Technological Change and Unions

An Intergenerational Conflict with Aggregate Impact

Leon Huetsch[†]

University of Pennsylvania

April 27, 2024

[Please click here for the latest version]

Abstract

How do labor adjustment costs, specifically in the form of unionization, shape the evolution of wages and employment of workers exposed to labor replacement by automation? I argue that, by raising adjustment costs, unions generate intergenerational redistribution by shifting the impact from existing, older to incoming, younger cohorts, and further generate aggregate effects by accelerating overall labor reallocation from automating to non-automating occupations. The reason is that labor adjustment costs incentivize firms to adjust through hiring rather than layoffs, and to reduce labor in anticipation of future adoption. Using variation across local labor markets in the U.S. since 1980, I document that unionization among exposed workers is associated with greater wage and employment decline among young relative to older workers, and with accelerated overall employment decline. I then develop an overlapping generations model of technological change and unionization that rationalizes the empirical findings through the impact of union-imposed labor adjustment costs on firms' choice how to transform their workforce over time when gradually adopting automation. Within automating occupations, unions reduce the welfare cost of automation of older workers along the transition by up to 4% of permanent consumption while raising the welfare costs of cohorts entering during the transition by up to 2%. Incoming workers endogenously respond to automation by entering non-adopting occupations. The union effect spills over into non-adopting occupations as the accelerated labor reallocation depresses wages there.

Keywords: Technological change, unions, intergenerational transfers, welfare cost of automation

JEL classification codes: E24, J24, J51, O33

[†]I am highly indebted to my advisors and committee members Dirk Krueger, Joachim Hubmer, Jesús Fernández-Villaverde as well as Alan Taylor and Alexander Ludwig for their invaluable support and guidance over the years. For helpful comments and discussions, I am also grateful to Víctor Ríos-Rull, Guillermo Ordoñez, Marko Mlikota, João Ritto, Priyanka Goonetilleke, Xincheng Qiu, Daniel Jaar, Marlon Azinovic, Min Kim as well as seminar and conference participants at University of Pennsylvania, the Chair of Public Finance and Macroeconomic Dynamics at Goethe University Frankfurt, the Philadelphia Fed and numerous other institutions. All errors are my own.

Email: lhuetsch@sas.upenn.edu; website: www.leonhuetsch.com.

1 Introduction

Technological change improves productivity and standards of living but creates winners and losers among workers. One of the most prominent technological changes in recent decades is the adoption of automation technologies, which boost productivity but can temporarily disrupt labor markets for transitional generations through worker displacement and reduced earnings (Graetz and Michaels (2018), Humlum (2020), Acemoglu and Restrepo (2020)). Increased adoption of automation technologies, such as industrial robots and artificial intelligence, has spurred an active literature studying the labor market impact on workers and discussing appropriate policy responses, most notably taxing automation (Beraja and Zorzi (2022), Costinot and Werning (2018)).

The existing literature on the labor market impact of automation assumes that firms can freely adjust their workforce, abstracting from labor adjustment frictions that firms face. Yet, the adoption of automation technologies entails substantial labor adjustment, making such frictions especially potent during the transition. Moreover, it is well documented that labor market institutions and employment protection laws that generate labor adjustment frictions are empirically associated with reduced employment flows and increased capital deepening (Bassanini and Duval (2006)), as well as with raising unemployment particularly during times of economic turbulence (Ljungqvist and Sargent (1998, 2008)).

In this paper, I ask how labor adjustment costs that firms face shape the impact of labor-replacing technological change on workers across different generations by studying the effect of unions on the evolution of employment and wages of workers who are substitutable with automation technologies. I argue that, by raising labor adjustment costs, unions shape the transformation of workforces in two ways. First, labor adjustment costs incentivize firms to adjust through hiring rather than layoffs, thereby shifting the adverse labor market impact from existing, older to incoming, younger cohorts and generating intergenerational redistribution effects. Second, labor adjustment costs incentive firms to shrink their workforce in anticipation of future automation adoption in order to smooth labor adjustment over time, thereby accelerating overall labor reallocation and generating aggregate employment effects.

To support this argument, I start by providing empirical evidence showing that unionization among workers exposed to automation is associated with, first, larger employment and wage decline among young relative to older workers and, second, accelerated overall employment decline among these exposed workers. To further strengthen the empirical evidence, I build a simple theoretical model that illustrates the effect of labor adjustment costs on the evolution of the age composition of workers during a labor-replacing technological transition and show that the empirical findings are consistent with the predictions of the model. I then develop a quantitative dynamic equilibrium model of unionization and technological transitions to first show that the proposed mechanism can quantitatively account for the documented intergenerational redistribution and aggregate employment effects of unionization. I then use the model to study how unionization, and labor adjustment costs more broadly, shapes the intergenerational distribution effects of automation in terms of welfare.

In the empirical exercise, I focus on the wages and employment of workers in routine manual occupations, which have been particularly exposed to labor-replacing technologies, such as automation (Goos et al. (2014)). I combine data from several sources to exploit variation in unionization and the evolution of worker outcomes within such occupations across local labor markets in the U.S. since 1980.[1] First, I find that unionization is associated with a greater fall in employment and wages among young workers entering the labor market relative to older workers, consistent with insider-outsider dynamics.[2] In particular, comparing a low-unionized to a high-unionized labor market implies that the routine manual employment share of young workers below the age of 30 falls by an additional 11% and their wage falls by an additional 9% relative to the average routine manual wage during the first 10 years of the transition between 1980 and 1990.[3] As a result, the routine manual workforce in more unionized labor markets becomes older relative to less unionized labor markets, and I show that this relative aging persists throughout the technological transition. Second, unionization is associated with an accelerated decline in overall routine manual employment while not significantly affecting the long-run decline. In particular, I document a greater employment decline in high-unionized labor markets early in the transition, and a subsequent slow catch-up of employment decline in less unionized labor markets from 2000 onwards. By 2020, the gap has mostly closed.

Motivated by these findings, I develop a quantitative dynamic equilibrium model of unionization and endogenous technological change that demonstrates that the interaction of union-imposed convex labor adjustment costs and gradual technology adoption over time can jointly rationalize the documented distributional and aggregate effects of

[1] I define a local labor market in the baseline analysis as a Metropolitan Statistical Area (MSA) and check that results hold at the state level and across Consistent Public Use Microdata Areas (conspuma) which is a finer local labor market measure than MSAs.

[2] See e.g. Carruth and Oswald (1987) who theoretically introduce insider-outsider dynamics into a standard union model.

[3] Here I define a low-unionized and a high-unionized labor market as the average MSA at the 10th and 90th percentile of routine manual unionization, respectively.

unions. First, adjustment costs give rise to a static effect by incentivizing firms to replace their workforce through reduced hiring rather than through costly layoffs when adopting automation. Second, there is a dynamic effect. In the context of expected gradual technology adoption, firms smooth their labor adjustment along the transition by shrinking their workforce preemptively today in order to avoid adjustment costs in the future. Intuitively, any worker not hired in 1980 is a worker the firm will not have to lay off when more automation technologies are adopted later. This dynamic effect is not dependent on but particularly strong in the context of convex labor adjustment costs as convexity generates an additional incentive to smooth labor adjustments over time. As a result, the dynamic mechanism gives rise to accelerated overall employment decline in routine occupations in high relative to low unionized labor markets. Similar in spirit to the extensive literature on labor market institutions and economic turbulence, the mechanism here builds on the interaction of adjustment costs and expectations about future technology adoption in driving current labor demand.[4]

To answer my research question, I use the model as a measurement device to quantify the impact of automation and unionization on labor market outcomes and life-cycle consumption paths of different cohorts of exposed workers during the automation transition. At its core, the model combines three building blocks that make it a suitable quantitative framework for that objective: 1) firms combine routine and non-routine occupations, they endogenously automate the routine occupations over time; 2) routine workers are represented by a labor union that raises labor adjustment costs and endogenously sets their wages; and 3) overlapping generations of workers make occupational choices and accumulate occupation-specific human capital, which makes switching occupations costly, more so when old. The rate of unionization in the model is parameterized by the level of exogenous labor adjustment costs, consistent with empirical evidence that unions raise firing costs (CITE). Moreover, labor adjustment costs determine the union's ability to impose wage premia by reducing the elasticity of labor demand of firms, making them an intuitive measure of unionization in the model. The union sets separate wages for routine workers of different ages as they are imperfect substitutes in production due to skill accumulation, which allows the model to speak to the observed union effect on the wage ratio between young and old workers. The two-sector setup with occupational choice endogenizes the supply of workers. This allows me to decompose the documented overall employment decline in the routine occupations into a downward shift in demand driven by technology adoption and an endogenous supply response driven by incoming workers entering and switching to non-routine occupations instead

[4]See, for example, Blanchard and Summers (1986), Ljungqvist and Sargent (1998, 2008).

in order to avoid the automation impact. The model therefore captures the life-cycle paths of a group of workers that is affected by unionization but that is difficult to identify empirically: workers who would have entered routine occupations but, due to the combination of automation and unionization, enter non-routine occupations instead.

Firms trade off two opposing forces in their routine labor demand when deciding how to optimally adjust their workforce along the transition. First, adjustment costs incentivize adjustment through incoming workers as well as preemptively reducing the workforce in anticipation of automation adoption to avoid adjustment costs in the future. Second, routine workers of different ages are imperfect substitutes in production because they are finitely lived and accumulate occupation-specific skills on the job which allows older workers to complete different tasks in production. Firms therefore prefer a balanced age composition of routine workers, which constrains the incentive to adjust through young, incoming workers only.

I calibrate the model to U.S. local labor market data, targeting in particular life-cycle wage profiles, the routine employment share, and the aggregate labor share in 1980 and 2010. I model a technological transition through an unexpected fall in the path of automation prices from 1980 onward that matches the price path of capital goods in the U.S. The level of adjustment costs measures the degree of unionization in the model and is calibrated to match the relative decline in the routine employment share between 1980 and 1990 in high and low-unionized MSAs. Lastly, I connect the model with the empirical findings by validating that it matches the untargeted evolution of overall routine employment and the evolution of the age composition of routine workers along the transition.

I first evaluate the impact of automation adoption on routine workers in a low-unionized labor market. Automation is most costly for incumbent routine workers who made their occupational choice without anticipating the upcoming transition. These workers are caught by surprise, facing the option to either stay in a declining sector or switch into non-routine occupations at the cost of losing their occupation-specific human capital. Especially routine workers who entered between 1970 and 1980 experience the full automation impact over their entire life-cycle, resulting in large permanent earnings losses. As a result, the welfare cost of automation to these workers reaches 10% of permanent consumption in 2000, measured as the permanent percent decrease in consumption they would be willing to accept to avoid automation and remain at the 1980 steady state. Workers entering the labor market during the automation transition take the current and future impact of automation into account when making their occupational choice. As routine jobs become less desirable, only workers with a sufficiently large labor productivity in routine tasks relative to non-routine tasks still enter the routine occupation. As a result, average labor productivity, and in turn average life-cycle earnings and consumption paths, of incoming routine cohorts rise. This endogenous response to automation limits its impact on incoming cohorts. Nevertheless, entering routine workers would still pay up to 7% of permanent consumption to avoid automation.

I then study the automation impact in a high-unionized labor market to quantify to what extent unions reallocate the welfare cost of automation across generations. Unions protect incumbent routine workers by lowering their layoff risk and limiting their wage decline, which reduces the welfare cost of automation for incumbent cohorts by up to 4% of permanent consumption along the transition in the high relative to the low unionized labor market. However, the impact is shifted to incoming cohorts. As a result, the welfare cost of automation for incoming routine workers is up to 2% of permanent consumption larger in the high relative to the low unionized labor market, driven by falling routine entry wages. The difference in the welfare benefit for incumbent and the welfare cost for incoming cohorts reflects the ability of incoming workers to endogenously respond to automation by entering the non-routine occupation instead. Consistent with the empirical findings, high unionization causes a faster reallocation of employment from the routine to the non-routine occupation as firms in the high-unionized labor market preemptively reduce hiring in order to avoid future adjustment costs. The accelerated reallocation of labor means that non-routine wages fall relatively more in the high-unionized labor market early in the transition, resulting in a larger spillover of the automation impact from routine to non-routine occupations.

Lastly, motivated by the model findings, I empirically evaluate the political implications of the intergenerational conflict that unions magnify. An emerging political economy literature connects adverse economic outcomes to ideological realignment as well as a shift in political preferences and voting behavior (McCarty et al. (2016), Voorheis et al. (2015)). Autor et al. (2020) link trade-exposure to rising political support for strong-left and strong-right views as well as to pure rightward shifts across local labor markets in the U.S. My welfare analysis emphasizes that unionization has magnified the negative impact of automation on labor market experiences of less skilled cohorts of workers who entered routine and non-routine occupations after 1980. Cohorts of workers that have entered the labor market between 1980 and 2000 are in their 50s and 60s today, thus, the workers whose voting behavior has shifted (Pew Research Center (2014, 2017)). I test and find empirical support for the hypothesis that the union-induced employment decline among young routine manual workers in the 1980s across local labor markets is associated with a shift in voting from Democrats to Republicans in the 2016 and 2020 presidential elections relative to previous elections. This suggests that while unions protected incumbent workers from the adverse automation impact, this also induces worsening labor market experiences for incoming workers and a shift in political preferences among these workers today.

Literature. This paper contributes first and foremost to the extensive empirical and quantitative literature studying the labor market impact of automation and its contribution to rising inequality (Acemoglu (2002) Goldin and Katz (2008)). Several papers provide empirical evidence on the effect of industrial robots adoption on worker outcomes and productivity (Graetz and Michaels (2018), Acemoglu and Restrepo (2020), Humlum (2020), Koch et al. (2021), Bessen et al. (2023)), documenting that robot adoption raises productivity, output and the wage bill of skilled workers while reducing the wages and employment share of less skilled workers. Similarly, across U.S. commuting zones, Acemoglu and Restrepo (2020) find negative effects on wages and employment that are more pronounced in routine manual and blue-collar occupations. A large empirical body of work specifically documents the decline of routine employment since 1980.[5] I contribute to this literature by studying the role of labor adjustment costs, specifically in the form of unionization. I show that labor adjustment costs shape the intergenerational distribution effects of automation as well as the timing of aggregate labor reallocation and put forward as the underlying mechanism the dynamic effects of adjustment costs on firms' decision how to optimally adjust their workforce over time. I thereby emphasize that, in the context of labor adjustment costs, automation is a dynamic choice because expectations about future workforce adjustments drive current labor demand. Consistent with Acemoglu and Restrepo (2020) who document negative spillover effects of robot adoption on nontradable sectors such as construction and services, my model suggests that as unionization induces faster labor reallocation to non-adopting occupations, it also accelerates the negative spillover effects on wages in these occupations.

A second body of work studies the intergenerational distribution effects of structural changes, such as trade or automation, and the resulting aggregate transition dynamics (Guerreiro et al. (2022), Costinot and Werning (2018), Beraja and Zorzi (2022)). Similar to me, several papers emphasize the role of occupation-specific human capital (Adao et al. (2021), Dvorkin and Monge-Naranjo (2019), Guren et al. (2015), Traiberman (2019)). Through the lens of an estimated occupational choice model, Traiberman (2019) argues that the costs of switching occupations in response to trade shocks is largely driven by the resulting loss of occupation-specific human capital rather than other switching

[5] See, for example, Autor et al. (2003), Goos and Manning (2007), Autor (2010), Cortes et al. (2020).

costs, such as retraining. Related to me, Adao et al. (2021) and Guren et al. (2015) focus on the role of different generations of workers during technological transitions within overlapping generation frameworks. Guren et al. (2015) focus on the dynamics of labor markets in response to trade shocks and study the role of sector-specific human capital in driving worker decisions to reallocate across sectors, and, thus, labor mobility. Adao et al. (2021) focus on transitions triggered by the arrival of new technologies and argue that transitions are slower when innovations require learning new skills as labor adjustment is then more strongly driven by entering cohorts rather than labor reallocation of existing, older workers. Focusing on automation, I also emphasize the importance of occupation-specific human capital of existing workers and the occupational choices by incoming cohorts in determining the incidence of wage and employment decline across generations of exposed workers and for the transitional dynamics. I contribute to this strand of the literature by developing a structural model of the intergenerational distribution effects of automation to quantify the welfare consequences of automation across generations. Moreover, I incorporate imperfect labor market competition to study how unions shift the distributional consequences of technological change.

Lastly, this paper also contributes to the extensive body of research trying to reconcile the rise and persistence of high European unemployment since the 1970s compared to the U.S. labor market. The literature has identified the interaction between shocks and labor market institutions as a key factor for the transatlantic gap in unemployment.[6] In their seminal work Ljungqvist and Sargent (1998, 2008) argue that labor market institutions in Europe, particularly policies of employment protection that increase the cost of layoffs, reduced employment flows in the 1950s and 1960s when the economic environment was calm, thereby lowering frictional unemployment. However, economic turbulence starting in the 1970s limited the reemployment options for laid-off workers due to old human capital becoming obsolete and, as a result, labor market institutions in the form of generous unemployment benefits then increased unemployment by reducing the incentive of laid off workers to accept wage cuts in their new jobs. My paper is closely related to this broad literature, also emphasizing the interaction between labor market institutions, here in the form of labor adjustment costs, and economic change, here in the form of automation adoption. I contribute to this literature in two ways. First, I focus on the impact on employed workers and labor reallocation rather than unemployment. Second, I emphasize the expectation about future economic turbulence rather than current turbulence in driving current labor demand decision, thus, the dynamic effects of labor

[6] See, among others, Blanchard and Summers (1986), Lindbeck and Snower (1988), Ljungqvist and Sargent (1998, 2008), Haan et al. (2005), Baley et al. (2020, 2023).

market institutions.

The remainder of the paper is organized as follows: Section 2 documents the empirical findings. Section 3 develops the model and discusses its elements. Section 4 takes the model to the data by outlining the calibration strategy and validating the model output. Section 5 presents the main quantitative analysis. Section 6 briefly presents evidence on the political implications of the model findings and Section 7 concludes.

2 Empirical Analysis

This section documents the two main empirical findings on the effect of unions on the wages and employment of workers exposed to labor-replacing technologies. I start this section by describing the data sources and outlining the empirical approach.

Data. I exploit variation in unionization across local labor markets in the US. In the main analysis, a local labor markets is defined as a metropolitan statistical area (MSA).[7] I use public use micro data from the 1980, 1990, 2010, and 2019 American Community Survey (ACS) to construct population, employment and wage income measures at the MSA level as well as at the industry and occupation level within MSAs at those four dates.[8] I further use the US Current Population Survey (CPS) to compute unionization by MSA. I focus on workers in routine manual (RM) occupations and follow the recent literature in the classification of occupations based on their routine- and manual-task content.[9] These occupations focus on tasks that follow a well-defined set of instructions and, as a result, can more readily be performed by automation technologies. Routine employment, classified as such, has fallen significantly since 1980, and progress in labor-replacing automation technology has been identified as a main driver (Autor et al. (2003), Goos et al. (2014)). Throughout the empirical analysis, worker outcomes of interest as well as unionization are measured within routine manual occupations and refer to routine manual workers unless otherwise specified. Unionization is computed as the share of routine manual workers who are either a member of or covered by a labor union, averaged between 1995 and 2005. I use the average unionization rate over 10 years starting in 1995 to have better coverage across MSAs. I then validate that the results hold when using different unionization measures, and in particular earlier measures going back to 1986. Finally, I use exposure to robots estimates from Acemoglu and Restrepo

[7] I provide additional evidence for robustness across states.

