I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 18218

Intersecting Shocks: The Combined Labor Market Impacts of Automation and Immigration

Patrick Bennett
Julian Vedeler Johnsen

OCTOBER 2025

I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 18218

Intersecting Shocks: The Combined Labor Market Impacts of Automation and Immigration

Patrick Bennett
University of Liverpool, IFS and IZA

Julian Vedeler Johnsen
University of Bergen

OCTOBER 2025

ISSN: 2365-9793

IZA – Institute of Labor Economics

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

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

www.iza.org

ABSTRACT

Intersecting Shocks: The Combined Labor Market Impacts of Automation and Immigration*

We study how the labor market shocks of automation and immigration interact to shape workers' outcomes. Using matched employer–employee data from Norwegian administrative registers, we combine an immigration shock triggered by the European Union's 2004 enlargement with an automation shock based on the adoption of industrial robots across Europe. Although these shocks largely occur in separate industries, we show that automation reduces earnings not only in manufacturing but also in construction, where tasks overlap with robot-exposed sectors. Importantly, workers jointly exposed to automation and immigration suffer earnings losses greater than those facing either shock in isolation. These losses are driven by downward occupational mobility into low-wage services and re-sorting into lower-premium firms. Even within the Norwegian welfare system, the ability of social insurance to offset these long-run earnings declines is limited. Our findings underscore the importance of analyzing labor market shocks jointly, rather than in isolation, to fully understand their distributional consequences.

Keywords: automation, immigration, labor market shocks

Corresponding author:
Patrick Bennett
University of Liverpool
765 Brownlow Hill
Liverpool L69 7ZX
United Kingdom
E-mail: Patrick.Bennett@liverpool.ac.uk

* The authors would like to thank participants at several university seminars and conferences for their valuable feedback and suggestions. This research was supported by the Research Council of Norway through projects 315809 and 262675.

1 Introduction

The changing nature of work, driven by a confluence of labor market shocks such as rapid technological progress (Autor et al., 2003; Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Dauth et al., 2021), expanding global trade (Autor et al., 2013, 2015; Charles et al., 2018), and shifting migration patterns (Bratsberg and Raaum, 2012; Dustmann et al., 2017), has fueled public debate about job displacement and wage stagnation. While each of these forces has been studied extensively in isolation, much less is known about how they *interact* to shape labor market outcomes.

We address this gap by showing that two of the most salient contemporary labor market shocks—automation and immigration—not only coincide in time and space but also amplify each other's effects on vulnerable workers. Our objective is not to delineate all possible interactions among labor market shocks. Instead, we focus on these two contemporaneous shocks for which we have robust identification strategies. By doing so, we reveal underexplored interdependencies that meaningfully shape labor market outcomes for certain groups of workers. Although many studies have examined the effects of automation or immigration on earnings in isolation, the extent to which multiple shocks amplify or mitigate each other remains largely unknown. Our analysis also contributes to the broader labor economics literature by showing how seemingly unrelated shocks can generate spillover effects across industries.

Using worker-level data from Norwegian administrative registers, we show that automation and immigration jointly amplify their adverse impact on the earnings of non-college-educated workers. Our identification strategy combines an exogenous industry-level immigration shock—triggered by a sudden influx of migrants into specific sectors following the European Union's expansion in 2004 (as in Bratsberg and Raaum, 2012)—with an exogenous area-level automation shock tied to the rising adoption of industrial robots across Europe (as in Acemoglu and Restrepo, 2020). While the bulk of immigration and automation occur in distinct industries, we show that they interact via cross-industry spillovers of automation. Crucially, while these shocks overlap in time, we exploit the fact that two workers facing the same automation exposure can differ in their exposure to immigration due to licensing requirements that bar migrants from specific occupations. Finally, we document how the joint effects of automation and immigration manifest through shifts in occupations (from blue-collar to lower-paid service jobs) and downward progression along the employer pay distribution, increasing mismatch between workers and firms.

We present several sets of main results. First, we uncover strong between-industry spillover effects of automation. In line with previous studies from other countries, we confirm that automation negatively impacts the earnings of non-college-educated workers in the affected manufacturing sector. Specifically, our estimates indicate that each additional robot per 1,000 workers reduces the earnings of manufacturing employees in the local labor market by 3.2%.[1] However, construction workers also experience similarly

[1] For a discussion of immigration's effect on earnings in Norway, see Bratsberg and Raaum (2012), who find that a 10%

large earnings declines, despite robot adoption being concentrated in manufacturing. This strong spillover effect can be explained by worker substitutability: earnings losses are concentrated among construction workers performing tasks directly transferable to manufacturing, whereas those who perform tasks less transferable to manufacturing experience no significant earnings declines.

Second, and most importantly, we establish that the shocks of automation and immigration amplify each other, and that this interaction effect is sufficiently large to alter the interpretation of each shock individually. Workers simultaneously exposed to both shocks experience earnings losses that are 2.5 percentage points greater than those exposed only to automation. This sizable interaction effect is quantitatively important in interpreting the labor market consequences of each shock. In the absence of one shock, the negative effect of the other is attenuated: workers exposed to only one shock see no significant decline in earnings. Moreover, while earnings fall for jointly exposed workers, their employment levels remain stable, suggesting that the effect operates primarily through lower wages rather than job loss. Our findings are robust to a host of specification checks, including controlling for any potential differences in cognitive ability and for potential exposure to Chinese import competition. Results are also similar when comparing two workers who are differentially exposed to the immigration shock within the same local labor market. We also find no evidence of pre-treatment differences in earnings among workers who are later differentially exposed to both shocks, further supporting the validity of our identification strategy.

Third, we highlight the mechanisms through which joint exposure to automation and immigration reduces earnings. Workers exposed to both shocks get stuck at the bottom: they move from higher-paid blue-collar jobs to lower-paid service occupations, transition to plants with lower employer-specific wage premia, suffer declines in match quality, and are significantly more likely to claim disability insurance. Jointly exposed workers are 18% more likely to work in service occupations and transition to plants with employer wage premia that are 1 log point lower. Together, downward job mobility and transitions to lower-paying employers account for 50% of the total earnings decline among jointly exposed workers, with employer shifts alone explaining 35%. Defining match quality as the overlap between estimated worker and employer-specific fixed effects, exposure to both shocks leads to significant increases in mismatch. Joint exposure increases the distance between employer and worker fixed effects by 3% relative to those exposed only to automation, driven by greater mismatch among workers with initially poor matches. Finally, we quantify the social insurance margin in a generous welfare state. While taxes and transfers play a limited role in offsetting earnings losses, joint exposure significantly increases the likelihood of claiming disability insurance.

Our findings contribute to a growing yet underexplored literature on the combined impacts of multiple labor market shocks. While Autor et al. (2015) and Faber et al. (2022) underscore the importance

increase in immigration decreases within-industry earnings of native workers by 0.6%.

of jointly considering technology and trade—showing that Chinese import competition and automation affect labor markets in distinct ways—we demonstrate that immigration and automation can similarly interact in ways that change the interpretation of each shock. On the theoretical side, Goos et al. (2014) develop a model that jointly considers routine-biased technical change and offshoring, while Mandelman and Zlate (2022) propose a structural framework to jointly analyze offshoring, automation, and immigration. However, empirical studies examining the joint effects of automation and immigration remain limited. By analyzing how these two shocks coincide in time and interact across industries, we provide reduced-form evidence that their joint impact depresses earnings beyond the sum of their individual effects.[2]

Our results also contribute to the literature on automation. Prior work shows that computerization and robots displace routine and mid-skill tasks, while raising productivity and complementing non-routine work, thus broadening disparities between college- and non-college-educated workers or polarizing labor markets (Autor et al., 2003, 2006; Goos et al., 2009; Autor and Dorn, 2013; Jurkat et al., 2023). More recent studies highlight the distinct role of industrial robots. Acemoglu and Restrepo (2020) show that U.S. commuting zones more exposed to robots experienced significant declines in both wages and employment, whereas Graetz and Michaels (2018) find that robot adoption across 17 countries boosted productivity and average wages but displaced low-skill jobs. For Germany, Dauth et al. (2021) document no net job loss, but major compositional shifts: robots eliminated many manufacturing jobs yet were offset by growth in services. Building on this literature, we show that automation shocks extend beyond directly exposed sectors: robot adoption in manufacturing significantly reduces earnings in construction, particularly among workers performing tasks that are transferable to manufacturing. This finding extends the task-based framework of Autor et al. (2003) and Acemoglu and Restrepo (2022), which emphasizes how technology substitutes for routine tasks and complements non-routine ones, by showing that task substitutability operates not only within industries but also across them, generating significant between-industry spillovers. We further show that these adverse effects are magnified when automation coincides with immigration, as jointly exposed workers are pushed down the job ladder into lower-paid occupations and lower-premium firms.

We also contribute to the literature on immigration and labor markets. While earlier studies reach divergent conclusions (Card, 1990; Borjas, 2003), more recent work reconciles these findings by harmonizing empirical specifications (Dustmann et al., 2016) or emphasizing task specialization and heterogeneity. Immigrants often concentrate in manual-intensive work, allowing natives to shift toward communication-

[2] A related literature studies how immigration shapes automation. In U.S. manufacturing, Lewis (2011) shows that cities receiving larger inflows of less-skilled immigrants adopted automation and computerization more slowly, consistent with immigrant labor substituting for capital. Conversely, when low-skill immigrant labor was abruptly curtailed in the farming sector, growers mechanized to offset the shortfall, with little evidence of native wage gains (Clemens et al., 2018). Our paper does not focus on this mechanism, and—crucially—it does not appear to confound our estimates. Our empirical strategy exploits two orthogonal sources of variation, and there is no systematic relationship between changes in immigration and automation across local labor markets.

intensive tasks or upgrade into more complex jobs (Peri and Sparber, 2009; Foged and Peri, 2016). Distributional studies often find (modest) losses for low-wage workers but gains for higher-wage groups (Dustmann et al., 2013; Glitz, 2012; Manacorda et al., 2012; Ottaviano and Peri, 2012). Our contribution to this literature is to demonstrate that the effects of immigration cannot be fully understood in isolation. When immigration coincides with automation, the two shocks amplify one another's effects: opportunities for nearby job transitions diminish as automation simultaneously reduces demand in manufacturing, thereby magnifying occupational downgrading and employer re-sorting among non-college-educated workers. Whereas prior studies emphasize average wage effects or task specialization between natives and immigrants, we show that the impact of immigration critically depends on the broader technological environment in which workers are embedded.

Our results also clarify the mechanisms through which occupational and employer shifts drive earnings declines under joint exposure to automation and immigration. This aligns with previous research emphasizing the importance of occupations (Edin et al., 2023) and employers (Abowd et al., 1999; Card et al., 2018) in shaping labor market outcomes. We confirm that occupations—already shown to matter for automation-related effects (as in Acemoglu and Restrepo, 2020)—also respond to joint exposure, consistent with prior findings for each shock individually (Peri and Sparber, 2009; Dauth et al., 2021). Crucially, however, the downward re-sorting across employers is large enough to explain a substantial share of the overall earnings decline. By highlighting how sorting to lower-paying firms and occupations amplifies workers' losses, our findings underscore the importance of studying employer and occupation dynamics together when assessing the full impact of labor market shocks.

