-2006)	-0.9826	(1.7760)
EU07 Empl (2002-2006)	0.3375	(0.2347)
EU07Wage (2002-2006)	-0.2511	(0.2907)

References

Acemoglu, D. and Autor, D. H. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Guerrieri, V. (2008). Capital deepening and nonbalanced economic growth. *Journal of Political Economy*, 116(3):467–498.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Alesina, A., Miano, A., and Stantcheva, S. (2018). Immigration and redistribution. NBER Working Paper 24733.

Altonji, J. G. and Blank, R. M. (1999). Race and gender in the labor market. In *Handbook of Labor Economics*, volume 3, pages 3143–3259. Elsevier.

Amior, M. and Manning, A. (2020). Monopsony and the wage effects of migration. CEP Discussion Paper 1690.

Aneja, A. and Avenancio-Leon, C. F. (2019). The effect of political power on labor market inequality: evidence from the 1965 Voting Rights Act. Unpublished.

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. 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., Katz, L. F., and Kearney, M. S. (2008). Trends in US wage inequality: revising the revisionists. *Review of Economics and Statistics*, 90(2):300–323.

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.

Autor, D. H. and Price, B. (2013). The changing task composition of the US labor market: an update of Autor, Levy, and Murnane (2003). Mimeo.

Batistich, M. K. and Bond, T. (2019). Stalled racial progress and Japanese trade in the 1970s and 1980s. IZA Discussion Paper 12133.

Baumol, W. J. (1967). Macroeconomics of unbalanced growth: the anatomy of urban crisis. *American Economic Review*, 57(3):415–426.

Bayer, P. and Charles, K. K. (2018). Divergent paths: A new perspective on earnings differences between black and white men since 1940. *The Quarterly Journal of Economics*, 133(3):1459–1501.

Beerli, A., Ruffner, J., Siegenthaler, M., and Peri, G. (2021). The abolition of immigration restrictions and the performance of firms and workers: evidence from Switzerland. *American Economic Review*, 111(3):976–1012.

Berman, E., Bound, J., and Griliches, Z. (1994). Changes in the demand for skilled labor within US manufacturing: evidence from the annual survey of manufactures. *The Quarterly Journal of Economics*, 109(2):367–397.

Berman, E., Bound, J., and Machin, S. (1998). Implications of skill-biased technological change: international evidence. *The Quarterly Journal of Economics*, 113(4):1245–1279.

Black, S. E. and Spitz-Oener, A. (2007). Explaining women's success: technological change and the skill content of women's work. NBER Working Paper 13116.

Black, S. E. and Spitz-Oener, A. (2010). Explaining women's success: technological change and the skill content of women's work. *Review of Economics and Statistics*, 92(1):187–194.

Bond, T. N. and Lang, K. (2013). The evolution of the Black-White test score gap in grades K–3: the fragility of results. *Review of Economics and Statistics*, 95(5):1468–1479.

Bond, T. N. and Lang, K. (2019). The sad truth about happiness scales. *Journal of Political Economy*, 127(4):1629–1640.

Bound, J. and Freeman, R. B. (1992). What went wrong? The erosion of relative earnings and employment among young black men in the 1980s. *The Quarterly Journal of Economics*, 107(1):201–232.

Cavounidis, C., Dicandia, V., Lang, K., and Malhotra, R. (2021). The evolution of skill use within and between jobs. NBER Working Paper 29302.

Cavounidis, C. and Lang, K. (2020). Ben-Porath meets Lazear: microfoundations for dynamic skill formation. *Journal of Political Economy*, 128(4):1405–1435.

Coate, S. and Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review*, 83(5):1220–1240.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and STEM careers. *The Quarterly Journal of Economics*, 135(4):1965–2005.

Derenoncourt, E. and Montialoux, C. (2021). Minimum wages and racial inequality. *The Quarterly Journal of Economics*, 136(1):169–228.

Dickens, W. T. (1990). Error components in grouped data: is it ever worth weighting? *Review of Economics and Statistics*, 72(2):328–333.

Dorn, D. (2009). *Essays on inequality, spatial interaction, and the demand for skills*. PhD thesis, University of St. Gallen.

Dustmann, C., Fasani, F., and Speciale, B. (2017). Illegal migration and consumption behavior of immigrant households. *Journal of the European Economic Association*, 15(3):654–691.

Dustmann, C. and Glitz, A. (2015). How do industries and firms respond to changes in local labor supply? *Journal of Labor Economics*, 33(3):711–750.

Elias, F., Monras, J., and Vàzquez Grenno, J. (2019). Understanding the effects of legalizing undocumented immigrants. CEPR Discussion Paper 12726.

Fang, H. and Moro, A. (2011). Theories of statistical discrimination and affirmative action: A survey. In *Handbook of Social Economics*, volume 1, pages 133–200. Elsevier.

Firpo, S. P., Fortin, N. M., and Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77(3):953–973.