[8] See Ruggles et al. (2010).

[9] The literature goes back to Berman et al. (1994), Levy and Murnane (1996), Autor et al. (1998). See Katz and Autor (1999) for a summary of the early literature and Cortes et al. (2020) for a classification.

(2020).

The empirical strategy is to study if routine manual workforces across MSAs, who differ in their rate of unionization but are similar otherwise, experience differential changes in the overall level and composition of their employment and wages from 1980 onwards when automation technologies increasingly became available in all MSAs. Thus, I estimate to what extent differential decline in employment and wages among routine manual workers across MSAs can be explained by variation in unionization among routine manual workers, controling for the ex-ante exposure to labor-replacing technologies within MSAs. First, in order to account for the ex-ante exposure to automation, that is, the expected amount of automation which is unrelated but potentially correlated with unionization, I construct a rich set of controls at 1980, prior to the transition. In particular, I control for the industry composition at the MSA level, the industry composition within routine manual occupations in each MSA, and the demographic composition of routine manual workers in the MSA. Lastly, I add the exposure to automation measure by Acemoglu and Restrepo (2020), which is a commuting zone level measure that combines the industry composition of a commuting zone with industry-level adoption of industrial robots between 1993 and 2014 at the national level. I aggregate the measure to the MSA level using 1980 population weights. The identification assumption is that the remaining variation in unionization among routine manual workers is exogenous to ex-ante exposure to technology adoption and changes in the age composition of workers conditional on adoption.

I use different measures of changes in employment and wages in routine manual occupations since the start of the transition as outcome variables. In particular, for the change in variable y after t years of the transition, I estimate the following model across MSAs i

$$\Delta y_{i,1980+t} = \beta_0 + \beta_1 \text{Unionization}_i + \gamma X_{i,1980} + u_{i,t}, \tag{1}$$

where $\Delta y_{\text{i},1980+t}$ is the realized change in y between in 1980 and $1980 + t$ (e.g. the decline in routine manual employment). The set of controls is constructed in 1980, prior to the transition, except for the exposure measure from Acemoglu and Restrepo (2020), which is based on adoption data between 1993 and 2014. I then run this model for different outcome variables y and at different stages of the transition, $t \in [10, 30, 40]$, to understand the effect of unionization on the level as well as timing of changes.

2.1 The Aggregate Effect of Unionization

In order to understand the aggregate effect of unionization, that is, the effect on the MSA-level routine manual employment share, I look at the timing and extent of overall employment decline. In particular, I regress the decline in the routine manual employment share in a MSAs since 1980 on its routine manual unionization rate and the set of controls. I do so for the decline until 1990, 2010 and 2019.

Table 1: The dependent variable is the change in the routine manual employment share since 1980, the independent variable is routine manual unionization. The regression further includes the set of controls outlined above.

	Change in RM share		
	1990	2010	2019
	(1)	(2)	(3)
Unionization	-0.080^{***}	-0.040^{**}	-0.036^{*}
	(0.018)	(0.018)	(0.020)
Mean dependent	-0.062	-0.11	-0.11
Observations	147	147	147
R^2	0.712	0.629	0.554
Adjusted R^2	0.684	0.592	0.510
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table 1 displays a negative effect of unionization on the change in employment, meaning employment falls more in high-unionized labor markets. Importantly, the effect is large between 1980 and 1990 and then falls off thereafter. In order to understand the size of the effect, I plot the change in the routine manual employment share when going from the 10th to the 90th percentile of unionization across MSAs. Thus, the graph below plots the estimated coefficient of the union effect, scaled by the difference in unionization between an MSA at the 90th and an MSA at the 10th percentile of unionization, which is a 29 percentage point difference in the unionization rate.

Figure 1: The graph shows the effect of going from the 10th to the 90th percentile of unionization on the RM employment share over time. The results hold for the 25th and the 75th percentile, see Appendix A.3.1 for details.

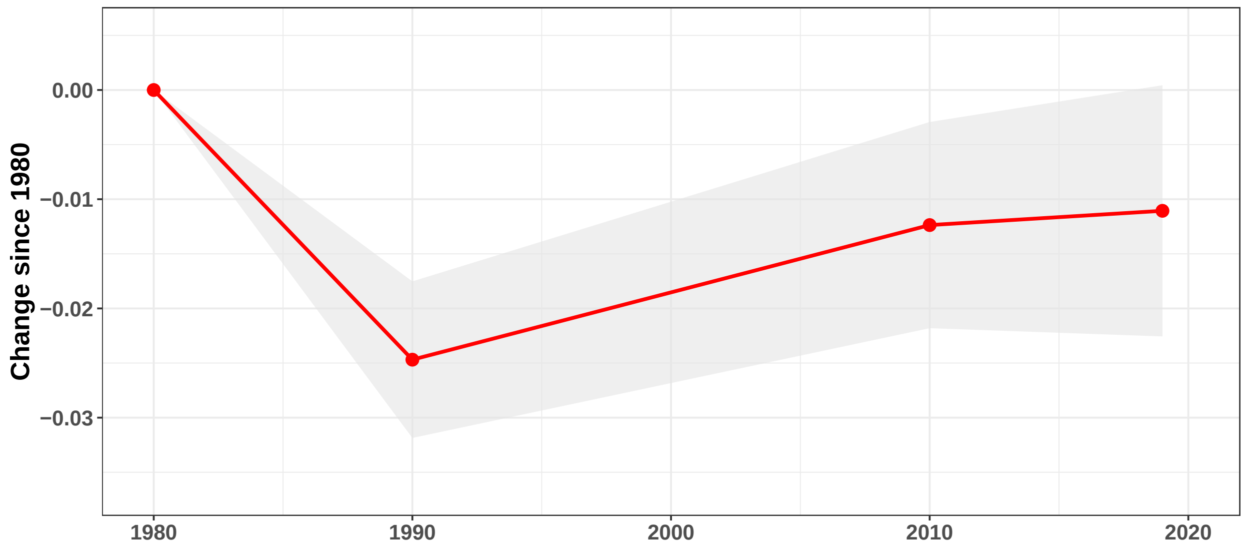

Figure 1 shows that unionization is associated with an accelerated decline in employment from 1980 onwards across MSAs. Going from the 10th to the 90th percentile in the rate of unionization leads to an additional decline of 2.5 percentage points in the employment share of exposed occupations between 1980 and 1990. Throughout the transition, the decline in the employment share then catches up in less unionized labor markets. In 2019, the union effect has fallen to roughly 1 percentage and is insignificant. Figure 2 relates the union effect to the average decline in the routine manual employment share across all MSAs, dividing the above union effect by the average decline across all MSAs in the same time period.

Figure 2: The graphs show the effect of going from the 10th to the 90th percentile of unionization on the RM employment share over time, relative to the mean decline in the RM employment share across all MSAs.

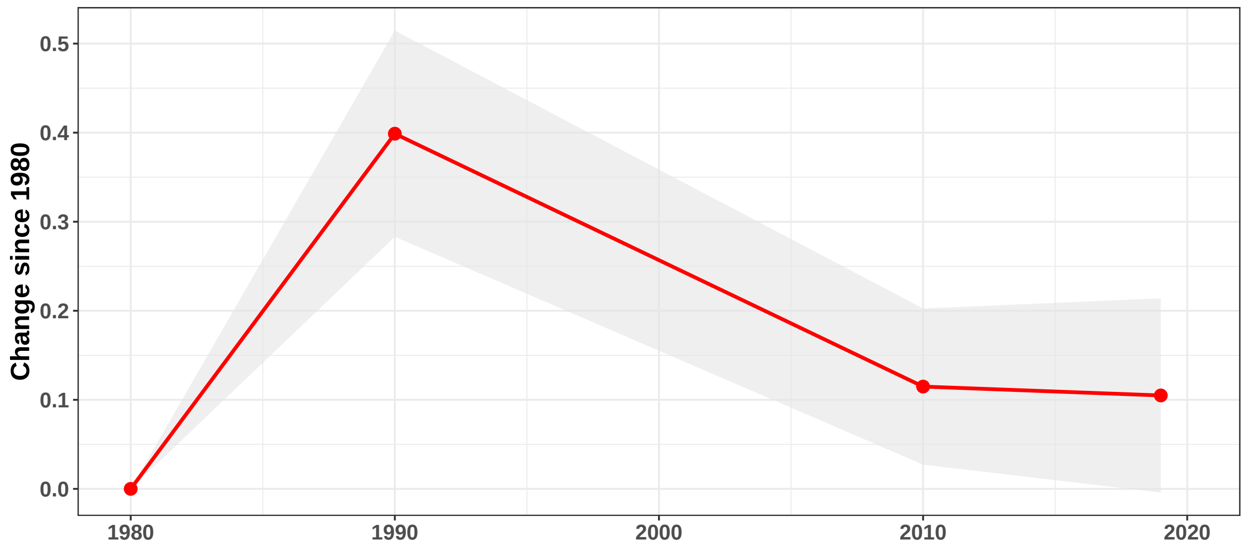

The union effect is large when relating it to the average decline across MSAs. In particular, comparing the 10th to the 90th percentile of unionization implies an additional fall in the employment share of 40% of the average decline across MSAs during the first 10 years. The union effect then falls to roughly 10% until 2019 and becomes insignificant, that is, employment decline in less unionized labor markets catches up over time. Thus, unionization among exposed workers is associated with a substantial acceleration of employment decline early in the transition. By 2019, the gap in employment decline by unionization is insignificant and has mostly closed with less unionized labor markets exhibiting a modestly smaller fall in their routine manual employment share.

2.2 The Distributional Effect of Unionization

An extensive literature has documented the insider-outsider dynamics of employment regulation and organized labor, as model by Carruth and Oswald (1987) and the literature thereafter. In particular, employment protection is associated with greater job security for older, incumbent workers, and reduced the employment opportunities and wage prospects for younger workers (Bassanini and Duval (2006), Botero et al. (2004)). This section documents to what extent the insider-outsider dynamics were prevalent during the decline of routine manual employment since the 1980s. Insider-outsider dynamics predict that unionization induces a downward shift in the demand for young, incoming workers. In turn, the downward shift in demand then translates into a fall in the price and quantity of young workers in high relative to low-unionized places. Guided by this intuition, I test whether unionization among routine manual workers is associated with a larger decline in wages and employment among young relative to older routine manual workers since 1980.

Throughout the analysis, I control for the decline in the routine manual employment share in order to measure how the composition of routine manual workers changed conditional on a decline in employment.

2.2.1 Employment Effect for Young and Old Workers

To quantify the impact of unionization on the relative employment of young and old workers, I measure the union effect on the age composition of routine manual workers along the transition. To build intuition and derive precise predictions to test in the data, I start by developing a simple model that isolates the effect of the hiring and layoff margin, and thereby illustrates how the age composition of workers evolves over time if the fall in employment is to a larger extent driven by reduced inflow (hiring) of young workers rather than increased outflow (layoffs) across the age distribution.

There are two labor markets, A and B, which are initially in steady state with an identical and uniform age distribution of homogeneous workers aged 20 to 60. Each year the 60 year old retire and are replaced by an inflow of 20 year old. In labor market A firms face zero labor adjustment costs, representing low or no unionization. By contrast, labor market B is unionized and firms face infinite labor adjustment costs. In 1980, an unexpected shock hits both labor markets which forces firms to shrink their workforce, firms in labor market A respond with uniform layoffs across the age distribution while firms in labor market B respond by lowering their hiring rate as layoffs are infinitely costly. Figure 3 shows the cdf of the age distribution of workers in both labor markets along the transition from a simple simulation.

Figure 3: CDF of age distribution in high and low-unionized labor markets in simple model.

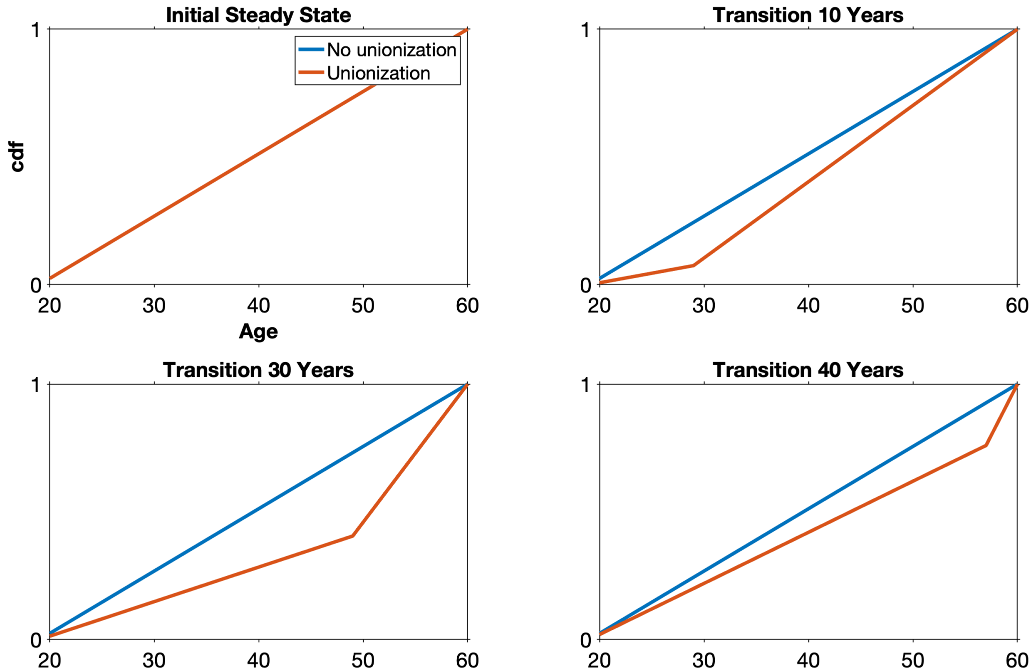

While the age composition in labor market A never changes, the reduction in hiring in labor market B leads to a fall in the share of young workers and a slow transition as the age composition adjusts. The workforce in labor market B ages relative to A as more old workers remain, which results in a downward shift of its CDF compared to labor market A. 10 years into the transition, the downward shift is largest at age 30, which is the first cohort that experienced reduced hiring. All cohorts between the age of 20 and 30 experienced reduced hiring while all cohorts above the age of 30 did not, and thus the share of workers below age 30 has fallen most relative to the steady state. The downward shift then evolves along the transition. In particular, the largest downward shift moves up the age ladder with the first cohort to experience reduced hiring as it ages over the course of the transition. 30 years into the transition the share of workers below the age of 50 has shifted down the most as all cohorts younger than age 50 have experienced reduced hiring. Thus, the simple model makes two detailed predictions about the effect of labor adjustment costs that drive relative changes in hiring and layoffs to test in the data. First, the routine manual workforce in more unionized MSAs becomes relatively older during the transition, measured as a relative downward shift in the CDF across all ages. Second, the downward shift is largest for the cohorts who entered around 1980 and moves up the age ladder with that cohort over time.

To account for the fact that routine manual workforces in all MSAs experience a mixture of reduced inflow and increased outflow in the data, I estimate the union effect on the downward shift in the CDF relative to 1980. Figure 4 shows the downward shift in the age distribution in each labor market relative to their initial steady state levels in the thought experiment.

Figure 4: Change in CDF of age distributions relative to steady state in simple model.

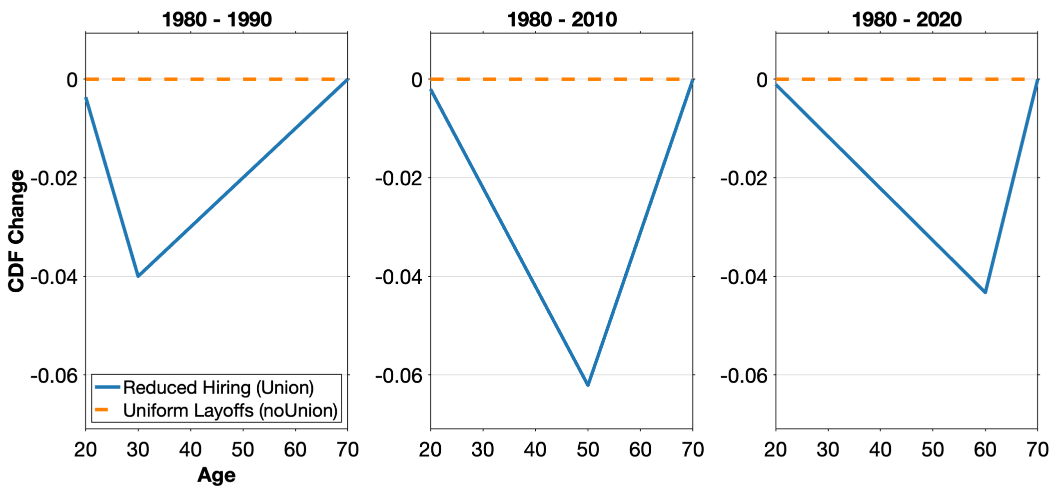

To test the predictions for the age composition of routine manual workers, I estimate the gap between the orange and blue line with the following model:

$$\mathrm{CDF}(\mathrm{a})_{\mathrm{i},t} - \mathrm{CDF}(\mathrm{a})_{\mathrm{i},1980} = \Delta\mathrm{CDF}(\mathrm{a})_{i,t} = \beta_0 + \beta_1^{\mathrm{a},t} \cdot \mathrm{Unionization}_{\mathrm{i}} + \gamma X_{\mathrm{i},t} + u_{\mathrm{i},t}. \quad (2)$$

The coefficient $\beta_1^{\mathrm{a},t}$ estimates the gap between the orange and blue line t years into the transition at age a. To account for differences in the initial age distributions, I control for the 1980 age composition among routine manual workers. Moreover, to isolate the insider-outsider dynamic from the aggregate effect, I further control for the decline in the routine manual employment share between 1980 and t ($\Delta \mathrm{RM}_{i,t}$).

Table 2: The table shows the effect of unionization on the change in the age distribution of routine manual workers between 1980 and different stages of the transition (1990, 2010, 2019). See A.3.2 for the full regression tables.

	Dependent variable: Change in CDF across Ages				
	Age 20	Age 30	Age 40	Age 50	Age 60
	(1)	(2)	(3)	(4)	(5)
CDF Change 1980-1990	-0.043^{***}	-0.126^{***}	-0.114^{***}	-0.062^{***}	-0.023^{**}
	(0.012)	(0.027)	(0.026)	(0.020)	(0.011)
CDF Change 1980-2010	-0.044^{***}	-0.106^{***}	-0.119^{***}	-0.142^{***}	-0.046^{***}
	(0.014)	(0.035)	(0.028)	(0.030)	(0.017)
CDF Change 1980-2019	-0.037^{**}	-0.026	-0.067^{**}	-0.087^{***}	-0.084^{***}
	(0.017)	(0.029)	(0.033)	(0.031)	(0.024)
Note:	*p<0.1; **p<0.05; ***p<0.01.				

Table 2 displays the coefficient $\beta_1^{a,t}$ measuring the union effect across the age distribution (columns) at different times in the transition (rows). First, the union effect is significant and negative throughout. Consistent with the first prediction, unionization is associated with a downward shift in the CDF of workers relative to 1980 across all ages. That means, more unionized workforces have become older relative to less unionized workforces along the transition as the share of young workers has fallen and more older workers have remained in unionized workforces. Second, between 1980 and 1990 the downward shift is largest at young ages and peaks at age 30. Over time, the downward shift moves up the age ladder with the cohorts who entered around 1980, consistent with the second prediction. Figure 5 constructs the graphs from the simple model from the regression estimates to directly compare the results with the prediction and to understand the magnitude of the effects. In particular, it again plots the union effect when comparing the average MSA at the 10th with the average MSA at the 90th percentile of routine manual unionization.

