



THEMA

théorie économique,
modélisation et applications

THEMA Working Paper n°2021-05
CY Cergy Paris Université, France

**Routine-biased technological change
and wages by education level:
Occupational downgrading and
displacement effects**

Clément Bosquet, Paul Maarek, Elliot Moiteaux



February 2021

Routine-biased technological change and wages by education level: Occupational downgrading and displacement effects*

Clément Bosquet[†] Paul Maarek[‡] Elliot Moiteaux[§]

February 19, 2021

Abstract

Taking advantage of geographic (and time) variation in the proportion of routine occupations in the US, we study the impact of this variation on the wage rate of workers by educational group. Using individual data and a Bartik-type IV strategy, we show that not only non-college-educated workers but also, in the same proportion, workers with fewer than four years of college are negatively impacted by this routine-biased technological change. The latter skill group currently represents 30% of the US population. We show that only 10% to 20% of the impact on both educational groups is related to occupational and industrial downgrading (the composition effect) and that most of the wage impact occurs within industries and occupations, including manual service occupations. This is consistent with the displacement effect described in the theoretical literature on task-biased technological change and automation.

JEL Classification: I24, J23, J24, J31, O33

Key words: job polarization, routine occupations, wages, education

*We are particularly grateful to Ghazala Azmat, Laurent Gobillon, Ariell Reshef, Thepthida Sopraseuth and various seminar participants for useful comments.

[†]Centre d'Économie de la Sorbonne, Université Paris 1 Panthéon-Sorbonne. Email: clement.bosquet@univ-paris1.fr.

[‡]*Corresponding author.* Université Paris 2, LEMMA. paul.maarek@u-paris2.fr

[§]ThEMA, UMR 8184, CY Cergy Paris Université, CNRS, F-95000 Cergy-Pontoise. 33 Boulevard du Port, 95011 Cergy-Pontoise Cedex, France.

1 Introduction

In recent decades, the labor market has become increasingly polarized. Technological progress and globalization have progressively made routine occupations disappear. Technology has substituted human labor in routine tasks such as bookkeeping. The corresponding jobs require some skills and are not the lowest paid in the labor market but are rather intermediate-pay jobs (Goos and Manning, 2007). The proportion of routine occupations in the US went from 42.9% in 1970 to 26.1% in 2017.¹

These intermediate occupations in the wage distribution have been replaced by less well-paid manual occupations, and, at the top of the distribution, abstract occupations. Many employees in these routine jobs have moved into manual occupations. Cortes (2016) shows that the least able workers tend to reallocate to manual occupations. Competition to reallocate toward abstract occupations is very high, especially for less able workers. Alongside the strong increase in the number of abstract occupations, the number of college graduate workers has drastically increased in the United States during the same period. Autor (2019) shows that many non-college-educated workers have moved from intermediate (routine) occupations requiring specific skills to low-paid occupations requiring generic skills. This has decreased the proportion of non-college-educated workers who hold intermediate-skill occupations.

We study how this process has impacted the wages of US workers by education/skill group over the period 1970-2010.² The literature generally distinguishes between college graduates (those with at least four years of college) and non-college-educated workers, and routine-biased technological change seems clearly unfavorable to the latter (Acemoglu and Autor, 2011; Autor, 2019). We distinguish not only low-skilled individuals without a college degree

¹The literature uses the terms ‘job polarization’, ‘routine-biased technological change’ or ‘task-biased technological change’ to describe this phenomenon.

²Like Michaels et al. (2014), “we follow the literature by referring to ‘education’ and ‘skills’ interchangeably; thus, ‘high skilled’ refers to ‘highly educated’, ‘middle skilled’ refers to those with intermediate levels of education, and ‘low skilled’ refers to those with lower levels of education.” For more details on how the variables are constructed, see section 2.

but also two different types of college-educated workers: those who went to college for three years or fewer and possibly obtained a diploma (middle skilled) and those who completed college, generally in four years or more (high skilled). Middle-skilled individuals are an important group of the US working-age population, with this group's share increasing from 9.9% in 1970 to 28.7% in 2017. Some of them occupy routine jobs and do not necessarily have the capacity to reallocate into abstract occupations: 45.1% of middle-skilled workers held a routine occupation in 1970 against only 32.8% in 2017. Conversely, only 18.8% of middle-skilled workers held a manual occupation in 1970 against 32.5% today. The proportion of middle-skilled workers holding an abstract occupation even decreased slightly during the period (from 36.1% to 34.6%), despite the strong increase in the proportion of abstract occupations (from 25.8% to 43.1%) among all jobs. The latter are mainly held by high-skilled individuals. It is thus very likely that the disappearance of routine jobs also affected medium-skilled workers. On the other hand, approximately 80% of high-skilled workers hold an abstract occupation, and this proportion has been very stable over time. Our intuition is that compared to medium-skilled workers, high-skilled workers should not be affected by routine-biased technological change given that the occupations that they overwhelmingly hold have tended to expand rather than decrease.

Using a local labor market approach and an instrumental variable strategy, we show that this current occupational change in the labor market has significantly decreased the wage rate of individuals with low and medium skills over the long run (1970-2010). Individuals with a four-year degree or more do not seem affected by the change in occupational structure, which is consistent with the job polarization literature. The (abstract) occupations that they occupy are not affected by automation. The decline of routine jobs affects low-skilled workers in similar proportion to middle-skilled workers.

The use of individual data on wages and occupations allows us to disentangle the mechanisms behind the impact of occupational change on the wage structure: the average wage of an occupation in a given industry can remain constant, but low- and medium-skilled individ-

uals may experience some occupational downgrading or/and industry switches (composition effects). Many workers whose routine occupation disappeared have reallocated into manual occupations that are characterized by lower wages. This occupational downgrading is highlighted by Autor (2019) for low-skilled workers and corresponds to the composition effects in the task-biased workhorse model of Acemoglu and Restrepo (2019). This composition effect should also hold for medium-skilled workers, as argued earlier, since such workers largely hold routine occupations and do not reallocate into abstract jobs. This effect is very important in the narrative in the job polarization literature explaining the decrease in economic opportunities of low-skilled workers (see, for instance, Autor, 2019). Alternatively, wages can change within occupations for a skill group (general equilibrium effect) due to a displacement effect as in Acemoglu and Restrepo (2018)'s task-biased model of technological change. The workers occupying the jobs that disappeared and who are reallocated into existing occupations create downward wage pressure within those occupations, in which there is an overflow of labor supply. Even workers who do not experience occupational downgrading are affected. In the case of medium-skilled workers who cannot reallocate into abstract jobs and therefore compete with low-skilled workers for manual jobs (and the remaining routine jobs), they should experience a decrease in wages relative to their productivity within their (manual and routine) occupations, as low-skilled workers. If the displacement effect is sufficiently high, the wages of medium-skilled workers may even decrease in absolute terms. This is an important implication of recent works on task-biased technological change and automation: workers are not only affected because they reallocate to some occupations that pay less (the composition or between effect) but also because the displacement effect and overflow of labor supply imply a decrease in wages within occupations.³

³The polarization process may also be related to or result in structural change. In Goos and Manning (2007), routine-biased technological change implies that some workers move between industries when changing occupations (the between component of job polarization). In Barany and Siegel (2018), some sector-level asymmetric productivity shocks (not task biased) that affect routine-intensive industries and lead to reallocation across industries generate structural change and job polarization between industries. This implies that workers also reallocate across industries. The wage impact for workers should be partly captured by an occupation-specific composition effect (occupational downgrading), but wages could also be affected by an industry effect. Occupations do not necessarily have the same wage rate across industries.

In our empirical strategy, using a Bartik shift-share instrument, we identify local labor market (commuting zone or CZ) shocks that alter the occupational structure of these areas. In the same spirit as Autor and Dorn (2013) or Acemoglu and Restrepo (2020a), we use the evolution of the proportion of routine jobs by sector at the national level to measure how each sector is affected by job polarization. We then interact those proportions with the initial specialization of the areas (sector shares), which allows us to determine how exposed each zone is, according to its initial specialization, to a global technological progress shock that eliminates routine occupations. Identifying an exogenous shock is crucial given that (i) many variables simultaneously influence the proportion of routine jobs as well as local wages and (ii) a change in the sectoral and occupational structure may be the consequence of wage changes. Next, we take into account the concerns raised by Goldsmith-Pinkham et al. (2018) to verify that it is not only a few industries that drive the variation in the instrument and provide several robustness tests and discussions of our results.

Our results suggest that both low- and medium-skilled workers are negatively affected in similar proportions by the disappearance of routine occupations. Consistent with theoretical priors, we find no effect (or only a very reduced effect) on the high-skilled group. The effects on the low- and medium-skilled groups are economically sizeable: a 5.7 percentage point increase in the initial proportion of routine occupations in the CZ, which corresponds to one standard deviation in the distribution observed in local labor markets over the 1970-2010 period, implies a decrease in average wages of 3.3% for the low-skilled group and 2.5% for the medium-skilled group over a 10-year period.

Our results also suggest that the two channels mentioned above are both important in explaining the wage impact of job polarization for the different skill groups. However, the within effect related to the displacement effect seems to be dominant since it explains more than 80% (90%) of the wage impact for medium-skilled (low-skilled) workers. This result is in line with the recent findings of Hunt and Nunn (2019), who argue that wage inequalities are only weakly related to occupations. Wage dispersion within occupations is

much more important than that between occupations. We show that the routine occupation wage premium (relative to the wage in manual occupations) for low- and medium-skilled workers is indeed very low. It is thus not surprising that the wage impact of the disappearance of routine jobs does not occur through the composition effect. This contradicts the traditional narrative in the job polarization literature positing (i) that wage distributions are highly related to occupations and (ii) that automation has been detrimental to low-skilled workers because it replaced well-paid routine occupations with low-paid manual ones. Instead, we show that routine-biased technological change has impacted low-skilled (and medium-skilled) workers mainly through the downward wage pressure it implies within occupations.

Furthermore, we show that the negative effects for low- and medium-skilled workers is observed for individuals in routine or manual occupations, including low-skilled (manual) service occupations. This is at odds with the traditional narrative on job polarization, according to which the highest wage growth should be observed at the top and bottom of the wage distribution. According to Autor and Dorn (2013), this is especially true for manual low-skilled service occupations, in which wages should increase due to strong complementarities with abstract high-paid occupations. We show that the disappearance of routine jobs decreases wages for both low- and medium-skilled workers in all manual occupations, including low-skilled manual service occupations. This finding is more consistent with the displacement effect of Acemoglu and Restrepo (2019) discussed earlier.⁴

Moreover, we find that the negative effects are more important for young workers, which is not surprising, as they face fewer job opportunities when they enter the labor market. Older low- and medium-skilled workers had the opportunity to climb the occupational ladder when they entered the labor market, at a time when more opportunities existed.

Our results may explain part of the wage and income stagnation recently observed in the data. (Since the 1980s, real income has not increased for around half of the American

⁴Hunt and Nunn (2019), among others, also note that the positive effect on manual service occupations in Autor and Dorn (2013) could be offset by an excess supply of labor in those occupations, which could create downward wage pressure. We show that the displacement effect clearly does offset the positive complementarity effect.

population.) We show that job polarization creates downward wage pressure for an important part of the population. (The low- and middle-skilled groups represented 70% of the US population in 2017.)

Our paper is related to several strands of the literature. It is first connected to the literature on polarization and task-biased technological change. In a seminal contribution Autor et al. (2003) and then Acemoglu and Autor (2011), Autor and Dorn (2013), Goos and Manning (2007), Goos et al. (2009), and Dustmann et al. (2009) show that middle-pay (routine) occupations tend to disappear and that this decline is related to ICT diffusion. Traditionally, the literature on polarization has examined the evolution of intermediate wages by looking at the middle of the wage distribution, sometimes referred to as the medium-skill segment. However, this does not tell who medium-wage individuals are and how different educational groups are affected by the disappearance of routine jobs. Generally, this disappearance is seen as detrimental for low-skilled workers who have not been to college, as they tend to hold routine occupations. (see, for instance, Acemoglu and Autor, 2011; Autor, 2019). We show that many workers with some college (fewer than four years) are also strongly affected by this process and that manual workers, including those in service occupations, also suffer from the disappearance of routine jobs, consistent with the literature on automation and task-biased technological change. The impact on wages occurs within occupations rather than between occupations. This casts some doubt on the hypothesis that wage inequality and the wage impact of occupational change are mostly driven by individual occupational switches (the composition effect).