Finally, our findings relate to a large literature on how social insurance shapes labor supply. In particular, disability insurance (DI) has become an increasingly important income source for workers hit by adverse shocks, with DI expenditures accounting for a growing share of public transfers in Norway and other advanced economies (OECD, 2010). Rising DI recipiency has been linked to declining employment rates, with Abraham and Kearney (2020) identifying higher DI take-up—among other factors—as a key driver of falling labor force participation in the 2000s. Within the broader DI literature, French and Song (2014) emphasize its negative effects on labor supply, especially among non-college workers, Kostøl and Mogstad (2014) show that recipients often retain substantial work capacity, and Autor et al. (2019) highlight the inherent insurance–work trade-off in DI systems. Against this backdrop, we find that while DI cushions some of the losses faced by workers jointly exposed to automation and immigration, it offsets only a small share of their long-run earnings decline. Taken together, these findings underscore the limits of social insurance in protecting workers when shocks push them down the job ladder into persistently lower-paying firms and occupations.

The remainder of the paper is organized as follows. Section 2 describes the Norwegian registry data and industrial robot adoption measures used in our analysis, and briefly describes relevant features of the institutional setting. Section 3 outlines the empirical framework and details the automation and immigration shocks leveraged for identification. Section 4 discusses the results: the direct impact of automation on non-college-educated workers in manufacturing, spillover effects on construction workers, and the joint impact of automation and immigration. It also examines mechanisms and the role of social insurance. We conclude the paper with some final remarks in Section 5.

2 Data and Institutional Setting

We use detailed administrative registers from Norway, which contain annual, individual-level data on demographic characteristics, employment, earnings, and education. We merge these registers to create a panel that follows workers over time. In addition, we incorporate data from the International Federation of Robotics (IFR) to measure exposure to automation at the local labor market level. We define exposure to the automation and immigration shock based on *pre-shock* area of residence and industry of employment. While this Section provides information on the data and a brief overview of the Norwegian institutional setting, Section 3 provides more information on the dynamics of the automation and immigration shocks and our strategy for identifying their (joint) effects. Throughout the paper, we will focus on male workers aged 18 to 49 in 2000.

2.1 Norwegian Register Data

We get demographic characteristics from the population register which covers all residents in Norway. For each individual, it records date of birth, municipality of residence, country of birth, and gender. We use country of birth to distinguish between native-born and immigrant workers. Our mapping from municipalities to local labor markets follows Gundersen and Juvkam (2013). This classification groups Norway's municipalities into 160 local labor markets defined by commuting patterns, travel distances, and centrality.[3]

To separate between college and non-college educated workers, we use the education registers. These registers tracks an individual's highest completed degree as well as current enrollment in formal education. We define "college educated" as having any post-secondary credential (for example, a degree from a university, university college, or technical college). "Non-college educated" workers, therefore, include both high school graduates and high school dropouts. We also measure cognitive ability from compulsory military testing performed at the age of 18, which is an aggregate measure of tests in arithmetic, word similarities, and figures.[4] Such data provides a measure of ability which allows us to account for any

[3]In their approach, Gundersen and Juvkam first establish a "center structure" by ranking municipalities according to population, employment, service availability, public institutions, and a centrality index. Based on these characteristics, municipalities are grouped into "functional" labor markets known as BA-regions, which reflect shared commuting patterns and local economies.

[4]Tests in arithmetic and word similarities are similar to a Wechsler Adult Intelligence Scale (WAIS) test while the test in figures resembles a Raven Progressive Matrix test.

potential differences in cognitive ability over and above education.

Finally, we observe annual employment spells from the employer register, which contains a plant identifier, firm identifier, industry classification, and detailed occupation codes (available from 2003). We group occupations into four broad categories: 1) Professional (managerial, professional, or technical occupations), 2) blue-collar (craft, trade, and machinery jobs), 3) services (sales jobs, protective services, personal care), and 4) elementary occupations (office cleaners, low-skilled maintenance and construction). The employer and income registers also provide annual earnings and income, which include both wages and any reported business income. Throughout our analysis, we focus on gross annual labor market earnings, which includes any business income earned throughout the year.

2.2 IFR Data

We measure industrial robot adoption using data from the International Federation of Robotics (IFR), covering the years 1993–2015. The IFR compiles information on robot shipments and installations across multiple industries. Our approach to measuring exposure to robots closely follows Acemoglu and Restrepo (2020), linking IFR data at a broad industry level—roughly two-digit NACE—with 19 distinct industries. Since Norway's own robot imports could potentially be endogenous to local labor market conditions, we instead use data from five other European countries—Denmark, Finland, France, Italy, and Sweden—to predict average yearly changes in robot installations at the local labor market level. The detailed construction of our industry- and local-level measures of robot exposure is presented in Section 3.

2.3 Labor Market Institutions

For our study, the most important institutional settings of the Norwegian labor market are coordinated wage-setting, open labor mobility for workers from within the European Economic Area (EEA), and a generous, work-compatible social insurance system. These features matter for our design and interpretation: while sectoral wage floors and compressed wage structures might shape the incidence of shocks, the 2004 EEA enlargement generated plausibly exogenous immigration inflows.

Norway has no economy-wide statutory minimum wage. Instead, minimum pay in a set of (mainly lower-wage) sectors is established by the general application of collective bargaining agreements. Collective bargaining coverage is high in international comparison, and the wage structure is compressed. Importantly, immigrant workers are generally covered by the same applicable minimum terms as natives in affected firms, limiting outright undercutting but not removing quantity pressures. The mix of sectoral wage floors and firm-level bargaining coverage should channel shocks into wage levels and firm premia.

The Norwegian welfare state provides generous insurance: unemployment insurance; sickness benefits with high initial replacement; work assessment allowance (AAP) during rehabilitation; and disability insurance (DI). As in many other countries, DI has become an increasingly important source of income: in 2015, nearly 10% of the overall population received DI.[5] DI can be partial and combined with earnings, subject to income thresholds, which facilitates income smoothing without full labor-force exit.

3 Empirical Framework

Understanding the causal impact of automation and immigration on labor market outcomes requires isolating exogenous variations in these shocks. In this section, we outline our identification strategy for capturing the exogenous exposure to both automation and immigration, drawing on established natural experiments and methodologies. Our strategy exploits two different sources of variation: shifts in automation at the local labor market level and shifts in immigration at the industry level. By combining these shifts, we can disentangle the independent and joint effects of each shock on worker outcomes.

3.1 Automation Shock

To isolate the impact of automation on labor market outcomes, we follow Acemoglu and Restrepo (2020) and estimate the predicted penetration of robots at the local labor market level. Specifically, this approach predicts the local labor market penetration of industrial robots based on its initial (1996) industry composition and the industry-specific robot adoption in other countries. Local labor markets with a higher concentration of industries experiencing rapid automation internationally are predicted to have higher exposure to automation. Because we rely on international (rather than Norwegian) trends in robot adoption, we mitigate endogeneity concerns and isolate exogenous variation in automation that is not driven by local economic conditions.

We begin by calculating the adjusted robot penetration at the industry level using import data from Denmark, Finland, France, Italy, and Sweden. By averaging the change in robot adoption per worker across multiple countries, we obtain an industry-level measure of automation that reflects international trends rather than domestic factors. The formula for the adjusted robot penetration in industry j is given by:

$$\overline{adjusted\ robot\ penetration_j} = \frac{1}{5}\sum_k \frac{M^k_{j,2015} - M^k_{j,1996}}{L^k_{j,1996}} \tag{1}$$

where: $M^k_{j,2015}$ is the number of robots in industry j in country k in 2015 (the end of our period of observation) and $M^k_{j,1996}$ is the corresponding figure in 1996 (the start of our period of observation). $L^k_{j,1996}$ is the baseline employment in industry j in country c in 1996. We illustrate the industry-level robot penetration over time in Figure 1a. As is clear from the figure, the growth in robot adoption in

[5]See "Table 11714: Recipients of disability benefit, by sex, age and education level 2015 - 2024", https://www.ssb.no/en/statbank/table/11714.

the period 1996 to 2015 mainly occurred in the manufacturing sector.

Because local labor markets in Norway differ in their initial industry composition, they face varying degrees of exposure to this international trend. We follow the literature (Acemoglu and Restrepo, 2020) and assume that local industries are similarly affected by international industry-level robot shocks. Specifically, we calculate ΔAuto_m, the predicted change in robot penetration in labor market m, by assigning each industry's adjusted robot penetration to a local labor market in proportion to that industry's 1996 employment share::

$$\Delta\text{Auto}_m = \sum_j \left(\frac{L_{j,m,1996}}{L_{m,1996}} \times \overline{\textit{adjusted robot penetration}}_j \right), \tag{2}$$

where: $L_{j,m,1996}$ is the 1996 employment in industry j within local labor market m and $L_{m,1996}$ is total 1996 employment in local labor market m. This design ensures that differences in initial industrial specialization drive observed variation in automation exposure.

Figure A.1 shows the spatial distribution of predicted automation shocks across Norway from 1996 to 2015. While there are some pockets of areas with larger increases in industrial robots, the shock is fairly spread across different regions of Norway. While all areas see increases in industrial robots, some local labor markets see smaller increases and some see much larger increases. This variation provides a basis for identifying the causal impact of automation on worker outcomes.

While our focus is on the importance of joint shocks rather than the direct labor market consequences of automation, it remains informative to understand the factors which are correlated with our exogenous measure of automation. Table A.1 performs a balance test which reports the factors at local labor market level which are correlated with the measured created in equation (2) as in Goldsmith-Pinkham et al. (2020); Borusyak and Hull (2020). Factors such as unemployment and education are not significantly correlated with future automation. Important for our empirical strategy, factors such as the share of workers employed in unlicensed construction industries (which are exposed to the immigration shock) and the general level of natives/immigrants in an area are also uncorrelated with the automation shock. In contrast, the initial manufacturing share in an area is positively and significantly correlated with the automation shock. Such a correlation is unsurprising, given the concentration of robots in the manufacturing industries (see Figure 1a). Given the potential importance of manufacturing areas for our automation shock, we include, as a robustness check, an additional control for pre-shock manufacturing share. In addition, results are unchanged when including LLM fixed effects, which account for any area-specific factor, including all the variables in Table A.1.