Figure 5: The plot shows the shift in the age distribution (CDF) relative to 1980 when going from the average MSA at the 10th percentile of routine manual unionization to the average MSA at the 90th percentile of routine manual unionization. The difference in unionization is 29 percentage points. The results hold for the 25th and the 75th percentile, see Appendix A.3.3 for details.

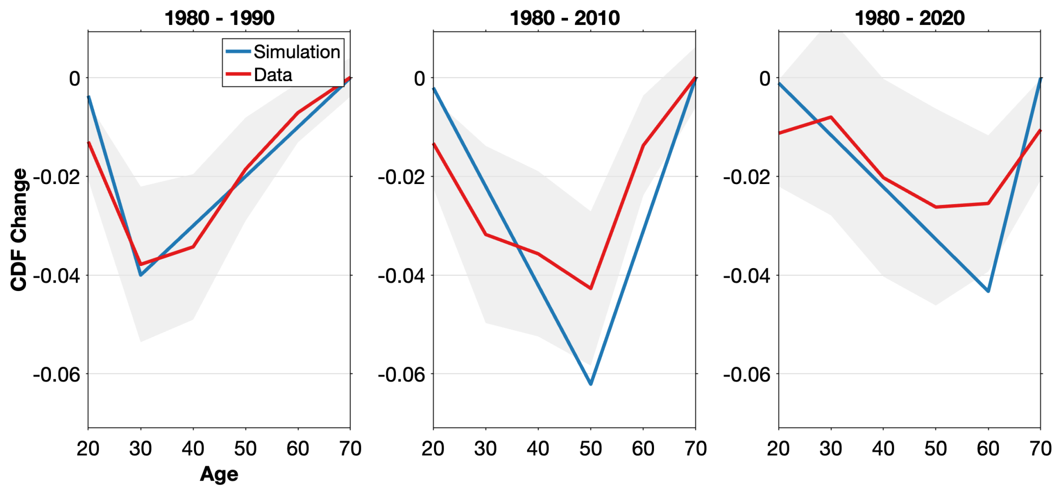

The plot closely resembles the prediction from the model. To understand the magnitude, note that the share of workers below the age of 30 falls by roughly 4 percentage points in the high-unionized MSA relative to the low-unionized MSA during the first 10 years of the transition. This translates into an additional 11% decline relative to the 1980 share of workers below the age of 30.

To summarize, conditional on reducing employment, unionzation is associated with a larger reduction in employment of young workers and a smaller reduction in employment of old workers, resulting in more unionized workforces becoming older relative to less unionized workforces throughout the transition. This is consistent with labor adjustment costs on incumbent workers imposed by unionization, which incentivize firms to adjust relatively more through young and incoming workers from 1980 onwards. The initial decline in the employment share of young workers in more unionized workforces has then translated into persistent changes in the age composition over time as the middle and right panel show.

2.2.2 Wage Effect for Young and Old Workers

A downward shift in the demand for young workers driven by unionization should further lead to a fall in the price of young workers, that is, their wage. To quantify the differential effect of unions on wages of young, incoming and older, incumbent workers, I look at the changes in the wage ratio between young to old routine manual workers.[10] Table 3 displays the results of regressing the change in the wage ratio between 1980 and 1990 on unionzation as well as the set of controls. The wage ratio is measured as the average wage of workers below the age of 30 divided by the average wage of workers over the age of 30 in the first two columns, and divided by the average wage of workers over the age of 50 in the last two columns. Columns 2 and 4 additionally control for the overall decline of the routine manual employment share between 1980 and 1990.

Table 3: Table shows the effect of unionization on the change in the wage ratio between young and older routine manual workers between 1980 and 1990. The first two columns define the wage ratio as the average wage of workers below the age of 30 divided by the average wage of workers over the age of 30. In the last two columns, the wage ratio is measured as the average wage of workers below the age of 30 divided by the average wage of workers over the age of 50.

	Change in Wage Ratio 1980-1990			
	$\Delta\frac{\text{Wage age}\leq 30}{\text{Wage age}>30}$		$\Delta\frac{\text{Wage age}\leq 30}{\text{Wage age}>50}$	
	(1)	(2)	(3)	(4)
Unionization	-0.184^{***}	-0.175^{**}	-0.307^{***}	-0.289^{***}
	(0.069)	(0.069)	(0.096)	(0.096)
Change RM-share 1980s		0.250		0.525
		(0.290)		(0.374)
Mean dependent	0.032		0.024	
Observations	200	200	200	200
R^2	0.282	0.285	0.281	0.289
Adjusted R^2	0.244	0.243	0.243	0.247
Note:	*p<0.1; **p<0.05; ***p<0.01.			

Unionization is negatively correlated with a change in the wage ratio, that is, in more unionized workforces the wages of young workers have declined more relative to wages of older workers compared to less unionized workforces. Looking at the first two columns, a 1 percentage point increase in unionzation is associated with a 0.18 percentage point decline in the wage ratio between workers below and above age 30. The effect rises to a roughly 0.3 percentage point decline for the wage ratio between workers below age 30 and above age 50, shown in the last two columns. To put this into perspective, going from the 10th to the 90th percentile of unionization is then associated

[10]Note, since I display a ratio, there is no need to deflate the wages by prices.

with a 6 and 9 percentage point decline in the wage ratio between 1980 and 1990 in the first and last two columns, respectively. This effect is quantitatively large as the average unconditional change in the wage ratio across local labor markets is even slightly positive with 0.032 and 0.024 percentage point increases, respectively. Thus, while the wage ratio between young and older routine manual workers has been stable on average across MSAs, wages of young workers have declined significantly relative to wages of older workers in unionized local labor markets. The effect grows larger when conditioning the group of older workers to higher ages, shown by the last two columns relative to the first two columns, which is consistent with union protection rising in tenure and age.

To the extent that the effects measured above are driven by a union induced fall in the demand for less protected young workers during the transition, the relative decline in wages of young workers should be a temporary effect during the transition rather than a persistent change in the wage structure. Table 4 looks at the union effect on changes in the wage ratio between 1980 and 2010. Over a longer horizon of the transition, between 1980 and 2010, the effect vanishes and unionization is not associated with an additional decline in the wage ratio between young and old workers.

Table 4: Table shows the effect of unionization on the change in the wage ratio between young and older routine manual workers between 1980 and 2010. The first two columns define the wage ratio as the average wage of workers below the age of 30 divided by the average wage of workers over the age of 30. In the last two columns, the wage ratio is measured as the average wage of workers below the age of 30 divided by the average wage of workers over the age of 50.

	Change in Wage Ratio 1980-2010			
	$\Delta\frac{\text{Wage age}\leq 30}{\text{Wage age}>30}$		$\Delta\frac{\text{Wage age}\leq 30}{\text{Wage age}>50}$	
	(1)	(2)	(3)	(4)
Unionization	−0.037	−0.029	0.030	0.035
	(0.096)	(0.094)	(0.117)	(0.116)
Change RM-share 1980-2010		0.837*		0.437
		(0.434)		(0.550)
Mean dependent	0.023		-0.062	
Observations	200	200	200	200
R^2	0.125	0.147	0.167	0.170
Adjusted R^2	0.069	0.088	0.114	0.112
Note:	*p<0.1; **p<0.05; ***p<0.01.			

To summarize, unionization is associated with a larger fall in employment and wages of young workers during the first decades of the transition. These effects are temporary and vanish by 2019, consistent with a fall in demand for young workers during the initial adjustment phase of the transition. Importantly, while the employment share of young workers in more unionized routine manual labor markets recovers by 2019, the initial fall in their employment moves up the age ladder over time and thereby has a persistent effect on the age composition of routine manual workers.

3 The Model

Motivated by the empirical findings, I develop a quantitative dynamic equilibrium model that interprets the documented distributional and aggregate effect of unions jointly through the lens of union-imposed labor adjustment costs interacting with gradual and endogenous technology adoption by firms over time. After validating that the model can replicate the different transitions observed in high and low unionized labor markets, I use it as a measurement device to quantify the welfare cost of automation for routine workers and the intergenerational transfer that unions give rise to during technological transitions. I first outline the model and provide a more detailed discussion of the model choices and properties in section 3.7.

3.1 Overview

Time is discrete and one period corresponds to 10 years. The model is a small open economy without aggregate uncertainty, combining three core elements. First, firms produce the final good by combining output from non-routine and routine occupations while endogenously and gradually adopting automation in routine production as capital prices fall. Second, overlapping generations of workers make an occupational choice between routine and non-routine occupations based on their expected life-cycle wage paths in each occupation. Third, a monopoly union represents incumbent routine workers by posting the wage schedule for routine workers of different skills, and thus ages, each period, taking labor demand of firms into account. The level of labor adjustment costs parameterizes the rate of unionization as the union derives its ability to impose wage premia from labor adjustment costs which limit how much firms can reduce employment in response.

3.2 Job levels

In the routine occupation, firms, the union, and workers interact through job levels. In practice, job levels describe the specific task requirements of each job. An extensive literature on the internal labor markets (ILMs) and the production process of firms has documented the importance of job levels in the design of production processes.[11] In particular, job levels are a key input in production, and progression across job level accounts for the majority of life-cycle wage growth of workers (Bayer and Kuhn (2023), Pierce (1999)). Moreover, unions directly bargain for wages at different job levels (Bayer and Kuhn (2023)). Based on these findings, routine production in this economy is organized around job levels, and firms decide how many workers to employ at each job level. Young workers entering the routine occupation are hired at the lowest job level and progress in job levels by accumulating human capital on the job. Lastly, the union sets the job level wage profile in the routine occupations.

3.3 Production

Technology. There is a continuum of perfectly competitive firms that produce the final consumption good by combining the output y_t^i from two occupations i, the non-routine and the routine occupations, with a CES production technology G according to:

$$y_t = G(y_t^R, y_t^N) = \left[\phi(y_t^R)^\nu + (1-\phi)(y_t^N)^\nu\right]^{\frac{\theta}{\nu}}, \tag{3}$$

where ϕ is the share of automatable (routine) occupations and $(1-\phi)$ is the share of non-automatable (non-routine) occupations in the economy. $\nu < 1$ is the elasticity of substitution between routine and non-routine occupations. To accommodate decreasing returns to scale from convex adjustment costs while abstracting from firm heterogeneity, I assume that firms need to use land as another input in production which is in fixed and limited supply L, as in Huo and Ríos-Rull (2020). Without loss of generality, I assume there is a total of one unit of land $L = 1$ and there is a firm operating each unit of land. $\theta < 1$ measures the returns to scale and the land is then priced by the value of the representative firm.

The non-routine occupations use homogenous labor input N_t to operate a constant returns to scale technology given by:

$$y_t^N = N_t. \tag{4}$$

[11] See, for instance, Doeringer and Piore (1985), Baker et al. (1994), Pierce (1999), Strub et al. (2008) Bayer and Kuhn (2023). Baker and Holmstrom (1995) provide an overview of the early literature.

Routine Production. The routine occupations use automation α_t, and workers at $J = 5$ different job levels $l_t = (l_{t,1}, ..., l_{t,J})$ to produce routine output with a CES production technology F given by:

$$y_t^R = F_t(l_{t,1}, ..., l_{t,J}, \alpha_t) = \left[\sum_{j=1}^{J} \eta_j l_{(t,j)}^{\varphi} + \eta_\alpha \alpha_t^{\varphi} \right]^{\frac{1}{\varphi}}, \tag{5}$$

where $\eta = (\eta_1, ..., \eta_J, \eta_\alpha)$ governs the share of automation and job level input, and $\varphi < 1$ is the elasticity of substitution between routine inputs.[12]

Workers accumulate the necessary skills to produce the tasks required at higher job levels on the job. In particular, I assume the technology is such that it takes workers one period (10 years) to learn the skills to work the next job level. That is, a worker on job level j in period t can be promoted to job level $j+1$ in period $t+1$. I restrict attention to employment contracts between firms and routine entrants that commit the firm to compensate the rising human capital path of workers or otherwise terminate the contract at firing cost c_f. Thus, when hiring a routine worker, firms commit to progressing workers one job level per period consistent with their accumulated skill or firing them otherwise. As a result, age becomes a sufficient statistic for job level which makes the firm and worker problem tractable.

Optimization. Taking the path for the non-routine wage and the routine wages across job levels, $(w_t^N, \{w_{t,j}^R\}_j)_t$, as well as the price of automation $(p_t)_t$ as given, the firm chooses period t automation α_t and labor demands $(\{l_{t,j}\}_j, N_t)$ to maximize the discounted sum of future profits:

$$\begin{aligned} W_t(\{l_{t-1,j}\}_{j=1}^J) = \max_{\alpha_t, \{l_{t,j}\}_j, N_t} \; & G(y_t^R, y_t^N) - \sum_{j=1}^{J} w_{t,j}^R l_{t,j} - p_t \alpha_t - w_t^N N_t - \sum_{j=1}^{J} c_f(f_{t,j}) \\ & + \frac{1}{1+r_t} W_{t+1}(\{l_{t,j}\}_j), \end{aligned} \tag{6}$$

$$\begin{aligned} \text{s.t.} \qquad f_{t,j} &= l_{(t-1,j-1)} - l_{t,j} \quad \forall j \geq 2, \\ c_f(f_{(t,j)}) &= c \cdot f_{(t,j)}^2, \end{aligned}$$

[12] Bayer and Kuhn (2023) document five possible job levels in a German Employment Survey (BIBB/BAuA) and 15 job levels for the United States using the National Compensation Survey (NCS), consistent with Pierce (1999).

where $f_{t,j} = l_{t-1,j-1} - l_{t,j}$ denotes fired workers at job level j. Firing costs are specified the same way across job levels and parameterized by c. Thus, labor adjustment costs in this model are governed by c. I assume firing is a lottery, thus, firms decide how many workers to fire at each job level but which workers at job level j are fired is random.

3.4 Households

Agents and Preferences. The economy is populated by overlapping generations of households. Each period a measure one of young households is born who live for 5 periods, from age 20 to 70. Thus, in every period there is a total of 5 generations alive. Young workers choose which occupation to work in and spend resources on consumption and saving while supplying labor inelastically and accumulating human capital on the job. Workers born in period t maximize expected lifetime utility U_t given by:

$$U_t = \sum_{a=1}^{5} \beta^{a-1} E\left[u(c_{t+a-1,a})\right], \tag{7}$$

where the period utility function $u(c)$ is at least twice continuously differentiable with $u'(c) > 0$ and $u''(c) < 0$, and satisfies the lower Inada condition, thus $\lim_{c\to 0} u'(c) = \infty$.

Human capital accumulation process. Workers are born and enter the labor market with initial routine labor productivity z^R and non-routine labor productivity z_1^N. Initial non-routine labor productivity z_1^N is ex-ante identical across workers while routine labor productivity z^R differs across workers and is drawn from distribution f_z. Human capital is occupation-specific and deterministically accumulates on the job in both occupations.

Households working in the non-routine occupation accumulate human capital each period in the form of labor productivity. They move up a discrete labor productivity ladder $(z_1^N, z_2^N, z_3^N, z_4^N, z_5^N)$, which is calibrated to match average life-cycle wage paths of non-routine workers in the data.

Workers in the routine occupation accumulate human capital on the job through job level progression which drives their life-cycle wage growth. In particular, routine workers who are not laid off move up one job level per period as specified in their employment contract. A worker's routine labor productivity z^R is a permanent type that applies to all job levels. As a result, human capital in the routine occupation follows a step function. Job level progression captures steps over the life-cycle which are common across workers and give rise to wage dispersion across age. z_R captures the overall level of the step function which differs across workers and gives rise to wage dispersion within age groups. Routine workers have perfect information about the endogenous probability of being laid off in the routine occupations.

Occupational Choice. At labor market entry, workers only differ in their permanent routine labor productivity type z^R which determines their initial occupational choice. They take into account the expected life-cycle path of earnings in each occupation. There is no aggregate uncertainty, thus, workers have full information about the future path of wages in each occupation $(\{w^R_{t,j}\}_j, w^N_t)$. They face individual uncertainty in the form of firing risk when working in the routine occupations. The probability of being fired at each job level is endogenously chosen by firms but is fully known by workers. In each consecutive period, workers choose whether to switch or stay in their current occupation. Routine workers who are laid off switch to the non-routine occupations and stay there for the remainder of their life. I assume they cannot reenter the routine occupations after being laid off.

Assets. Financial markets are incomplete, in particular routine workers cannot trade contingent assets against the risk of being laid off. Households have access to risk-free bonds at world interest rate R which is exogenous and constant. In the baseline model, the land, and thus firms, is owned by risk-neutral capitalists who receive the firm dividends. In practice, equity participation is limited, especially for low and medium skilled workers who usually are less educated (Mankiw and Zeldes (1991)), and 90% of firm equity is held by the 10% wealthiest households in the U.S. (Survey of Consumer Finances, (2022)). Since the model is designed to capture the impact of automation on the subset of less skilled workers who are most exposed to displacement by automation technologies, I take as baseline an economy in which these workers do not hold equity and are therefore impacted through wages but not through profits. In appendix B.1, I show results for the case when workers hold fixed equity shares in the firms.

The Worker Problem. At the beginning of period t, the state of a worker is her age a, wealth b, labor productivities (z^R, z^N), and previous occupation, s. There is no incentive for a worker to initially enter the non-routine occupation and then switch to the routine occupation later. Thus, the problem of a worker previously employed in the non-routine occupation, $s = 1$, can be simplified and solved by imposing the occupational choice to stay in the non-routine occupation. I verify in equilibrium that non-routine workers do not want to switch. The problem then consists only of the consumption-savings decision for the remainder of her life given by:

$$V_t^{hh}(k, z^R, z_i^N, a, s=1) = \max_{c,k'} u(c) + \beta V_{t+1}^{hh}(k', z^R, z_{i+1}^N, a+1, s'=1), \tag{8}$$
$$\text{s.t.} \quad c + k' = w_t^N z_i^N + (1+r_t)k,$$
$$k' \geq 0.$$

The problem of a worker previously employed in the routine occupation, $s = 0$, is more complicated as it involves the discrete choice about whether to stay a routine worker or switch into the non-routine occupation. It can be formulated as a two-stage problem. In the first stage, the household decides on the occupation s', in the second stage the household makes a consumption-savings decision conditional on the realization of the occupational choice. Note that since routine workers face an endogenous, possibly positive probability of being fired, a worker may decide to stay in the routine occupation but is fired and nevertheless forced to switch occupations. Conditional on working in occupation s' after stage 1, the stage 2 problem of a worker previously employed in the routine occupation is then given by:

$$v_t^{hh}(k, z^R, z^N, a, s=0|s') = \max_{c,k'} u(c) + \beta V_{t+1}^{hh}(k', z^R, z^N, a+1, s'), \tag{9}$$
$$\text{s.t.} \quad c + k' = s' w_t^N z^N + (1-s') w_{t,j(a)}^R z^R + (1+r_t)k,$$

where $w_{t,j(a)}^R$ is the routine wage at job level $j(a) = \frac{a}{10} - 1$ which maps the age of workers $a = (20, 30, 40, 50, 60)$ into their job level $j = (1, 2, 3, 4, 5)$.