Several other papers distinguish multiple classes of college workers. First, Valletta (2018) tries to explain the stagnation in the high-skill premium since 2000. He distinguishes between high-skilled workers (those with four-years degrees) and very high-skilled workers (those with postgraduate degrees). and looks at national-level factors, including polarization, to explain the slowdown in demand for cognitive skills.⁵ Another important exception that is closely

⁵Lindley and Machin (2016) also distinguish two types of college educated workers: graduates and post-graduates. They document a rise in the postgraduate relative to the graduate wage premium.

related to our paper is Michaels et al. (2014), who look at how ICT diffusion at the sector level has polarized skill demand. More specifically, they look at the wage share of high-skilled (college graduates), medium-skilled (some college) and low-skilled (no college) workers. in the total wage bill at the industry level for 11 countries and find a clear negative impact of ICT diffusion on the wage share of medium-skilled individuals. Our approach differs significantly. We make use of geographical disparities in industry specialization and the destruction of routine jobs within industries at the national level to identify task-biased technological shocks at the local level. We focus on individual-level data, which allows us to study the sources of wage variation (within versus between occupations and industries) among groups with different levels of education.

The rest of the paper is organized as follows. Section 2 presents the data and some stylized facts regarding the evolution of the skill supply (education) and occupational structure over time. The empirical/identification strategy is developed in section 3, while section 4 displays our results. Section 5 concludes.

2 Data and stylized facts

Census/American Community Survey

The main data source we use in this paper is the Census/American Community Survey from IPUMS. We mostly use the 1960, 1970, 1980, 1990, 2000 and 2010 waves for regressions and the harmonized variables provided by IPUMS, which are consistently coded across years.⁶ We construct three samples for the different exercises. For all of them, we restrict the dataset to nonmilitary individuals aged between 16 and 64 years old, following Acemoglu

⁶We also use the 1950 and annual waves from 2005 to 2017 in our discussion of the stylized facts below. IPUMS provides two different samples for 1970: the state sample and the metro sample. We use the state sample for the stylized facts below due to its very slightly higher number of observations, but the two datasets are so similar that the statistics are exactly the same if we use the metro sample. We also use the state sample for the robustness regressions at the state level, and we use the metro sample for the regressions at the commuting zone level, as the smallest identifiable area in the state sample is the state. Both samples are composed of two forms each. We combine both forms for each sample and divide the weights by two for each observation.

and Autor (2011). Skill shares, for instance, are calculated with this working-age population. To calculate industry and occupational shares at the sectoral or geographical level, we also drop nonworking individuals.⁷ Finally, again following Acemoglu and Autor (2011), we also exclude self-employed, part-time and part-year workers from the regression sample.⁸ Our wage variable is the annual labor income of workers during the previous year.⁹ Following Autor and Dorn (2013)’s methodology, we create local labor markets at the CZ level. This method uses probabilistic matching to map substate geographic units in the US Census to CZs. Each observation in each CZ is then weighted according to the fraction of the county group/Public Use Microdata Area (PUMA) that maps to the given CZ. There are 722 CZs in the mainland USA and 741 when Alaska and Hawaii are included.¹⁰ Finally, to interact industry shares at the CZ level with industry routine shares at the national level for each census year, we need to have a consistent industry classification over time. Based on the major industry categories, we unify the classification and end up with 34 industries. Following Goldsmith-Pinkham et al. (2018), we separate industries with a high Rotemberg weight in the overall instrument and end up with 43 industries. (See Appendix A for a detailed discussion on the Bartik shift-share instrument.)¹¹ Finally, we also complement the census data with the IPUMS March Current Population Survey, mainly to derive the stylized facts documented below.¹²

Measure of routineness

There are two methods generally found in the literature to classify occupations according to their degree of routineness. The first method is to proxy for job tasks by directly working

⁷Our results are robust to calculating the local shares of routine workers over the full working-age population. (See the online appendix.)

⁸Our results are robust to the inclusion of part-time, part-year or self-employed workers. (See the online appendix.)

⁹According to IPUMS, this income is composed of “wages, salaries, commissions, cash bonuses, tips, and other money income received from an employer. Payments-in-kind or reimbursements for business expenses are not included”.

¹⁰To construct CZs for 1960, we use a similar methodology to that provided by Rose (2018).

¹¹Our results are robust to using the original classification with 34 industries. (See the online appendix.)

¹²A description of this dataset is presented in the online appendix.

with broad occupational categories (see Acemoglu and Autor (2011) or Verdugo and Allegre (2020) among others). Another standard method, used, for instance, by Autor and Dorn (2013), is to use data from the Dictionary of Occupational Titles (DOT) or Occupational Information Network (O*NET).¹³

As our main measure of routineness, we rely on the first method and use the broad occupational groupings defined by Acemoglu and Autor (2011): managerial, professional and technical occupations are specialized in abstract, nonroutine cognitive tasks; sales, clerical and administrative support occupations are specialized in routine cognitive tasks; production, craft, repair, and operative occupations are specialized in routine manual tasks; and service occupations are specialized in nonroutine manual tasks. As is usually done in the literature, we group the routine manual and routine cognitive occupations into a single category. The main advantage of this measure is that it does not rely on a rigid task-based classification of occupations made in 1968, as is the case for the DOT. It is very likely that the task content of occupations was very different at that time from what it is today. Relying on the DOT also requires adopting criteria to weight the different task scores of each occupation. Nevertheless, in several robustness checks, we also use the classification of occupations based on the DOT using the criteria of Autor and Dorn (2013).¹⁴

Skills

Skill levels are coded using the education information available in the database. Educational attainment is measured by the highest year of school or degree completed. In this paper, the low-skilled category refers to individuals who have completed or dropped out before completing high school. The medium-skilled category corresponds to individuals who went

¹³These databases give the task content of each occupation. Each task of each occupation receives a score between 0 and 5 depending on its importance for the occupation. Then, those task scores are aggregated by occupation into three categories: routine, manual and abstract. Next, for each occupation, a routine task intensity measure, equal to $\ln(\text{Routine score}) - \ln(\text{Manual score}) - \ln(\text{Abstract score})$, is computed. Finally, routine occupations are defined as occupations belonging in the top third of the routine task intensity distribution. Goos et al. (2014) show that the two measures are highly correlated.

¹⁴We directly use the data from Autor and Dorn (2013), which classify each occupation as routine, manual or abstract and are available on the AEA website.

to college but completed fewer than four years. Finally, the high-skilled category corresponds to individuals who have at least a four-year college degree. Overall, in 2017, the three skill groups represented a sizable share of the population: the medium- and high-skilled groups represented 30.8% and 28.7% of the working-age population, respectively. For a more detailed analysis, we also divide each skill group into two categories and end up with six categories: lower low-skilled workers, who do not have a high school degree; upper low-skilled workers, who have a high school degree; lower medium-skilled workers, who have completed fewer than three years of college; upper medium-skilled workers, who have completed three years of college; lower high-skilled workers, who have completed four years of college; and upper high-skilled workers, who have completed more than four years of college.¹⁵

Stylized facts: occupations and skills

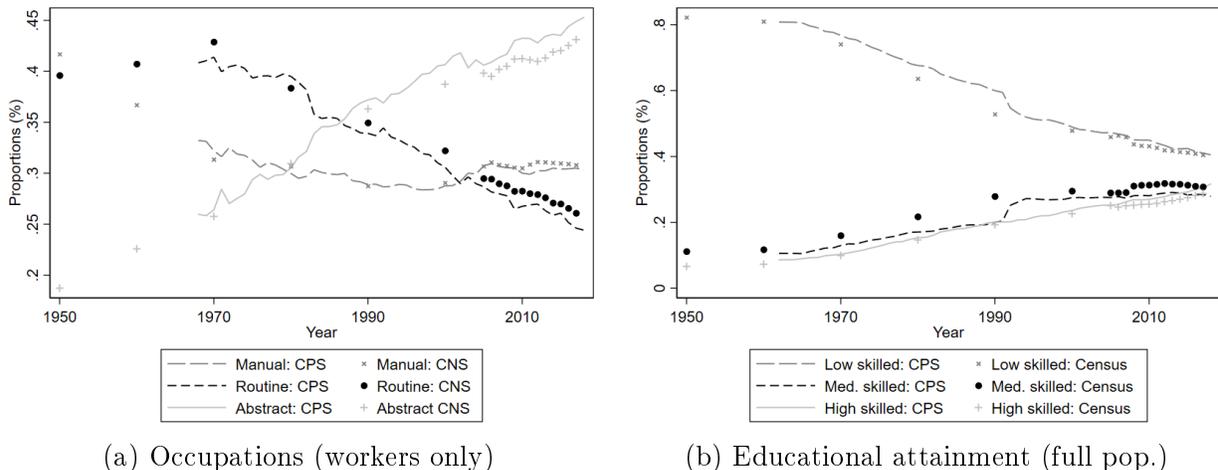
In this section, we document both the global evolution of occupations and skills (Figure 1) and the evolution of occupations for workers in each skill group (Figure 2).

The evolution of workers' broad occupations is depicted in Figure 1a. As is well known in the literature, the proportion of routine occupations has sharply decreased. In the US, it went from 42.9% in 1970 to 26.1% in 2017. These occupations have been replaced by abstract occupations and, to a lesser extent, by manual ones. Another important well-known fact is the sharp increase in the share of the working-age population that has been to college, as shown in Figure 1b. This share goes from 25.9% in 1970 to 59.5% in 2017. More precisely, the proportion of individuals who have completed a four-year program (high skilled) rose from 9.9% in 1970 to 28.7% in 2017, and the proportion of individuals with fewer than four years of college education (medium skilled) went from 15.9% in 1970 to 30.8% in 2017.

To clarify the reallocation pattern for individuals working in routine occupations in the 1970s, we next show that the share of low- and medium-skilled workers in routine jobs substantially decreased over the 1970-2017 period. First, Figure 2 shows that 45.1% (47.8%)

¹⁵These categories correspond to IPUMS codes 1 to 50, 60 to 64, 65 to 71, 80 to 90, 100 to 101 and 110 to 116, respectively.

Figure 1: Evolution of workers' occupations and educational attainment for the overall population



Sources: March CPS data for years 1963-2017, Census IPUMS 5% samples for 1950, 1960, 1980, 1990 and 2000, Census IPUMS 1% state sample for 1970 and Census American Community Survey for 2005-2017. The data used in figure (b) include all persons aged 16-64, excluding those employed by the military. In figure (a), non-working individuals are also excluded.

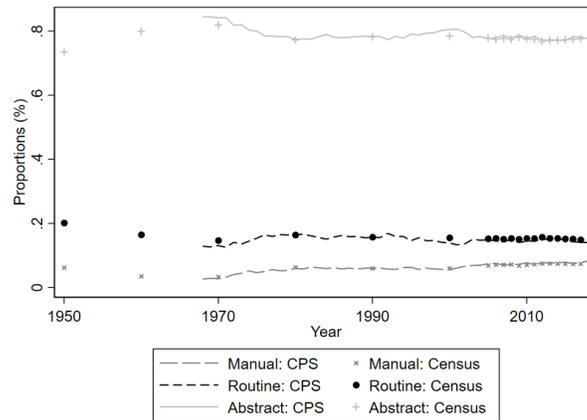
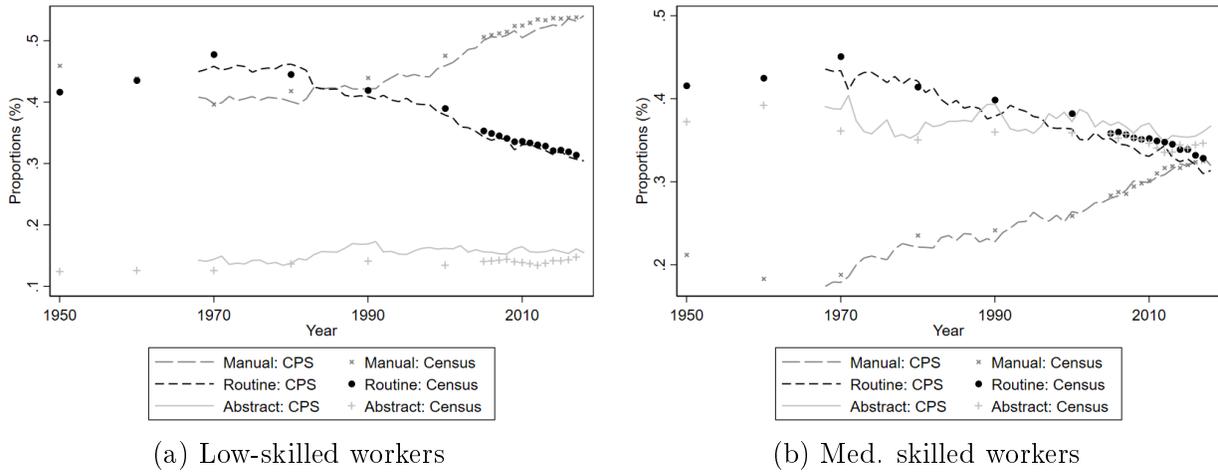
of medium-skilled (low-skilled) workers had a routine occupation in 1970 but only 32.8% (31.4%) did so in 2017. Second, the share of abstract occupation holders did not increase in this group (and even slightly decreased from 36.1% in 1970 to 34.6% in 2017 for medium-skilled workers) despite the strong increase in the proportion of abstract jobs in the economy. Conversely, the share of manual occupation holders among these individuals rose substantially, from 18.8% in 1970 to 32.5% in 2017 for medium-skilled workers and from 39.7% to 53.8% for low-skilled workers, suggesting an occupational downgrading pattern for those skill groups.¹⁶ At the same time, the occupational structure of high-skilled workers has been very stable over time: approximately 80% of these individuals hold abstract jobs.¹⁷

These stylized facts indicate that due to the sharp increase in the supply of high-skilled workers, competition for abstract jobs has increased. As a consequence, medium-skilled workers who have lost their routine jobs have tended to reallocate into manual occupations

¹⁶In the online appendix, we also provide graphs of the evolution of broad occupation shares among the full working-age population.

