Working Paper Series 821
(ISSN 2788-0443)

Engines of Change: Automation, Wages, and Labour Demand in Slovakia's Automotive-Centred Economy

Matěj Bělín

CERGE-EI
Prague, June 2026

Engines of Change: Automation, Wages, and Labour Demand in Slovakia's Automotive-Centred Economy

Matěj Bělín*

Abstract

This paper estimates the effects of robot adoption on wages and employment in Slovakia, one of the most automation-exposed economies in Europe. Using matched employer-employee administrative data combined with survey-based measures of occupational task exposure, we identify the impact of automation through an instrumental-variables strategy that exploits variation in robot adoption from major non-European adopters. Across specifications, we find no evidence of economically meaningful effects of robot adoption on labour-market outcomes. The estimated impacts on wages and job separations are small and precisely estimated around zero. The Slovak estimates point to increased productivity among robot-adopters, which, in turn, creates new employment opportunities rather than displacing human labour.

1 Introduction

Automation has emerged as a central force shaping contemporary labour markets, raising fundamental questions about the future of work, wage formation, and the allocation of labour across sectors. At the core of this debate lies a longstanding concern whether technological progress displaces human labour

*CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences, Politickych veznu 7, 111 21 Prague, Czech Republic; Inštitút sociálnej politiky, Ministerstvo práce, sociálnych vecí a rodiny Slovenskej republiky. The work was supported from ERDF/ESF project New technologies and changes in education, research and the labor market (No. CZ.02.01.01/00/23_025/0008693). Insightful feedback on an earlier version of this paper from participants of the EUBA Economic Research Seminar on 6 February 2025 and stimulating discussions with Štěpán Jurajda, Ján Komadel, and Vladimír Novák are gratefully acknowledged. Any remaining errors are my own.

or instead enhances productivity in ways that sustain, or even expand, employment. This question has gained renewed urgency as advances in robotics and artificial intelligence diffuse across manufacturing and services, particularly in highly specialised economies where sectoral concentration may amplify exposure to technological shocks.

A large empirical literature has examined the labour-market effects of robotics, yet the evidence remains mixed and, in several respects, incomplete. Early contributions documented negative associations between robot adoption and employment in the United States (Acemoglu and Restrepo, 2017), with related findings in France (Acemoglu et al., 2020) and China (Giuntella and Wang, 2019), particularly for low-skilled workers.[1] At the same time, other studies find little or no effect on employment, including firm-level evidence from the Czech Republic (Jurajda and Bělín, 2020) and the Netherlands (Acemoglu et al., 2023), the latter also documenting positive wage effects. Cross-country analyses point to modest positive wage effects on average (Graetz and Michaels, 2018), while evidence from Germany suggests that automation may primarily induce reallocation of workers across sectors rather than aggregate displacement (Dauth et al., 2021).

These divergent findings highlight several open questions. First, the extent to which automation affects labour demand appears to depend on institutional context, industrial structure, and the stage of technological diffusion. Second, existing studies point to heterogeneity in outcomes (Sun et al., 2023; Graetz and Michaels, 2018; Boddin and Kroeger, 2022; Ge and Zhou, 2020), including evidence that displacement effects may materialise only once technologies reach sufficient maturity (Jaccoud et al., 2024).

Slovakia provides a particularly informative setting to address these open questions, as it combines exceptionally high exposure to automation with the characteristics of a small, open economy. Its production structure is heavily concentrated in the automotive sector, which is one of the most robot-intensive industries worldwide. These traits make it a natural testing ground for theories linking technological adoption, sectoral specialisation, and labour-market outcomes. Slovakia is the world's largest car producer in per-capita terms (Laurenson, 2026), and consequently ranks among the most exposed economies to automation technologies (Nedelkoska and Quintini, 2018). Despite this, there

[1] At the same time, it should not be taken for granted that the displacement of human labour is inherently undesirable. Automation can substitute for workers in physically demanding or hazardous tasks, potentially improving job quality and worker safety. This dimension became particularly salient during the COVID-19 pandemic, when a range of technological solutions were deployed to reduce physical contact and mitigate contagion risks (Aymerich-Franch and Ferrer, 2022; Shen et al., 2021).

is limited evidence on how automation has affected labour-market outcomes in Slovakia, with existing work largely confined to descriptive accounts of global adoption trends (e.g. Müller, 2023, pp. 406–415). This paper addresses this gap by providing causal evidence on the effects of robot adoption on wages and employment using micro-level data.