For routine workers who either decide to switch into the non-routine occupation or who are fired, the above stage 2 problem is the same as the beginning of period problem of workers who were previously employed in the non-routine occupation:

$$v_t^{hh}(k, z^R, z^N, a, s=0|s'=1) = V_t^{hh}(k, z^R, z^N, a, s=1). \tag{10}$$

Given the value in stage 2, one can solve for the occupational choice in stage 1 of a routine worker. I assume workers face choice specific taste shocks to smooth the discrete occupation choice, $\sigma_s \epsilon_t(s)$. The taste shocks are additively separable and follow an extreme value distribution, as in McFadden (1973) and the literature thereafter. Due to firing, the realized occupation s' can differ from the chosen occupation $\tilde{s}'$, which solves the stage 2 problem given by:

$$V_t^{hh}(k, z^R, z^N, a, s=0) = \max_{\tilde{s}'} \Bigg\{ \tilde{s}' \Big(V_t^{hh}(k', z^{R\prime}, z^{N\prime}, a+1, s'=1) + \sigma_s \epsilon_t(\tilde{s}'=1) \Big) \tag{11}$$
$$+(1-\tilde{s}') \Big(\mu_{t,j(a)} V_t^{hh}(k', z^{R\prime}, z^{N\prime}, a+1, s'=1) + \sigma_s \epsilon_t(\tilde{s}'=1)$$
$$+ (1-\mu_{t,j(a)}) V_t^{hh}(k', z^{R\prime}, z^{N\prime}, a+1, s'=0) + \sigma_s \epsilon_t(\tilde{s}'=0) \Big) \Bigg\},$$

where $\mu_{t,j(a)}$ denotes the probability of being fired from the routine occupation as a worker at job level $j(a)$. The probability of deciding to stay a routine worker is given by the discrete choice policy function, $\mathcal{P}_t(\tilde{s}'|k, z^R, z^N, a)$, which is equal to the standard logit choice probability with an extreme value distributed taste shock:

$$\mathcal{P}_t(\tilde{s}'|k, z^R, z^N, a) = \frac{\exp\left(V_t^{hh}(k', z^{R\prime}, z^{N\prime}, a+1, \tilde{s}')/\sigma_s\right)}{\exp\left(V_t^{hh}(k', z^{R\prime}, z^{N\prime}, a+1, =0)/\sigma_s\right) + \exp\left(V_t^{hh}(k', z^{R\prime}, z^{N\prime}, a+1, 1)/\sigma_s\right)}. \tag{12}$$

Based on the discrete choice policy function, the probabilities for the realization of the occupational choice, $P_t(s'=1|k, z^R, z^N, a)$, that accounts for endogenous firing is then given by:

$$P_t(s'=1|k, z^R, z^N, a) = \mathcal{P}_t(\tilde{s}'=1|k, z^R, z^N, a) + \mu_{t,j(a)} \mathcal{P}_t(\tilde{s}'=0|k, z^R, z^N, a), \tag{13}$$
$$P_t(s'=0|k, z^R, z^N, a) = (1-\mu_{t,j(a)}) \mathcal{P}_t(\tilde{s}'=0|k, z^R, z^N, a). \tag{14}$$

3.5 The Union

All incumbent workers in the routine occupations are represented by a labor union. The union acts within a standard monopoly union framework by setting wages as a monopolist while firms choose labor demand in response.[13] However, I extend the basic monopoly union model by allowing the union to set the full job level wage profil to account for the fact that workers at different job levels are imperfect substitutes in production and therefore can have different wages. The union then chooses wage growth across job levels by setting the full job level wage profile in the current period, and seeks

[13] The basic monopoly union framework goes back to Fellner (1949) and Cartter (1959).

to maximize the total wage bill paid to its current members:

$$\Theta_t = \max_{\{w^R_{t,j}\}_j} \sum_{j=2}^{J} w^R_{t,j} l_{t,j}. \tag{15}$$

Thus, the union is constrained by the employment response of firms $\{l_{t,j}\}_j$ which is endogenous to the wage schedule, $\{w^R_{t,j}\}_j$, posted by the union.

The rate of unionization in the model is measured by the level of firing costs c. The union's ability to extract rents from firms by raising wage premia is determined by the elasticity of labor demand of firms. If the elasticity is high, firms respond to wage premia by reducing their labor input which drives down the wage bill. Firing costs reduce the demand elasticity and thereby increase the ability of the union to increase the wage bill through wage premia, thus, providing a measure of the degree of unionization in the model. In practice, unions obtain higher firing costs for their members, consistent with the union model here.[14]

3.6 Competitive Equilibrium

I focus on perfect foresight equilibria in which there is no aggregate uncertainty.

Definition 1 (Competitive Equilibrium). Given a path for automation prices p_t and interest rates r_t, and an initial worker distribution Φ_0, a competitive equilibrium consists of paths for non-routine wages w^N_t, routine wages $(w^R_{t,1}, ..., w^R_{t,J})$, firm policies $(l_{t,1}, ..., f_{t,J}, \alpha_t, N_t)$, worker policies $V_t, c_t, k_{t+1}, \mathcal{P}_t$, and the worker distribution Φ_t that satisfy for all $t \geq 0$:

1. Given the paths for prices $\{r_t, w^N_t, (w^R_{t,1}, ..., w^R_{t,J})\}_{t\geq 0}$, and the firm implied firing probabilities $\{(\mu_{t,1}, ..., \mu_{t,J})\}_{t\geq 0}$, $V_{t\geq 0}$ solves the optimization problem of workers and $\{c_t, k_{t+1}, \mathcal{P}_t\}_{t\geq 0}$ are the corresponding decision rules.

2. Given the paths for prices $\{r_t, w^N_t, (w^R_{t,1}, ..., w^R_{t,J})\}_{t\geq 0}$, $W_{t\geq 0}$ solves the optimization problem of the firm and $\{l_{t,1}, ..., f_{t,J}, \alpha_t, N_t\}_{t\geq 0}$ are the corresponding policies.

3. Given the paths for prices $\{r_t, w^N_t, (w^R_{t,1}, ..., w^R_{t,J})\}_{t\geq 0}$, the period t routine wage schedule $\{w^R_{t,1}, ..., w^R_{t,J}\}_t$ solves the period t union problem.

[14] See, for example, Parsons (2005c,b,a), Millward et al. (1992) and Colonna (2008).

4. The non-routine wage clears the labor market for non-routine workers:

$$N_t = \int z^N \mathrm{d}\Phi_t(k, z^R, z^N, a, s=1) + \int z^N \mathrm{d}\Phi_t(k, z^R, z^N, a, s=0) P_t(s'=1|k, z^R, z^N, a, s=0). \quad (16)$$

5. The labor market for incoming routine workers clears:

$$l_{t,1} = \int z^R \mathrm{d}\Phi_t(k, z^R, z^N, a=1, s) P_t(s'=0|k, z^R, z^N, a=1, s). \quad (17)$$

6. The law of motion of the worker distribution is induced by the optimal decisions of the firm and workers.

3.7 Discussion

This section discusses the two key model elements, the job-level based routine production process and the union model.

Job Levels. Routine production in this model is organized around job levels. In practice, job levels categorize jobs by explicitly describing specific task requirements of jobs along the dimensions of responsibilities, complexity, and autonomy. This builds on an extensive literature that studies internal labor markets (ILMs) of firms and career dynamics. It emphasizes the role of the organizational structure of firms, and, in particular, the importance of job levels in the design of the production process.[15] One of the main insights going back to Doeringer and Piore (1985) and confirmed by the subsequent literature is that "in many jobs in the economy, wages are not attached to workers, but to jobs." (Doeringer and Piore (1985, p. 77)). Based on that idea, the literature documents two findings with respect to the determinants of wages and wage growth that make job levels a suitable modeling choice in this context. First, life-cycle wage growth is largely driven by job level progression over time (Baker et al. (1994), Dohmen et al. (2004), Bayer and Kuhn (2023)). Second, unions bargain for wages and benefits at the level of job levels (Bayer and Kuhn (2023)).

In the model, routine workers accumulate skill on the job and, as a result, move up in job levels. Thus, this yields a standard process of human capital accumulation as wages rise in response to skill accumulation. Consistent with the empirical evidence on how unions operate, the union sets wages at the level of job levels in the model, taking into

[15] See the seminal works of Doeringer and Piore (1985) and Baker et al. (1994), as well as the literature thereafter.

account labor demand. As a result, the job level model yields life-cycle wage growth that reflects a standard human capital accumulation process as well as an endogenous and time-varying union wage premium.

From the firm's perspective, the job level model allows for workers with different experience levels to be imperfect substitutes in production because they perform different tasks, which gives rise to an endogenous and time-varying wage ratio between younger and older routine workers that can be mapped to the data. This aligns with the task-based approach (Autor et al. (2003)), which builds on the idea that wages are determined by the tasks that a worker performs in their job. As a result, the firm optimally produces with a range of workers of different experience levels and therefore of different age. I restrict the analysis to employment contracts with full utilization of human capital, meaning firms commit to progressing workers to the next job level each period as their experience accumulates. This makes the model tractable as age becomes a sufficient statistic for job levels, and it allows me to abstract from the managerial decision about which workers to promote and when to promote. In practice, unions similarly bargain for wage growth in employment contracts, and, thus, for contractual commitments by the firm to compensate rising human capital over time or terminate the contract otherwise.

The Union. The model of the labor union here follows the literature initiated by Dunlop (1944) whose starting point is the microeconomic theory of firms. The labor union is modeled as an economic entity that maximizes an economic objective, such as the wage bill, while facing constraints, in particular the labor demand of the firm. The literature thereafter largely uses two frameworks. First, monopoly union models in which the union acts as a monopolist and imposes its wage policy while the firm chooses employment in response. Second, since the 1980s game-theoretic bargaining frameworks were developed. These frameworks were developed with the intend to properly model the sources of bargaining power, such as strikes, and the bargaining process. In both cases, the union generally imposes a wage premium, resulting in an inefficient outcome in which the wage is above and employment is below their market clearing levels. While in the bargaining framework variation in the bargaining power offers a direct model analog to the empirically observed variation in the rate of unionization, the difficulty of introducing bargaining here stems from the fact that it requires specifying the value of the disagreement outcome. Therefore, I abstract from the bargaining process and model the union within a monopoly union framework. Importantly, the solution to bargaining frameworks in which the firm and union bargain over wages coincides with the monopoly union model when the union has all the bargaining power. I then use the level of firing costs as a measure of unionization instead of an explicit bargaining power. Empirically, unions obtain higher firing costs for members (see, for instance, Parsons (2005c,b,a), Millward et al. (1992), Colonna (2008)), through bargaining for higher severance pay and the ability to impose strike costs. In the model, the union maximizes the wage bill of its members by extracting rents from firms. Since the union acts as a monopolist, its ability to extract rents is driven by the elasticity of labor demand of firms. Firing costs reduce the demand elasticity and, thus, increase the ability of the union to expand the wage bill through wage premia. Thus, firing costs drive the strength of the union and thereby provide an intuitive measure of the rate of unionization in the model.

4 Quantitative Evaluation

In this section, I outline the calibration strategy before evaluating the quantitative behavior of the model. I then connect the model back to the empirical findings by validating that it replicates the untargeted distributional and aggregate union effects along the transition.

4.1 Calibration

I calibrate the initial steady state of the model to MSA-level data in the U.S. in 1980. I then explore the response of the economy to an unexpected fall in the path of automation prices from 1980 to 2010 that matches the decline in capital prices observed in the U.S., as measured by Hubmer (2023). In particular, agents in the economy learn in 1980 about the complete future path of automation prices, thus, there is no aggregate uncertainty. The timing is motivated by the fact that existing measures of capital prices show a more rapid decline from the 1980s onward (Hubmer (2023)), and the adoption of industrial robots has picked up from 1990 onwards (Acemoglu and Restrepo (2020)). This is also consistent with the observed fall in the routine employment share and the manufacturing labor share from 1980 onwards (Hubmer (2023), Cortes et al. (2020)).

I take the low-unionized labor market as an economy that is characterized by low firing costs, and calibrate that economy to the average MSA at the 10th percentile of unionization. The high-unionized labor market corresponds to a MSA at the 90th percentile of unionization with high firing costs. I calibrate the common parameters in the low-unionized labor market, which is the baseline economy. A subset of parameters is calibrated exogenously, either following direct empirical observation or the existing literature. The remaining parameters are estimated in the model using the method of simulated moments. Since the objective is to use the model as a measurement device to quantify the impact of automation and unionization on the consumption paths of different workers, the calibration aims in particular at matching three sets of targets: First, the 1980 and 2010 routine manual employment share to capture the amount of workers that are exposed to displacement by automation technologies as well as the amount of employment loss they experience along the transition. Second, the 1980 and 2010 aggregate labor share to capture the fact that while automation raises productivity and thereby output, the share of output that accrues to labor has declined over time. Third, life-cycle wage profiles of workers in both occupations.

Data. I estimate the targeted data moments using the same data as in the empirical section. Thus, I construct MSA-level estimates by combining public use micro data from the American Community Survey (ACS), and the Current Population Survey (CPS).[16] I take the remaining targets from the existing literature and indicate when I do so.

Share paramaters in the production technology. The share parameters in the production function are calibrated to match the employment and labor shares. In particular, I calibrate the share parameter ϕ for routine output to match an initial routine manual employment share in 1980 of 27%. The share parameter of automation in the routine technology, η_α, is calibrated to match an initial aggregate labor share in 1980 of 64%, which I take from the U.S. Bureau of Labor Statistics (BLS). Lastly the share parameters of job level inputs in the routine technology, $(\eta_1, ..., \eta_5)$, are calibrated to match life-cycle wage profiles in 1980, the estimation of life-cycle wage profiles is outlined below.

Substitution elasticities in the production technology. The substitution elasticities are calibrated to match the change in the routine employment and aggregate labor share between 1980 and 2010. I calibrate the substitution elasticity across the two occupations, ν, to match the routine manual employment share in 2010 of 16%. The substitution elasticity across routine inputs, φ, is calibrated to match an aggregate labor share in 2010 of 56% (BLS).

Firing costs. The firing cost schedule is pinned down by one parameter c. Recall that the low-unionized economy is characterized by a low level of firing costs, c_l, which I calibrate to match the change in the average age of routine manual workers in MSAs at the 10th percentile of unionization between 1980 and 1990. The level of firing costs in the high-unionized economy, c_h, is then calibrated to match the documented union effect on the decline in the routine manual employment share between 1980 and 1990. That is, I

[16] See Ruggles et al. 2010

target that the routine employment share in the high-unionized economy declines by 2.5 percentage points more between 1980 and 1990 than in the low-unionized economy.

Human capital accumulation. Using repeated cross-sectional data from the CPS, I estimate life-cycle wage profiles by decomposing earnings growth into cohort, experience and time effects for cohorts born between 1940 and 1980, following Heckman et al. (1998), and more recently Lagakos et al. (2018) and Fang and Qiu (2021). The estimated experience effects capture the component of life-cycle wage growth that is driven by human capital accumulation. I then estimate experience effects separately for routine manual and non-routine occupations to calibrate life-cycle wage paths in both sectors in the model.[17] Workers in the non-routine occupation all enter with the same labor productivity z_1^N, and accumulate labor productivity every period on the job. I calibrate the life-cycle path of non-routine labor productivity, $z^N = (z_1^N, z_2^N, z_3^N, z_4^N, z_5^N)$, to match the estimated experience effect, and normalize mean labor productivity, $\bar{z}^N = 1$.

Preferences. The period utility function of households is given by Constant Relative Risk Aversion (CRRA) utility:

$$u(c_t) = \frac{c^{1-\sigma}}{1-\sigma}$$

where $\sigma = 2$ to get an intertemporal elasticity of substitution of 0.5. I set the discount factor to $\beta = 0.75$; recall that one model period corresponds to 10 years, and thus, the annualized discount factors of households is $\beta_{\text{annualized}} = 0.97$.

Remaining parameters. The world interest rate is set to 3% annually, based on estimates of the natural rate of interest for the U.S. from Davis et al. (2024). I abstract from the fact that real rates started to decline particularly from 2000 onward and keep the interest rate constant over the transition. Note that savings play a minor role in this model as households have rising life-cycle wage paths, do not face retirement, and face permanent income risk in the form of a small layoff probability in the routine occupation. The initial automation price $p_{1980} = 0.12$ is fixed in a first-stage calibration. The automation share η_α is then calibrated conditional on $p_{1980} = 0.12$ as the two are not separately identified.

[17] See appendix A.1 for details on the estimation.

Table 5: Externally calibrated parameters.

Parameter	Description	Value	Target
		Preferences	
$1/\sigma$	IES	0.5	Standard
β	Discount factor	0.75	$\beta_{\text{annualized}} = 0.97$
σ_s	Taste shocks	0.05	Small - smooth occ choice
		Human capital	
z^N	Non-routine labor productivity		Life-cycle wage profile
$\bar{z}^N$	Mean labor productivity	1	Normalization
		Small open economy	
r	Rate of return	0.34	3% annual Davis et al. (2024)
p_{1980}	Automation price 1980	0.12	Normalization
g_p	Growth rate of price	-0.06	Hubmer (2023)

Table 6: Internally calibrated parameters.

Parameter	Description	Value	Target
		Production and technology	
f_z	Routine labor productivity	$\mathcal{U}(0.2, 1.8)$	Routine Wage Dispersion
ϕ	Share of automatable occupations	0.75	1980 RM employment share
θ	Returns to scale	0.8	Land-output ratio
η_l	Job level shares		Life-cycle wage profile
η_α	Automation share	0.3	1980 labor share
ν	Substitution elasticity: sectors	0.75	2010 RM employment share
φ	Substitution elasticity: routine inputs	0.85	2010 labor share
		Union	
(c_l, c_h)	Firing costs	$(0.01, 0.05)$	Agg. union effect 1980-1990

4.2 Model Mechanism

Figure 6 shows the fall in capital prices which induces firms to gradually adopt automation over time and thereby triggers the transitional dynamics.

Figure 6: Price of automation along the transition matches the decline in measured capital prices in the U.S. between 1980 and 2010 from Hubmer (2023).

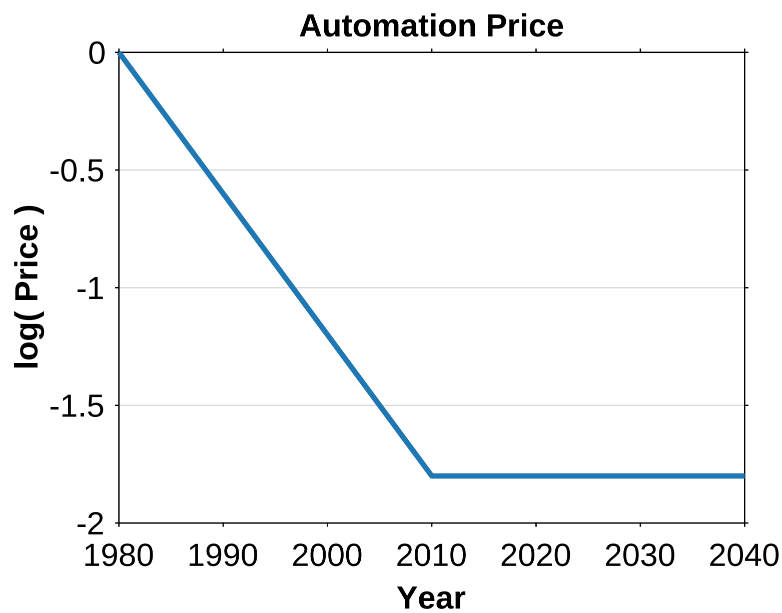

In 1980, the economy is in steady state. In 1990, agents in the economy wake up and learn about the new path of automation prices. Thus, they learn that the automation price has already fallen in 1990 and will further decline until 2010. The price decline matches the price decline of capital goods in the U.S. since 1980. The fall in automation prices triggers a fall in routine employment in the high as well as low-unionized economy as routine workers are imperfectly substitutable with automation. However, higher firing costs in the high-unionized economy increase the cost of layoffs and incentive firms to adjust to a larger extent through the hiring margin. As a result, demand for young workers entering the routine occupation falls relatively more in the high-unionized economy. Figure 7 shows hiring of young workers in the routine occupations relative to the 1980 steady state in both economies.