¹⁷In the online appendix, we also provide graphs that display the shares of each skill group within each occupational category.

Figure 2: Evolution of occupations' shares for low, medium, and high-skilled workers



Sources: See Figure 1(a).

and have been unable to reallocate into abstract jobs. Given those patterns, we expect this change in the occupational structure to have negatively affected low- and medium-skilled workers, first, because of occupational (and possibly industry) downgrading, which pushes low- and medium-skilled workers into lower-ranked occupations (and industries), and second, because of the general equilibrium effect, which should affect wages within the remaining occupations.

3 Identification strategy

To estimate the impact of the decrease in the proportion of routine occupations on the wages of low-, medium- and high-skilled workers, we adopt a local labor market approach. As Census data are individual level data whereas our variable of interest is at the CZ level, we proceed in two steps. In the first step, we regress individual annual log wages ω_i on individual demographic characteristics for each year y and education group g separately:

$$\omega_{i \in gy} = X_i \beta_{gy} + \epsilon_i \quad (1)$$

where X_i includes age and its square, gender and race dummies.¹⁸

In a second step, we use the average residual over each CZ c for each year and education group $w_c = \frac{1}{N_c} \sum_{i \in c} \hat{\epsilon}_i$.¹⁹ Variations in w_c reflect local variations in average wages net of individual characteristics to take into account both the spatial sorting of workers and the sampling of IPUMS. We regress these variations $\Delta w_{ct} = w_{ct_1} - w_{ct_0}$ on the initial share of routine jobs in the CZ for each skill group separately:

$$\Delta w_{ct} = \gamma \text{Routine}_{ct_0} + \Delta Z_{ct} \theta + \lambda_t + \varepsilon_{ct} \quad (2)$$

where Routine_{ct_0} is the share of routine jobs in CZ c at date t_0 , Z_{ct} is a vector of time-varying area controls such as the shares of the different skills and other local labor market variables, and λ_t is a time fixed effect.

As standard in this literature, we first estimate the model in stacked differences over 10-year periods (see Verdugo and Allegre, 2020; Autor and Dorn, 2013; Acemoglu and Restrepo,

¹⁸Thus, we allow the return of individual characteristics to vary over time and across education groups. Those returns are provided in the online appendix. We weight these first stage regressions by the Census weights (i.e. how many persons in the US population are represented by a given person in an IPUMS sample) in order to ensure their representativeness with respect to the US population.

¹⁹To compute these average residualized wages, each observation is weighted by its Census weight multiplied by the CZ specific weight (i.e. the fraction of the county group/PUMA that maps to this given CZ, as mentioned in section 2). This ensures that the local variations in average wages for the IPUMS sample reflect the variations for the US workers in each CZ.

2020b). Thus, t refers to decades 70's 80's 90's and 00's and Δw_{ct} correspond to wage variations over the decades. We also consider long run differences over the 1970-2010 period or stacked differences over 20-year periods. As in Verdugo and Allegre (2020) or Autor and Dorn (2013), we consider as our main regressor the initial share of routine jobs in t_0 of each decade t , which is a very strong predictor of the decrease in routine jobs over the period. It can be interpreted as the exposition of a CZ to task-biased technological change and to the decrease in routine occupations.²⁰

γ , our main coefficient of interest, is then identified with the time variations of the local shares of routine jobs. The variations in the proportions of medium- and high-skilled individuals at the local level are introduced to control for the supply factors of skills as it could affect the local equilibrium wages of the different skill groups and the occupational structure of the workforce. As often done in the literature, we also control for the variations in the proportion of foreign-born individuals in a CZ as those individuals may have lower bargaining power and often work in the manual service sector (see Autor and Dorn, 2013).

Despite the introduction of individual and local control variables in the first- and second-stage regressions, respectively, the initial share of routine occupations might still be endogenous. For instance, the rate of routine job destruction and then the proportion of routine jobs could respond to the wage adjustment. If wages sufficiently adjust downward, this could limit the technology induced destruction of routine jobs and our OLS coefficient could be downward biased. Also, it is still possible that an unobserved/omitted variable affects simultaneously wages and occupation composition at the local level.²¹ As a consequence, we instrument the initial proportion of routine jobs in the CZ ($Routine_{ct_0}$) by the following arguably exogenous measure of exposure to task-biased technological change, based on the

²⁰As a robustness, we also directly consider the variation in the share of routine occupations as our main regressor. Results are qualitatively similar across specifications.

²¹For instance, a tax policy which affects more routine intensive sector than other sectors.

past industrial specialization of each CZ:

$$Exposure_{ct} = \sum_{j=1}^N \varphi_{cj} Routine_{jt-1} \quad (3)$$

where φ_{cj} is the employment share of sector j in area c in 1960 and $Routine_{jt-1}$ is the 10-year lag of period t initial share of routine occupations at the national level in sector j .²² $Routine_{jt-1}$ summarizes a variety of shocks that affect the proportion of routine jobs at the sector level and which should be only marginally related to what happens in a given area. This should be the result of task-biased technological change or the possibility to offshore some tasks abroad. φ_{cj} summarizes how those shocks at the national level translate into shocks at the local level. We follow the literature (see for instance Verdugo and Allegre, 2020; Autor and Dorn, 2013) and consider the employment shares φ_{cj} ten years before the start of our first period as it should better satisfy their exogeneity with respect to the evolution of wages.²³ For the national sector shares of routine occupations, we consider a 10-year lag of period t initial shares (t_{-1}) as Autor and Dorn (2013).²⁴

Following Goldsmith-Pinkham et al. (2018), we discuss extensively the variations that drive the identification of our coefficient of interest and the exogeneity of the instrument in Appendix A. As in most empirical settings using Bartik shift-share instruments, our identification relies on the industry shares rather than on specific shocks at the sector level given that automation and ICT affected many sectors at the same time.²⁵ It is thus important that the initial industry shares are exogenous to the local development of CZ that affects the wage rate of individuals. In some IV regressions, we then control for the variations in the

²²Alternatively, we estimate a reduced form model using our exposure variable directly as a regressor: $\Delta w_{ct} = \gamma Exposure_{ct} + Z_{ct} \theta + \lambda_t + \varepsilon_{ct}$, for each skill group separately.

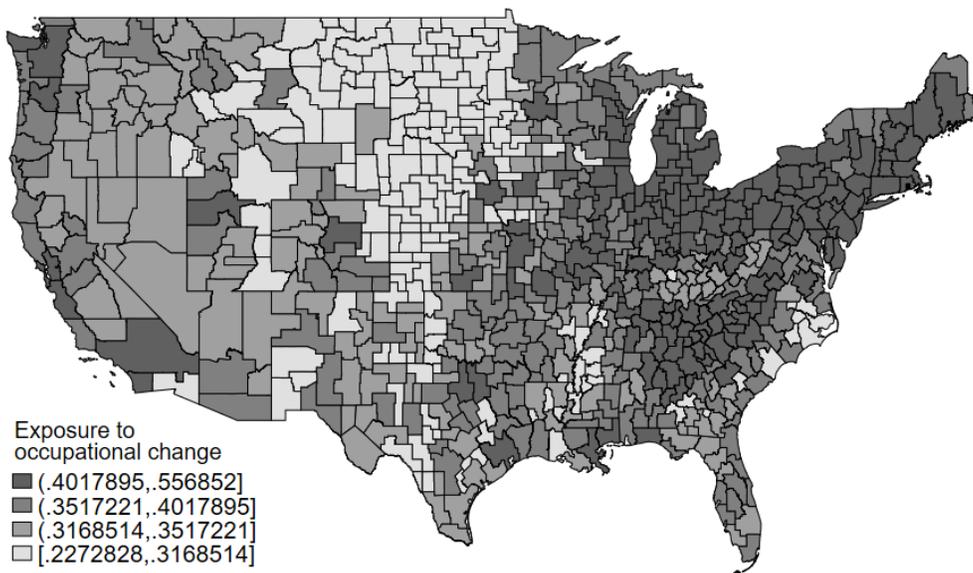
²³Note however that many papers that use shift-share (Bartik) instrument consider the shares at the initial period of their sample. In our case, it would imply to use the local sector shares of 1970. We provide estimates using the 1970 shares in the online appendix. Results are unchanged.

²⁴Results in the online appendix show robustness to the use of national sector shares of routine jobs in t_0 instead of t_{-1} .

²⁵Appendix A shows that for the main specification, only 1% (0.107²) of the identification relies on sector shocks of the share of routine jobs. In some specifications presented in the paper, sector shocks never account for more than 16% of the identifying variations.

shares of the broad manufacturing and service sectors, as suggested by Goldsmith-Pinkham et al. (2018), given that the initial industry shares in our instrument could be endogenous to the variations in the aggregate sectoral composition in CZ which could affect the wage rate. Autor and Dorn (2013) also control for the share of the manufacturing sector. Results are robust to the inclusion of these local variations of aggregate shares as additional controls.^{26,27}

Figure 3: Distribution of the instrument at the CZ level, 1970



Source: Census IPUMS 1% metro sample for 1970. The value of the Bartik shift-share instrument is computed at the commuting zone \times year level using 43 industries.

Figure 3 plots $Exposure_{ct}$ for 1970, the beginning of our estimation period. We can see that CZs had very different expositions to shocks that lead to a decrease in the proportion of routine jobs given their initial industry specialization.

To study the mechanisms at work behind the impact of the routine-biased technological change on wages by education group, we then decompose the wage variations into a between and a within components. The between component corresponds to the composition effect induced by the switch of workers from occupations that disappear to some other occupations

²⁶In the online appendix, we also control for the evolution of industry shares using a more detailed classification of 13 industries.

²⁷For $Exposure_{ct}$ to be a valid instrument, the 10-year lag of period t initial proportion of routine jobs in an industry at the national level, $Routine_{jt-1}$, should also be exogenous to local development of CZ. However, as stated earlier, most of the identification relies on sector shares rather than sector shocks.

which have a different wage rate. The within term, at the heart of the displacement effect of Acemoglu and Restrepo (2018), corresponds to the wage impact within occupations due to higher labor supply in remaining occupations.