This combination of sectoral concentration and rapid technological adoption makes Slovakia a particularly informative setting for studying the labour-market effects of automation. The dominance of a single, highly robot-intensive sector generates substantial variation in exposure across firms and workers within a relatively small and integrated economy. At the same time, the openness of the Slovak economy implies that firms face strong competitive pressures to adopt cost-saving technologies. These features jointly provide a stringent test of the labour-displacement hypothesis in an environment where automation incentives are strong.

We focus on *robotic* technologies, a well-defined and empirically tractable subset of automation. Following the standard definition (e.g. Madesäter, 2005), consistent with ISO 8373, robots are programmable, multi-axis machines with a degree of autonomy capable of performing a range of tasks. This definition excludes purpose-built machinery and sector-specific equipment that cannot be reprogrammed, such as conventional assembly lines or agricultural machinery. Typical examples include industrial welding robots, widely used in automotive production (Müller and Kutzbach, 2020, p. 69). Their defining feature is the ability to substitute for human labour across a range of tasks rather than perform a single dedicated function.

This paper contributes to this literature along three dimensions. First, we provide evidence from a highly exposed economy, offering a stringent test of the labour-displacement hypothesis. Second, we combine administrative matched employer-employee data with survey data containing detailed occupational information. This combination allows us to exploit the strengths of both sources: the administrative data enable precise identification using worker-firm matches, while the survey data allow us to construct task-based measures of automation exposure. Third, we document that the effects of robot adoption on wages and employment are not only statistically insignificant but also economically negligible. Across specifications, we obtain estimates that are tightly centered around zero, with confidence intervals sufficiently narrow to exclude effects of meaningful magnitude. We refer to these as *precise zero* effects.

These findings are consistent with the predictions of task-based models of automation. In such frameworks, automation reallocates tasks between labour and capital, generating opposing forces: displacement from automated tasks and productivity gains in remaining and/or new tasks. In environments where the productivity gains dominate, the net effect on wages and employment is theoretically neutral or positive. Our results provide empirical support for this case.

Empirically, we estimate the effects of robot exposure on worker-level wages and employment outcomes using a panel of matched employer-employee data. Our primary measure of automation is the stock of industrial robots from the International Federation of Robotics (IFR), complemented by firm-level R&D-based proxies to capture within-industry variation. The use of matched data allows us to control for firm-worker heterogeneity and identify the effects of automation more precisely than aggregate approaches (Graetz and Michaels, 2018). Across all specifications, we find no evidence of substantial labour displacement or wage effects.

The remainder of this paper is organised as follows: a discussion of the theory is followed by a brief survey of the data with notes on historical trends. We then proceed with an exposition of the instrumental-variable identification strategy and its alternatives and variants. Results are discussed with robustness checks and the paper closes with a summary of conclusions and directions for future work.

2 Theory and empirical predictions

We estimate the effects of firms' possibly endogenous automation choices on workers' labour-market outcomes within a task-based framework of production. In the task approach associated with Acemoglu and Restrepo (2018), production is understood as a collection of tasks, some of which can be reassigned from labour to capital when automation technologies become sufficiently cost-effective. Robotisation therefore affects labour-market outcomes not simply through capital deepening, but through a reorganization of the task content of production itself.

Nakamura and Nakamura (2019) provide a useful bridge between this task-based perspective and the classic aggregate production function framework. They show that a task-based model with endogenous automation can be represented as a neoclassical production function in reduced form, thereby allowing empirical specifications that relate wages and labour demand to automation exposure in a theoretically coherent manner. Their argument (which was subsequently developed in Nakamura et al., 2026) is summarised below, highlighting its empirical implications.