Figure 7: Hiring falls more in the high-unionized labor market as firing cost generate insider-outsider dynamics.

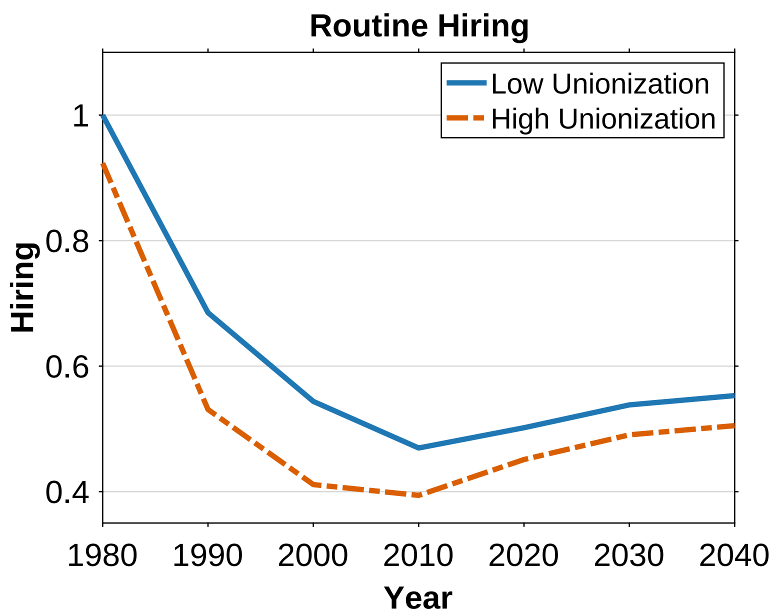

In the initial steady state, hiring of routine workers is 8% lower in the high-unionized than in the low-unionized economy since wage premia imposed by the union increase wages but reduce the level of routine employment across all age groups. After automation becomes available, routine hiring falls in both labor markets, however, it falls substantially more in the high-unionized labor market early in the transition. The greater fall in hiring in the high-unionized labor market in 1990 is driven by larger adjustment through incoming workers in response to current automation adoption as well as by a further preemptive reduction in hiring in anticipation of future adoption to avoid adjustment costs along the transition.

Figure 8 shows the union effect on wages and paints a similar picture. Early in the transition, unionization further depresses entry wages in the routine occupation by lowering demand for incoming workers. In particular, routine entry wages decline 3.5% more until 1990 and 5% more until 2000 in the high-unionized relative to the low-unionized labor market. The union effect spills over into the non-routine occupation, resulting an additional 0.3% decline in non-routine wages until 2000 in the high-unionized labor market. The spillover is driven by the accelerated routine employment decline in the high-unionized labor market, which in turn accelerates the reallocation of workers to the non-routine occupations and reduces wages there.

Figure 8: The routine entry and non-routine wage temporarily fall in the high-unionized relative to the low-unionized economy during the transition.

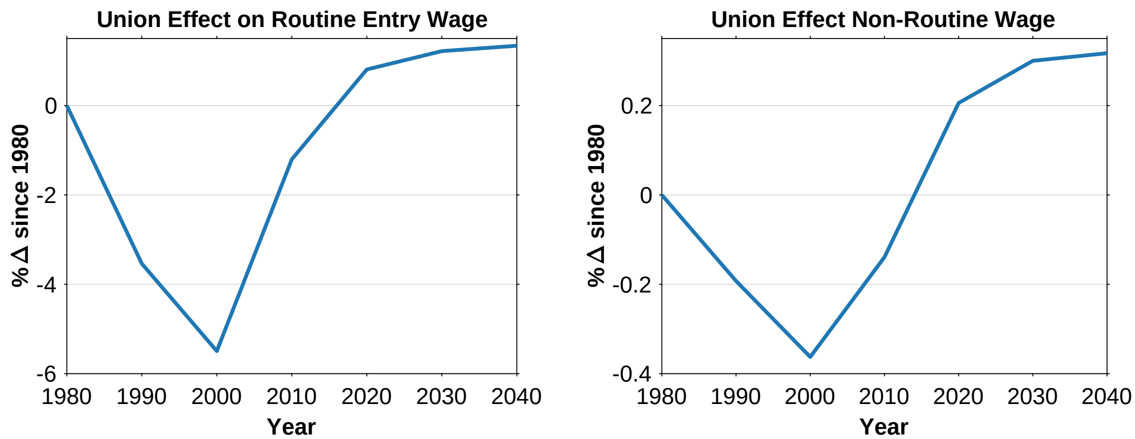

4.3 Model Validation

The evolution of aggregate routine employment and the evolution of the age distribution of routine workers along the transition are not targeted by the calibration. To validate the model and connect it with the empirical findings, I test whether it matches the evolution of routine employment and the age composition, and thus, whether it matches the aggregate and distributional union effect documented in the data.

Figure 9: Effect of unionization on routine employment over time. The red line shows data estimates for going from the 10th to the 90th percentile of unionization. The blue line shows the model output, comparing the low and high-unionized economy.

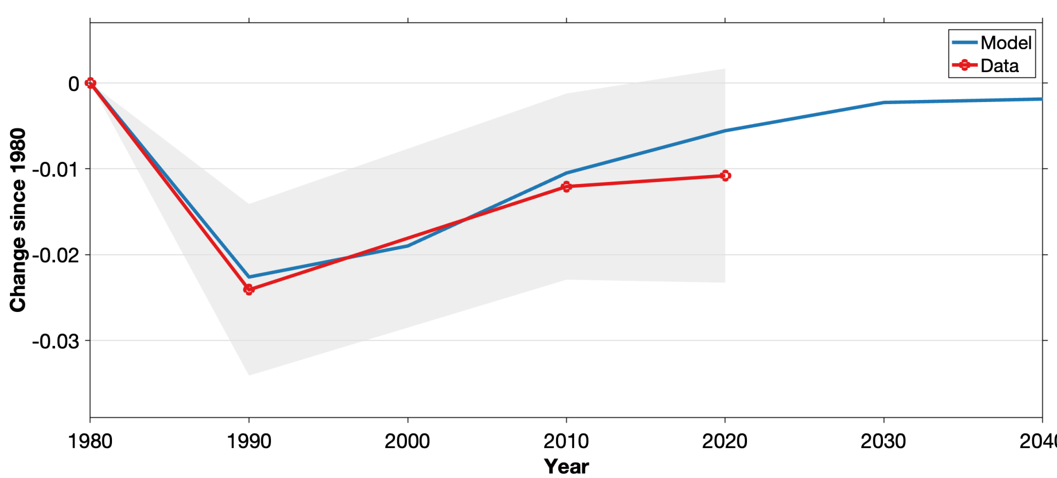

How does unionization shape the timing and extent of the overall routine employment decline in the model? Figure 9 displays the union effect on the routine employment share as documented in the data as well as in the model. In particular, for the model it shows the difference between the change in the routine employment share in the low and high-unionized economy. The model captures the accelerated routine employment decline in the high relative to the low-unionized economy between 1980 and 1990, driven by a preemptive reduction in routine employment in anticipation of future adoption to avoid firing costs along the transition. Recall, the difference in firing costs between the two economies is calibrated to match the 1990 data point and, thus, the close match between data and model in 1990 is no surprise. Without being targeted, the model matches well that after 1990 the routine employment decline in the low-unionized economy catches up as automation adoption increases and routine workers are being displaced in the low-unionized economy. By 2020, most of the gap has closed as routine employment in the high-unionized economy has declined by roughly 0.5 percentage points more, which is slightly larger than in the data. I now turn to the underlying distributional effect to see if the model can match how unions shape the distribution of transitional costs between young and old.

Figure 10: Effect of unionization on the age composition of routine workers over time. The red line shows data estimates for going from the 10th to the 90th percentile of unionization. The blue line shows the model output, comparing the low and high-unionized economy.

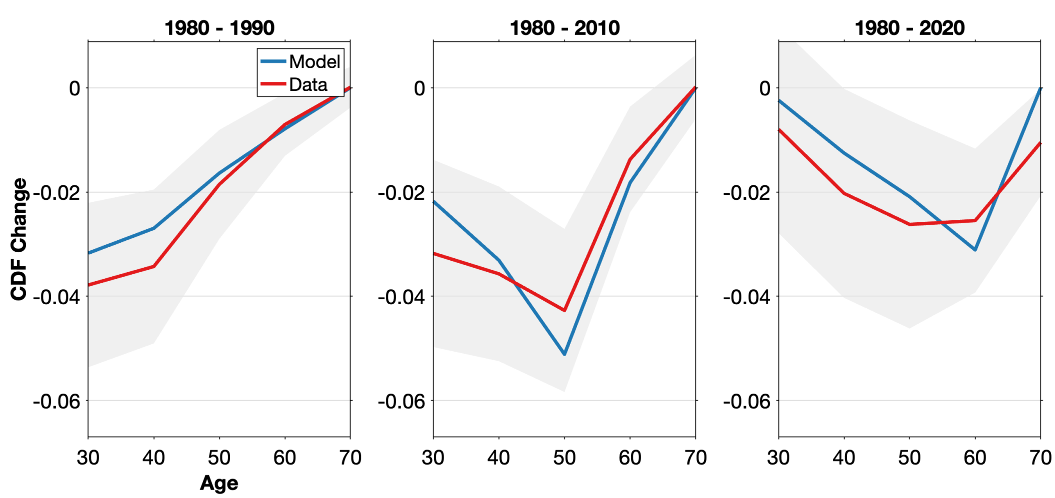

The model matches well the untargeted downward shift in the cdf of the age distribution of routine workers as displayed by Figure 10. Consistent with the data, unionization results in a downward shift in the cdf that is driven by a fall in the share of young workers. By 1990, the share of incoming workers below the age of 30 declines by 3.2 percentage points more in the high-unionized economy. This initial reduction in new hires then moves up the age ladder over time as the 1980 cohort ages. As a result, unionization induces an aging of the routine workforce, and these differences between the age composition of routine workers in the low and high-unionized economy persist throughout the transition.

5 The Welfare Cost of Automation

5.1 Measurement of the Welfare Cost of Automation

To measure the individual-specific welfare impact of automation, I calculate the permanent percent decrease in consumption a worker would be willing to accept to return to the 1980 steady state and thereby avoid automation, keeping her individual states fixed. I then compute this consumption equivalent variation at different times during the transition for different cohorts of workers to understand the evolution of the cross-sectional automation impact over time. Let $x_t(s, k, z^R, z^N, a)$ be the required compensation in consumption to be indifferent to automation for a worker with individual state (s, k, z^R, z^N, a) in period t. It is given by

$$x_t(s, k, z^R, z^N, a) = \left(\frac{V_{1980}(s, k, z^R, z^N, a)}{V_t(s, k, z^R, z^N, a)} \right)^{\frac{1}{1-\sigma}}.$$

My primary interest is in understanding to what extent welfare costs differ across routine workers of different cohorts and how unionization impacts these welfare costs. In particular, I focus on the difference between workers who entered the routine occupation before the automation shock hit and are caught by surprise, and workers who entered during the transition and therefore anticipate the current and future impact of automation when making their occupational choice.

5.2 The Welfare Cost of Automation for Routine Workers

Before studying how unionization affects the welfare cost of automation, I start by looking at the welfare cost of automation for routine workers in the low-unionized labor market to understand how automation shapes life-cycle wage paths of routine workers from 1980 onwards.

Figure 11: Welfare cost of automation for routine workers in 1990 in the low-unionized economy.

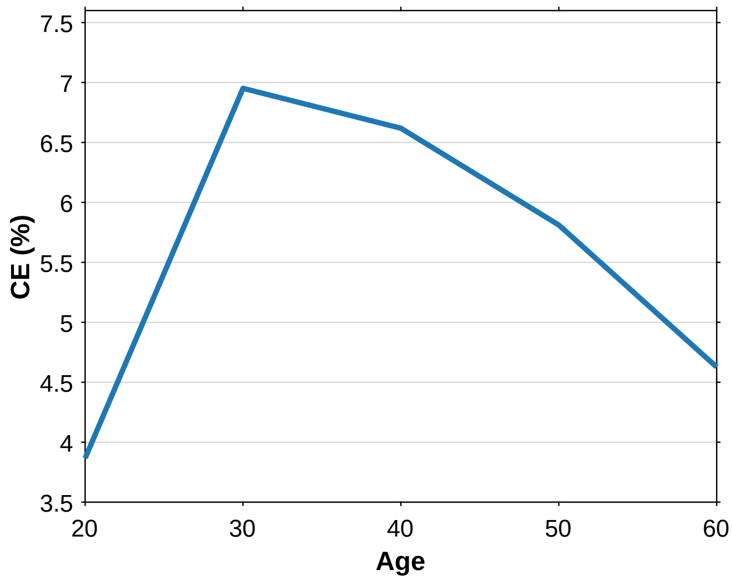

Figure 11 displays the welfare impact of automation for routine workers of different ages, averaged across all routine workers in that age group, in the low unionized labor market in 1990. The welfare costs are postitive for all age groups, meaning automation in 1990 is costly for all existing routine workers. For these workers, automation is costly for two reasons: displacement risk in the form of layoff risk, and permanent earnings losses as current and future routine wages fall. The welfare costs are large and range from 4% of permanent consumption for the incoming workers to 7% for workers aged 30.

What drives the inverted u-shape of the welfare costs by age? As the shock hits in 1990, incoming workers fully anticipate the negative impact of current and future automation on their entire life-cycle wage path and expected layoff risk in the routine occupation. In turn, a routine career becomes less desirable and the required routine labor productivity z^R that justifies entering the routine occupation rises. The average labor productivity of the incoming routine cohort in 1990 is therefore higher than for the older, incumbent cohorts. More productive workers have a higher life-cycle wage and consumption path which limits the welfare impact of automation for them. By contrast, incumbent routine workers are less productive on average as they made their occupational choice prior to the automation shock. Their consumption paths on average are lower which increases the welfare impact of automation. The costs are particularly driven by the subset of workers who would not have entered the routine occupation if they had anticipated the automation transition but are now stuck as switching occupations comes at the cost of losing their accumulated occupation-specific human capital. Among incumbent cohorts, the welfare costs are highest for the youngest workers aged 30 as they still have 40 years of routine work ahead of them and will experience the full automation impact over time. The horizon of older, incumbent routine workers is shorter which limits the welfare impact of automation for them to its short and medium-term effects, resulting in falling welfare costs by age.

Figure 12: Welfare cost of automation for routine workers along the transition in the low-unionized labor market.

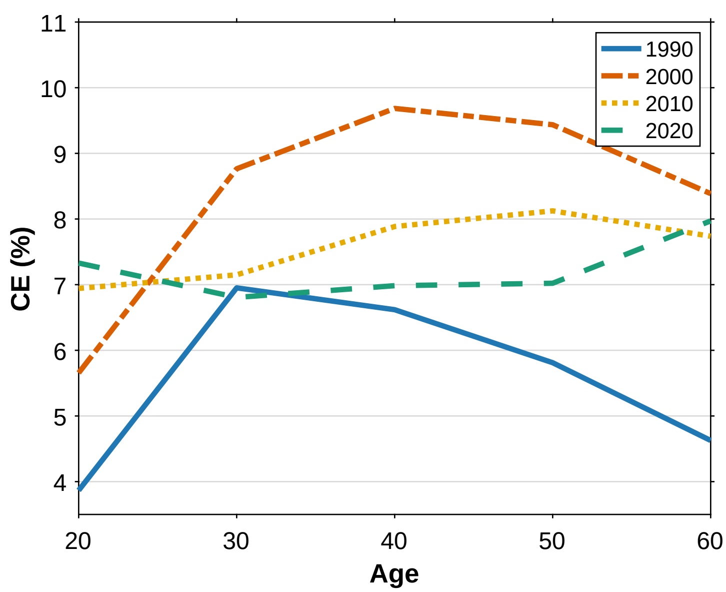

Figure 12 displays the welfare cost to routine workers along the transition. As capital prices keep falling, firms increase automation adoption in 2000 which raises the level of welfare cost relative to 1990. Note that the routine workers aged 40 in 2000 are, abstracting from layoffs and occupational switches, the same workers who were 30 years old in 1990. Along the transition, the welfare cost remains highest for the cohort of workers that entered the routine occupations in 1980, right before the automation shock. These workers experience large permanent earnings losses over their entire life-cycle and would be willing to give up almost 10% of permanent consumption in 2000 to avoid automation. Automation prices fall until 2010 and cohorts entering from 2010 onward enter either at the end or after the transition. As a result, the welfare costs in 2010 and

2020 are driven by the long-term impact of automation on routine workers which is similar across cohorts, leading to a flattening of the curve.

5.3 The Union Induced Transfer of Automation Costs

How does unionization shape the welfare cost of automation for routine workers along the transition? To answer this question, I compare the welfare cost of automation as computed above in the low and high unionized labor market.

Figure 13: Union effect on the welfare cost of automation in 1990.

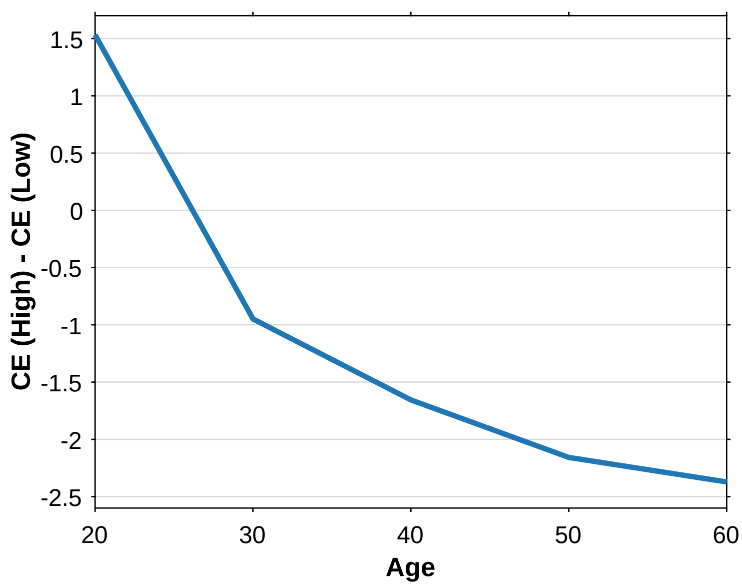

Figure 13 shows the difference in the welfare impact of automation between the low and high-unionized economy in 1990. Unionization shifts the welfare costs from incumbent workers to young, incoming workers who enter the routine occupation. Driven by the negative union effect on routine entry wages shown in Figure 8, the cost of automation for incoming routine workers is 1.5% of permanent consumption larger in the high-unionized relative to the low-unionized economy. By contrast, the union protects incumbent routine workers by stabilizing their current wages and limiting their layoff risk, reducing the cost of automation to them by up to 2.5% of permanent consumption.