Figure 1: Describing the dynamics of the Automation and Immigration Shocks

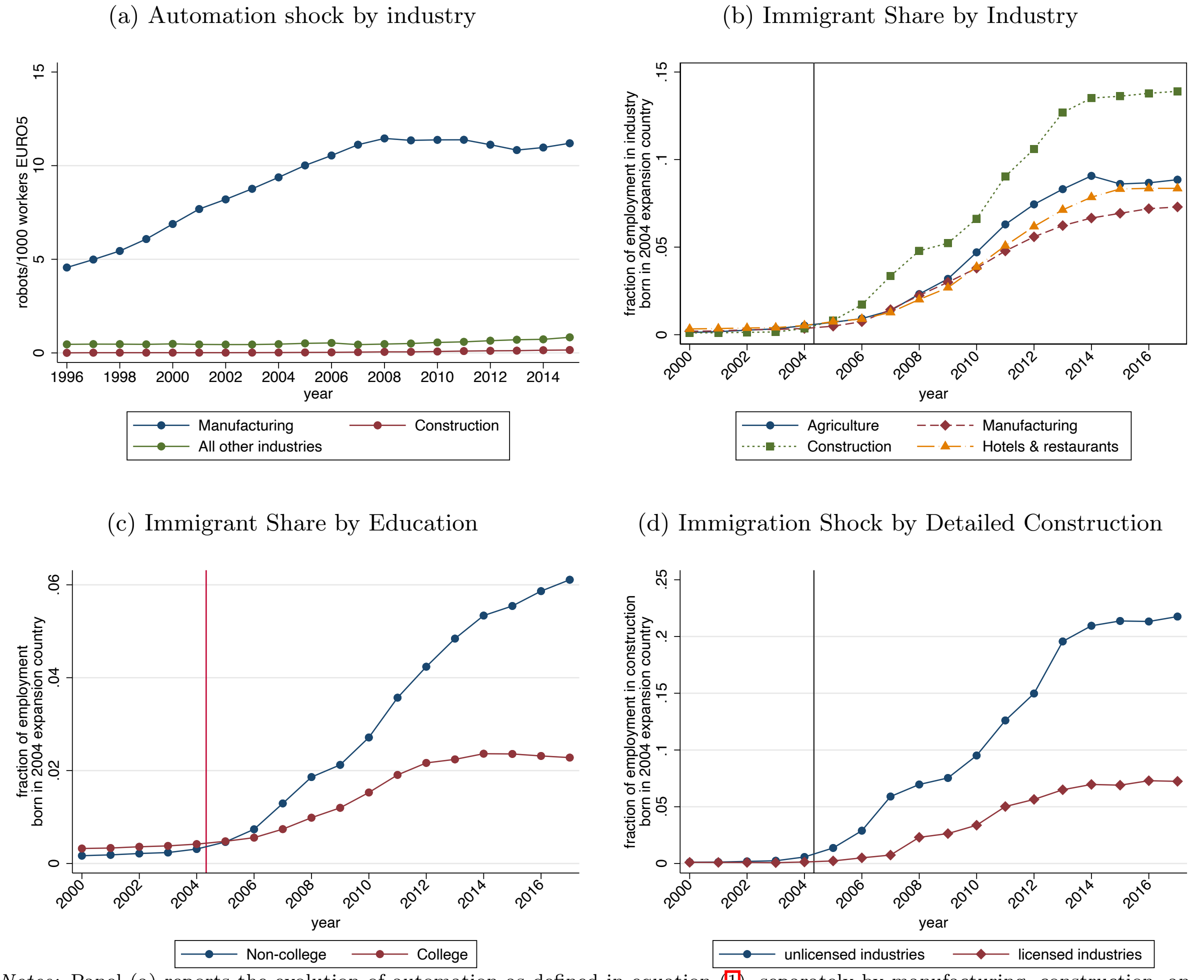

Notes: Panel (a) reports the evolution of automation as defined in equation (1), separately by manufacturing, construction, and all other industries. Panel (b) reports the share of immigrants employed in the top 4 industries affected by immigration. Panel (c) reports the share of immigrants employed in any industry, separately by college/non-college educated. Panel (d) reports, within the construction industry, the share of immigrants who are employed in either unlicensed or licensed industries.

3.2 Immigration Shock

During the same period in which automation advanced in Norway, immigration also surged. To isolate exogenous variation in immigration, we exploit the 2004 expansion of the European Union (EU), a policy event that significantly increased migration flows into Norway. Although Norway is not an EU member, it participates in the European Economic Area (EEA), which permits the free movement of goods, services, capital, and, importantly, people between EU and EEA countries. The 2004 enlargement admitted eight Central and Eastern European countries (the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia, and Slovenia) into the EU. While immigration from these countries to Norway was relatively restricted prior to 2004 due to stringent immigration policies, the 2004 EU enlargement enabled citizens from these new EU member states to access the Norwegian labor market more easily (Bratsberg and Raaum, 2012). This policy event resulted in a swift and sudden influx of immigrant workers.

The post-2004 inflow of immigrant workers tended to concentrate in specific industries, as seen in Figure 1b. Construction experienced the largest proportional increase in immigrant labor, with workers from the 2004 expansion countries comprising roughly 15% of the industry by 2017. Although immigration was most pronounced in construction, there was also considerable inflow into manufacturing, which had the fourth-largest increase in immigrant workers (Figure 1b). Taken together, these post-2004 immigrant inflows mainly affected blue-collar, non-college-intensive sectors. As seen in Figure 1c, the vast majority of new arrivals did not hold a college degree—non-college immigrants outnumbered college-educated immigrants by nearly four to one.

In our main analysis, we will focus on the joint impact of automation and immigration among construction workers. This enables us to leverage an important institutional feature that allows for a clean identification of immigration's impact: licensing requirements. Certain construction jobs in Norway have strict licensing requirements that are tied to Norwegian credentials and therefore that pose time-consuming and costly barriers to entry for foreign workers. This limits immigrant entry into "licensed" construction segments (Bratsberg and Raaum, 2012). Immigrants thus tended to cluster in unlicensed construction roles (e.g., general labor or non-specialized trades), leaving licensed segments relatively protected from this surge.

Figure 1d reveals the important role of licensing requirements in restricting immigration within the construction sector: while immigrants comprise 7% of employment in licensed construction industries, they make up 22% of unlicensed construction workers by 2017. Both industries had similar immigrant shares close to zero prior to the expansion of the EU, underlying the importance of the EU's expansion for immigration. In our main analysis, we will compare outcomes of those workers employed in unlicensed construction industries in 2000 (i.e., before the enlargement), and to outcomes of workers in licensed construction industries. This distinction allows us to compare labor market outcomes between workers directly affected by the immigration shock and those who were not, within the same geographic and economic environment. Because both sets of workers are in the same sector and geographic context, this distinction helps us cleanly isolate the causal effect of immigration—and its interaction with automation—on labor market outcomes.[6]

3.3 Regression Models

To investigate the impact of these automation and immigration shocks on individual labor market outcomes, we employ regression models that relate changes in individual outcomes to shock exposure. We start by identifying the effect of exposure to one shock—automation—in isolation. Specifically, we estimate the following baseline model:

$$\Delta y_i = \beta_0 + \beta_1 \Delta Auto_{m(i)} + \gamma_{c(i)} + \lambda_{e(i)} + \varepsilon_i \tag{3}$$

where: Δy_i is the change in a given outcome for individual i from 2000 to 2015 and $\Delta \mathrm{Auto}_{m(i)}$ is the predicted change in automation exposure in the local labor market m where individual i resides, as measured in equation (1). $\gamma_{c(i)}$ represents cohort fixed effects for the birth cohort c of individual i, while $\lambda_{e(i)}$ denotes education fixed effects for the education level e of individual i. ε_i is the error term. Throughout, standard errors are clustered at the initial local labor market level. Equation 3 captures the average effect of automation exposure on changes in outcomes across all individuals, controlling for cohort and education effects.

We then contrast these effects of exposure to automation with estimates of joint exposure to both automation and the intersecting shock of immigration. To estimate the differential impact of automation on workers affected by the immigration shock, we extend the model by introducing an interaction term between automation exposure and an indicator variable for being treated by the immigration shock. This model compares unlicensed and licensed construction workers within the same area. Specifically, we estimate the following model:

$$\Delta y_i = \delta_0 + \delta_1 \Delta Auto_{m(i)} + \delta_2 ImmExp_i + \delta_3 \Delta Auto_{m(i)} \times ImmExp_i + \gamma_{c(i)} + \lambda_{e(i)} + \varepsilon_i \tag{4}$$

where: ImmExp_i is an indicator variable equal to 1 if individual i was employed in unlicensed construction industries in 2000 (thus directly affected by the immigration shock), and 0 if employed in licensed construction industries. Unlicensed workers are more likely to have been affected by the influx of immi-

[6] In Norway, vocational education follows a 2 + 2 model (two years school-based + two years apprenticeship), culminating in a trade certificate. This shared pipeline means licensed and unlicensed construction workers are closely comparable on formal schooling, reinforcing the credibility of our within-sector comparison.

grants following the 2004 EU expansion, as unlicensed industries faced fewer barriers to entry for foreign workers. In contrast, licensed construction industries required specific certifications, limiting the impact of immigration on those sectors. In our final sample, there is an approximately equal split between workers in unlicensed and licensed industries.

In Equation 4, δ_1 measures the effect of automation exposure on earnings changes for licensed construction workers who are not exposed to immigration. δ_2 captures earnings changes between those exposed vs. not exposed to immigration due to licensing protection. Both δ_1 and δ_2 are interpreted as the effect of one shock, in the absence of the other shock. For instance, δ_1 estimates the change in earnings among workers who are affected by automation, but are insulated from the immigration shock. Similarly, δ_2 estimates the change in earnings among those exposed to the immigration shock, absent any change in the automation shock.

By including the interaction term between both shocks, we assess whether the impact of automation differs for workers who were also exposed to the immigration shock. The coefficient on the interaction term $\Delta Auto_{m(i)} \times ImmExp_i$, δ_3, captures the additional effect of automation for those exposed to immigration compared to those not exposed to immigration due to licensing requirements. When looking at earnings as the outcome, a negative and significant δ_3 would indicate that automation has a more adverse effect on earnings for those simultaneously exposed to immigration.

Our empirical strategy exploits two different sources of variation: shifts in automation at the local labor market level and shifts in immigration at the industry level. While these shocks occur at different levels, it is useful to understand the overlap between the two of them at the area-level. Figure A.2 presents the correlation between the change in immigration and automation at the local labor market level. There is no systematic relationship between the immigration shock and automation shock. Areas differentially exposed to automation experience similar immigration inflows, suggesting that the automation shock alone does not influence local immigration patterns.

4 Results

4.1 The Area-Level Impact of Industrial Robots

We begin by presenting the estimates from equation (3) in Table 1, showing the impact of predicted automation in the local labor market on the earnings among the population of male workers aged 18 to 49 in 2000.[7] Column 1 of Panel A shows that the negative consequences of automation are concentrated among non-college-educated workers, whose earnings decline by 2.4% for each additional robot per 1,000 workers. Similarly, the negative consequences of automation are concentrated among the manufacturing and construction industries, two industries which are dominated by non-college educated workers. For

[7] As shown in Table B.1, the impact of automation is concentrated among men. There is no effect on the earnings or employment of women.

college educated workers and those employed outside of manufacturing and construction, there are no significant effects of automation on earnings. The declines in earnings are not accompanied by declining employment rates for any workers (Panel B), suggesting that the overall earnings decline is at least partially due to reductions in relative wages over time. The size of this effect is notable compared to previous findings: Acemoglu and Restrepo (2020) find that men with less than college education experience declines in hourly wages of 0.75–1.1% as a result of a similar change.

Table 1: The Impacts of Automation by Education/Industry

	Education		Industry		
	(1) Non-College Educated	(2) College Educated	(3) Manufacturing	(4) Construction	(5) All Other Industries
Panel A: Δ Log Earnings					
Δ Auto	-0.024*	-0.004	-0.032**	-0.029***	-0.009
	(0.014)	(0.014)	(0.014)	(0.010)	(0.014)
N	521189	242265	139337	92305	531813
Panel B: Δ Employed					
Δ Auto	-0.001	0.003	-0.003	-0.003	0.004
	(0.007)	(0.003)	(0.004)	(0.005)	(0.007)
N	573218	252686	152419	99452	574034

Notes: Sample of men aged 18–49 initially employed in 2000. Figure presents estimates of equation (3), estimated separately for workers with different education levels (columns 1–2) and workers initially employed in different industries columns (3–5). Outcome variables correspond to the change in log earnings (panel a) and the change in employment (panel b). Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

4.2 Between-Industry Spillover Effects of Automation

Next, we investigate the impact of industrial robots by industry. We focus on manufacturing and construction versus all other industries. As shown in column 3 of Table 1, we see that there is unsurprisingly a strong effect of industrial robots on the earnings among men employed in the affected industry of manufacturing. However, as shown in column 4, there is a similarly large 2.9% earnings loss among workers initially employed in construction. Contrastingly, there is no statistically significant impact on earnings for workers in all other industries.