The starting point is a continuum of tasks indexed by i such that

$$i \in [0, 1]. \tag{1}$$

Each task can be performed either by capital or labour, with task-specific efficiencies $\theta(i)$ and $\lambda(i)$ for capital or labour respectively. Tasks are ordered such that capital has a comparative advantage in lower-indexed tasks, while labour has a comparative advantage in higher-indexed tasks. There exists a threshold $a_t \in [0,1]$ such that tasks $i \leq a_t$ are automated and tasks $i > a_t$ are performed by labour.

Production is assumed to exhibit a bottleneck structure across tasks, implying a Leontief-type technology. Aggregating across tasks yields

$$Y_t = \min\{\Theta(a_t)K_t, \Lambda(a_t)L_t\}, \tag{2}$$

where

$$\Theta(a_t) = \left[\int_0^{a_t} \theta(i)^{-1} di\right]^{-1}, \tag{3}$$

$$\Lambda(a_t) = \left[\int_{a_t}^{1} \lambda(i)^{-1} di\right]^{-1}. \tag{4}$$

The term $\Theta(a_t)$ represents the effective productivity of capital in automated tasks, while $\Lambda(a_t)$ captures the effective productivity of labour in non-automated tasks. As automation expands, capital is applied to tasks that are progressively more difficult to automate, implying $\Theta'(a_t) < 0$, whereas labour becomes concentrated in tasks where it has comparative advantage, implying $\Lambda'(a_t) > 0$.

Cost minimization implies

$$\Theta(a_t)K_t = \Lambda(a_t)L_t, \tag{5}$$

which yields

$$\frac{K_t}{L_t} = \frac{\Lambda(a_t)}{\Theta(a_t)} \equiv \Omega(a_t). \tag{6}$$

Since $\Omega'(a_t) > 0$, the automation threshold can be written as a function of the capital-labour ratio:

$$a_t = \Omega^{-1}\left(\frac{K_t}{L_t}\right). \tag{7}$$

Substituting back into into the Leontief technology gives an aggregate production function

$$Y_t = \Lambda\left(a\left(\frac{K_t}{L_t}\right)\right)L_t \equiv F(K_t, L_t), \tag{8}$$

which defines a neoclassical production function in reduced form. Automation is thus embedded in the production technology through the endogenous task threshold.

While not strictly necessary for the empirical task ahead, it is worth noting that a specific functional form of the production function can be derived under an additional restriction on the trade-off between capital and labour efficiency across tasks,

$$1 = b\Theta(a_t)^{\rho} + (1-b)\Lambda(a_t)^{\rho}, \tag{9}$$

such that $0 < b < 1,\ -1 \leq \rho$. This restriction can be interpreted as assuming that automation reallocates tasks in a balanced way, namely, as capital takes over more complex, high-skill tasks and becomes less efficient on average, labour simultaneously concentrates in tasks where it becomes more efficient (cf. Caselli and Coleman, 2006, for similar argumentation). Condition (9) ensures that this trade-off is sufficiently regular to generate a constant elasticity of substitution between the two factors, thus reducing (8) to a constant-elasticity-of-substitution (CES) form:

$$Y_t = \left[bK_t^{-\rho} + (1-b)L_t^{-\rho}\right]^{-1/\rho}, \tag{10}$$

with elasticity of substitution

$$\sigma = \frac{1}{1+\rho}. \tag{11}$$

This framework yields direct implications for labour demand and wages. Under perfect competition, wages equal the marginal product of labour:

$$w_t = \frac{\partial Y_t}{\partial L_t}. \tag{12}$$

Using the CES representation, this can be written as

$$w_t = (1-b)\Lambda(a_t)^{1+\rho}. \tag{13}$$

Wages therefore depend on automation through $\Lambda(a_t)$, the effective productivity of labour in non-automated tasks. As automation increases, labour is reallocated toward tasks in which it has higher comparative advantage, raising $\Lambda(a_t)$ and exerting upward pressure on wages. However, automation simultaneously reduces the set of tasks performed by labour, generating a displacement effect.