More specifically the variation over period t of the average residualized wage in CZ c can be decomposed as follows:

$$\Delta w_{ct} = \sum_j ((\varphi_{cjt_0} + \varphi_{cjt_1})/2)(w_{cjt_1} - w_{cjt_0}) + ((w_{cjt_1} + w_{cjt_0})/2)(\varphi_{cjt_1} - \varphi_{cjt_0}) \quad (4)$$

where t_0 correspond to the initial date and t_1 the end of each period t . φ_{cjt_0} and φ_{cjt_1} are the local shares of employment in cell j in t_0 and t_1 . w_{cjt_1} and w_{cjt_0} are the average residualized wages in CZ c and cell j in t_0 and t_1 . Depending on the chosen decomposition, j can index occupations or occupation \times industry cells. The first component is the within term. It corresponds to the variation of average wages within each occupation and CZ, weighted by the share of this occupation in the CZ (averaged over t_0 and t_1). The second component is the between term and corresponds to changes in the local shares of occupations, weighted by the average residualized wage of each occupation (averaged over t_0 and t_1). It is the composition effect and it captures average wage variations due to occupation (and possibly industry) switches. We first consider a decomposition based on 4 occupation categories (routine, abstract, manual and not-classified) in order to capture the pure occupational downgrading effect described in the literature. We then consider a decomposition based on both 4 occupations and 13 industries.²⁸

²⁸ $4 \times 13 = 52$ bins. We restrict ourselves to few occupation and industry categories in the main specification because, for the decomposition to hold exactly, we need an average residualized wage for each bin in each CZ at each date. With a too detailed classification, we end up with many occupations(\times industry) \times CZ bins with no observation. As robustness we propose an alternative decomposition based on 77 occupations. In order to perform this robustness, we create a new classification of occupations. Similarly to the unification of the industry classification, we unify the occupation classification based on the major categories given by IPUMS. The only exception are occupation ‘258: Sales Engineers’ of the source classification, which has been separated from the ‘Sales Representatives, Commodities’ broad occupation because it is an abstract occupation, while the others are routine occupations. In addition, we combined ‘243: Supervisors and proprietors, sales occupations’ with ‘Sales Representatives, Finance and Business Services’, ‘415: Supervisors of guards’ with ‘Guards’, ‘503: Supervisors, mechanics and repairers’ with ‘Vehicle and Mobile Equipment Mechanics and Repairers’, and ‘558: Supervisors, construction’ with ‘Construction Trades’ for consistency

In some regressions, we also estimate the impact of the decrease in the share of routine occupations separately for each broad occupational category, to see if the wage impact within occupations is observed in each broad manual or routine occupation, in particular within the manual service sector which is at the heart of the polarization theory of Autor and Dorn (2013). In their analysis, the wage rate increases in the manual service sector, at the bottom of the wage distribution, because of complementarities with abstract jobs. The displacement effect of recent works modeling the task-biased technological change of Acemoglu and Restrepo (2018) goes in the opposite direction.

All the regressions are clustered at the State level and weighted by the share of CZ in the full population in t_0 .

Additionally, we also provide in Appendix C individual-level regressions which are consistent with our CZ-level regressions. We directly estimate a wage equation using individual log wages as dependant variable, which we regress on individual characteristics, CZ time-varying controls and the proportion of routine jobs at the CZ level. Regressions include CZ and year fixed effects. In some of these regressions, we also include occupation or occupation \times industry fixed effects to disentangle the mechanisms at work (composition and within effects). Whereas this strategy using directly individual wage data is slightly different, the identification still relies on spatial differences in variations of the share of routine occupations over time and results are qualitatively similar.²⁹ More details are provided in Appendix C.

Several robustness checks are provided in the online appendix (Bosquet et al., 2020): As mentioned before, our results are robust to calculating the local shares of routine workers over the full working-age population, to the inclusion of part-time, part-year or self-employed workers, to an alternative definition of routine occupations, and to various specifications of the instrument (using the original industry classification with 34 posts instead of 43, using the local sector shares of 1970 instead of 1960, or using the national sector shares of routine

over time. We end up with 77 occupations but we use only 76 of them as we exclude those in the military.

²⁹As the Census data does not follow individual over time, we do not estimate the model in first difference as in our main specification using data aggregated at the CZ level.

jobs in t_0 instead of t_{-1}). The online appendices also show that our results are robust to using different sets of weights (population count instead of population shares or using population fixed in 1970), with controls in levels instead of variations, including state fixed effects, running state-level regressions, and including more controls: the share of individuals born in another state or population density. In online appendix Table O20, we also re-estimate Table 1 changing the end date to 2013 instead of 2010 for the last period stacked difference (2000-2013) and also for the long differences estimates (1970-2013). Indeed, in 2010, the economy had not fully recovered from the great recession of 2008 which could have affected the proportion of routine jobs as well as wages of the different education groups. Finally, the last set of robustness checks requires to compute new residuals from equation 1: using pooled instead of year-specific regressions, trimming top and bottom 2.5% of wages, including the workers for which occupations or industries are missing and using CPS data.

4 Results

4.1 Baseline results

Table 1 presents our baseline results for the three education groups and three specifications: OLS, IV and IV with controls for local variation in broad sector shares.³⁰ All the regressions include time fixed effects as well as control variables at the CZ level. (Individual characteristics are controlled for in the first stage.) They all give qualitatively similar results of the impact of a reduction in the proportion of routine jobs on the wage structure in local labor markets. Only the magnitude differs between the OLS and IV estimates.

The impact of the local initial share of routine occupations is significantly negative for low- and medium-skilled workers in the OLS and IV estimates, meaning that a higher exposure to routine occupations decreases the wages of low- and medium-skilled workers. The estimated coefficients are three times larger in the IV than in the OLS specification, which

³⁰First-stage regressions of the 2SLS estimator are available in the online appendix.

Table 1: First-difference CZ average residualized wages on the local share of routine jobs, by skill levels, OLS and IV, stacked differences

Workers educ.:	OLS			IV					
	low	medium	high	low	medium	high	low	medium	high
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sh. of rout. occs ₀	-0.196 ^b (0.075)	-0.112 ^b (0.046)	-0.034 (0.047)	-0.581 ^a (0.114)	-0.437 ^a (0.089)	-0.167 ^b (0.080)	-0.496 ^a (0.149)	-0.372 ^a (0.124)	-0.046 (0.087)
Δ sh. of high skill	0.583 ^a (0.170)	0.657 ^a (0.135)	0.981 ^a (0.092)	0.749 ^a (0.195)	0.797 ^a (0.164)	1.039 ^a (0.111)	0.785 ^a (0.195)	0.835 ^a (0.167)	1.086 ^a (0.114)
Δ sh. of med. skill	-0.353 ^c (0.179)	-0.629 ^a (0.164)	-0.394 ^a (0.132)	-0.235 (0.166)	-0.529 ^a (0.150)	-0.353 ^a (0.127)	-0.180 (0.171)	-0.411 ^b (0.154)	-0.308 ^b (0.136)
Δ sh. of foreign	-0.386 ^a (0.066)	0.077 (0.101)	0.328 ^a (0.059)	-0.394 ^a (0.059)	0.070 (0.109)	0.325 ^a (0.061)	-0.400 ^a (0.059)	0.055 (0.098)	0.322 ^a (0.056)
Δ sh. of manuf.							0.104 (0.298)	-0.133 (0.224)	0.249 (0.153)
Δ sh. of services							-0.149 (0.260)	-0.449 ^b (0.192)	-0.056 (0.144)
R ²	0.05	0.09	0.20						
Kleibergen-Paap				323	323	323	183	183	183

Notes: Standard errors clustered by state between brackets. ^a, ^b, ^c indicate significance at the 1%, 5% and 10% level, respectively. Each regression is weighted by the share of the CZ in the national population. 2888 observations (722 CZ × 4 periods).

suggests that the OLS results are downward biased. This could be because the proportion of routine jobs reacts to the wage adjustment or because there is a potential measurement error problem. Reductions in the wage rate of low- and medium-skilled workers could prevent further destruction of routine jobs. The Kleibergen-Paap statistics are well above the threshold values provided by Stock and Yogo (2005), suggesting that exposure to routine occupations is a strong instrument for the local proportion of routine jobs. Columns 7-9 control for variation in the shares of the manufacturing (as in Autor and Dorn, 2013) and service sectors. These controls only marginally affect the results.³¹

The impact seems to be slightly higher for low-skilled than for medium-skilled workers. Given that the differences are neither large nor statistically significant, we conclude that low- and medium-skilled workers are affected in similar proportion by the disappearance of routine jobs. The effects are economically sizeable. According to the IV estimates, a 5.7

³¹In the online appendix, we also control for the evolution of industry shares using a more detailed classification of 13 industries.

percentage point increase in the initial proportion of routine jobs (which corresponds to one standard deviation in the distribution observed at the CZ level for the 1970-2010 period) implies a decrease in average wages of 3.3% for low-skilled and 2.5% for medium-skilled workers over a 10-year period.³² If we consider the variation in the initial proportion of routine jobs from the first to the top decile of CZs, i.e., 14.6 percentage points, the wages of low- and medium-skilled workers decrease by 8.1% and 6.2% over ten years on average.³³

High-skilled workers do not seem affected. The coefficient on the initial share of routine jobs is not significant and close to zero in the OLS regression (column 3) and marginally negative and significant in the IV estimates (column 6), but the effect totally vanishes when we include the industry shares (column 9). This is in line with the literature on task-biased technological change. The occupations in which high-skilled workers are concentrated (abstract) do not disappear. Only a few high-skilled individuals have routine occupations, and medium-skilled workers do not reallocate into abstract occupations and compete for them with high-skilled individuals. As a consequence, high-skilled individuals do not suffer the displacement effect of the disappearance of routine occupations.

Table B1 presents the estimates of the baseline IV model dividing each skill group into two subcategories. Two interesting patterns emerge. First, the lower low-skilled category seems to be the most affected (coefficient for routine jobs at -0.620). Then the lower medium-skilled (-0.358), higher medium-skilled (-0.367) and higher low-skilled (-0.368) groups are all affected in the same proportion.

Many alternative specifications to estimate the impact of local shocks can be found in the literature. We provide several additional robustness checks in the appendix. Table B2 shows the results with 20-year stacked differences and 40-year long differences (1970-2010) instead of our baseline 10-year stacked differences. Table B3 shows the results based on variation in the proportion of routine jobs instead of the initial level. Table B4 shows the reduced-form estimates directly using the instrument as a regressor. All these results are

³² $\exp(-0.581 \times 0.057) - 1$ and $\exp(-0.437 \times 0.057) - 1$.

³³ $\exp(-0.581 \times 0.146) - 1$ and $\exp(-0.437 \times 0.146) - 1$.

qualitatively very similar to those of our main specification. Finally, Appendix C shows the results of the individual-level regressions, which are remarkably consistent with those of the two-stage aggregate specification.

Moreover, Table B5 reproduces Table 1 but excludes the last decade 2000-2010 from the stacked differences. Interestingly, the impact for medium-skilled workers becomes lower, decreasing by 39%. The impact for low-skilled workers remains very similar. This suggests that task-biased technological change has become more detrimental to medium-skilled workers in the more recent period. This could be because technological change has increasingly affected complex occupations over time or because within an occupation, low-skilled workers are affected first, and then medium-skilled workers see an effect.

4.2 Mechanisms

4.2.1 Within vs. between effects

In this section, we try to understand the channels through which the occupational structure affects the wages of workers with different education levels. A common narrative in the polarization literature is that low-skilled workers are affected due to occupational downgrading, as routine jobs pay higher wages than manual jobs (see, for instance, Autor, 2019). This reasoning should extend to medium-skilled workers, as they do not seem to reallocate into abstract occupations but into manual ones. As explained earlier, industry switches can also have a wage impact for workers. On the other hand, the existence of the displacement effect highlighted by Acemoglu and Restrepo (2018) suggests that wages for low- and medium-skilled workers could decrease within occupations.

To disentangle the two effects, we perform regressions on the within and between terms of the decomposition presented in equation (4). The results are displayed in Table 2 for the IV specification. Panel A displays the results for the decomposition based on occupations only, and panel B displays the results for the decomposition based on occupation \times industry bins. Columns 1-3 and 4-6 display the results for low- and medium-skilled workers, respectively.

For each skill group, the first column corresponds to the total effect (i.e., the corresponding coefficient from Table 1), the second column to the between effect and the third column to the within effect.

Table 2: Decomposition of the IV coefficients from Table 1, columns 4 and 5

Workers educ.:	low			medium		
Effect:	total	between	within	total	between	within
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: based on occupations only					
Sh. of rout. occs ₀	-0.581 ^a (0.114)	-0.018 ^c (0.010)	-0.563 ^a (0.107)	-0.437 ^a (0.089)	-0.048 ^a (0.011)	-0.390 ^a (0.082)
Δ sh. of high skill	0.749 ^a (0.195)	0.025 (0.016)	0.724 ^a (0.189)	0.797 ^a (0.164)	-0.036 ^b (0.016)	0.833 ^a (0.162)
Δ sh. of med. skill	-0.235 (0.166)	0.089 ^a (0.014)	-0.325 ^c (0.171)	-0.529 ^a (0.150)	0.008 (0.008)	-0.537 ^a (0.148)
Δ sh. of foreign	-0.394 ^a (0.059)	-0.054 ^a (0.008)	-0.341 ^a (0.061)	0.070 (0.109)	-0.002 (0.018)	0.072 (0.093)
	Panel B: based on occupations×industries					
Sh. of rout. occs ₀	-0.581 ^a (0.114)	-0.072 ^a (0.024)	-0.503 ^a (0.098)	-0.437 ^a (0.089)	-0.096 ^a (0.016)	-0.342 ^a (0.077)
Δ sh. of high skill	0.749 ^a (0.195)	0.047 (0.034)	0.705 ^a (0.174)	0.797 ^a (0.164)	-0.009 (0.032)	0.789 ^a (0.149)
Δ sh. of med. skill	-0.235 (0.166)	0.029 (0.021)	-0.278 ^c (0.164)	-0.529 ^a (0.150)	-0.074 ^a (0.018)	-0.471 ^a (0.142)
Δ sh. of foreign	-0.394 ^a (0.059)	-0.056 ^b (0.024)	-0.340 ^a (0.065)	0.070 (0.109)	-0.014 (0.014)	0.073 (0.095)

Notes: See Table 1. Occupation classification is composed of 4 posts: manual, routine, abstract and not-classified. Industry classification composed of 13 posts, corresponding to most aggregated industries in the 1990 Census Bureau industrial classification. The Kleibergen-Paap rk Wald F statistic is equal to 323 in all regressions.