In the following, we provide evidence which suggests that this cross-industry spillover is primarily driven by worker substitutability and employment flows between manufacturing and construction. The similarities between workers in these two industries are detailed in Table C.1, showing comparable demographics, earnings, and educational profiles, especially relative to workers in all other industries. Specifically, around 80–85% of workers in construction and manufacturing lack college education while, in contrast, fewer than 65% of workers in other industries lack college education.

To further investigate whether worker substitutability and mobility drive these spillovers, we predict the likelihood of workers transitioning into manufacturing based on the specific tasks they perform and their individual characteristics.[8] Figure 2 illustrates the predicted probabilities of transitioning to manufacturing employment for non-construction (Panel a) and construction workers (Panel b), respectively. Non-construction workers' predicted likelihood of transitioning into manufacturing clusters at predominantly lower levels, with relatively fewer individuals exhibiting a probability of 0.5 or above. In contrast, construction workers generally have a higher probability of moving into manufacturing. Specifically, the probability distribution for construction workers is centered around 0.5 and exhibits considerable dispersion, indicating diverse yet overall higher likelihoods of industry transition. Furthermore, certain construction workers possess especially high probabilities of transitioning to manufacturing, driven by their specific characteristics and task profiles.

Building on these results, Panel (c) of Figure 2 further investigates the differential impact of automation on construction workers, distinguishing between those above and below the median probability of transitioning to manufacturing. We find that the negative earnings effects of automation are notably concentrated among construction workers with higher probabilities of shifting industries. This indicates that the spillover effects of automation predominantly affect workers in closely related industries who possess similar skills and perform comparable tasks. These results suggest that the underlying mechanism likely involves intensified competition for construction jobs as automation-displaced manufacturing workers seek employment opportunities elsewhere, thus expanding labor supply and placing downward pressure on wages. Additionally, this increased competition diminishes the outside employment options available to construction workers who share comparable skillsets with their manufacturing counterparts.

4.3 The Joint Effect of Automation and Immigration

The preceding analysis highlights the significant impact of automation on non-college-educated male workers, not only within the manufacturing industry but also through spillover effects to the construction industry. The results underscore the interconnectedness of industries and emphasize between-industry worker substitutability as a critical factor influencing the impact of labor market shocks.

In this section, we focus on the sample of construction workers. These workers were affected not only by the automation shock, but also a contemporaneous shock of a large inflow of non-college educated immigrants following the 2004 EU expansion. However, as detailed in Section 3, licensing requirements acted as employment barriers for immigrants, particularly affecting unlicensed industries and thereby shaping the interaction between immigration and automation shocks. This institutional feature of the Norwegian construction industry allows us to explore the interaction effects between automation and

[8] We employ four task dimensions: routine, service, math, and social tasks, derived from occupational data first available in 2003. We match occupations in Norway to these task measures following Deming (2017), utilizing task classifications from O*NET. Individual characteristics include age and age squared.

Figure 2: The Differential Impacts of Automation Across Different Types of Non-Manufacturing Workers

(a) Distribution of Probability of Shifting to Manufacturing, Non-Construction

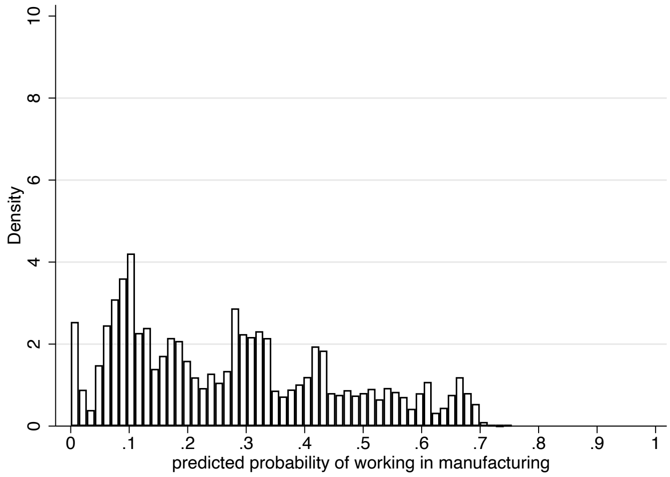

(b) Distribution of Probability of Shifting to Manufacturing, Construction

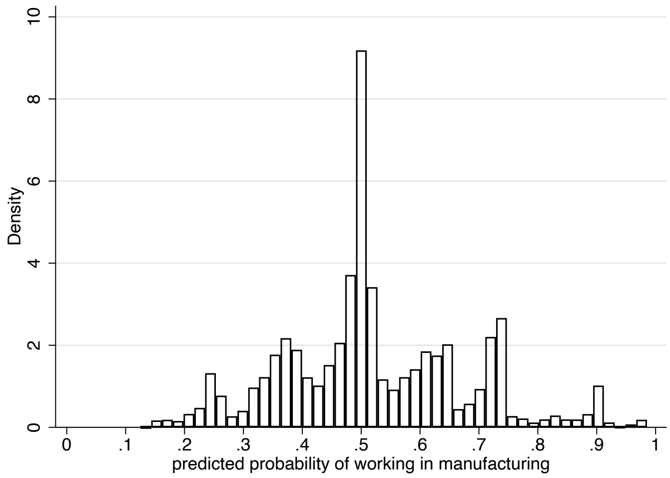

(c) Differential Impacts of Automation by Probability of Shifting to Manufacturing, Construction Workers

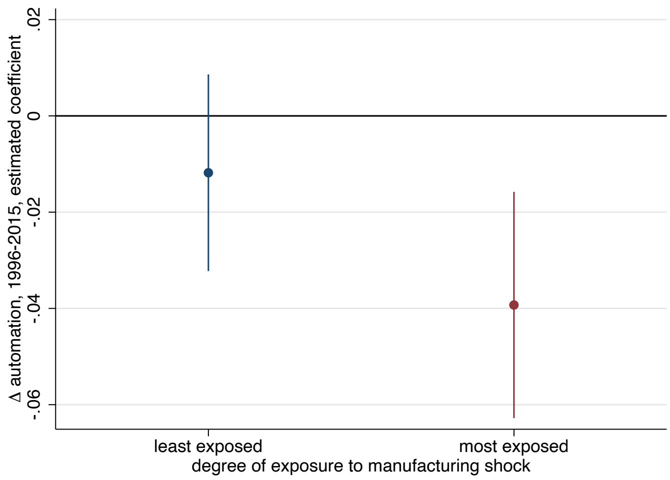

Notes: panel (a) plots the distribution of the probability of shifting to manufacturing among all non-construction employees. Panel (b) plots the distribution of the probability of shifting to manufacturing among all construction employees. Panel (c) plots the impact of automation among construction workers, separately by whether they are above/below the median probability of shifting to manufacturing. Sample of men aged 18–49 initially employed in 2000. Tasks are measured in 2003, the first year when data on occupations is available.

immigration to demonstrate the importance of considering multiple, concurrent labor market shocks to fully understand their combined impact on non-college educated workers. Specifically, within construction workers, we compare the differential impacts of automation among those working in unlicensed construction industries and those in licensed construction industries.

Table 2 estimates the differential impact of automation, comparing the effects of automation among construction workers exposed to the immigration shock to those insulated from immigration. Focusing on the interaction between the change in automation and exposure to labor immigrants reveals that the earnings losses resulting from automation are significantly worse among unlicensed construction workers who are also treated by the immigration shock (column 1). For these workers, earnings decline by an additional 2.8% for every additional robot per 1,000 workers compared to licensed construction workers, clearly demonstrating how immigration amplifies the negative consequences of automation.

Table 2: The Interaction Between Automation & Immigration

	Δ Log Earnings		
	(1) Full Sample	(2) Licensed	(3) Unlicensed
Δ Auto	-0.010	-0.010	-0.036***
	(0.014)	(0.014)	(0.011)
ImmExp	-0.022		
	(0.021)		
Δ Auto × ImmExp	-0.028**		
	(0.012)		
Education FE (1 digit)	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes
Individuals	91884	46110	45774

Notes: Column (1) reports estimates from equation (4). Columns (2)–(3) report estimates from equation (3), separately for licensed and unlicensed workers. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

In contrast to the joint effect of automation and immigration, we see no significant effect of each shock in isolation. In column 1, the non-interacted coefficient for automation indicates the earnings impact among construction workers not directly affected by immigration. For these workers, automation alone leads to a modest and statistically insignificant decline in earnings of approximately 1%. This suggests that absent the immigration shock, the effects of automation are considerably smaller in magnitude. Similarly, the non-interacted coefficient for unlicensed workers shows the earnings effect of exposure to immigration alone, absent automation. Like automation, immigration by itself results in a negative but statistically insignificant earnings impact, indicating that the observed declines primarily arise from workers simultaneously exposed to both shocks.

Further emphasizing this interaction, we compare unlicensed and licensed construction workers directly. Results shown in columns (2) and (3) reveal striking differences: while licensed workers experience no significant earnings changes due to automation, unlicensed workers suffer substantial earnings reductions of nearly 4%. These findings strongly suggest that the interplay of shocks exacerbates adverse earnings outcomes.[9] Figure 3 provides additional insight into these dynamics by illustrating the impacts of automation over time, comparing unlicensed and licensed construction workers separately. Prior to the immigration shock in 2004 (indicated by the vertical line), the estimated earnings effects for both groups are flat and statistically indistinguishable from zero. This similarity supports the parallel trends assumption and suggests that both groups of construction workers followed similar earnings trajectories before the onset of the combined shocks. After 2004, however, a clear divergence emerges: unlicensed workers experience significant and increasingly negative earnings effects due to automation, while licensed workers experience no effects. The timing of these observed effects closely aligns with the occurrence of the respective shocks, as depicted in Figure 1. Specifically, the automation shock peaks around 2008, whereas immigration flows substantially increase between 2004 and 2013 before stabilizing.[10]

Figure 3: The Dynamic Impacts of Automation for Unlicensed and Licensed Construction Workers

(a) Unlicensed (b) Licensed

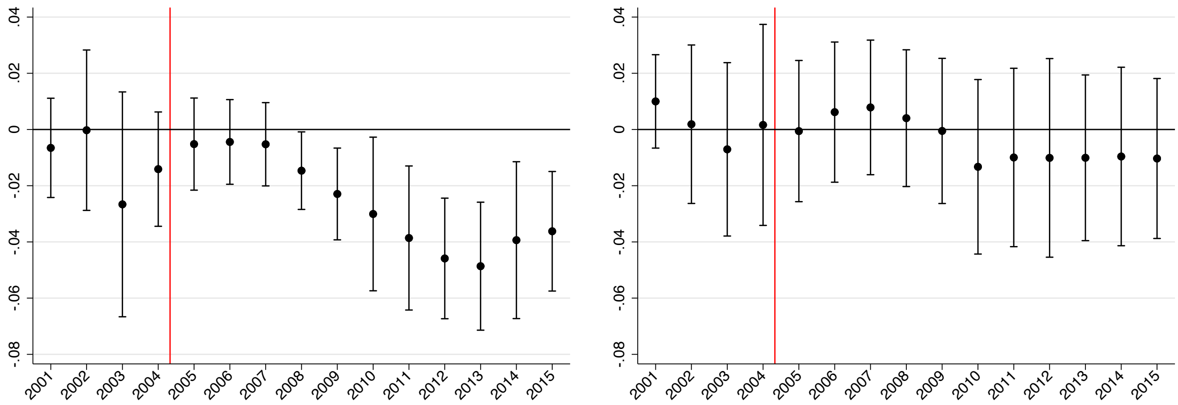

Notes: This figure shows the dynamic impact of automation for unlicensed and licensed construction workers. The specification uses a fixed automation shock from 1996–2015 while allowing earnings impacts to vary annually from 1998–2015. The vertical line represents the 2004 EU expansion that triggered the inflow of immigrant workers into unlicensed construction industries.