The wage implications of the model hinge on the assumption that the measure of tasks is *fixed* as required by eq. (1). Under this normalization, automation operates exclusively through the reallocation of existing tasks between labour and capital, so that improvements in labour productivity reflect only compositional effects. If, instead, the set of tasks expands over time (Acemoglu and Restrepo, 2019), automation is accompanied by an increase in total factor productivity. The creation of new tasks relaxes the displacement effect by generating additional demand for labour, thereby generating ambiguous associations with wages and employment.

Labour demand follows from

$$\Theta(a_t)K_t = \Lambda(a_t)L_t, \tag{14}$$

which implies

$$L_t = \frac{\Theta(a_t)}{\Lambda(a_t)}K_t. \tag{15}$$

Since $\Theta'(a_t) < 0$ and $\Lambda'(a_t) > 0$, the ratio $\Theta(a_t)/\Lambda(a_t)$ declines with automation. Holding capital fixed, an increase in automation reduces labour demand by reallocating tasks from labour to capital.

In sum, the task-based model combined with intuitive assumptions of technological bottlenecks, cost minimisation, and gradual automatisation of tasks that are increasingly difficult to automate, yields a well-behaved neoclassical production function with an additional automation parameter. This, in turn, implies that robotisation alters labour-market outcomes through an endogenous reallocation of tasks. This mechanism simultaneously affects labour demand (through task displacement) and wages (through changes in effective labour productivity), providing the theoretical foundation for our empirical analysis.

3 Data

The main source of information about the Slovak labour market is the panel of employer-worker data from the Slovak Social Security Agency. This is a monthly panel dataset of all employment contracts[2] in Slovakia identifying the worker,[3] employer and the specific employment contract.

[2] This dataset does not cover police and military forces. In addition, self-employed persons who are exempt from obligations to pay social contributions are not recorded. These exclusions are of a minor consequence to this study as robot adoption is concentrated mainly among large manufacturing firms. See Bělín and Veselková (2023) for further discussion of this dataset.

[3] For the purposes of this study, the term "employment" will be used to cover also persons engaged in work outside of an employment contract. These are independent contractors engaged

Table 1 reports the numbers of contracts recorded in the Social Security data. Following the financial crisis of 2008, the steady increase in the overall number of contracts until that point was broken and the Slovak economy hovered around 2.2 million contracts until 2012. From 2013, there was a fairly rapid growth in the number of contracts peaking around 3.2 million until the COVID years 2020-2021, which marked another major decline.

Tab. 1: Number of contracts recorded in the Social Security data (thousands)

Year	All	Empl.	Other	Auto.	Hires	Switch	Dism.
2004	2109	2012	97	47	524	122	394
2005	2211	2098	113	50	669	154	371
2006	2321	2195	126	62	688	183	367
2007	2419	2282	137	73	705	207	373
2008	2427	2280	148	72	650	198	385
2009	2243	2085	158	62	460	118	358
2010	2177	2016	162	62	476	121	292
2011	2246	2078	168	68	543	129	313
2012	2209	2037	173	70	468	125	310
2013	2839	2043	797	68	1124	261	503
2014	2971	2104	868	73	1011	298	525
2015	3109	2219	890	83	1084	334	556
2016	3155	2291	863	90	1065	339	541
2017	3197	2381	816	94	1051	331	544
2018	3231	2435	796	98	1035	313	548
2019	3234	2421	813	94	993	297	539
2020	3101	2329	772	86	833	246	535
2021	3066	2299	767	86	882	263	485
2022	3097	2340	756	84	905	215	409

Notes: Empl. = employee contracts; Auto. = Automotive industry (NACE 29); Switch = separations followed by another employment the next month; Dism. = dismissals, separations followed by longer periods without recorded employment.

Employee contracts constitute a large majority of all recorded work arrangements even though their share declined steadily from about 95% in 2004 to about 75% in 2022. By contrast, automotive employment held a steady share

to carry out a specific set of tasks for their employer. Further, we treat simultaneous contracts in cases when a single person is employed by multiple firms or a single person holds multiple contracts with the same firm as separate instances of employment.

of the overall employment, between 2% and 3%, throughout the observed period.