Figure 14: Union effect on the welfare cost of automation along the transition.

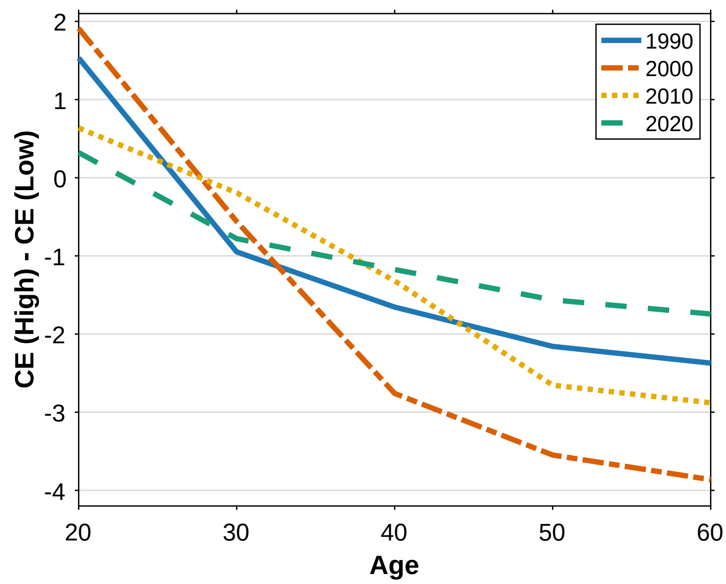

Figure 14 shows the union effect along the transition. As automation adoption increases in 2000, the productivity of routine workers falls which puts downward pressure on wages and increases the incentive of firms to lay off workers. As a result, the union protection becomes more valuable for incumbent routine workers. Especially for 60 year old routine workers layoffs would result in large earnings losses, meaning unionization reduces the welfare cost of automation for them by almost 4% of permanent consumption. While the positive union effect on older, incumbent workers increases substantially between 1990 and 2000, the negative impact on incoming workers rises only slightly, reaching close to 2% of permanent consumption. This reflects the fact that incoming workers in 1990 and 2000 endogenously respond to the impact of automation on their expected life-cycle earnings path in the routine occupation by not entering the routine occupation unless their routine labor productivity is sufficiently high. From 2010 onwards, capital prices stop falling and the economy transitions to its new long-run steady state. As a result, the union effect begins to flatten and decline in magnitude.

The welfare analysis emphasizes two things: First, among exposed routine workers who are substitutable with technology, automation is particularly costly for existing workers who made their occupational choice prior to the transition. These workers are caught by surprise and are stuck in a declining occupation, facing increased layoff risk and reduced earnings while switching occupations comes at the cost of losing their occupation-specific human capital. While automation also comes at substantial cost for workers who still enter the exposed occupations during the transition by reducing their expected life-cycle earnings path, these workers are on average more productive as they incorporate the current and future consequences of automation into their occupational choice which limits the welfare impact. Second, unionization shifts the welfare cost from incumbent routine cohorts to incoming workers. Workers who still enter the routine occupation during the transition experience declining entry wages due to unionization, and less productive young workers who would have entered the routine occupation in the past instead enter the non-routine occupation in order to avoid the automation impact.

6 Political Implications of the Conflict

An emerging political economy literature connects adverse economic shocks and outcomes to ideological realignment, which induces shifts in political preferences and economic policy. In particular, ideological polarization by race and education have widened among voters, most notably seen in a shift of less-educated Whites to the GOP (Pew Research Center (2014, 2017)). Mian et al. (2014) document a temporary increase in polarization in congressional voting outcomes following financial crises. Several studies document that the widening ideological polarization in Congress correlates with rising U.S. income inequality (McCarty et al. (2016), Voorheis et al. (2015)). Autor et al. (2020) find support for an ideological realignment in trade-exposed local labor markets in the form of rising support for strong-left and strong-right views, as well as pure rightward shifts. However, the causal relation between economic outcomes and shifts in voting behavior remains unclear.

How can the findings of this paper speak to and inform the narrative that puts economic factors at the center of political polarization? The model emphasizes that while unionization protects incumbent routine workers in response to an automation shock, it does so by shifting the cost of automation to young cohorts entering the labor market. In particular, unionization thereby causes a greater decline in routine entry wages for the subset of workers that still enter the routine occupation, as well as a larger reallocation of young workers to non-routine jobs, such as service sector jobs. As a result, unionization has intensified the deterioration of labor market experiences of less skilled incoming cohorts in routine and non-routine occupations since 1980. Cohorts of workers that have entered the labor market between 1980 and 2000 are in their 50s and 60s today, and precisely the workers whose voting behavior has shifted (Pew Research Center (2014, 2017)).

In order to test the importance of economic hardship as a driver of the shift in voting behavior, I test the hypothesis that union-induced employment decline of young routine manual workers between 1980 and 1990 across local labor markets is associated with a larger shift of voting to Republicans in the 2016 and 2020 presidential election.

I use data on county-level returns for presidential elections from 2000 to 2020 from the MIT Election Data and Science Lab (2017) and aggregate that data to the MSA level. The data shows only Democratic and Republican voter shares at the overall MSA level, not for the subset of routine manual workers. The dependent variable is the change in the Republican voter share in a MSA in the 2020 elections relative to the four previous elections in 2004, 2008, 2012 and 2016. I then regress the change in the voter share on the change in the share of routine manual workers that is below the age of 40 between 1980 and 1990, the unionization rate among routine manual workers, the interaction between unionization and the age shift, as well as controls. The interaction term is the coefficient of interest. As before, a greater fall in the share of young workers below the age of 40, that is, an aging of the workforce, indicates less employment prospects for young routine manual workers. The interaction with unionization then measures the additional fall in the routine manual employment share of young workers that is driven by unionization which is the variation of interest.

Table 7: The table shows the results from regressing changes in the republican voter share over time on the change in the share of young routine manual workers between 1980 and 1990, routine manual unionization and their interaction at the MSA level.

	Dependent variable: Change in Republican Voter Share			
	2020-2004	2020-2008	2020-2012	2020-2016
	(1)	(2)	(3)	(4)
1980s CDF Change Age 40	0.129	0.370^{***}	0.348^{***}	0.077
	(0.131)	(0.127)	(0.120)	(0.095)
Unionization	0.038	0.034	0.082^{**}	-0.021
	(0.051)	(0.042)	(0.041)	(0.026)
Interaction	-1.241	-1.620^{***}	-1.504^{***}	-0.440
	(0.611)	(0.587)	(0.535)	(0.378)
Mean dependent	-0.025	0.026	0.0071	0.0059
Observations	167	167	167	167
R^2	0.495	0.540	0.491	0.285
Adjusted R^2	0.433	0.484	0.429	0.198
Note:	*p<0.1; **p<0.05; ***p<0.01.			

Table 7 shows the regression results. The interaction terms are negative for all four models. Across MSAs, high unionization combined with a greater fall in the share of young workers in routine manual occupations between 1980 and 1990 is associated with an increase in the Republican voter share in the 2020 presidential election relative to previous elections. The effect is particular pronounced relative to the 2004, 2008 and 2012 election. Concretely, comparing a MSA at the 10th to a MSA at the 90th percentile in both independent variables, routine manual unionization and the fall in the share of young routine manual workers during the 1980s, is associated with a 3.9 percentage point and 3.6 percentage point additional increase in the Republican voter share in the 2020 relative the 2008 and 2012 presidential elections, respectively. The effect is weaker and not significant relative to the 2016 election, consistent with the 2016 presidential election already being polarized.

7 Conclusion

This paper argues that labor adjustment costs shape how the adverse labor market impact of labor-replacing technological change, such as automation, is distributed across different generations of workers as well as the timing of aggregate labor reallocation. To support the argument, I first document that unions have shifted the incidence of wage and employment declines among routine manual workers from older, incumbent to young, incoming cohorts since 1980. Moreover, unions have accelerated the decline in overall employment within routine manual occupations, resulting in a greater fall in employment in high-unionized labor markets between 1980 and 2000, and a subsequent catch-up of employment decline in less unionized labor markets.

I build a quantitative dynamic equilibrium model of endogenous technological change and unionization which demonstrates that the combination of gradual automation adoption over time and labor adjustment costs imposed by unions can jointly rationalize the two empirical observations. Labor adjustment costs incentivize firms to replace their workforce through reduced hiring rather than through costly layoffs in response to automation adoption. Moreover, when firms anticipate further automation in the future, they shrink their current workforce preemptively in order to avoid adjustment costs along the transition. This anticipatory adjustment channel is strong in the model and gives rise to an accelerated overall employment decline in routine occupations in high-unionized labor markets.

I use the model to quantify the effect of automation and unionization on the life-cycle consumption paths of routine workers across cohorts. The automation impact is most pronounced for incumbent routine workers who made their occupational choice prior to the transition, the welfare cost of automation to these workers reaches 10% of permanent consumption in 2000 in a low-unionized labor market. Workers entering the labor market during the transition endogenously adjust their occupational choice, and particularly less productive workers select into non-routine occupations in order to avoid the automation impact. Nevertheless, entering routine workers would still pay up to 7% of permanent consumption to avoid automation. In a high-unionized labor market, unions protect incumbent routine workers by lowering their layoff risk and wage decline which reduces the welfare cost of automation to these workers along the transition by up to 4% of permanent consumption relative to low-unionized labor markets. However, the cost is shifted to incoming cohorts. Routine entry wages and hiring falls relatively more in high-unionized labor markets, increasing the welfare cost of automation to incoming routine workers by up to 2% of permanent consumption along the transition. The difference in the welfare benefit to incumbent workers and welfare cost to incoming workers reflects the endogenous response of incoming workers to the automation shock. While older, incumbent workers are stuck in a declining occupation, having made their occupational choice not anticipating automation, incoming workers take the automation transition into account when making their occupational choice.

Finally, I provide suggestive evidence that the union induced shift of the adverse automation impact to young, incoming workers in the 1980s and 1990s has, through its persistent effects on the labor market experiences of these workers, implications for their voting behavior today.

References

Acemoglu, D. (2002). Technical change, inequality, and the labor market. *Journal of Economic Literature 40*(1), 7–72.

Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy 128*(6), 2188–2244.

Adao, R., M. Beraja, and N. Pandalai-Nayar (2021). Fast and Slow Technological Transitions. *Working Paper*.

Autor, D. (2010). The Polarization of Job Opportunities in the U . S . Labor Market. *Community Investments 23*(April), 360–361.

Autor, D., D. Dorn, G. Hanson, and K. Majlesi (2020). Importing political polarization? The electoral consequences of rising trade exposure. *American Economic Review 110*(10), 3139–3183.

Autor, D., F. Levy, and R. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Autor, D. H., L. F. Katz, and A. B. Krueger (1998). Computing Inequality: Have Computers Changed the Labor Market ? *The Quarterly Journal of Economics 113*(4), 1169–1213.

Baker, G., M. Gibbs, and B. Holmstrom (1994). The Internal Economics of the Firm : Evidence from Personnel Data. *Quarterly Journal of Economics 109*(4), 881–919.

Baker, G. and B. Holmstrom (1995). Internal Labor Markets: Too Many Theories, Too Few Facts. *The American Economic Review 85*(2), 255–259.

Baley, I., L. Ljungqvist, and T. J. Sargent (2020). Quit Turbulence and Unemployment. From the perspective of a particle physicist.

Baley, I., L. Ljungqvist, and T. J. Sargent (2023). Cross-phenomenon restrictions: Unemployment effects of layoff costs and quit turbulence. *Review of Economic Dynamics 50*, 43–60.

Bassanini, A. and R. Duval (2006). Employment Patterns in OECD Countries: Reassessing the Role of Policies and Institutions. *OECD Social, Employment and Migration Working Papers* (4), 1–129.

Bayer, C. and M. Kuhn (2023). Job levels and Wages. pp. 1–78.

Beraja, M. and N. Zorzi (2022). Inefficient Automation. *SSRN Electronic Journal*.

Berman, E., J. Bound, and Z. Griliches (1994). Changes in the Demand for Skilled Labor within U . S . Manufacturing: Evidence from the Annual Survey of Manufacturers. *The Quarterly Journal of Economics 109*(2), 367–397.

Bessen, J. E., M. Goos, A. Salomons, and W. Van den Berge (2023). Automatic Reaction - What Happens to Workers at Firms that Automate? *Review of Economics and Statistics*.

Blanchard, O. J. and L. H. Summers (1986). *Hysteresis and the European unemployment problem*. Number July.

Botero, J., S. Djankov, R. La Porta, F. Lopez-De-Silanes, and A. Shleifer (2004). The regulation of labor* j. *The Quarterly Journal of Economics* (November), 1339–1382.

Carruth, A. A. and A. J. Oswald (1987). On Union Preferences and Labour Market Models: Insiders and Outsiders. *97*(386), 431–445.

Cartter, A. (1959). *Theory of wages and employment*. Homewood; Ill.: R.D. Irwin.

Colonna, F. (2008). Does Union Membership increase Job Security? Evidence from British Panel Data. pp. 1–21.

Cortes, G. M., C. J. Nekarda, N. Jaimovich, and H. E. Siu (2020). The dynamics of disappearing routine jobs: A flows approach. *Labour Economics 65*(October 2019), 101823.

Costinot, A. and I. Werning (2018). Robots, trade, and luddism: A sufficient statistic approach to optimal technology regulation.

Davis, J., C. Fuenzalida, L. Huetsch, B. Mills, and A. M. Taylor (2024). Global Natural Rates in the Long Run: Postwar Macro Trends and the Market-Implied R* in 10 Advanced Economies. *Journal of International Economics* (March), 103919.

Doeringer, P. and M. Piore (1985). *Internal labor markets and manpower analysis*.

Dohmen, T. J., B. Kriechel, and G. A. Pfann (2004). *Monkey bars and ladders: The importance of lateral and vertical job mobility in internal labor market careers*, Volume 17.

Dunlop, J. T. (1944). *Wage Determination Under Trade Unions*. Reprints of economic classics. Macmillan Company.

Dvorkin, M. A. and A. Monge-Naranjo (2019). Occupation Mobility, Human Capital and the Aggregate Consequences of Task-Biased Innovations.

Fang, H. and X. Qiu (2021). "Golden Ages": A Tale of the Labor Markets in China and the United States. *SSRN Electronic Journal*.

Fellner, W. (1949). *Competition Among the Few: Oligopoly and Similar Market Structures*. A Borzoi book. A.A. Knopf.

Goldin, C. and L. F. Katz (2008). *The Race Between Education and Technology*. Cambridge, Mass. ; London :: Belknap,.

Goos, M. and A. Manning (2007). Lousy and Lovely Jobs : The Rising Polarization of Work in Britain. *The Review of Economics and Statistics 89*(1), 118–133.

Goos, M., A. Manning, and A. Salomons (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review 104*(8), 2509–2526.

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

Guerreiro, J., S. Rebelo, and P. Teles (2022). Should Robots Be Taxed? *Review of Economic Studies 89*(1), 279–311.

Guren, A., D. Hémous, and M. Olsen (2015). Trade dynamics with sector-specific human capital. *Journal of International Economics 97*(1), 126–147.

Haan, W. J. D., C. Haef, and G. Ramey (2005). Turbulence and Unemployment in a Job Matching Model. *Journal of the European Economic Association 3*(6), 1360–1385.

Heckman, J. J., L. Lochner, and C. Taber (1998). Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents. *Review of Economic Dynamics 1*, 1–58.

Hubmer, J. (2023). The Race Between Preferences and Technology. *Econometrica 91*(1), 227–261.

Humlum, A. (2020). Robot Adoption and Labor Market Dynamics.

Huo, Z. and J. V. Ríos-Rull (2020). Demand induced fluctuations. *Review of Economic Dynamics 37*, S99—-S117.

Katz, L. F. and D. H. Autor (1999). Chapter 26 Changes in the wage structure and earnings inequality. *Handbook of Labor Economics 3 PART*(1), 1463–1555.

Koch, M., I. Manuylov, and M. Smolka (2021). Robots and Firms. *Economic Journal 131*(638), 2553–2584.

Lagakos, D., B. Moll, N. Qian, and T. Schoellman (2018). Life-Cycle Human Capital Accumulation across Countries : Lessons from US Immigrants. *Journal of Human Capital 12*(2).

Levy, F. and R. J. Murnane (1996). With What Skills Are Computers a Complement? *The American Economic Review 86*(2), 258–262.

Lindbeck, A. and D. J. Snower (1988). Harassment , and Involuntary Unemployment: An Insider-Outsider Approach. *The American Economic Review 78*(1), 167–188.

Ljungqvist, L. and T. J. Sargent (1998). The European Unemployment Dilemma. *Journal of Political Economy 106*(3), 514–550.

Ljungqvist, L. and T. J. Sargent (2008). Two Questions about European Unemployment. *Econometrica 76*(1), 1–29.

Mankiw, N. G. and S. P. Zeldes (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics 29*(1), 97–112.

McCarty, N. M., K. T. Poole, and H. Rosenthal (2016). *Polarized America: the dance of ideology and unequal riches*. Cambridge, Mass. SE - xii, 240 pages : illustrations ; 24 cm.: MIT Press Cambridge, Mass.

McFadden, D. (1973). Conditional Logit Analysis of Qualitative Choice Behavior. *Drying Technology 33*(8), 907–914.

Mian, A., A. Sufi, and F. Trebbi (2014). Resolving Debt Overhang : Political Constraints in the Aftermath of Financial Crises. *American Economic Journal: Macroeconomics 6*(2), 1–28.

MIT Election Data and Science Lab, M. (2017). U.S. Presidential Elections Data 19762020.

Oxfam America (2018). THE BEST AND WORST STATES TO WORK IN AMERICA.

Parsons, D. O. (2005a). Benefit Generosity in Voluntary Severance Plans: The U.S. Experience. *SSRN Electronic Journal*.

Parsons, D. O. (2005b). Private Job Displacement Insurance in the United States, 1954-1979: Expansion and Innovation. *SSRN Electronic Journal*.

Parsons, D. O. (2005c). The Emergence of Private Job Displacement Insurance in the United States: Severance Pay Plans 1930-1954. *SSRN Electronic Journal*.

Pew Research Center (2014). Political Polarization in the American Public. Technical Report 1.

Pew Research Center (2017). The Partisan Divide on Political Values Grows Even Wider. Technical Report 1.

Pierce, B. (1999). Using the National Compensation Survey to Predict Wage Rates. *Compensation and Working Conditions* (Winter 1999), 8–16.

Rubinstein, Y. and Y. Weiss (2006). Chapter 1 Post Schooling Wage Growth: Investment, Search and Learning. *Handbook of the Economics of Education 1*(06), 1–67.

Ruggles, S., J. T. Alexander, K. Genadek, R. Goeken, M. B. Schroeder, and M. Sobek (2010). Integrated public use microdata series: Version 5.0 [Machine-readable database]. *Minneapolis: University of Minnesota 42*.

Strub, S., M. Gervin, and A. Bütikofer (2008). Vergleichende Analyse der Löhne von Frauen und Män- nern anhand der Lohnstrukturerhebungen 1998 bis 2006. Untersuchungen im Rahmen der Evaluation der Wirksamkeit des Gleichstellungsgesetzes.

Traiberman, S. (2019). Occupations and import competition: Evidence from Denmark. *American Economic Review 109*(12), 4260–4301.

Voorheis, J., N. McCarty, and B. Shor (2015). Unequal Incomes, Ideology and Gridlock: How Rising Inequality Increases Political Polarization.