The contrast between licensed and unlicensed workers is unlikely to be driven by underlying skill differences. In Norway, both groups follow the same standardized vocational track—two years of upper-secondary vocational coursework followed by a two-year apprenticeship—so their formal qualifications are comparable. Furthermore, all specifications in Table 2 absorb residual schooling heterogeneity by including education fixed effects, ensuring that differences in completed education cannot explain the patterns we observe. Finally, our unique data allows us to condition directly on cognitive ability: when we control for a full set of fixed effects for IQ score collected at conscription, the interaction coefficient

[9]These declines in earnings are not explained by significant employment changes over time (see Appendix Table D.1).

[10]A similar event-study illustration for the differential effects of automation between unlicensed and licensed construction workers is shown in Appendix Figure E.1, and, extending further back, in Appendix Figure E.2. In both figures, the pre-trends are flat and statistically indistinguishable from zero, reinforcing that differences observed post-2004 are indeed attributable to the interaction of automation and immigration rather than pre-existing trends.

between automation and immigration exposure remains virtually unchanged (see Table 3). Based on this, the larger earnings losses suffered by unlicensed construction workers likely reflect exposure to both the automation and immigration shocks rather than pre-existing differences in skill endowments.

Table 3: The Interaction Between Automation & Immigration

	Δ Log Earnings		
	(1) Full Sample	(2) Licensed	(3) Unlicensed
Δ Auto	-0.008	-0.009	-0.034***
	(0.015)	(0.015)	(0.010)
ImmExp	-0.014		
	(0.022)		
Δ Auto × ImmExp	-0.027**		
	(0.013)		
Education FE (1 digit)	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes
IQ FE	Yes	Yes	Yes
Individuals	91884	46110	45774

Notes: This table reports coefficients from specifications that include IQ fixed effects. Column (1) reports estimates from equation (4). Columns (2)–(3) report estimates from equation (3), separately for licensed and unlicensed workers. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

The joint effect of automation and immigration remains robust to several sensitivity checks, as reported in Table 4. First, omitting the 1% of local labor markets that experienced the largest rise in automation leaves the estimates essentially unchanged, showing that our results are not driven by extreme cases (column 1). Second, controlling for initial manufacturing employment shares confirms results are not driven by general manufacturing trends (column 2). Third, accounting for the rise in Chinese imports as in Balsvik et al. (2015) does not alter our findings, ruling out confounding from contemporaneous trade shocks (column 3). Fourth, results are also robust when controlling for demographic factors (initial shares of females, college-educated, Norwegian-born, and older workers) and population density (columns 4 and 5). Fifth, accounting for initial unionization rates has no impact on our conclusions (column 6). Including all these controls simultaneously yields consistent results (column 7), highlighting the robustness of the observed interaction effects. The robustness to the inclusion of an extensive set of controls confirms that differences across local labor markets in their preexisting trends do not drive the observed changes in earnings among those exposed to both the automation and immigration shocks. Finally, column 8 introduces local-labor-market fixed effects to net out *all* time-invariant market-specific factors. Although this specification precludes the inclusion of further local labor market-level covariates (as they are collinear with the local labor market fixed effects), it still recovers the joint effect of automation and immigration. Taken together, these exercises suggest that the large earnings penalties we document stem from the joint shock of automation and immigration rather than from pre-existing trends or other contemporaneous forces.

4.4 Mechanisms: Household and Worker-Level Responses

Our earlier findings highlight that construction workers jointly affected by automation and immigration experience notably worse outcomes compared to those facing only one shock. In this section, we explore the underlying mechanisms explaining why these combined shocks disproportionately harm earnings. Previous results indicate no differential changes in employment among these doubly impacted workers. Therefore, any worker adjustments in response to these shocks occur within employment, potentially affecting worker-level outcomes such as job types or employer characteristics.

4.4.1 Household responses

Table 5 investigates household responses to the joint automation–immigration shock. For context, note that at the start of our period (2000), 43% of workers in our sample were the sole earner in their household. Column 1 reveals that redefining the outcome as total household labor income leaves the interaction coefficient virtually unchanged relative to the individual-level estimates, implying that spouses neither offset nor amplify the earnings losses of the directly treated worker. Consistent with this interpretation, column 2 shows no significant effect on *added-family earnings*—the combined contribution of every household member except the focal worker. Column 3 further indicates that the probability that the household becomes (or remains) a single-earner household is unaffected. Turning to other potential margins of adjustment, columns 4 and 5 detect no discernible changes in marriage formation or fertility. Taken together, these patterns suggest that the earnings effect of the joint shock of automation and immigration on treated workers is not smoothed within the household.

4.4.2 Occupational shifts

The existing literature emphasizes that occupational factors are central to understanding the impacts of both automation and immigration shocks. Similar to our findings, Acemoglu and Restrepo (2020) demonstrate that the adverse effects of automation predominantly affect blue-collar occupations, which are particularly vulnerable to displacement by technological advances. Furthermore, Dauth et al. (2021) highlight significant occupational reallocation resulting from automation, observing a notable reduction in manufacturing employment offset by an equivalent increase in service-sector jobs. Occupational decline negatively impacts workers (Edin et al., 2023), aligning with a substantial body of research on skill-biased technological change documenting the erosion of mid-skilled occupations (Autor et al., 2006; Goos et al., 2009). Similarly, immigration shocks trigger occupational shifts among native workers. Peri and Sparber (2009) find that natives increasingly move into occupations emphasizing communication and language-based tasks in response to immigration influxes. Additionally, given direct competition from immigrants

Table 4: The Interaction Between Automation & Immigration, Robustness

	Outcome: Δ Log Earnings							
	(1) Excluding High Exposure	(2)	(3)	(4)	(5) Additional Controls	(6)	(7)	(8)
Δ Auto $\times$ ImmExp	-0.033**	-0.027**	-0.028**	-0.027**	-0.00269**	-0.027**	-0.026**	-0.025**
	(0.014)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Manufacturing share	No	Yes	No	No	No	No	Yes	No
Chinese Imports	No	No	Yes	No	No	No	Yes	No
Demographic Controls	No	No	No	Yes	No	No	Yes	No
Union Density	No	No	No	No	No	Yes	Yes	No
Population Controls	No	No	No	No	Yes	No	Yes	No
LLM FE	No	No	No	No	No	No	No	Yes
Observations	91081	91884	91884	91884	91884	91884	91884	91884

Notes: Table reports estimates from equation (4), where each column corresponds to a separate regression for each robustness check. Column (1) excludes the places with the highest rates of automation. Columns (2)–(6) add one control variable in each model, while column (7) includes all of the variables added from columns (2)–(6) in the same regression. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

Table 5: The Interaction Between Automation & Immigration, Family Earnings

	(1) Δ Log Family Earnings	(2) Δ Log Added Family Earnings	(3) Δ Sole Earner	(4) Δ Married	(5) Δ Children
Δ Auto × ImmExp	-0.026**	-0.006	-0.002	0.005	-0.011
	(0.011)	(0.013)	(0.006)	(0.004)	(0.010)
Education FE (1 digit)	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes
Individuals	94493	42796	94493	98922	98922

Notes: This table reports estimates from equation (4), for different outcomes. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

in sectors such as construction (Bratsberg and Raaum, 2012), native workers may strategically pursue employment in fields where they maintain a comparative advantage relative to immigrant labor.

Table 6 reports occupational changes among workers jointly affected by the automation and immigration shock.[11] Occupations are classified into four broad groups: professional (primarily white-collar), blue collar, service, and elementary occupations. Service occupations predominantly include roles such as shop salespersons and protective service workers, while elementary occupations mostly comprise office cleaners and support workers in maintenance and construction. We find that the joint shock causes construction workers to move down the occupational ladder, reflected by a substantial rise in the likelihood of employment in service occupations (column 1). The growth in service occupations mainly occurs at the expense of blue-collar positions (column 3) and, to a lesser extent, professional occupations (column 2). This shift toward service roles is notable: among workers exposed to both shocks, one additional robot per 1,000 workers corresponds to a 1 percentage point increase in the probability of employment in services. Given that only 5.5% of the sample were employed in services in 2003, this change represents an 18% relative increase.

Service occupations typically offer lower wages compared to blue-collar occupations; among all workers, men employed in service occupations earn approximately 35% less than those in blue-collar roles. Using this wage differential, we can quantify the earnings loss associated with shifting from blue-collar to service occupations. Specifically, the observed 1 percentage point shift toward service occupations from blue-collar positions can account for roughly 13% of the total earnings decline observed in column (1) of Table 2 ($\frac{0.01\times -0.35}{-0.028} = 0.13$). Thus, occupational downgrading represents a significant mechanism through which the joint shock of automation and immigration negatively affects workers' earnings.

[11] Occupations are only available from 2003. We therefore restrict the sample to those with an occupation defined in 2003 and report changes in occupations from 2003 to 2015.

Table 6: The Interaction Between Automation & Immigration, Occupational Shifts

	(1) Δ Service	(2) Δ Professional	(3) Δ Blue Collar	(4) Δ Elementary
Δ Auto $\times$ ImmExp	0.010**	-0.007	-0.018*	0.005
	(0.004)	(0.005)	(0.009)	(0.004)
Education FE (1 digit)	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Individuals	78891	78891	78891	78891

Notes: This table reports estimates from equation (4), with the outcome variable as the change in occupation. Changes in occupations are measured from 2003–2015, where 2003 is the first year which occupation data is available. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

4.4.3 Employer shifts

While workers affected by both shocks reallocate their labor supply to specialize in occupations performing service tasks, the majority of the earnings decline is unexplained by occupational shifts. In this Section, we focus on employer shifts as another important mechanism. The importance of employers in determining workers' labor market outcomes has been emphasized by an extensive literature (Abowd et al., 1999; Card et al., 2013; Sorkin, 2018). Lachowska et al. (2020); Schmieder et al. (2023) show that the loss of employer-specific earnings premiums are an important factor which explains the magnitude of earnings losses following job displacement shocks.

We examine the importance of changes in plant-level factors, estimating employer-specific premiums as in the AKM fixed effects literature. Our measure of employer-specific premium is estimated as follows:

$$ln(earnings)_{iey} = \alpha_i + \gamma_{e(i,y)} + \theta_y + \varepsilon_{iey} \tag{5}$$

where α_i and $\gamma_{e(i,y)}$ correspond to worker- and employer-specific fixed effects, our measure of the earnings premium attributed to each factor. Equation (5) is estimated among the universe of workers, taking the largest connected employer set, as is standard in the AKM literature (Card et al., 2013). We interpret $\gamma_{e(i,y)}$ an employer-specific earnings premium, measured net of the influence of fixed worker factors and year fixed effects (θ_y).