The fluctuations in the labour market (hires, changes, dismissals) follow the pattern of the business cycle exhibited by the overall number of contracts. Periods of growth (2004-2008 and 2013-2019) are associated with greater flux in the labour market in terms of more hires and also job separations.[4] Recession periods show higher numbers of dismissals at the beginning of the crisis followed by periods of few hires but also few separations. This is consistent with the standard economic intuition that firms shed non-essential personnel at the outset of a difficult period but retain their core employees to preserve accumulated expertise.

Despite their breadth, the data have several limitations. First, information on hours worked is not available, so earnings reflect total yearly compensation rather than wage rates. Second, job separations must be inferred from interruptions in observed employment relationships rather than directly observed separation events. While these limitations are standard in the literature, and to an extent addressed in an alternative version of the model using survey data, they should be kept in mind when interpreting the results.

Data on the labour market situation are combined with data on robot installations reported by the International Federation of Robotics (IFR, cf. Müller, 2023). This is a yearly dataset reporting the estimated levels of operational robot stocks in different industries. The classification of industries within IFR data is somewhat idiosyncratic, so concordance tables from Jurkat et al. (2022) were used to match the IFR data classification with 2-digit NACE classifications used in the Social Security data. The classification used is listed in Table 2. This step involves some loss of information as not all IFR codes can be directly associated with NACE groups. For instance, category 299 “Automotive unspecified" does not map to any NACE group and for this reason, we omit robot stocks in this category to avoid amplifying the measurement error caused by the coarse resolution of the IFR data with imprecise matching by industries. The loss of information on robot stocks due to the requirement of precise identification of the industry in which a given installation operates is relatively minor: in 2022, robots without precise identification of their industry amounted to about 6% of the total reported operational stock in Slovakia.

Following Dauth et al. (2021) we express robot usage intensity as number of robot installations per thousand workers because this relative measure lends

[4] The Social Security data do not record dates or reasons for contract terminations. Thus, job switches and dismissals have to be inferred from the interruptions in observed relationships between employer and worker within a specific type of a work contract.

Tab. 2: ISIC classification of industries used

Code	Description
10-12	Food and beverages
13-15	Textiles, wearing apparel
16+31	Wood and furniture
17-18	Manufacture of paper and paper products + Printing and reproduction of recorded media
19-22	Manufacture of coke and refined petroleum products + Manufacture of chemicals and chemical products Manufacture of pharmaceuticals, medicinal chemical and botanical products + Manufacture of rubber and plastics products
23	Manufacture of other nonmetallic mineral products
24	Manufacture of basic metals
25	Manufacture of fabricated metal products, except machinery and equipment
26	Manufacture of computer, electronic and optical products
27	Manufacture of electrical equipment
28	Manufacture of other machinery and equipment
29	Manufacture of motor vehicles, trailers, and semi-trailers
30	Manufacture of other transport equipment
32-33	Other manufacturing + Repair and installation of machinery and equipment
A	Agriculture, forestry, and fishing
B	Mining and quarrying
D-E	Utilities (electricity, gas, steam, air conditioning, water, waste collection and remediation)
F	Construction
P+M	Education + Scientific research and development
G-U\PM	All other non-manufacturing branches (wholesale and retail trade, transportation, accommodation, financial and insurance etc.)

Notes: Categories 26 and 27 were constructed from codes 260+261+262+263+265 and 271+275+279 respectively.