A Further Empirical Evidence

A.1 Cohort Effects Across US States

I estimate cohort effects of low-skilled and routine workers over time across US states in order to test whether unionization has specifically affected the price of human capital of incoming cohorts during the automation transition. Following Heckman et al. (1998), and more recently Lagakos et al. (2018) and Fang and Qiu (2021), I decompose earnings growth into cohort, experience and time effects for cohorts born between 1940 and 1980 at the state and education (and occupation) level. Experience effects measure human capital growth over the life cycle, cohort effects measure the relative human capital level of a cohort at labor market entry, and time effects capture growth of the price of human capital over time. Thus, cohort effects essentially measure the value of human capital of each cohort at labor market entry in units of wages. Therefore, comparing the cohort component of earnings growth across states with high and low unionization in routine occupations quantifies by how much more the value, or the marginal product of labor, of incoming routine workers, measured in units of entry wages, has declined in states with high routine unionization relative to states with low routine unionization. It is well known that experience, cohort, and time effects cannot be separately identified without further assumptions due to perfect collinearity. In order to solve the identification issue, I closely follow the literature, and in particular Fang and Qiu (2021), by adopting the standard identification strategy first used by Heckman et al. (1998). The identifying assumption is that there is no experience growth in the final years of a worker's career which is based on theories of life cycle wage growth.[18] To see this, denote log wage $w_{i,c,t}$ of individual i from cohort c at time t as

$$w_{i,c,t} = p_t + h_{c,t} + \epsilon_{i,c,t}, \qquad \text{where } E_i[\epsilon_{i,c,t}] = 0.$$

Further decompose the cohort component into entry level human capital $s_c = h_{c,c}$ and return to e years of experience $r_{c,e} = r_e$ according to

$$w_{c,t} = p_t + s_c + r_e,$$

where p_t reflect time effects, s_c reflect cohort effects, and r_e reflect experience effects. It is

[18] See, e.g., Rubinstein and Weiss (2006)

straight forward to see that perfect collinearity $e = t - c$ now results in non-identification

$$w_{c,t+\tau} - w_{c,t} = p_{t+\tau} - p_t + s_c - s_c + r_{e+\tau} - r_e.$$

However, with the identifying assumption that there is no experience growth in final years of a worker's career, $r_e = 0$ for cohorts with $e \geq \bar{e}$, the above equation reduces to the following for workers of cohorts with $e \geq \bar{e}$:

$$w_{c,t+\tau} - w_{c,t} = p_{t+\tau} - p_t.$$

Thus, the assumption allows to identify common time effects through older cohorts. Since those time effects by definition are common across cohorts, this then allows for the identification of experience and cohort effects for all other cohorts. I apply the above estimation to decompose repeated cross-sectional annual earnings (total income) profiles from CPS non-parametrically into experience, cohort, and time effects at the state level.

Figure 15 displays the average cohort effects for highschool dropouts for states in the bottom and top quartile of unionization. In particular, the x-axis shows cohorts by birth year, the y-axis shows the cohort component of entry wages for each cohort relative to the cohort born in 1940. For the top row of plots states are weighted equally when ranked by percentile of unionization while the bottom row shows results when weighing states by their overall routine employment decline between 1980 and 2010. In the left column of plots, unionization percentiles are computed based on routine unionization in 1987 while routine employment in 2003 is used in the right column.

Figure 15: The panels compare the average estimated cohort effect in states with high and low routine manual unionization.

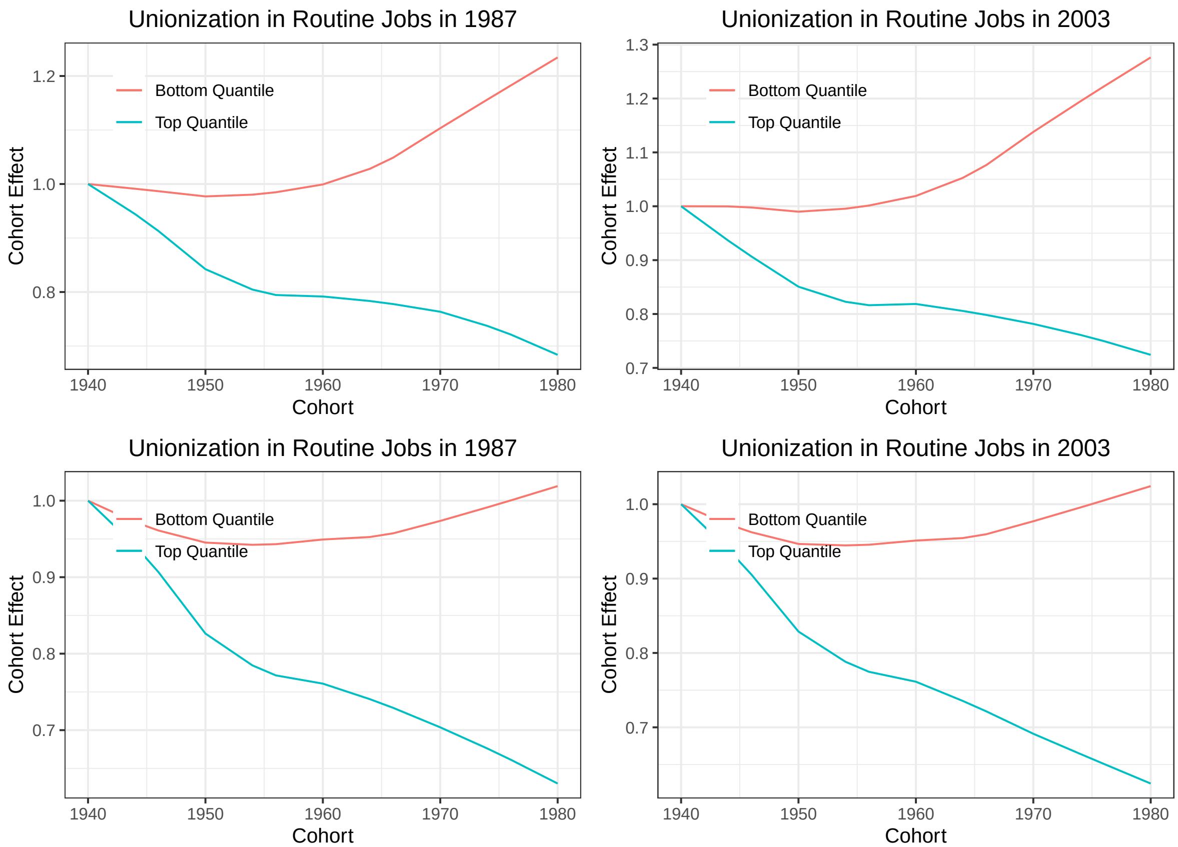

The top right panel shows that entry wages in states with high routine unionization decline relative to states with low routine unionization from 1940 onwards. However, entry wages diverge more strongly from 1960 onwards, that is, for cohorts that entered the labor market after 1980. Thus, labor market entry conditions for highschool dropouts have deteriorated particularly from 1980 onwards in states with high routine unionization. Moreover, the results are very similar when using the 2003 routine unionization. This is an important robustness test as it bolsters the MSA level analysis which uses average routine manual unionization in a MSA between 1995 and 2005 as explanatory variable since CPS coverage at the MSA level is not sufficient before 1995.

Figure 16 displays the cohort effect for highschool dropouts, when grouping states not only by their level of routine manual unionization, but also by other measures of employment protection. In particular, I use the Union Coverage, Right to Organize, Real Minimum Wage, and Total Employment Protection Score meausure from Oxfam America from 2018 (Oxfam America (2018)). In each case, the plot shows the average cohort effect for states in the top and bottom half of the corresponding measure. The top right panel groups states by a total employment protection score, the bottom left pannel groups states by their right-to-organize laws, and the bottom right panels uses the minimum wage at the state level. Across measures, highschool dropouts experienced declining entry wages, and the effect is most pronounced for cohorts entering the labor market from 1980 onwards. This bolsters the view taken in the model, that the mechanism underlying the documented union effect works through firing costs.

Figure 16: The panels compare the average estimated cohort effect in states with high and low employment protection measures.

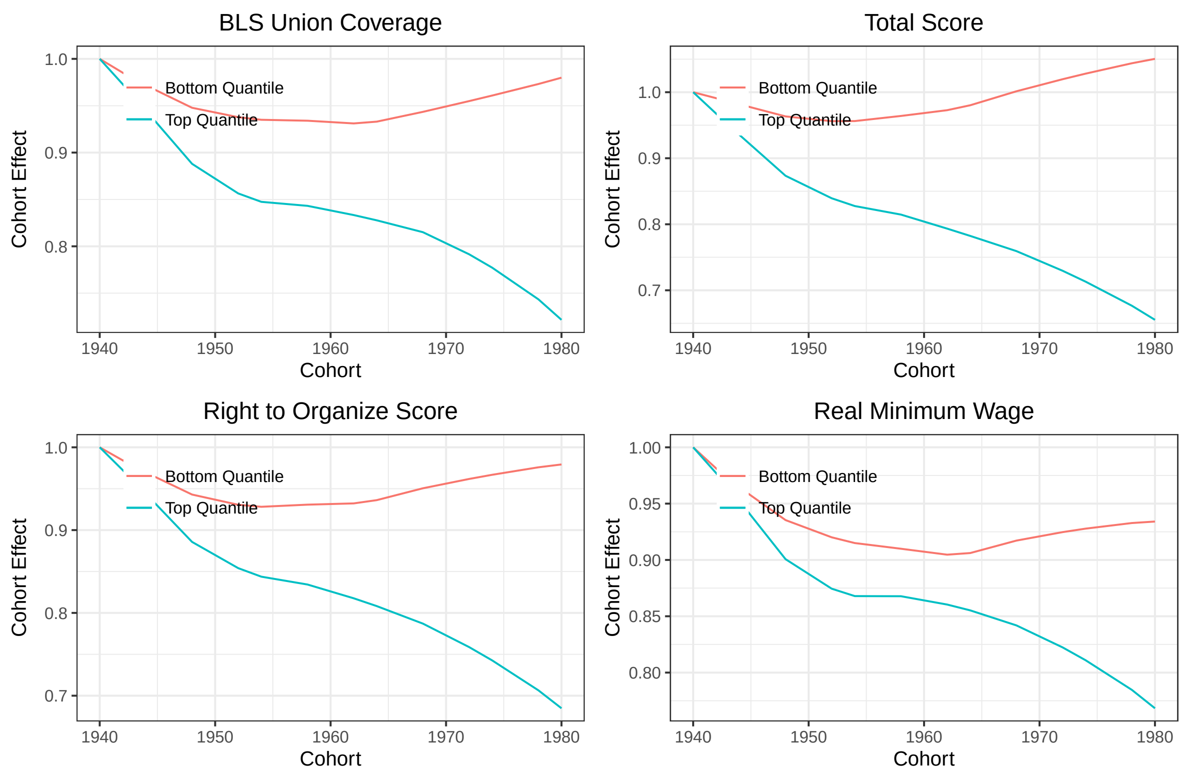

Figure 17 displays the cohort effects for all individuals in the sample, not just highschool dropouts. The differential evolution of entry wages from 1980 onwards by the level of employment protection in a state disappears. Thus, the relative decline in entry wages for workers entering the labor market from 1980 onwards in high protection relative to low protection states holds only for highschool dropouts, meaning less skilled workers, but is absent for other workers.

Figure 17: The panels compare the average estimated cohort effect in states with high and low employment protection measures.

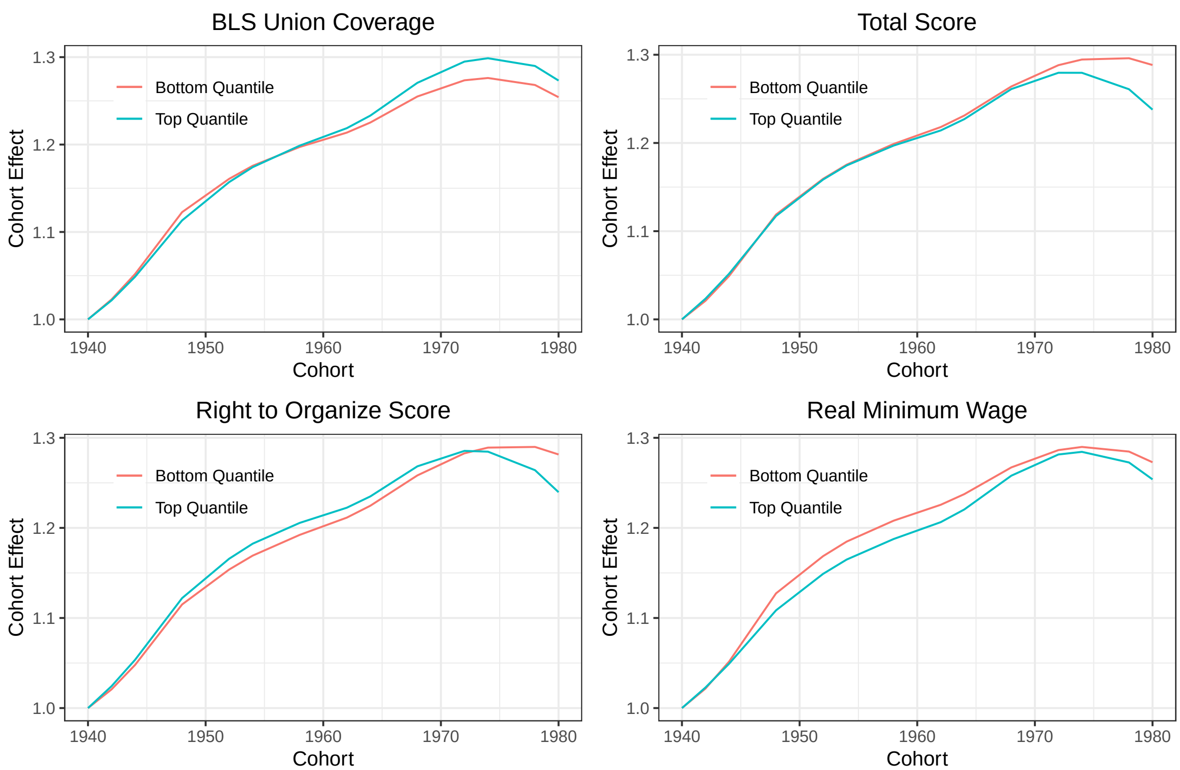

A.2 Inflow-Outflow Decomposition

Following Cortes et al. (2020), I use monthly individual-level matched CPS data, and classify every individual observation into 9 mutually exclusive employment states: Non-routine cognitive (NRC), routine cognitive (RC), routine manual (RM), non-routine manual (NRM) (employed or unemployed), and not in the labor force (NLF). I then compute monthly labor market flows from 1986 to 2012 for each U.S. state between these employment states, thus, transition rates between employment states over time. In order to quantify how much the outflow rate from routine employment (ERM) out of the labor force contributed to the reduction in overall routine employment (ERM), I construct counterfactual routine employment paths as Cortes et al. (2020) in the following way:

1. Fix the outflow rate from ERM to NLF at 1986 level:

$$\hat{\mu}_t(NLF, ERM) = \mu_{1986}(NLF, ERM)\ \forall t.$$

2. Leave other transition rates as in data, only rescale such that transition rates add up to 1 ($\sum_j \mu_t(NLF, i) = 1$).

3. Construct counterfactual employment shares over time:

$$S_{t+1} = \hat{\mu}_t S_t.$$

4. Compare the counterfactual decline in ERM to the realized one:

$$F(\text{ERM} \rightarrow \text{NLF}) \equiv 1 - \frac{\Delta E\hat{R}M_{cf}}{\Delta ERM}.$$

$F(\text{ERM} \rightarrow \text{NLF})$ measures how much decline in ERM would have been avoided if the transition rate from ERM to NLF stayed at its 1986 level $\mu_{1986}(NLF, ERM)$.

I then regress $F(\text{ERM} \rightarrow \text{NLF})$ on routine manual unionization and a set of controls including the 1980 industry composition and demographics. In particular, I run the following model across US states s:

$$F(\text{ERM} \rightarrow \text{NLF}) = \beta_0 + \beta_1 U_s + \gamma X_s + u_s.$$

Table 8 shows the results for three regressions using as independent variable the percentile of unionization for each state. The first column uses a dummy variable that measures whether a states is above or below the median of unionization, columns 2 and 3 use categorial variables that measure the quartile and quintile of a state's unionization.

Table 8: The table shows the results of regressing the contribution of outflow from routine manual employment into non-employment for routine manual employment decline on different measures of unionization.

	Dependent variable: $F(ERM \rightarrow NLF)$		
	Union Coverage Q2	Union Coverage Q4	Union Coverage Q5
	(1)	(2)	(3)
	-0.064^{*}	-0.038^{**}	-0.031^{***}
	(0.038)	(0.018)	(0.012)
Observations	51	51	51
R^2	0.103	0.157	0.169

The results shows that the statistical contributions of outflow rates to the overall decline in routine employment is negatively correlated with routine manual unionization.

As expected, the coefficient falls in magnitude from the left to the right when going to smaller percentiles while becoming better identified. This is because the change in unionization between two quintiles in column 3 is smaller than between the bottom and top half of unionization in column 1. Moreover, the estimates are economically meaningful. Going from the 1st (lowest) to the 5th (highest) quintile of unionization is associated with a 15pp increase in the share of routine manual employment decline accounted for by outflow from routine manual employment out of the labor force.

A.3 Additional Material for Empirical Analysis

A.3.1 Additional Robustness: Effect on RM Employment Decline

Figure 18: The graphs show the effect of going from the 25th to the 75th percentile of unionization on the RM employment share over time.