Table 7 shows that workers affected by both shocks shift to lower quality employers who have lower earnings premiums. As a result of both shocks, affected workers step down the pay scale and move to a plant which pays 1 log point lower (column 1). This implies that the shift to lower quality employers explains around 35% of the observed 2.8% decline in earnings. As we examine a sample of workers who are all initially employed, and the estimated employer-specific premiums are time-invariant, the observed changes in the estimated values of $\gamma_{e(i,y)}$ correspond to changes in job to job mobility.

While the measure of employer-specific premium in column (1) is our preferred measure of employer quality, columns (2) and (3) show that workers move to plants which are lower ranked in the pay distribution. Consistent with the results using AKM fixed effects, workers affected by both shocks move to plants which are 1.3 percentiles lower paid. As employers at different points in the pay scale differ in a number of ways over and above what is measured by the employer-specific fixed effect, the downward shift in the employer pay distribution corresponds to an even larger fraction of what employer factors can explain.[12] Workers treated by both shocks also experience absolute declines in the rank of their employer, where the probability of moving down the employer pay distribution increases by 3.1 percentage points, an 8% increase relative to the average.

Columns (4)–(6) combine the estimated worker- and employer-specific premiums to show that workers treated by both shocks experience declines in match quality. We define match quality as the absolute value of the distance between the estimated worker and employer fixed effects, $|\alpha_i - \gamma_{e(i,y)}|$. Large positive values indicate a divergence between the estimated measures of worker and employer quality, while small positive values indicate a better quality match. As before, both measures are time invariant, and the measure of match quality can only change when workers change employers through job to job mobility.

As a result of being exposed to both the automation and immigration shocks, treated workers experience significant declines in match quality: the distance between worker and employer-specific premiums increases by 0.013, corresponding to a 3% increase in this gap. Columns (5) and (6) split the sample to examine whether those with an initial bad or good match drive the overall shifts in match quality. Those with bad initial matches drive the overall results, where those with good initial matches experience no significant changes in match quality. As such, being treated by both shocks makes those with initial bad matches shift to an employer which results in an even worse match.

Table F.1 shows that while employer-specific pay premiums and match quality decline, there are no corresponding shifts in other measures of plant quality such as education levels of coworkers or the fraction of natives employed in the plant. As such, the decline in employer-specific earnings is related to pay rather than changes in the composition of coworkers. Taken in tandem, results suggest that getting stuck at lower paid employers, and to some extent lower paid service occupations, are important mechanisms behind why workers are worse off following the combined immigration and automation shocks. Jointly, these factors explain approximately 50% of the observed decline in earnings among workers affected by both shocks.

[12]Moving one percentile rank down the pay distribution corresponds to roughly a 2% decline in median earnings, such that the downward shift to lower paid employers is sufficiently large to explain a large portion of the overall decline in earnings.

Table 7: The Interaction Between Automation and Immigration, plant-level outcomes

				Entire Sample	Bad Initial Match	Good Initial Match
	(1) Δ Plant Specific Premium	(2) Δ Plant Rank	(3) Decline in Plant Rank	(4)	(5) Δ Match Quality	(6)
Δ Auto$\times$ ImmExp	-0.010**	-1.289**	0.031**	0.013**	0.025***	0.002
	(0.005)	(0.523)	(0.013)	(0.006)	(0.009)	(0.004)
Education FE	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Individuals	81956	84888	84889	81956	40997	40958

Notes: Table reports estimates from equation (4), with outcome variables measured at the plant-level. Column (1) measures the employer-specific earnings premium, calculated in equation (5) as in the AKM fixed effects literature. Column (2) measures the change in the percentile rank of plants, for the median earnings of all workers. Column (3) measures a variable equal to 1 if an individual experiences a decline in the percentile rank of their employing plant from 2000–2015. Columns (4)–(6) define a measure of match quality, as $|\alpha_i - \gamma_{e(i,y)}|$ from equation (5), estimating changes in this measure for the entire sample, those with initial bad matches, and those with initial good matches respectively. Sample sizes of columns (1) and (4) differ as the employer-specific earnings premium in column (1) is estimated only among the largest connected set of plants and workers. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

4.5 Policy Implications and the Role of Social Insurance

Intersecting labor market shocks widen inequality not just between the non-college and college educated, but also within the non-college educated. Our evidence shows that exposure to automation and immigration in combination moves workers down the job ladder—toward lower-paying firms and into lower-wage service jobs. However, we see no impact on overall employment. As a result, while affected workers see wage declines, the social insurance of the welfare system is unlikely to fully compensate for the persistent earnings losses.

Table 8 illustrates the modest mitigation effects of the social insurance in the Norwegian welfare system. Relative to an earnings decline of 2.8% for jointly exposed workers, pre-tax income falls by 2.5% (column 1) and after-tax income by 2.3% (column 2). A 4% reduction in taxes paid (column 3) offsets a small share of the earnings loss, consistent with progressivity. The main transfer margin is disability insurance (DI): the probability of claiming DI rises by 0.5 percentage points (column 4), a 7.7% increase relative to the average among those exposed only to automation in 2015.[13] This statistically meaningful uptick in DI receipt offsets only a small fraction of earnings losses.

The DI responses are also consistent with partial labor-force attachment. DI in Norway is awarded on medical and non-medical criteria (Kostøl and Mogstad, 2014), and benefits can be combined with some employment while maintaining eligibility. Given the magnitude of the increase in DI is slightly larger than the non-significant declines in employment (see Table D.1), this suggests that those who claim DI as a result of both shocks also remain in some employment alongside. This is consistent with prior work in Norway, which shows those on DI have considerable capacity for work (Kostøl and Mogstad, 2014).

[13] There are no corresponding changes in the probability of receiving other benefits/transfers.

Table 8: The Interaction Between Automation & Immigration, The Role of Transfers

	(1) Δ Log Pre-Tax Income	(2) Δ Log After Tax Income	(3) Δ Log Assessed Tax	(4) Δ Claim DI
Δ Auto × ImmExp	-0.025***	-0.023***	-0.040***	0.005**
	(0.007)	(0.006)	(0.011)	(0.002)
Education FE (1 digit)	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Individuals	97994	97994	96996	97994

Notes: Columns (1) reports estimates from equation (4), for income, which is the sum of earnings and any transfers. Column (2) defines a variable equal to one if an individual is claiming disability insurance (DI). Standard errors are clustered at the initial local labor market level. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

Taken together, three implications emerge from our results. The central margin of adjustment is reallocation rather than non-employment: workers exposed to both automation and immigration shocks move to lower-paying employers and into lower-paid service occupations, with roughly half of the total earnings loss explained by these two channels, and employer downgrades accounting for the larger share. This underscores the importance of policies that improve job matching and wage progression, not simply employment per se. At the same time, standard insurance mechanisms attenuate but do not neutralize losses, despite higher DI uptake and lower taxes. Finally, the outside option of switching into manufacturing weakens precisely when automation reduces demand there as well, rendering "nearby" job moves less remunerative and magnifying mismatch.

These findings point to policy directions that emphasize the quality of reallocation rather than the quantity of employment. Because employer downgrading explains a large share of losses, placement and counseling services could be tied more directly to transitions into higher-premium firms, with active labor market programs incentivized to improve wage progression. Retraining and upskilling efforts should focus on tasks and sectors less exposed to automation or immigration. The observed increase in DI take-up, consistent with partial work capacity, suggests a role for strengthening work-oriented elements of disability insurance. Geographically targeted active labor market programs may also be warranted in high-exposure regions, while reforms to licensing could ease upward mobility into better-paying, less-exposed occupations. More broadly, our evidence highlights the need for policy that insures short-run losses in ways that accelerates productive reallocation into firms and occupations with higher pay premia, so that workers are not left stuck at the bottom of the job ladder.

5 Conclusion

This paper has shown that two of the most prominent labor market shocks of recent decades—automation and immigration—not only coexist in time and space but also amplify one another's effects on vulnerable workers. Our results demonstrate that automation and immigration should not be analyzed in isolation, as their interaction alters both the magnitude and the interpretation of each shock's impact. Automation reduces earnings not only in directly exposed manufacturing but also in construction through cross-industry spillovers, and immigration exacerbates these effects by limiting workers' outside options. Workers facing both shocks suffer substantially larger earnings losses than those facing either in isolation.

Our findings shed light on the mechanisms behind these outcomes. Workers exposed to both shocks move down the job ladder, away from higher-paid blue-collar occupations and into lower-paid service jobs, while simultaneously transitioning into employers with systematically lower wage premia. Crucially, the employer margin explains a large fraction of the observed earnings decline, underscoring the importance of firm heterogeneity in shaping distributional consequences.

Turning to social insurance, we find that disability insurance provides partial compensation but offsets only a modest share of long-run earnings losses. This is consistent with evidence that DI systems insure against adverse shocks but often at the cost of weaker attachment to high-paying jobs. These results imply that cushioning income losses requires more than transfer programs: policies that enhance reallocation quality—through targeted retraining, task-specific upskilling, and support for transitions into higher-premium firms—are critical to prevent workers from becoming stuck at the bottom of the job ladder.

Our results focus on non-college educated incumbent workers, for whom we can identify the labor market consequences of two exogenous shocks. Such a focus is necessary to isolate workers for whom we possess robust identification strategies for both shocks. Both the automation and immigration shocks lead to changes in demand for certain types of labor, and our findings show that college educated are more insulated from the shocks which afflict the non-college educated. Non-college educated incumbent workers may be more disposed to the negative displacement effects of automation relative to other workers and, as such, the average effects on local labor markets themselves may differ from our focus on non-college educated workers who are treated by both shocks. For instance, employers in treated industries may benefit from lower wage bills as a result of increasing immigration, but understanding the dynamic general equilibrium consequences of the interaction between both shocks remains beyond the scope of this paper. However, Jaimovich et al. (2021) highlight the importance of the general equilibrium effects of automation and Mandelman and Zlate (2022) show that the joint effects of automation and immigration shocks also matter in general equilibrium.

The objective of this paper is not to examine the potentially exhaustive set of interactions among all possible shocks which affect workers differently; instead, our results highlight the need for an integrated approach to comprehensively understand the key patterns which shape the labor market. While extensive literatures on automation and immigration present mixed results, we demonstrate that the interaction between different shocks can play an important role in understanding the divergence in findings. A unified approach is needed to inform any policy response to mitigate these negative labor market shocks, and an understanding of how different shocks interact with each other is imperative.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica 67*(2), 251–333.

Abraham, K. G. and M. S. Kearney (2020). Explaining the decline in the us employment-to-population ratio: A review of the evidence. *Journal of Economic Literature 58*(3), 585–643.

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

Acemoglu, D. and P. Restrepo (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica 90*(5), 1973–2016.

Autor, D., D. Dorn, and G. H. Hanson (2013). The china syndrome: Local labor market effects of import competition in the united states. *The American Economic Review 103*(6), 2121–2168.

Autor, D., A. Kostøl, M. Mogstad, and B. Setzler (2019). Disability benefits, consumption insurance, and household labor supply. *American Economic Review 109*(7), 2613–2654.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the u.s. labor market. *American Economic Review 103*(5), 1553–1597.