Considering occupations alone (panel A), the between term accounts for a very small share of the overall effect for the low- and medium-skilled groups (3.1% and 11.0%, respectively).³⁴ Most of the wage impact of the disappearance of routine occupations for each skill group occurs within occupations. This is at odds with the traditional narrative of the wage impact of polarization on low-skilled individuals. Indeed, the literature often posits that these workers are mainly affected by occupational downgrading, switching from well-paid routine jobs to low-paid manual occupations.

³⁴0.018/0.581 = 0.0309 and 0.048/0.437 = 0.1098.

Table 3: Individual regressions of residualized wages, OLS

Workers educ.:	low		medium		high	
	(1)	(2)	(3)	(4)	(5)	(6)
Routine	0.122 ^a (0.003)	0.114 ^a (0.004)	0.088 ^a (0.004)	0.081 ^a (0.004)	0.281 ^a (0.007)	0.264 ^a (0.008)
Abstract	0.301 ^a (0.007)	0.293 ^a (0.008)	0.295 ^a (0.006)	0.284 ^a (0.006)	0.531 ^a (0.013)	0.515 ^a (0.014)
CZ FE	no	yes	no	yes	no	yes

Notes: All regressions include year fixed effects. 722 CZ fixed effects are also included in some regressions, as indicated in the bottom of the table. Standard errors clustered by commuting zone between brackets. ^a, ^b, ^c indicate significance at the 1%, 5% and 10% level, respectively. All regressions are weighted by individual Census weights multiplied by CZ specific weights. The numbers of observations for each skill are 9,110,818, 4,831,197 and 3,901,517, respectively.

This result is not so surprising when we take a closer look at the data. Table 3 shows estimates of the premium from having a routine or an abstract occupation, compared to having a manual occupation, for each education group. The residuals of our individual-level first-stage wage regressions (equation 1), which correspond to wages net of individual characteristics, are regressed on dummies for routine and abstract occupations. The coefficients of Table 3 thus capture the mean differences in log wages between broad occupational categories.³⁵ All regressions include year fixed effects, and including or not including CZ fixed effects does not make much difference. The routine occupation premium is quite small for both low- and medium-skilled workers (columns 1-4). It is equal to 11.4% and 8.1% according to columns 2 and 4, respectively. This is quite low in comparison to wage dispersion in the US or to the abstract premium, which is close to 30% for low- and medium-skilled workers and greater than 50% for high-skilled workers, as reported in Table 3.³⁶ The routine occupation premium. is also much lower than the medium (23.7%) and high (61.4%) skill premiums over the 1970-2010 period.³⁷ This finding confirms the important result of Hunt and Nunn (2019), who also contradicts the polarization narrative by showing that

³⁵Full individual wage regressions with both individual controls and occupational dummies yield very similar results.

³⁶The difference between the lowest and highest wage deciles corresponds to 155% during the 1970-2010 period.

³⁷These premiums are obtained from a simple wage regression with controls for the same individual demographic characteristics.

occupations are a poor determinant of wage differences, given that wage dispersion within occupations is much larger than the average wage differences across occupations. It is thus not surprising to find that the between effect, related to occupational switching, is quite modest in the overall impact. Most of the wage impact occurs within occupations. This is in line with the displacement effect in Acemoglu and Restrepo (2019)’s model of automation in which destruction of occupations due to technical change displaces labor to other existing occupations, which creates downward pressure on wages. The between effect also exists in their model but depends on the wage differential between the occupations that are destroyed and the occupations into which workers reallocate. This differential is low in our case, which explains the modest role of the between effect. Online appendix Figure O3 shows that the routine premiums tend to decrease over time.

When we take into account industries in the decomposition (occupation×industry bins), panel B of Table 2 shows that the magnitude of the between effect significantly increases to 12.4% of the total impact for the low-skilled group and 22.0% for the medium-skilled group.³⁸ This is consistent with the finding of Barany and Siegel (2018), who suggest that occupational structure changes are related to industrial shocks and structural transformation. The decrease in the proportion of routine occupations goes together with a reallocation of individuals across industries (within occupations). This has a negative wage impact, which means that in the structural change process that makes routine jobs disappear, many individuals do not experience occupational change but rather switch to industries that pay less.

Appendix Table B6 reproduces the results of Table 2 panel A with 77 occupations.³⁹ The between effect accounts for 9.1% of the total effect for low-skilled and 18.5% for medium-

³⁸ $0.072/0.581 = 0.1239$ and $0.096/0.437 = 0.2196$.

³⁹We prefer to keep the more aggregated classification in the main text for two reasons. First, the literature on polarization distinguishes three broad occupational groups (manual, routine and abstract) and argues that reallocation occurs across those three groups. Second, a more detailed classification makes the decomposition more imprecise since it increases the number of occupation×industry×CZ bins that can be empty. See section 3.

skilled workers.⁴⁰ The between effect is more important than the within effect in the 4 occupation decomposition. This suggests that some reallocation occurs within the broad occupation categories.⁴¹

Appendix Table C2 distinguishes the between and within effects directly using micro data. To do so, we introduce occupation, industry and industry×occupation fixed effects in individual wage regressions. This strategy allows us to control for very detailed occupation (77 posts) and industry (43 posts) classifications. The between effect accounts for 13.8% (occupations only) or 24.6% (industry×occupation) of the total wage impact for the low-skilled group and 17.4% (occupations only) or 29.8% (industry×occupation) for the medium-skilled group.⁴² The results are consistent with those of the aggregate specification: most of the wage impact occurs within occupations.

The results suggest that the most important part of the wage impact from the destruction of routine occupations occurs within occupations and industries. This within effect accounts for more than 70% of the total impact for low- and medium-skilled workers when we look at the detailed classification in the micro regressions and between 80% and 90% when we look at broader occupation categories. This is very consistent with the task-biased technological change and automation model of Acemoglu and Restrepo (2018) and their displacement effect. The traditional occupational downgrading narrative explaining the wage impact of polarization on non-college educated workers and the industry composition effects are still in evidence in the data but are clearly not the main drivers of the wage impact. The fact that occupational downgrading also threatens medium-skilled individuals (and in even higher proportion than low-skilled workers) is consistent with the stylized facts previously discussed: the proportion of medium-skilled workers holding a routine occupation has sharply decreased and has been entirely compensated for by an increase in manual occupations among

⁴⁰ $0.053/0.581 = 0.0912$ and $0.081/0.437 = 0.1853$.

⁴¹For instance, the share of low-paid occupations could increase within the routine category. The same kind of phenomenon could occur within the manual occupation category.

⁴² $1 - (1.880/2.182) = 0.1384$, $1 - (1.645/2.182) = 0.2461$, $1 - (1.626/1.970) = 0.1746$, $1 - (1.382/1.970) = 0.2984$.

this population. The proportion of medium-skilled workers who have an abstract job has even slightly decreased, suggesting occupational downgrading for this population rather than upgrading. A quite similar pattern is observed for the low-skilled group, but the wage impact of this occupational shifting is much lower than that for the medium-skilled group.

4.2.2 Wage impact on each occupational category

This section tests whether the large within impact is observed in both manual and routine occupations, consistent with the task-biased technological change literature and the displacement effect. The important within impact could hide some disparities across occupations. An important prediction at the heart of Autor and Dorn (2013) is that wages should increase in manual occupations and, more specifically, manual occupations in services. It is possible that the overall within effect is negative but can remain positive in some specific occupations, notably manual service occupations. We thus perform the previous analysis separately for two manual and three routine broad occupation categories of low- and medium-skilled workers: (1) transport and construction and (2) service occupations (both manual) and (3) production, craft and repair, (4) operative and (5) clerical and administrative support occupations (routine). We do not report results for abstract occupations since only very few low- and medium-skilled individuals work in these occupations.

Panels A and B of Table 4 report the results for the low- and medium-skilled workers, respectively. The disappearance of routine occupations has a negative impact on wages within each of the five occupational categories, including manual service occupations. The overall negative impact within occupations does not hide important disparities across occupational categories. Again, this finding seems to be highly consistent with the displacement effect of Acemoglu and Restrepo (2018, 2019). The disappearance of routine jobs makes wages decrease in all the occupations into which workers reallocate. Low- and medium-skilled workers have massively reallocated into manual occupations and not into abstract ones. According to the task-biased technological change models of Acemoglu and Restrepo (2018, 2019) and

Table 4: First-difference CZ average residualized wages on the local share of routine jobs, by broad occupation categories, IV, stacked differences

Occupations:	Transport, construction and misc. occ.	Service occ.	Production, craft, repair, occ.	Operative occ.	Clerical and administrative support occ.
	(1)	(2)	(3)	(4)	(5)
Panel A: low-skilled workers					
Sh. of rout. occs ₀	-0.721 ^a (0.128)	-0.516 ^a (0.120)	-0.504 ^a (0.132)	-0.605 ^a (0.109)	-0.337 ^a (0.092)
Δ sh. of high skill	0.768 ^a (0.221)	0.783 ^a (0.231)	0.671 ^a (0.198)	0.636 ^a (0.197)	0.838 ^a (0.180)
Δ sh. of med. skill	-0.264 (0.203)	-0.412 ^b (0.199)	-0.372 ^c (0.199)	-0.378 (0.265)	-0.108 (0.137)
Δ sh. of foreign	-0.517 ^a (0.076)	-0.333 ^a (0.085)	-0.363 ^a (0.078)	-0.669 ^a (0.155)	0.039 (0.082)
Panel B: medium-skilled workers					
Sh. of rout. occs ₀	-0.436 ^a (0.098)	-0.399 ^b (0.152)	-0.223 ^b (0.103)	-0.347 ^a (0.115)	-0.311 ^a (0.084)
Δ sh. of high skill	0.697 ^a (0.198)	0.611 ^b (0.231)	0.299 (0.213)	0.871 ^a (0.205)	0.984 ^a (0.166)
Δ sh. of med. skill	-0.547 ^b (0.244)	-0.757 ^a (0.257)	-0.802 ^a (0.214)	-0.732 ^b (0.320)	-0.394 ^a (0.129)
Δ sh. of foreign	-0.119 (0.131)	-0.158 (0.149)	-0.039 (0.122)	-0.298 ^b (0.146)	0.183 ^b (0.082)

Notes: See Table 1. The Kleibergen-Paap rk Wald F statistic is equal to 323 in all regressions.

their exposition of a displacement effect, wages should decrease within each occupation into which workers reallocate. This is precisely what we find. This is at odds with the traditional polarization literature, which suggests stronger wage growth at the bottom of the wage distribution, i.e., for manual occupations.

In Appendix Table C3, we directly use micro data and an interaction term between the proportion of routine jobs in a CZ and dummy variables for routine, manual and abstract occupations. The results are consistent with those of the aggregate-level regressions and show that the wage impact for low- and medium-skilled workers occurs within all (routine and manual) occupation categories in which they remain.

4.3 Heterogeneous effects by age

In this section, we look at the heterogeneity in the impact of task-biased technological change on wages by age category.

There are several reasons to think that the impact may be more important for young individuals who recently entered the labor market. First, wage rigidities could exist, and the wages of existing employees may have been fixed some time ago. Second, the margin of adjustment may be job creation and the vacancy rate of routine jobs rather than an increase in the job destruction rate. A decrease in the vacancy rate should impact new entrants more than insiders. Finally, individuals who entered the labor market when routine employment was flourishing had time to build human capital and experience while working to climb the occupational ladder and perform less routine-intensive tasks. The conditions when workers enter the labor market should matter. Beaudry et al. (2016) argue that the impact of the slowdown in demand for cognitive tasks should mostly affect younger workers who entered the labor market at the time of the slowdown. These authors' empirical analysis accordingly focuses on young workers.