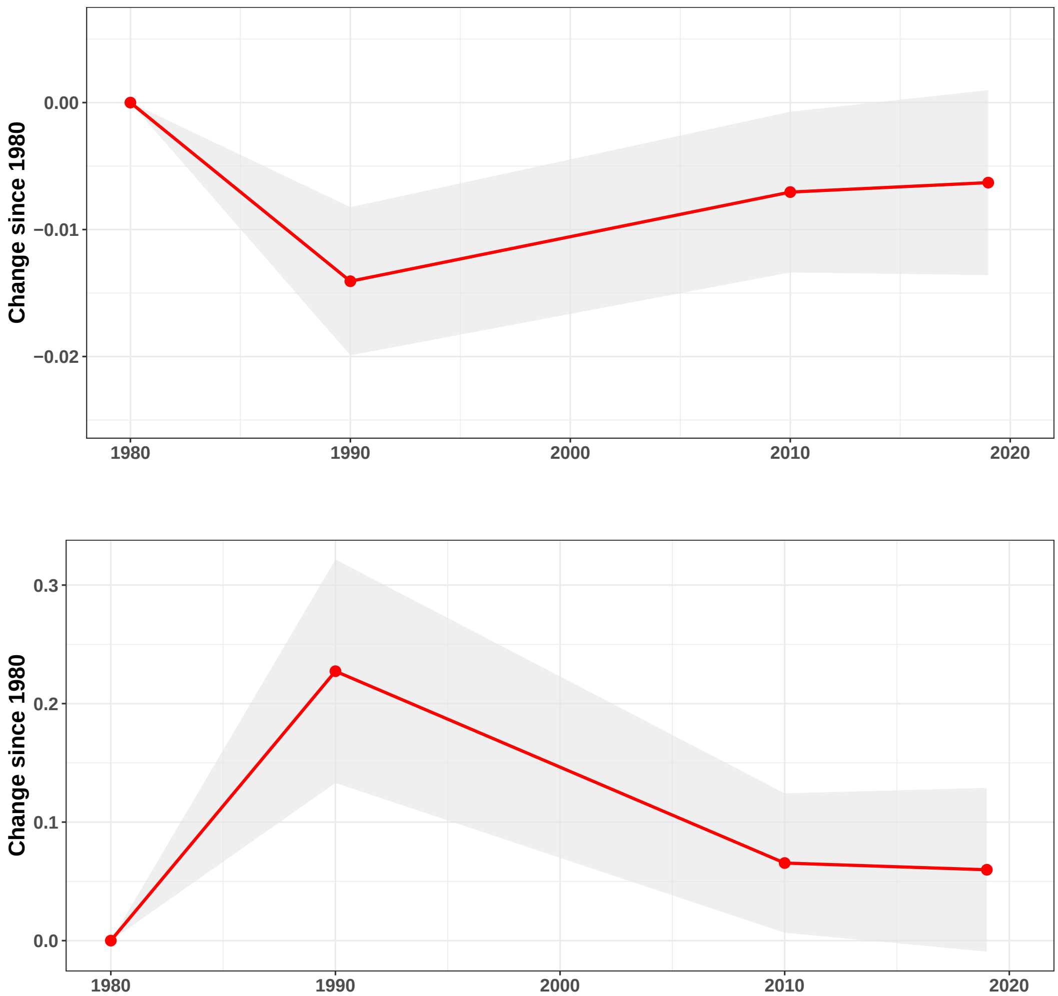

Figure 18 repeats the main exercise of regressing the change in the routine manual employment share on routine manual unionization and the set of controls, and it plots the effect of going from a MSA at the 25th percentile to a MSA at the 75th percentile of unionization. As expected, the magnitude of the effect falls relative to the main result that compares the 10th to the 90th percentile of unionization. However, the union effect remains significant and large. The routine manual employment share falls significantly more in high-unionized MSAs between 1980 and 1990, after which employment decline in low-unionized MSAs starts to catch up. The union effect is large, reaching almost 25% of the mean routine manual employment decline across MSAs between 1980 and 1990.

Figure 19: The graph shows the effect of going from the 10th to the 90th percentile of unionization in 1986 at the MSA level on the RM employment share over time. Unionization is measured at the MSA level in 1986.

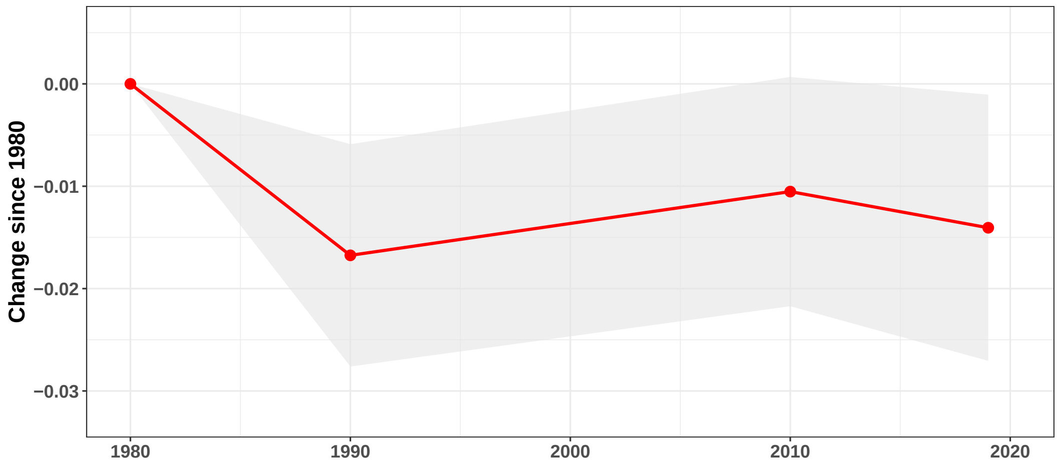

Figure 20: The graph shows the effect of going from the 10th to the 90th percentile of unionization on the RM employment share over time. The set of controls additionally includes the change in the age composition at the MSA level as a proxy for migration.

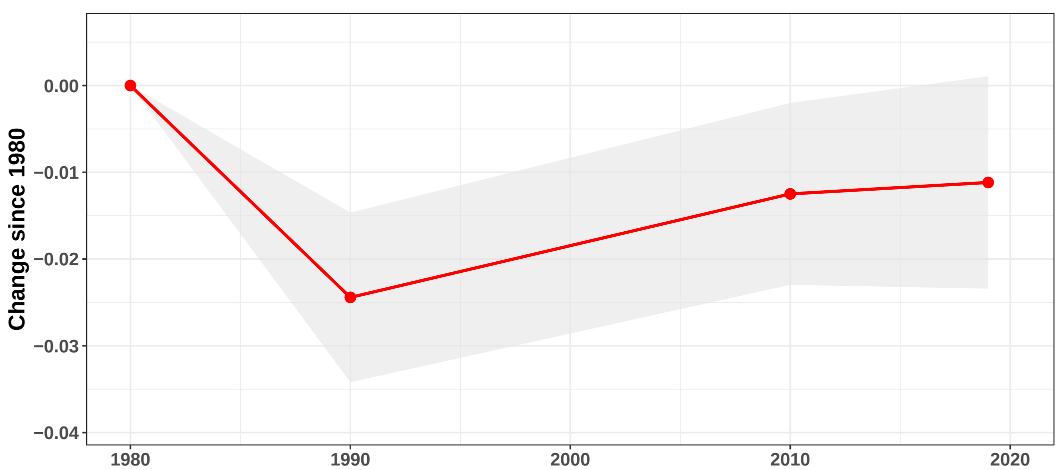

A.3.2 Main Results: Union Effect on Age Composition

Table 9: Effect of unionization on the change in the age distribution of routine manual workers between 1980 and 1990.

	Dependent variable: Change in CDF across Ages				
	Age 20	Age 30	Age 40	Age 50	Age 60
	(1)	(2)	(3)	(4)	(5)
Unionization	-0.043^{***}	-0.126^{***}	-0.114^{***}	-0.062^{***}	-0.023^{**}
	(0.012)	(0.027)	(0.026)	(0.020)	(0.011)
Change RM 1980-1990	0.165^{***}	0.687^{***}	0.405^{***}	0.111	0.088^{*}
	(0.059)	(0.156)	(0.141)	(0.091)	(0.052)
Mean dependent	-0.072	-0.099	-0.017	0.026	0.012
Observations	200	200	200	200	200
R^2	0.314	0.474	0.367	0.261	0.260
Adjusted R^2	0.262	0.434	0.319	0.205	0.204
Note:	*p<0.1; **p<0.05; ***p<0.01				

Table 10: Effect of unionization on the change in the age distribution of routine manual workers between 1980 and 2010.

	Dependent variable: Change in CDF Gap across Ages				
	Age 20	Age 30	Age 40	Age 50	Age 60
	(1)	(2)	(3)	(4)	(5)
Unionization	-0.044^{***}	-0.106^{***}	-0.119^{***}	-0.142^{***}	-0.046^{***}
	(0.014)	(0.035)	(0.028)	(0.030)	(0.017)
Change RM 1980-2010	0.051	0.580^{***}	0.551^{***}	0.272	0.055
	(0.065)	(0.189)	(0.155)	(0.175)	(0.094)
Mean dependent	-0.1	-0.2	-0.18	-0.083	-0.0063
Observations	200	200	200	200	200
R^2	0.260	0.334	0.331	0.364	0.256
Adjusted R^2	0.204	0.284	0.280	0.316	0.199
Note:	*p<0.1; **p<0.05; ***p<0.01				

Table 11: Effect of unionization on the change in the age distribution of routine manual workers between 1980 and 2019.

	Dependent variable: Change in CDF Gap across Ages				
	Age 20	Age 30	Age 40	Age 50	Age 60
	(1)	(2)	(3)	(4)	(5)
Unionization	-0.037^{**}	-0.026	-0.067^{**}	-0.087^{***}	-0.084^{***}
	(0.017)	(0.029)	(0.033)	(0.031)	(0.024)
Change RM 1980-2019	0.115	0.226^{*}	0.349^{**}	0.302^{**}	0.237^{***}
	(0.078)	(0.135)	(0.153)	(0.143)	(0.086)
Mean dependent	-0.094	-0.17	-0.16	-0.11	-0.051
Observations	147	147	147	147	147
R^2	0.253	0.327	0.379	0.287	0.291
Adjusted R^2	0.174	0.256	0.313	0.211	0.216
Note:	*p<0.1; **p<0.05; ***p<0.01				

A.3.3 Additional Robustness: Union Effect on Age Composition

Figure 21: The graphs show the effect of going from the 25th to the 75th percentile of unionization on the change in the routine manual age composition over time.

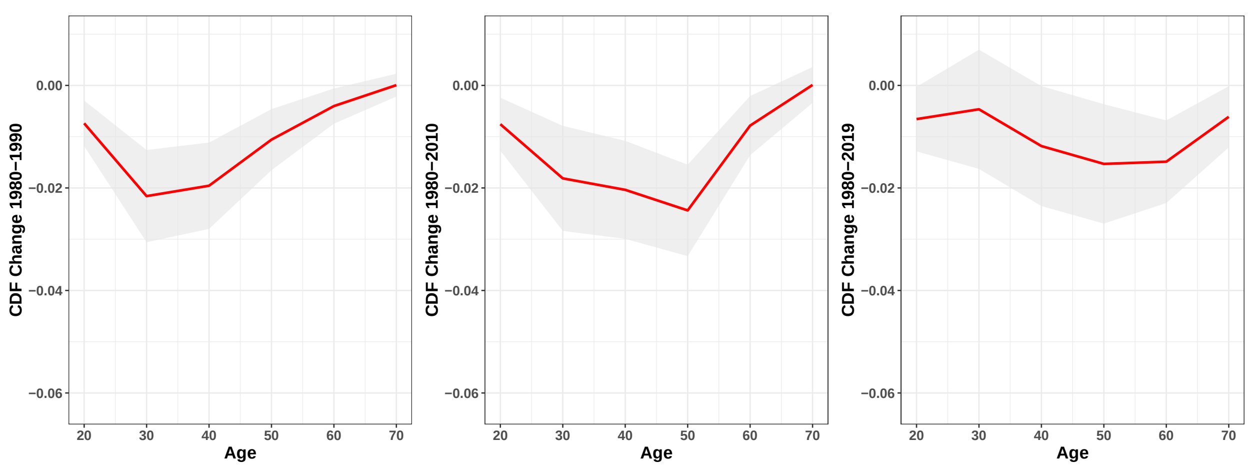

Figure 21 repeats the main exercise of regressing the change in the age composition (cdf) of the routine manual workforce in a MSA on routine manual unionization and the set of controls, and it plots the coefficient scaled by the difference in unionization between a MSA at the 25th percentile and a MSA at the 75th percentile of unionization. As expected, the magnitude of the effect falls relative to the main result that compares the 10th to the 90th percentile of unionization. However, the union effect remains significant and large.

Figure 22: The graphs show the effect of going from the 10th to the 90th percentile of unionization in 1986 at the MSA level on the change in the routine manual age composition over time. Unionization is measured at the MSA level in 1986.

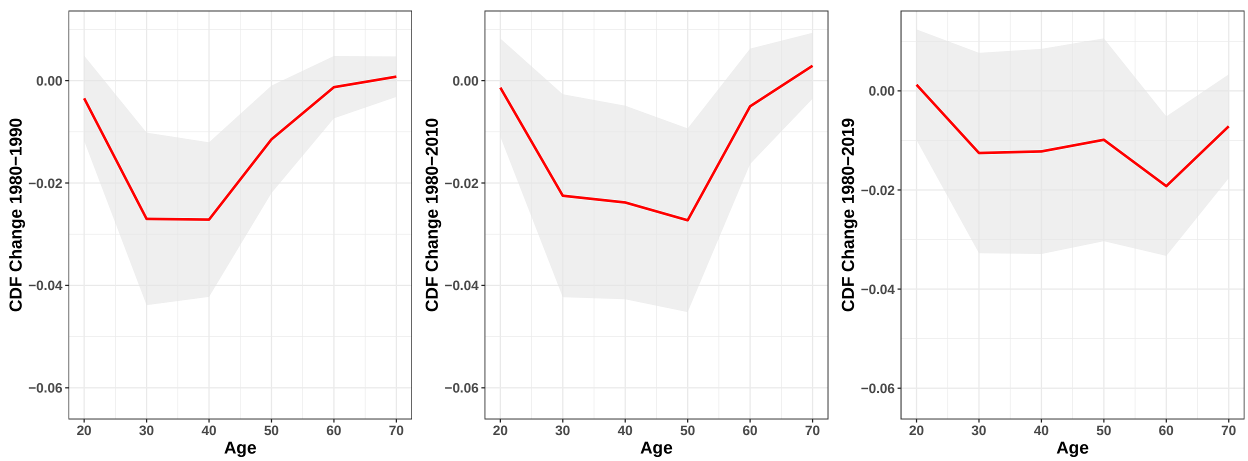

Figure 22 repeats the main exercise of regressing the change in the age composition (cdf) of the routine manual workforce in a MSA on unionization and the set of controls. Instead of using the average share of unionization among routine manual workers between 1995 and 2005, I here use unionization in 1986. However, in order to have sufficient coverage, unionization is measured among all employed workers at the overall MSA level. Again, the magnitude of the effect falls relative to the main result as expected, but the union effect remains.

Figure 23: The graphs show the effect of going from the 10th to the 90th percentile of unionization on the change in the routine manual age composition over time. The set of controls additionally includes the change in the age composition at the MSA level as a proxy for migration.

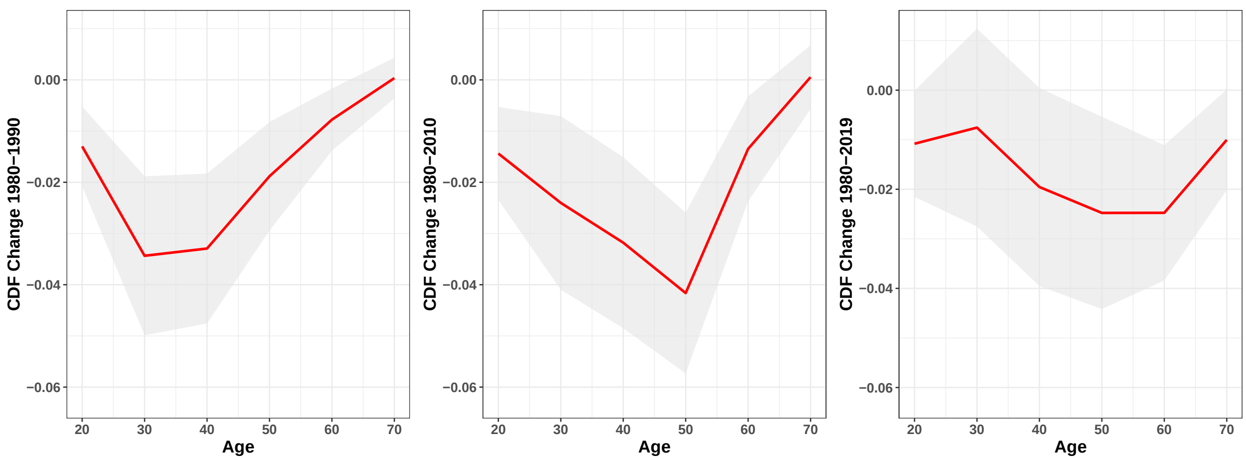

One concern may be that young workers in highly unionzed MSAs respond to bad employment prospects in routine manual occupations by migrating to less unionized MSAs. Figure 23 shows results when additionally controlling for changes in the age composition of all employed workers at the MSA level as a proxy for in and out migration for the corresponding time periods. The union effect remains basically unchanged, indicating that the main results are not driven by migration.

A.3.4 Additional Robustness: Union Effect on Age Composition

Table 12: Robustness: Effect of unionization on the change in the age distribution of routine manual workers between 1980 and different stages of the transition (1990, 2010, 2019). Regression uses routine manual employment share in 1980 for each MSA as weights.

	Dependent variable: Change in CDF across Ages				
	Age 20	Age 30	Age 40	Age 50	Age 60
	(1)	(2)	(3)	(4)	(5)
CDF Change 1980-1990	-0.042^{***}	-0.121^{***}	-0.110^{***}	-0.060^{***}	-0.023^{*}
	(0.012)	(0.028)	(0.027)	(0.021)	(0.012)
CDF Change 1980-2010	-0.043^{***}	-0.104^{***}	-0.113^{***}	-0.138^{***}	-0.042^{**}
	(0.014)	(0.034)	(0.029)	(0.031)	(0.018)
CDF Change 1980-2019	-0.036^{**}	-0.022	-0.063^{*}	-0.080^{**}	-0.079^{***}
	(0.017)	(0.029)	(0.034)	(0.032)	(0.025)

Table 12 reports the results when estimating the effect of unionization on the change in the age distribution of routine manual workers between 1980 and different stages of the transition using the routine manual employment share in 1980 for each MSA as regression weights. The results are robust to reweighting.

B Model Appendix

B.1 Workers hold fixed equity shares

Figure 24 displays the welfare cost to routine workers along the transition in the low-unionized labor market if workers hold fixed and equal equity shares in the firms. While the overall distribution and evolution of welfare costs to routine workers is similar to the baseline economy in which workers do not own equity, the level of welfare costs is lower. If workers hold equity, they benefit from automation due increased profits, which partially offsets the earnings losses they incur.

Figure 24: The graph shows the welfare cost of automation for routine workers along the transition in the low-unionized labor market when workers hold fixed and equal equity shares.

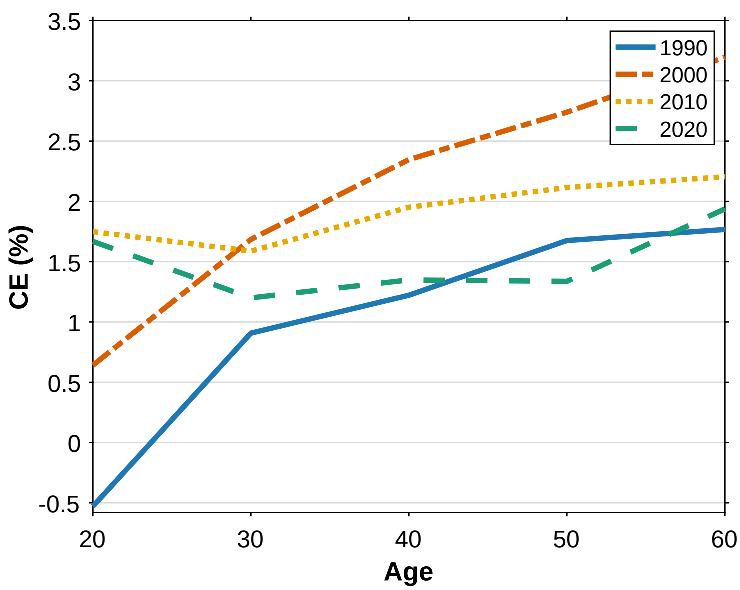

Figure 25 displays the union effect on welfare cost to routine workers along the transition for the fixed equity case. Again, the shape is similar to the baseline economy, unions shift the cost from older, incumbent cohorts to young workers. However, the union effect falls as wage income becomes a smaller component of workers' income, which in turn reduces the importance of lower layoff risk and limited earnings losses due to high unionization.

Figure 25: The graph shows the union effect on the welfare cost of automation along the transition.