Firpo, S. P., Fortin, N. M., and Lemieux, T. (2018). Decomposing wage distributions using recentered influence function regressions. *Econometrics*, 6(2):28.

Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., and Westberry, M. (2021). Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]. IPUMS, Minneapolis, MN. DOI : 10.18128/D030.V9.0.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–26.

Hurst, E., Rubinstein, Y., and Shimizu, K. (2021). Task-based discrimination. NBER Working Paper 29022.

Jones, C. I. (2011). Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 3(2):1–28.

Juhn, C. (1999). Wage inequality and demand for skill: evidence from five decades. *ILR Review*, 52(3):424–443.

Juhn, C., Murphy, K. M., and Pierce, B. (1991). Accounting for the slowdown in black-white wage convergence. *Workers and their wages*, pages 107–43.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Lang, K. and Lehmann, J.-Y. K. (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature*, 50(4):959–1006.

Larrimore, J., Burkhauser, R. V., Feng, S., and Zayatz, L. (2008). Consistent cell means for topcoded incomes in the public use March CPS (1976–2007). *Journal of Economic and Social Measurement*, 33(2-3):89–128.

Lemieux, T. (2006). Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? *American Economic Review*, 96(3):461–498.

Lundberg, S. J. and Startz, R. (1983). Private discrimination and social intervention in competitive labor market. *American Economic Review*, 73(3):340–347.

Mastrobuoni, G. and Pinotti, P. (2015). Legal status and the criminal activity of immigrants. *American Economic Journal: Applied Economics*, 7(2):175–206.

Mitaritonna, C., Orefice, G., and Peri, G. (2017). Immigrants and firms' outcomes: evidence from France. *European Economic Review*, 96:62–82.

Naidu, S., Nyarko, Y., and Wang, S.-Y. (2016). Monopsony power in migrant labor markets: Evidence from the United Arab Emirates. *Journal of Political Economy*, 124(6):1735–1792.

National Academy of Sciences. Committee on Occupational Classification and Analysis (1977). Dictionary of Occupational Titles (DOT): Part I - Current Population Survey, April 1971, Augmented With DOT Characteristics and Dictionary of Occupational Titles (DOT): Part II - Fourth Edition Dictionary of DOT Scores for 1970 Census Categories. ICPSR7845 - v2, Inter - University Consortium for Political and Social Research, Ann Arbor , MI. DOI : 10.3886/ICPSR07845.v2.

National Research Council (2014). The growth of incarceration in the United States: exploring causes and consequences. Technical report, Washington, DC.

Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth. *American Economic Review*, 97(1):429–443.

Ottaviano, G. I. P. and Peri, G. (2012). Rethinking the effect of immigration on wages. *Journal of the European Economic Association*, 10(1):152–197.

Ottaviano, G. I. P., Peri, G., and Wright, G. C. (2018). Immigration, trade and productivity in services: evidence from UK firms. *Journal of International Economics*, 112:88–108.

Pinotti, P. (2017). Clicking on heaven's door: The effect of immigrant legalization on crime. *American Economic Review*, 107(1):138–168.

Porto, E. D., Martino, E. M., and Naticchioni, P. (2018). Back to black? The impact of regularizing migrant workers. CSEF Working Papers 517.

Ruggles, S., Flood, S., Foster, S., Goeken, R., Pacas, J., Schouweiler, M., and Sobek, M. (2021). IPUMS USA: Version 11.0 [dataset]. IPUMS, Minneapolis, MN. DOI : 10.18128/D010.V11.0.

Sabety, A. and Spitzer, A. K. (2021). Missing black men? The impact of underreporting on estimates of black male labor market outcomes. Unpublished.

Signorelli, S. (2020). Do skilled migrants compete with native workers? Analysis of a selective immigration policy. Job market paper.

Smith, J. P. and Welch, F. R. (1986). Closing the gap: Forty years of economic progress for Blacks. Technical report, Rand Corporation.

Solon, G., Haider, S. J., and Wooldridge, J. M. (2015). What are we weighting for? *Journal of Human Resources*, 50(2):301–316.

United States Department of Labor (1965). *Dictionary of Occupational Titles (DOT): Third Edition*, volume 1 and 2. US Government Printing Office, Washington, DC.

United States Department of Labor. United States Employment Service, and the North Carolina Occupational Analysis Field Center (1991). Dictionary of Occupational Titles (DOT): Revised Fourth Edition. ICPSR6100 - v1, Inter - University Consortium for Political and Social Research, Ann Arbor , MI. DOI : 10.3886/ICPSR06100.v1.

U.S. Bureau of Economic Analysis (2021). Personal Consumption Expenditures [PCE]. Federal Reserve Economic Data, Federal Reserve Bank of St. Louis.

Western, B. and Pettit, B. (2005). Black-white wage inequality, employment rates, and incarceration. *American Journal of Sociology*, 111(2):553–578.

CURRICULUM VITAE

[redacted]

[redacted]

[redacted]