Table 5: First-difference CZ average residualized wages on the local share of routine jobs, by age categories, IV, stacked differences

Workers educ.:	low			medium		
	≤ 29	30-49	≥ 50	≤ 29	30-49	≥ 50
Age categories:	(1)	(2)	(3)	(4)	(5)	(6)
Sh. of rout. occs ₀	-0.729 ^a (0.152)	-0.522 ^a (0.106)	-0.582 ^a (0.105)	-0.649 ^a (0.116)	-0.361 ^a (0.083)	-0.338 ^a (0.082)
Δ sh. of high skill	0.819 ^a (0.268)	0.745 ^a (0.187)	0.670 ^a (0.177)	0.934 ^a (0.219)	0.786 ^a (0.145)	0.600 ^a (0.156)
Δ sh. of med. skill	-0.527 ^b (0.213)	-0.149 (0.165)	0.072 (0.173)	-0.716 ^a (0.241)	-0.362 ^a (0.125)	-0.268 ^c (0.135)
Δ sh. of foreign	-0.460 ^a (0.118)	-0.349 ^a (0.058)	-0.355 ^a (0.062)	-0.129 (0.183)	0.166 ^c (0.090)	0.168 ^a (0.057)

Notes: See Table 1. The Kleibergen-Paap rk Wald F statistic is equal to 323 in all regressions.

To test for a differentiated impact by age group, we estimate our IV model for each skill group on three different age classes: workers aged 29 years and below, workers aged between

30 and 49 and workers aged between 50 and 64. The results are displayed in Table 5 in columns 1-3 for the low-skilled group and columns 4-6 for the medium-skilled group. Several interesting results emerge. The heterogeneity of the impact is much stronger for medium- than for low-skilled workers. This could be due to the fact that more experienced medium-skilled workers have more stable jobs given their seniority in firms, whereas new entrants are more affected, as they face more adverse labor market conditions. For the low-skilled group, the effect is also heterogeneous but much less so than for the medium-skilled group.

We also split the sample by gender, but there is no significant differences across both subsamples. Men and women appear to be affected in the same proportion in the low- and medium-skilled groups. The results are reported in Table O3 in the online appendix.

5 Conclusion

This paper studies the impact on wages of the massive change in the occupational structure that is occurring in most OECD countries. The most common narrative in the literature is that job polarization should be detrimental to non-college educated workers who hold routine occupations with specific skills (in the middle of the wage distribution) and experience an occupational downgrading to manual occupations that require generic skills. We show that the impact on wages extends beyond non-college educated workers only and that this process also impacts in similar proportion some workers who have been to college and potentially even obtained an intermediate postsecondary degree. Altogether, low- and medium-skilled individuals represent 70% of the US population in 2017. We show that, consistent with the literature on task-biased technological change and the displacement effect highlighted by Acemoglu and Restrepo (2018), more than 70% of this impact occurs within industries and occupations. The rest of the wage loss corresponds to occupational downgrading and some workers (non-college educated workers and workers with intermediate postsecondary education) having to accept work in (manual) occupations and industries that pay less. This

composition effect explains less than 30% of the total impact of routine-biased technological change on wages when we use more detailed categories in micro regressions and less than 10% when we look at the broader occupational categories that match the polarization literature narrative.

Our results may contribute to the literature on the wage stagnation affecting an important share of US workers and on wage–productivity decoupling. It is now well documented that the median wage in the US has not evolved for decades (see, for instance, Pessoa and Van Reenen, 2013). Overall, the labor share of income seems to be decreasing in the US economy (Autor et al., 2020): while median wages tracked productivity until 1973, this trend then stopped. Productivity kept growing, but wages remained stagnant. According to Pessoa and Van Reenen (2013), 40% of this wage decoupling is attributable to inflation measurement errors. Still, it seems that wages have not caught up with productivity. This is a prediction of the task-biased technological change model à la Acemoglu and Restrepo (2018) and the displacement effect that occurs when tasks disappear. We empirically show that the change in the occupational structure of the economy at work in OECD countries is likely to explain part of this phenomenon.

Recently, researchers have also suggested that the Phillips curve may have disappeared. Despite low unemployment in the US, wages do not seem to react accordingly. A potential explanation is the recent modification in the occupational structure, which has accelerated in the decade since the subprime crisis (Jaimovich and Siu, 2020). The combination of low unemployment and stagnant wages may be due to downward wage pressure induced by the task-biased technological change that occurred simultaneously.

References

- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics, Volume 4. Amsterdam: Elsevier-North*, pages 1043–1171.
- Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares and employment. *American Economic Review*, 108(6):1488–1542.
- Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.
- Acemoglu, D. and Restrepo, P. (2020a). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Acemoglu, D. and Restrepo, P. (2020b). Unpacking skill bias: Automation and new tasks. *American Economic Association Papers and Proceedings*, 110:356–361.
- Autor, D. H. (2019). Work of the past, work of the future. *American Economic Association Papers and Proceedings*, 109:1–32.
- Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597.
- Autor, D. H., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, 135(2):645–709.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.
- Barany, Z. and Siegel, C. (2018). Job polarization and structural change. *American Economic Journal: Macroeconomics*, 10(1):57–89.

- Beaudry, P., Green, D. A., and Sand, B. M. (2016). The great reversal in the demand for skill and cognitive tasks. *Journal of Labor Economics*, 34(S1):S199–S247.
- Bosquet, C., Maarek, P., and Moiteaux, E. (2020). Online appendices to ‘Routine-biased technological change and wages by education level: Occupational downgrading and displacement effects’. Available on <https://sites.google.com/site/clementbosquet/>.
- Cortes, G. M. (2016). Where have the middle-wage workers gone? A study of polarization using panel data. *Journal of Labor Economics*, 34(1):63–105.
- Dustmann, C., Ludsteck, J., and Schönberg, U. (2009). Revisiting the German wage structure. *Quarterly Journal of Economics*, 124(2):843–881.
- Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2018). Bartik instruments: What, when, why, and how. *NBER Working Paper*, 24408.
- Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *Review of Economics and Statistics*, 89(1):118–133.
- Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in Europe. *American Economic Review: Papers & Proceedings*, 99(2):58–63.
- Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.
- Hunt, J. and Nunn, R. (2019). Is employment polarization informative about wage inequality and is employment really polarizing? *NBER Working Paper*, 26064.
- Jaimovich, N. and Siu, H. E. (2020). Job polarization and jobless recoveries. *Review of Economic and Statistics*, 102(1):129–147.
- Lindley, J. and Machin, S. (2016). The rising postgraduate wage premium. *Economica*, 83(330):281–306.

- Michaels, G., Natraj, A., and Van Reenen, J. (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics*, 96(1):60–77.
- Pessoa, J. P. and Van Reenen, J. (2013). Decoupling of wage growth and productivity growth? Myth and reality. *CEP Discussion Papers*, 1246.
- Rose, E. K. (2018). The rise and fall of female labor force participation during World War II in the United States. *The Journal of Economic History*, 78(3):673–711.
- Stock, J. H. and Yogo, M. (2005). Testing for Weak Instruments in Linear IV Regression. In Andrews, D. W. and Stock, J. H., editors, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pages 80–108. Cambridge University Press, Cambridge.
- Valletta, R. G. (2018). Recent flattening in the higher education wage premium: Polarization, skill downgrading, or both? *Working Paper*.
- Verdugo, G. and Allegre, G. (2020). Labour force participation and job polarization: Evidence from Europe during the Great Recession. *Labour Economics*, 66.

Appendix A Discussion on the Bartik instrument

In this section we open the black box of the Bartik shift-share instrument we use in the paper. We follow the recommendations of Goldsmith-Pinkham et al. (2018) who provide a methodology to discuss the sources of identification used in such settings. Goldsmith-Pinkham et al. (2018) show that the Bartik estimator can be decomposed into a weighted sum of individual estimators, each using a single industry share and industry shock at a time. From this decomposition, Goldsmith-Pinkham et al. (2018) propose several diagnostic tests in order to discuss if the exclusion restrictions are plausibly satisfied. In particular, the exogeneity of the industries that have the higher (Rotemberg) weights in the estimator, i.e. which drive its identification, needs to be discussed.

Table A1 presents various statistics about the Bartik estimator we use in Table 1 columns 4-6, for low, medium and high skilled workers. $\hat{\alpha}_j$ is the sum of Rotemberg weights for industry j , g_j is the weighted average national share of routine jobs in industry j over the full period,⁴³ $\hat{\gamma}_j$ is the weighted average just-identified instrumental variable estimator that use the share of industry j as an instrument and \hat{F}_j is the first step F-statistic for the instrument based on the single industry j . Finally, φ_j is the national share of industry j in 1960.

Panel A reports the sum of the negative and positive Rotemberg weights (column 1), their mean value (column 2) and their respective proportion (column 3). 76.1% of the Rotemberg weights are positive. According to Goldsmith-Pinkham et al. (2018), the Rotemberg weights “can be interpreted as sensitivity-to-misspecification elasticities”. In other words, the value of $\hat{\alpha}_j$ represents the potential bias introduced by the instrument of industry j if it is misspecified. If $\hat{\alpha}_j$ is relatively small, then the bias introduced by the the corresponding instrument on the Bartik estimate will be relatively small.

Panel B reports correlations of the various industry statistics. The low correlation coefficient between the Rotemberg weights $\hat{\alpha}_j$ and the average national share of routine jobs in

⁴³ $g_j = \frac{1}{T} \sum_{t=1}^T \hat{\alpha}_{jt} \text{Routine}_{j,t-1}$.

Table A1: Summary of Rotemberg weights for the IV specification reported in columns 4-6 of Table 1 for each skill

Panel A: Negative and positive weights						
	Sum	Mean	Share			
Negative	-0.459	-0.027	0.239			
Positive	1.459	0.056	0.761			
Panel B: Correlations						
	$\hat{\alpha}_j$	g_j	$\hat{\gamma}_j$	\widehat{F}_j	$Var_j(\varphi_{cj})$	
$\hat{\alpha}_j$	1					
g_j	0.107	1				
$\hat{\gamma}_j$	-0.059	-0.997	1			
\widehat{F}_j	-0.049	0.093	-0.123	1		
$Var_j(\varphi_{cj})$	-0.031	-0.373	0.340	0.449		1
Panel C: Variation across years in $\hat{\alpha}_j$						
	Sum	Mean				
1980	0.495	0.012				
1990	0.289	0.007				
2000	0.131	0.003				
2010	0.085	0.002				
Panel D: top 5 Rotemberg weight industries						
	$\hat{\alpha}_j$	g_j	$\hat{\gamma}_{j,low}$	$\hat{\gamma}_{j,med}$	$\hat{\gamma}_{j,hi}$	φ_j
Yarn, thread, and fabric mills	0.189	0.754	0.560	0.258	0.135	1.005
Apparel and accessories, except knit	0.151	0.884	0.337	-0.042	0.027	1.846
Machinery, except electrical, n.e.c. and Machinery, n.s.	0.141	0.717	-0.747	-0.517	-0.155	1.694
Electrical machinery, equipment, and supplies	0.149	0.673	-0.736	-0.579	-0.377	2.282
Motor vehicles and motor vehicle equipment	0.140	0.679	-1.135	-0.623	-0.145	1.109
Relative weight of top 5 industries:	0.527					
Panel E: Estimates of γ_j for positive and negative weights						
	$\hat{\alpha}$ -weighted sum	Share of overall $\hat{\gamma}$	Mean			
Low-skill						
Negative	0.076	-0.131	59.858			
Positive	-0.657	1.131	-0.817			
Medium-skill						
Negative	0.112	-0.255	41.574			
Positive	-0.549	1.255	-0.631			
High-skill						
Negative	0.063	-0.378	16.682			
Positive	-0.230	1.378	-0.283			

Notes: This table reports statistics about the Rotemberg weights. In all cases, we report statistics about the aggregated weights, where we aggregate a given industry across years. Panel B reports correlations between the weights ($\hat{\alpha}_j$), the national share of routine jobs (g_j), the just-identified coefficient estimates ($\hat{\gamma}_j$), the first-stage F-statistic of the industry share (\widehat{F}_j), and the variation in the 1960 industry shares across locations ($Var_j(\varphi_{cj})$). Panel C reports variation in the weights across years. Panel D reports the top five industries according to the Rotemberg weights. φ_j is the national industry share in 1960 (multiplied by 100 for legibility). Panel E reports statistics about how the values of $\hat{\gamma}_j$ vary with the positive and negative Rotemberg weights.

industry j g_j (0,107) indicates that the identification of our Bartik estimator relies mainly on the industry shares rather than the industry shocks.⁴⁴ Goldsmith-Pinkham et al. (2018) argue that it is very unusual in empirical frameworks to be able to rely on both sources of identification. Our framework makes no exception. The reason is quite intuitive: the technology shocks that make routine occupations disappear have affected many industries at the same time. The identification relies on industry shares which distribute this shock across CZ.