Autor, D. H., D. Dorn, and G. H. Hanson (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal 125*(584), 621–646.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006, May). The polarization of the u.s. labor market. *American Economic Review 96*(2), 189–194.

Autor, D. H., F. Levy, and R. J. Murnane (2003, 11). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Balsvik, R., S. Jensen, and K. G. Salvanes (2015). Made in china, sold in norway: Local labor market effects of an import shock. *Journal of Public Economics 127*, 137–144.

Borjas, G. J. (2003). The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market. *The Quarterly Journal of Economics 118*(4), 1335–1374.

Borusyak, K. and P. Hull (2020, September). Non-random exposure to exogenous shocks: Theory and applications. Working Paper 27845, National Bureau of Economic Research.

Bratsberg, B. and O. Raaum (2012). Immigration and wages: Evidence from construction. *The Economic Journal 122*(565), 1177–1205.

Card, D. (1990). The impact of the mariel boatlift on the miami labor market. *Industrial and Labor Relations Review 43*(2), 245–257.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics 36*(S1), S13–S70.

Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics 128*(3), 967–1015.

Charles, K. K., E. Hurst, and M. Schwartz (2018, March). The transformation of manufacturing and the decline in u.s. employment. Working Paper 24468, National Bureau of Economic Research.

Clemens, M. A., E. G. Lewis, and H. M. Postel (2018). Immigration restrictions as active labor market policy: Evidence from the mexican bracero program. *American Economic Review 108*(6), 1468–1487.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2021, 05). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association 19*(6), 3104–3153.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The quarterly journal of economics 132*(4), 1593–1640.

Dustmann, C., T. Frattini, and I. P. Preston (2013). The effect of immigration along the distribution of wages. *Review of Economic Studies 80*(1), 145–173.

Dustmann, C., U. Schönberg, and J. Stuhler (2016). The impact of immigration: Why do studies reach such different results? *Journal of Economic Perspectives 30*(4), 31–56.

Dustmann, C., U. Schönberg, and J. Stuhler (2017). Labor supply shocks, native wages, and the adjustment of local employment. *The Quarterly Journal of Economics 132*(1), 435–483.

Edin, P.-A., T. Evans, G. Graetz, S. Hernnäs, and G. Michaels (2023). Individual consequences of occupational decline. *The Economic Journal 133*(654), 2178–2209.

Faber, M., A. P. Sarto, and M. Tabellini (2022). Local shocks and internal migration: The disparate effects of robots and chinese imports in the us. NBER Working Paper 30048, National Bureau of Economic Research.

Foged, M. and G. Peri (2016). Immigrants and native workers: New analysis using longitudinal employer–employee data. *American Economic Journal: Applied Economics 8*(2), 1–34.

French, E. and J. Song (2014). The effect of disability insurance receipt on labor supply. *American economic Journal: economic policy 6*(2), 291–337.

Glitz, A. (2012). The labor market impact of immigration: A quasi-experiment exploiting immigrant location rules in germany. *Journal of Labor Economics 30*(1), 175–213.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020, August). Bartik instruments: What, when, why, and how. *American Economic Review 110*(8), 2586–2624.

Goos, M., A. Manning, and A. Salomons (2009). Job polarization in europe. *American economic review 99*(2), 58–63.

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

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

Gundersen, F. and D. Juvkam (2013). Inndelinger i senterstruktur, sentralitet og ba-regioner. NIBR-rapport 2013:1, Norsk institutt for by- of regionforskning.

Jaimovich, N., I. Saporta-Eksten, H. Siu, and Y. Yedid-Levi (2021). The macroeconomics of automation: Data, theory, and policy analysis. *Journal of Monetary Economics 122*, 1–16.

Jurkat, A., R. Klump, and F. Schneider (2023). Robots and wages: A meta-analysis.

Kostøl, A. R. and M. Mogstad (2014). How financial incentives induce disability insurance recipients to return to work. *American Economic Review 104*(2), 624–655.

Lachowska, M., A. Mas, and S. A. Woodbury (2020). Sources of displaced workers' long-term earnings losses. *American Economic Review 110*(10), 3231–3266.

Lewis, E. G. (2011). Immigration, skill mix, and capital–skill complementarity. *The Quarterly Journal of Economics 126*(2), 1029–1069.

Manacorda, M., A. Manning, and J. Wadsworth (2012). The impact of immigration on the structure of wages: Theory and evidence from britain. *Journal of the European Economic Association 10*(1), 120–151.

Mandelman, F. S. and A. Zlate (2022). Offshoring, automation, low-skilled immigration, and labor market polarization. *American Economic Journal: Macroeconomics 14*(1), 355–389.

OECD (2010). *Sickness, disability and work: breaking the barriers: a synthesis of findings across OECD countries*. OECD Publishing.

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

Peri, G. and C. Sparber (2009). Task specialization, immigration, and wages. *American Economic Journal: Applied Economics 1*(3), 135–169.

Schmieder, J. F., T. Von Wachter, and J. Heining (2023). The costs of job displacement over the business cycle and its sources: evidence from germany. *American Economic Review 113*(5), 1208–1254.

Sorkin, I. (2018). Ranking firms using revealed preference. *The quarterly journal of economics 133*(3), 1331–1393.

A Understanding the Sources of the Automation Shock

Table A.1: The Importance of Pre-Shock Factors for Future Automation Levels

	(1) Unemployment Rate	(2) Share Low Educ	(3) Mfg. Emp. Share	(4) Construction Unlicensed Emp. Share	(5) Share Native	(6) Share Imm. Expansion Country
Δ Auto	0.008	0.007	0.067***	0.003	0.024	-0.001
	(0.005)	(0.006)	(0.010)	(0.002)	(0.015)	(0.000)
LLMs	160	160	160	160	160	160

Notes: Table tests the importance of pre-shock factors (measured in 1996) for the future change in automation from 1996–2015. Factors are calculated for each local labor market using data on all individuals aged 18–64. Each column reports separate regression of the stated variable on the change in automation as measured in equation (2), weighted by the population size. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

Figure A.1: Describing the Local Labor Market Variation in Automation

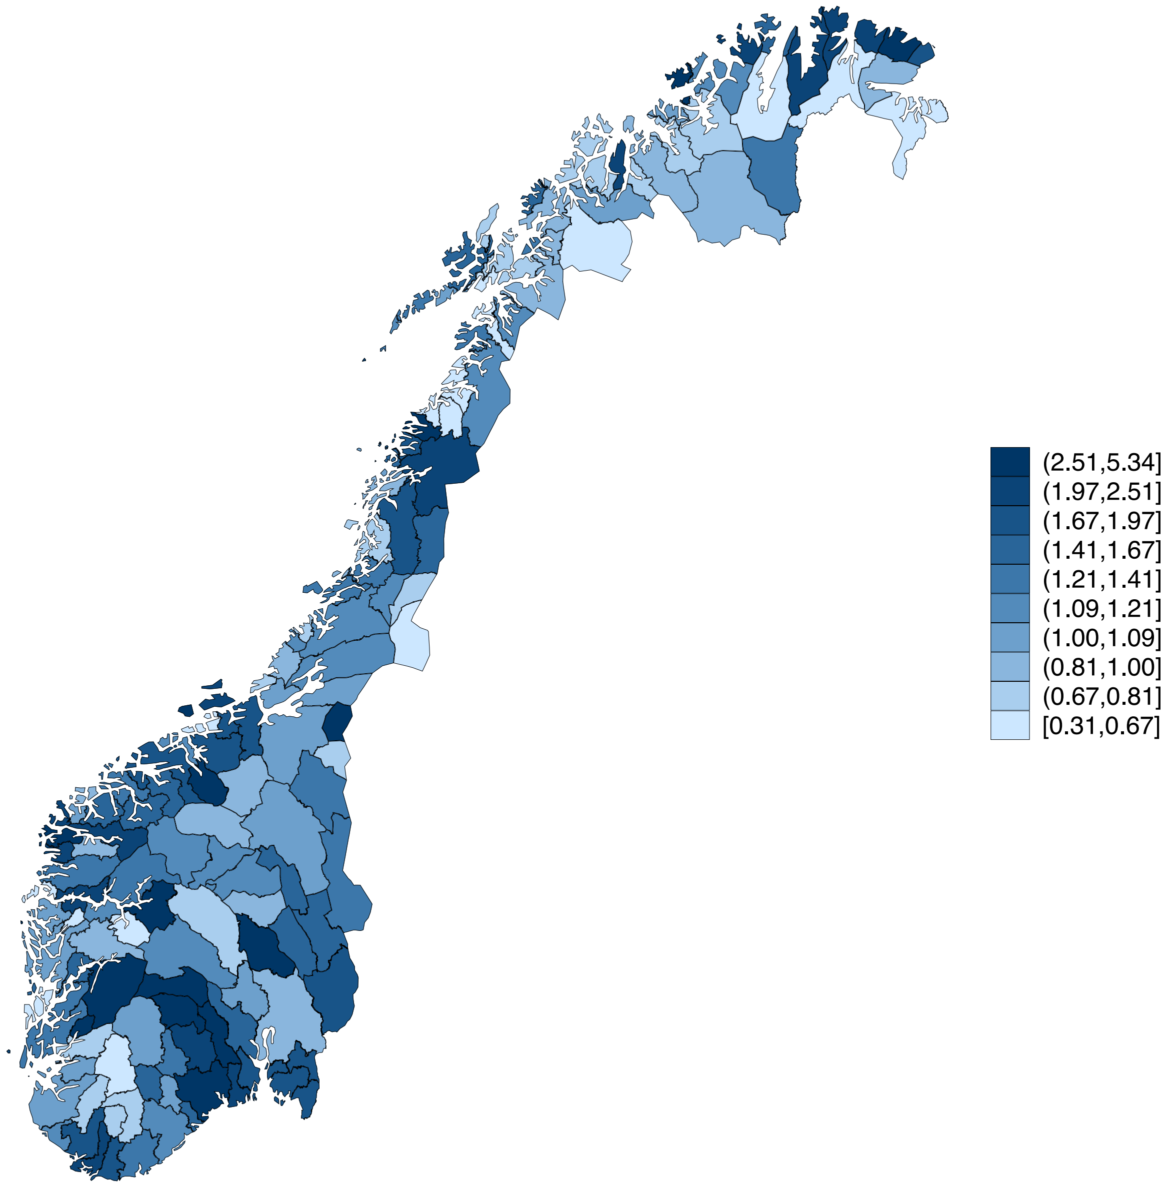

Notes: Figure reports the exposure to robots (robots per 1000 workers) as defined as in Acemoglu and Restrepo (2020) from 1996–2015.