Panel C displays the sum of Rotemberg weights for each decade. The first decades have more weights in the total Bartik estimate. This is not surprising: as we use the local industry shares in 1960, Rotemberg weights are higher for the first periods. This is standard when using stacked differences with 10-year variations (see Goldsmith-Pinkham et al., 2018, for examples). The estimates using long run variations over 40 years avoid this problem as we only use one time variation. Results are qualitatively very consistent across the stacked and the long run differences.

Panel D reports the five industries which have the highest Rotemberg weights, their Rotemberg weight $\hat{\alpha}_j$ (column 1), their weighted average national share of routine jobs g_j (column 2), the just-identified IV coefficient $\hat{\gamma}_j$ using the single industry j for the three types of skills (columns 3-5) and the corresponding 1960 national industry shares (φ_j) multiplied by 100. These five industries together account for 52.7% of the Rotemberg weight. This means that 52.7% of the identification comes from those five industries. It is quite standard in the literature using Bartik shift-share instruments that few industries account for an important share of the total weights (see Goldsmith-Pinkham et al., 2018, for examples). It is thus important to ensure that the initial industry shares for those five industries are exogenous to future development in the CZ.

Panel E shows descriptive statistics of $\hat{\gamma}_j$ for positive and negative weights. Column 1 reports the sum of coefficients $\hat{\gamma}_j$ weighted by the Rotemberg weights $\hat{\alpha}_j$ (for the negative and

⁴⁴This correlation means that the g_j explains about 1.1% (0.107²) of the variance of the Rotemberg weights.

positive Rotemberg weights separately) for the three skill groups. Column 2 reports their share in the Bartik estimate $\hat{\gamma}$ and column 3 presents the mean of $\hat{\gamma}_j$ for the positive and negative Rotemberg weights. The weighted sum of coefficients is very low for the negative weights and much higher and negative for the positive weights. The overall coefficient $\hat{\gamma}$ is thus mostly explained by the industries with positive weights. The unweighted sum of $\hat{\gamma}_j$ is very high for the negative weights. This is due to an industry which is a clear outlier and exhibits a very high $\hat{\gamma}_j$ coefficient (Business and repair services). The weighted sum of coefficients for negative Rotemberg weights is very low because the Rotemberg weight associated with this coefficient is very low (1.93×10^{-5}) and hence, this industry does not matter for the overall Bartik estimate. If this high coefficient is due to a misspecification, this should not invalidate the overall identification strategy.

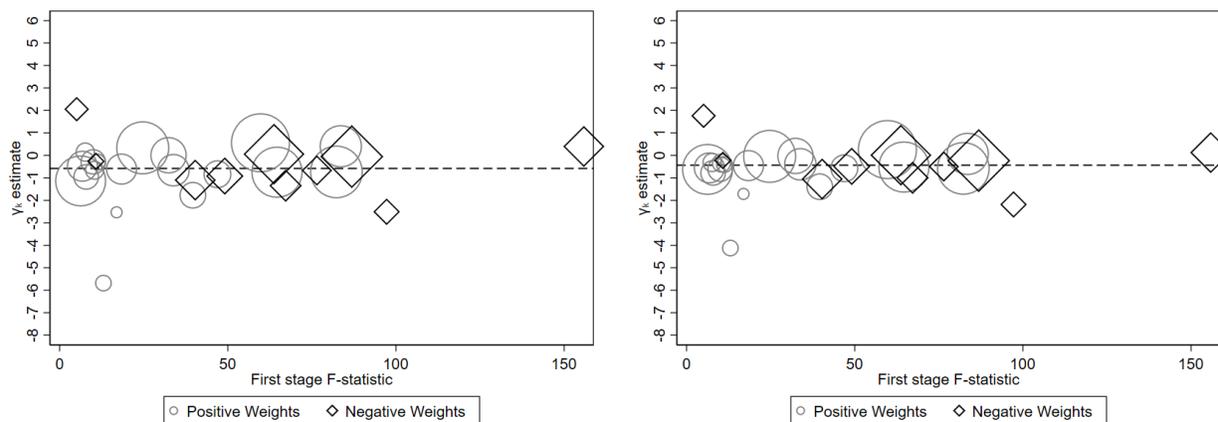
We now try to convince that the local industry shares in 1960, and more specifically the ones of the five industries that drive the Bartik estimator (i.e. the sum of their Rotemberg weights is greater than 50%), are exogenous to future development that could affect wages in a CZ. As argued in Goldsmith-Pinkham et al. (2018), the local industry shares in 1960 do not have to be exogenous to wages in levels, but rather to variations of wages since we estimate our model in variation. We first take a 10-year lag relative to the first year of our panel (1970). In Table 1, we then control for various factors that can affect the evolution of wages and could be affected by the initial industry shares. First, in all regressions we control for the evolutions of the skills supply and of the share of foreign born workers, which could be related to the initial industry composition. Second, as suggested by Goldsmith-Pinkham et al. (2018), we provide some estimates controlling for the evolution of broad industry shares. The local industry composition in 1960 could lead to major changes in the former industry composition such as variations in the manufacturing sectors which could explain the wage dynamics beyond the decrease in the proportion of routine jobs. On the other hand, as argued by Barany and Siegel (2018), structural change and industry shocks could explain itself the decrease in routine occupations. By controlling for structural change,

we could capture part of the polarization phenomenon. Autor and Dorn (2013) control for the share of manufacturing using similar arguments. We can see in columns 7-9 of Table 1 that controlling for the variations of broad sector shares only marginally affects the results. Online appendix Table O12 shows regressions controlling for the evolutions of industry shares using a more detailed 13 industry classification. Coefficients are still significant at 1% for the low and medium skilled and insignificant for the high skilled.

Finally, we perform an overidentification test when using the single industry instruments separately. It rejects the null hypothesis of instrument exogeneity. As suggested by Goldsmith-Pinkham et al. (2018), a plausible reason in those kind of settings is that our IV estimate measures a local average treatment effect rather than an average treatment effect. It is very likely that some local labor markets are affected by the instrument very differently than others and that the estimated coefficient $\hat{\gamma}_j$ of a variation in routine may be heterogeneous. Removing one instrument will affect the overall $\hat{\gamma}$ even if the model is well specified. In order to see if this interpretation is correct, we graph on Figure A1 the estimated coefficient $\hat{\gamma}_j$ (y-axis) according to their F statistics (x-axis). The figure only includes instruments with first-stage F-statistics above 5. The circles represent positive Rotemberg weights and the squares represent negative weights. The size of the circles and squares corresponds to the value of the Rotemberg weights.

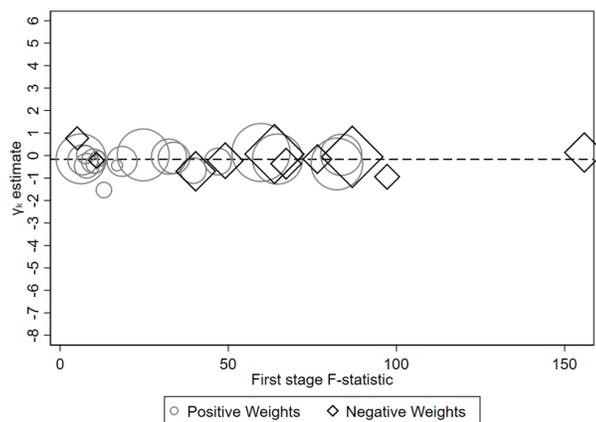
We can see that the dispersion of the estimated coefficients is not very high except for a few outliers which have some very low Rotemberg weights. According to Goldsmith-Pinkham et al. (2018), this gives support for the LATE interpretation of the IV estimate rather than a misspecification. If the model was misspecified, the dispersion of the estimated coefficients would have been higher, and industries with high negative Rotemberg weights would have been associated with misspecified coefficients, meaning that those coefficients would account for an important part of the overall Bartik estimator. Here, the few outlier coefficients for which we could suspect a misspecification of the model have very low Rotemberg weights.

Figure A1: Distribution of Rotemberg Weights for each skill



(a) Low-skilled workers

(b) Med. skilled workers



(c) High-skilled workers

Notes: Those figures represent the relationship between each just-identified coefficient $\hat{\gamma}_j$, first stage F-statistics \hat{F}_j and the Rotemberg weights $\hat{\alpha}_j$. Each point corresponds to a separate instruments' estimates (industry share). The figure plots the estimated $\hat{\gamma}_j$ for each instrument on the y-axis and the estimated first-stage F-statistic \hat{F}_j on the x-axis. The size of the points are scaled by the magnitude of the Rotemberg weights, with the circles denoting positive Rotemberg weights and the diamonds denoting negative weights. The horizontal dashed line is plotted at the value of the overall $\hat{\gamma}$ reported in columns 4-6 in Table 1 for each skill. The figure excludes instruments with first-stage F-statistics below 5.

Appendix B Additional results

B.1 Robustness of Table 1

Table B1: First-difference CZ average residualized wages on the local share of routine jobs, by *detailed* skill levels, IV, stacked differences

Workers educ.:	lower low	upper low	lower med.	upper med.	lower high	upper high
	(1)	(2)	(3)	(4)	(5)	(6)
Sh. of rout. occs ₀	-0.620 ^a (0.076)	-0.368 ^a (0.069)	-0.358 ^a (0.062)	-0.367 ^a (0.070)	-0.254 ^a (0.076)	0.006 (0.057)
Δ sh. of up. high skill	1.111 ^a (0.301)	1.599 ^a (0.245)	1.716 ^a (0.242)	1.216 ^a (0.230)	1.311 ^a (0.211)	1.551 ^a (0.215)
Δ sh. of low. high skill	-1.942 ^a (0.440)	-1.392 ^a (0.325)	-1.507 ^a (0.355)	-1.159 ^a (0.367)	-0.176 (0.336)	-0.370 (0.363)
Δ sh. of up. med. skill	-1.111 ^a (0.246)	-0.908 ^a (0.248)	-1.233 ^a (0.253)	-0.998 ^a (0.252)	-0.872 ^a (0.188)	-0.349 ^c (0.202)
Δ sh. of low. med. skill	-0.507 (0.303)	-0.317 (0.255)	-0.668 ^b (0.262)	-0.742 ^a (0.258)	-0.602 ^a (0.195)	-0.214 (0.197)
Δ sh. of up. low skill	-1.473 ^a (0.153)	-0.903 ^a (0.152)	-0.961 ^a (0.172)	-0.833 ^a (0.176)	-0.713 ^a (0.142)	-0.472 ^a (0.140)
Δ sh. of foreign	-0.270 ^b (0.120)	0.157 (0.102)	0.355 ^a (0.121)	0.322 ^a (0.105)	0.590 ^a (0.103)	0.464 ^a (0.092)

Notes: Standard errors clustered by state between brackets. ^a, ^b, ^c indicate significance at the 1%, 5% and 10% level, respectively. Each regression is weighted by the share of the CZ in the national population. 2888 observations (722 CZ × 4 periods).

Table B2: First-difference CZ average residualized wages on the local share of routine jobs, by skill levels, OLS and IV, *long differences and 20-year stacked differences*

Workers educ.:	OLS			IV					
	low	medium	high	low	medium	high	low	medium	high
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: long differences									
Sh. of rout. $occs_0$	-0.919 ^a (0.222)	-0.737 ^a (0.150)	-0.207 (0.131)	-1.174 ^a (0.275)	-0.849 ^a (0.182)	-0.275 ^c (0.151)	-2.062 ^a (0.408)	-1.199 ^a (0.296)	-0.203 (0.245)
Δ sh. of high skill	0.196 (0.215)	0.199 ^b (0.099)	0.644 ^a (0.122)	0.325 ^c (0.197)	0.256 ^b (0.100)	0.678 ^a (0.119)	0.243 (0.221)	0.222 ^b (0.103)	0.674 ^a (0.124)
Δ sh. of med. skill	0.248 (0.364)	-0.159 (0.203)	-0.212 (0.157)	0.300 (0.361)	-0.136 (0.203)	-0.198 (0.159)	-0.224 (0.341)	-0.349 (0.216)	-0.237 (0.160)
Δ sh. of foreign	-0.313 ^c (0.171)	0.200 ^b (0.093)	0.416 ^a (0.120)	-0.304 ^c (0.168)	0.205 ^b (0.094)	0.418 ^a (0.119)	-0.241 (0.156)	0.230 ^a (0.088)	0.419 ^a (0.118)
Δ sh. of manuf.							-0.190 (0.430)	-0.061 (0.318)	0.179 (0.291)
Δ sh. of services							1.020 ^b (0.423)	0.423 (0.308)	0.175 (0.239)
R ²	0.29	0.29	0.47						
Kleibergen-Paap				512	512	512	89	89	89
Panel B: 20-year stacked differences									
Sh. of rout. $occs_0$	-0.320 ^b (0.129)	-0.142 (0.091)	0.029 (0.075)	-0.804 ^a (0.177)	-0.556 ^a (0.119)	-0.125 (0.107)	-1.032 ^a (0.277)	-0.801 ^a (0.226)	-0.076 (0.178)
Δ sh. of high skill	0.238 (0.171)	0.269 ^b (0.114)	0.655 ^a (0.076)	0.463 ^a (0.163)	0.462 ^a (0.128)	0.726 ^a (0.084)	0.405 ^b (0.157)	0.416 ^a (0.110)	0.727 ^a (0.084)
Δ sh. of med. skill	0.495 ^b (0.199)	0.136 (0.132)	0.001 (0.092)	0.626 ^a (0.196)	0.249 ^c (0.142)	0.042 (0.102)	0.522 ^b (0.221)	0.226 (0.173)	0.003 (0.109)
Δ sh. of foreign	-0.186 ^c (0.099)	0.316 ^b (0.150)	0.480 ^a (0.077)	-0.188 ^c (0.103)	0.314 ^c (0.158)	0.479 ^a (0.081)	-0.097 (0.113)	0.369 ^b (0.167)	0.489 ^a (0.085)
Δ sh. of manuf.							0.100 (0.303)	-0.228 (0.285)	0.211 (0.199)
Δ sh. of services							0.657 ^a (0.233)	0.174 (0.209)	0.227 (0.158)
R ²	0.12	0.08	0.29						
Kleibergen-Paap				376	376	376	136	136	136

Notes: See Table B1. 722 observations in panel A, 1444 (722 CZ \times 2 years) observations in panel B.

Table B3: First-difference CZ average residualized wages on the *variations* of the local share of routine jobs, by skill levels, OLS and IV, stacked differences

Workers educ.:	OLS			IV		
	low	medium	high	low	medium	high
	(1)	(2)	(3)	(4)	(5)	(6)
Δ sh. of rout. occs	0.337 ^c (0.181)	0.136 (0.130)	0.011 (0.140)	1.896 ^a (0.335)	1.428 ^a (0.255)	0.546 ^b (0.255)
Δ sh. of high skill	0.678 ^a (0.193)	0.681 ^a (0.164)	0.972 ^a (0.114)	1.508 ^a (0.265)	1.369 ^a (0.218)	1.257 ^a (0.183)
Δ sh. of med. skill	-0.448 ^b (0.178)	-0.677 ^a (0.161)	-0.405 ^a (0.128)	-0.610 ^a (0.220)	-0.811 ^a (0.194)	-0.461 ^a (0.155)
Δ sh. of foreign	-0.362 ^a (0.079)	0.087 (0.101)	0.329 ^a (0.063)	-0.275 ^a (0.087)	0.159 ^c (0.088)	0.359 ^a (0.062)
R ²	0.05	0.09	0.20			
Kleibergen-Paap				245	245	245

Notes: See Table B1.

Table B4: First-difference CZ average residualized wages on the local share of routine jobs, by skill levels, *reduced form*, stacked differences

Workers educ.:	low	medium	high
	(1)	(2)	(3)
Exposure	-0.336 ^a (0.061)	-0.253 ^a (0.048)	-0.097 ^b (0.045)
Δ sh. of high skill	0.924 ^a (0.210)	0.929 ^a (0.175)	1.089 ^a (0.126)
Δ sh. of med. skill	-0.439 ^b (0.181)	-0.683 ^a (0.165)	-0.412 ^a (0.137)
Δ sh. of foreign	-0.403 ^a (0.061)	0.063 (0.106)	0.322 ^a (0.061)
R ²	0.11	0.13	0.21

Notes: See Table B1.

Table B5: First-difference CZ average residualized wages on the local share of routine jobs, by skill levels, OLS and IV, stacked differences, *excluding 2010*

Workers educ.:	OLS			IV		
	low	medium	high	low	medium	high
	(1)	(2)	(3)	(4)	(5)	(6)
Sh. of rout. occs ₀	-0.159 ^c (0.082)	-0.056 (0.062)	0.027 (0.061)	-0.479 ^a (0.105)	-0.266 ^a (0.094)	-0.071 (0.091)
Δ sh. of high skill	0.810 ^a (0.194)	0.903 ^a (0.165)	1.116 ^a (0.123)	0.969 ^a (0.233)	1.008 ^a (0.203)	1.164 ^a (0.145)
Δ sh. of med. skill	-0.736 ^a (0.250)	-0.998 ^a (0.237)	-0.505 ^a (0.168)	-0.679 ^a (0.248)	-0.961 ^a (0.233)	-0.487 ^a (0.166)
Δ sh. of foreign	-0.523 ^a (0.109)	-0.046 (0.150)	0.318 ^a (0.070)	-0.533 ^a (0.102)	-0.052 (0.156)	0.315 ^a (0.071)
R ²	0.09	0.15	0.26			
Kleibergen-Paap				381	381	381

Notes: See Table B1. 2166 (722 CZ × 3 years) observations.

B.2 Robustness of Table 2

Table B6: Decomposition of the IV coefficients from Table 1, columns 4 and 5 *on 77 occupation posts*

Workers educ.:	low			medium		
	total	between	within	total	between	within
	(1)	(2)	(3)	(4)	(5)	(6)
Sh. of rout. occs ₀	-0.581 ^a (0.114)	-0.053 ^b (0.022)	-0.517 ^a (0.095)	-0.437 ^a (0.089)	-0.081 ^a (0.018)	-0.368 ^a (0.077)
Δ sh. of high skill	0.749 ^a (0.195)	0.080 ^a (0.029)	0.667 ^a (0.172)	0.797 ^a (0.164)	-0.043 (0.028)	0.816 ^a (0.153)
Δ sh. of med. skill	-0.235 (0.166)	0.103 ^a (0.019)	-0.351 ^b (0.160)	-0.529 ^a (0.150)	-0.049 ^b (0.021)	-0.503 ^a (0.147)
Δ sh. of foreign	-0.394 ^a (0.059)	-0.071 ^b (0.027)	-0.317 ^a (0.065)	0.070 (0.109)	-0.001 (0.012)	0.061 (0.095)

Notes: See Table B1. Occupation classification composed of 77 posts, more details available in Section 3. The Kleibergen-Paap rk Wald F statistic is equal to 323 in all regressions.

Appendix C Individual-level regressions

To estimate the impact of the decrease in the proportion of routine occupations on the individual wages of the three skill groups we have identified, we also apply a local labor market approach on individual data. We regress individual wages on the share of routine occupations in a CZ by skill group:

$$w_{it} = \gamma \text{Routine}_{ct} + X_{it} \beta + Z_{ct} \theta + \alpha_c + \lambda_t + \varepsilon_{it} \quad (5)$$

where w_{it} is the log wage of worker i at time t , Routine_{ct} is the share of routine jobs in area (CZ) c at time t , X_{it} is a vector of individual characteristics including gender, age and its square, race and possibly occupation and industry dummies, Z_{ct} is a vector of time-varying area controls and α_c and λ_t are area and time fixed effects, respectively. γ , our main coefficient of interest, is then identified with the time variations of the local shares of routine jobs. As in our baseline specification, we control for the supply of skills and foreign-born individuals in Z_{ct} as it could affect the local equilibrium wages of the different skill groups and the occupational structure of the workforce.

As for our CZ-level regressions, the local share of routine occupations (Routine_{ct}) might be endogeneous. Thus, we also instrument it by the exposure to occupational structure changes. As a reminder:

$$\text{Exposure}_{ct} = \sum_{j=1}^N \varphi_{cj} \text{Routine}_{jt-1} \quad (6)$$

In order to study the mechanisms at work behind the impact of the routine-biased technological change on wages by education groups, we introduce additional fixed effects in some specifications. As stated earlier, the effect of a change in the occupational structure on wages could be related to a between industry/occupation effect and/or a within one. In order to disentangle between the two effects, we introduce occupation, industry and

occupation×industry fixed effects in Table C2. In Table C3 we introduce interactions between the share of routine jobs and individual-level dummies for manual, routine and abstract occupations to test if the effect is heterogenous across the broad occupational categories.

All the regressions are clustered at the commuting zone level and weighted. The weights are equal to the Census weight (i.e. how many persons in the US population are represented by a given person in an IPUMS sample) multiplied by the CZ specific weight (i.e. the fraction of the county group/PUMA that maps to this given CZ, as mentioned in Section 2).

Table C1: Individual (log) wage regressions on the local share of routine jobs, by skill levels

Workers educ.:	OLS			IV		
	low	medium	high	low	medium	high
	(1)	(2)	(3)	(4)	(5)	(6)
Sh. of rout. occs	0.273 ^b (0.129)	0.351 ^a (0.099)	-0.083 (0.115)	2.182 ^a (0.364)	1.970 ^a (0.249)	0.699 ^b (0.280)
Sh. of high skill	0.459 ^b (0.192)	0.464 ^a (0.123)	0.627 ^a (0.123)	1.662 ^a (0.305)	1.403 ^a (0.166)	1.078 ^a (0.201)
Sh. of med. skill	0.100 (0.142)	-0.294 ^a (0.107)	-0.578 ^a (0.145)	-0.011 (0.143)	-0.350 ^a (0.116)	-0.591 ^a (0.150)
Sh. of foreign	-0.173 ^c (0.095)	0.310 ^a (0.074)	0.359 ^a (0.076)	-0.049 (0.108)	0.460 ^a (0.097)	0.444 ^a (0.093)
R ²	0.57	0.55	0.52			
Kleibergen-Paap				216	174	154

Notes: Standard errors clustered by CZ between brackets. ^a, ^b, ^c indicate significance at the 1%, 5% and 10% level, respectively. The numbers of observations for each skill are 10,472,235, 5,341,197 and 4,130,037, respectively.

Table C2: Individual (log) wage IV regressions on the local share of routine jobs, by skill levels, with occupation and industry fixed effect

Workers educ.:	low			medium			high		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sh. of rout. occs	1.880 ^a (0.336)	1.656 ^a (0.310)	1.645 ^a (0.307)	1.626 ^a (0.230)	1.419 ^a (0.215)	1.382 ^a (0.211)	0.365 (0.249)	0.173 (0.249)	0.155 (0.242)
Sh. of high skill	1.517 ^a (0.286)	1.460 ^a (0.263)	1.452 ^a (0.260)	1.307 ^a (0.151)	1.238 ^a (0.141)	1.223 ^a (0.138)	0.825 ^a (0.177)	0.711 ^a (0.180)	0.696 ^a (0.175)
Sh. of med. skill	-0.153 (0.130)	-0.142 (0.120)	-0.142 (0.118)	-0.322 ^a (0.104)	-0.300 ^a (0.098)	-0.300 ^a (0.096)	-0.614 ^a (0.126)	-0.615 ^a (0.121)	-0.598 ^a (0.118)
Sh. of foreign	-0.062 (0.101)	-0.082 (0.089)	-0.080 (0.088)	0.387 ^a (0.090)	0.370 ^a (0.088)	0.372 ^a (0.086)	0.359 ^a (0.083)	0.335 ^a (0.082)	0.331 ^a (0.079)
Occ. FE	yes	yes	no	yes	yes	no	yes	yes	no
Ind. FE	no	yes	no	no	yes	no	no	yes	no
Occ. × Ind. FE	no	no	yes	no	no	yes	no	no	yes
Kleibergen-Paap	217	218	218	174	174	174	154	154	154

Notes: See Table C1. The numbers of observations are 10,472,235 (10,472,141 in column 3), 5,341,197 (5,341,077 in column 6) and 4,130,037 (4,129,864 in column 9) for low, medium and high-skill workers, respectively.

Table C3: Individual (log) wage IV regressions on the local share of routine jobs, by skill levels, with the routine share interacted with broad occupation dummies

Workers educ.:	low	medium	high
	(1)	(2)	(3)
Sh. of rout. occ. × man.	2.124 ^a (0.337)	1.790 ^a (0.249)	1.133 ^a (0.341)
Sh. of rout. occ. × rout.	1.634 ^a (0.337)	1.709 ^a (0.228)	0.900 ^a (0.298)
Sh. of rout. occ. × abst.	1.609 ^a (0.317)	1.207 ^a (0.217)	0.224 (0.247)
Sh. of high skill	1.504 ^a (0.282)	1.272 ^a (0.145)	0.836 ^a (0.183)
Sh. of med. skill	-0.170 (0.128)	-0.294 ^a (0.102)	-0.612 ^a (0.124)
Sh. of foreign	-0.028 (0.100)	0.423 ^a (0.090)	0.366 ^a (0.081)
Kleibergen-Paap	73	59	53

Notes: See Table C1. The numbers of observations for each skill are 9,110,818, 4,831,197 and 3,901,517, respectively.