Figure A.2: The Relationship Between Automation and Immigration Shocks at the Area-Level

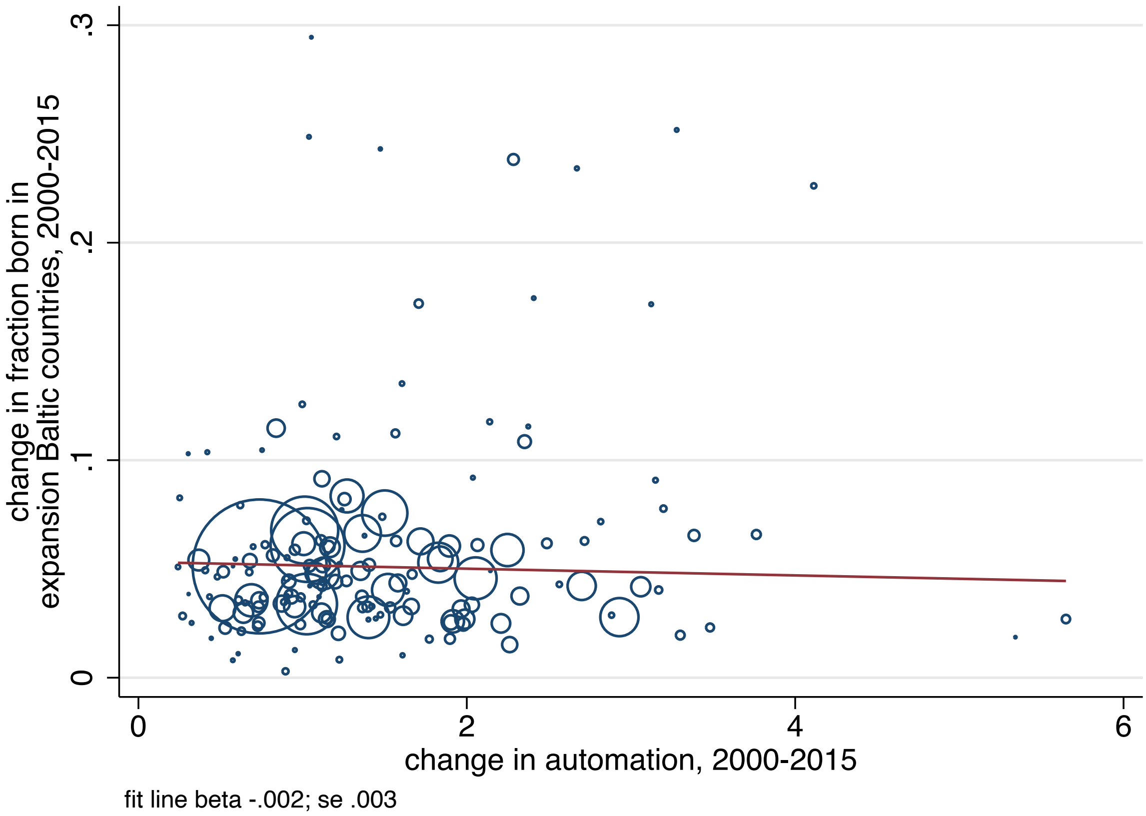

Notes: Figure plots the relationship between realized immigration (y-axis) and the change in automation (x-axis), both measured at the area-level. Each circle corresponds to a local labor market, where the size corresponds to the population. Both scatter plot and regression fit line (estimated at the area-level) are weighted by population.

B The Effects of Automation on Women by Education / Industry

Table B.1: The Impacts of Automation on Earnings by Industry/Education, from 2000–2015

	Education		Industry		
	(1) Non-College Educated	(2) College Educated	(3) Manufacturing	(4) Construction	(5) All Other Industries
Panel A: Δ Log Earnings					
Δ Auto	0.010	0.003	-0.020	0.011	0.007
	(0.018)	(0.010)	(0.016)	(0.032)	(0.013)
N	423286	250321	48336	8182	617090
Panel B: Δ Employed					
Δ Auto	-0.001	-0.000	-0.006	-0.011	-0.004
	(0.006)	(0.002)	(0.006)	(0.008)	(0.003)
N	480996	264295	54942	9116	681234

Notes: sample of men aged 18–49 initially employed in 2000. Figure presents estimates of equation (3), estimated separately for workers with different education levels (columns 1–2) and workers initially employed in different industries (columns 3–5). Outcome variables correspond to the change in log earnings (panel A) and the change in employment (panel B). Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

C Comparing Workers in Manufacturing, Construction, and Other Industries

Table C.1: Comparing Construction, Manufacturing, and Other Industries

	(1) construction	(2) manufacturing	(3) all other industries
Demographics:			
male	0.917	0.737	0.461
	(0.276)	(0.440)	(0.498)
norwegian	0.952	0.923	0.921
	(0.214)	(0.266)	(0.270)
age	33.99	34.71	34.18
	(8.768)	(8.401)	(8.831)
Labor Market:			
log earnings	12.37	12.39	12.16
	(0.606)	(0.637)	(0.797)
Education:			
HS grad	0.602	0.581	0.661
	(0.489)	(0.493)	(0.474)
College grad	0.139	0.206	0.364
	(0.346)	(0.405)	(0.481)
Individuals	112653	215516	1303627

Notes: sample of workers in one of three industries, age 18–49. All variable measured in year 2000.

D The Interacted Impacts of Automation and Immigration on Employment

Table D.1: The Interaction Between Automation & Immigration, Employment

	Δ Employed		
	(1) Full Sample	(2) Licensed	(3) Unlicensed
Δ Auto	0.001	0.001	-0.002
	(0.005)	(0.005)	(0.007)
ImmExp	-0.024***		
	(0.005)		
Δ Auto × ImmExp	-0.004		
	(0.004)		
Education FE (1 digit)	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes
Individuals	98922	48903	50019
Average x	1.326		

Notes: Columns (1) reports estimates from equation (4), for employment. Columns (2)–(3) report estimates from equation (3), separately for licensed and unlicensed workers, for employment. Standard errors are clustered at the initial local labor market level. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.

E Robustness

Figure E.1: The Interaction Between Automation & Immigration, Over Time

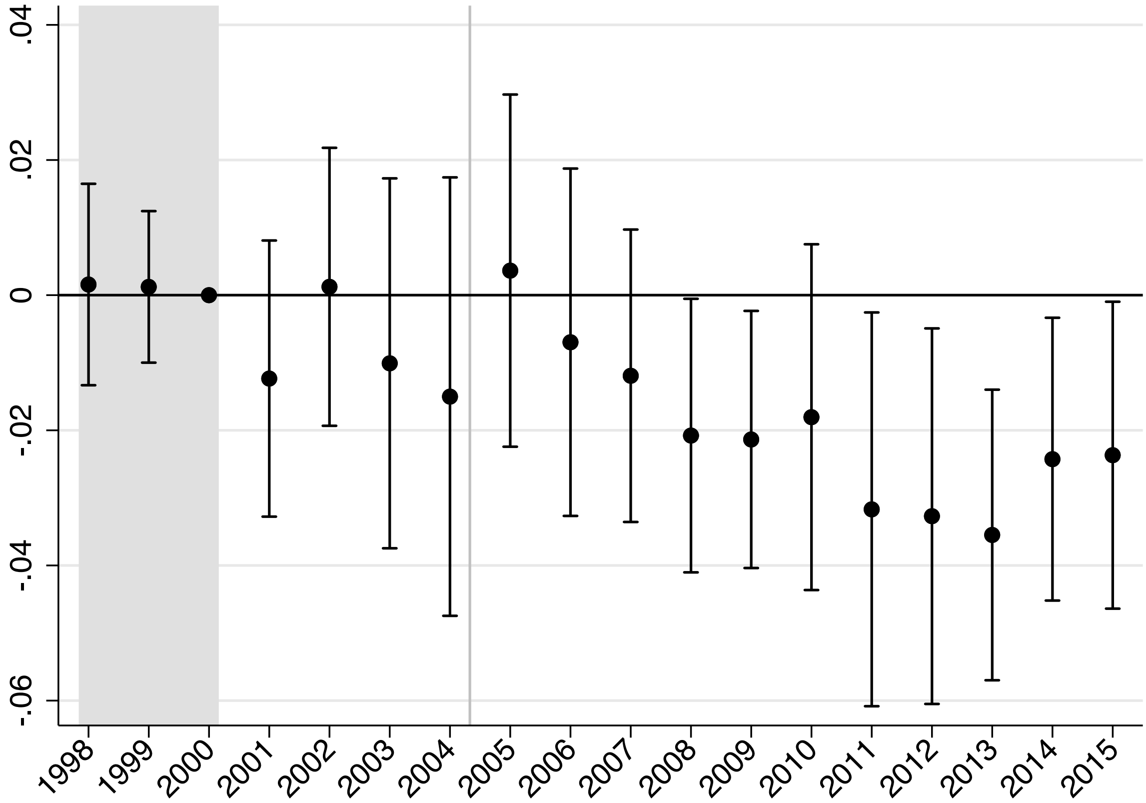

Notes: Figure reports estimates from equation (4), with outcome variables measured at the change in earnings for each year relative to 2000. Regressions are estimated separately year-by-year, such that the value for 2000 is not estimated. Vertical line at 2004 represents when EU expansion occurs. Change in automation is measured as the change in automation from 1996–2015 for all years, such that the value of the automation shock is fixed over time.

Figure E.2: The Interaction Between Automation and Immigration, Prior to Onset of Both Shocks

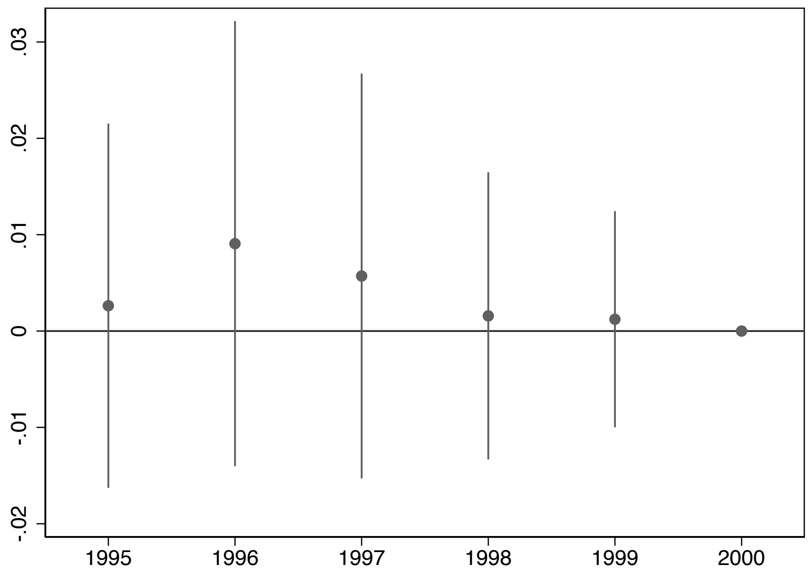

Notes: Figure plots the estimated interaction coefficients between automation and immigration shocks, prior to the onset of both shocks. Sample size becomes increasingly smaller going back further in time, as younger individuals employed in 2000 have not yet transitioned into the labor market 5 years earlier.

F The Interaction Between Automation and Immigration, detailed plant-level outcomes

Table F.1: The Interaction Between Automation and Immigration, detailed plant-level outcomes

	(1) Δ Plant Rank Natives	(2) Δ Frac. Native	(3) Δ Frac. College Educ.	(4) Δ Frac. Comp. Educ.
Δ Auto× ImmExp	-1.375**	0.001	-0.003	-0.006
	(0.581)	(0.003)	(0.004)	(0.004)
Education FE (1 digit)	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Individuals	84207	84889	84889	84889

Notes: This table reports coefficients from separate regressions at the plant level, relating changes in workforce composition and relative plant rank to exposure to automation and import competition shocks. Column (1) reports the estimated change in the plant's percentile rank (based on average earnings) of native workers relative to plants in the same industry–region cell. Column (2) reports the change in the fraction of native workers employed at the plant. Column (3) reports the change in the fraction of college-educated workers, and column (4) the change in the fraction of workers with only compulsory education. Standard errors are clustered at the initial local labor market level. ***, **, and * correspond to significance at the 1%, 5%, and 10% levels respectively.