Fig. 1: Robot adoption

(a) All sectors

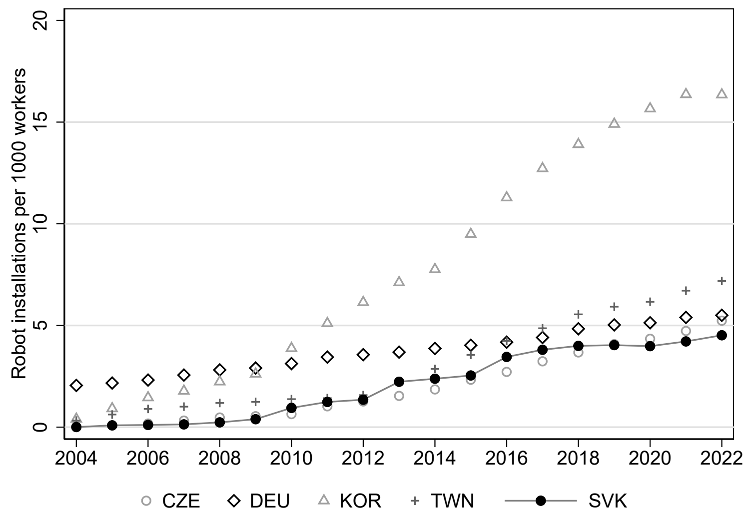

(b) Manufacturing only

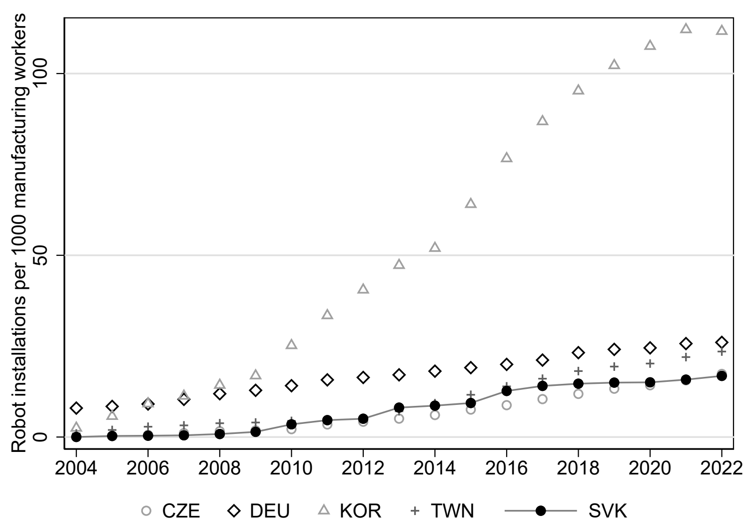

itself to easier international comparisons. To this end, we utilise data from the International Labour Organisation (ILO) for the number of workers in different industries.[5] To supplement missing entries in the ILO data, most notably for China and Taiwan, which are both major robot adopters, we add estimates of the distribution of Chinese workers by sector extracted from the China Industrial Productivity (CIP) Database 3.0 (Wu and Ito, 2015) and Taiwanese workers from the Asia KLEMS database (Pyo et al., 2011). Due to the older vintage of these estimates, information about the workers' distribution was extrapolated[6] forward in time using using population data reported in the Penn World Tables 10.01 (Feenstra et al., 2015).

Figure 1 shows the robot stocks in Slovakia compared with a selection of other robot-adopting economies. From virtually zero robot installations in 2004, very little growth in Slovak robotisation was recorded until the recession of 2008, at which point the concentration of robots started to rise. Toward the end of the observed period, Slovakia had roughly 4.5 robot installations per 1000 workers, which is broadly on a par with the Czech Republic, and lags slightly behind those in Germany (under 6 robots per 1000 workers). Virtually identical image emerges when looking at the manufacturing sector in isolation (Figure 1b) as opposed to the whole economy (Figure 1a) bar the scale.

However, we caution against over-interpretation of these robot densities as they are subject to errors in the estimated size of the labour force and inconsistencies in reporting practices in different countries (Müller, 2023, p. 73). For instance, Dauth et al. (2021) report just under 8 robots per 1000 workers in Germany in 2014 using Institute for Employment Research (IAB) data. However, when using Labour Force Survey (LFS) estimates of German employment at that time, we obtain only about 4.7 robots per 1000 workers.[7] This number is further reduced to 3.9 in Figure 1a when robots in problematic industries are removed. However, neither the precise robot density, nor the country rankings are important to our substantive conclusions because using absolute number of robot installations leads to very similar estimates (see section 5.3.3).

Finally, we turn to our last major data source, namely, the R&D tax credit data extracted from Financial Administration database. These data are available on the firm-year level for the period 2015-2021 and will be used as an

[5] ILOSTAT Labour Force Statistics: Employees by sex and economic activity [EES_TEES_SEX_EC2_NB_A]

[6] See Müller (2023, pp. 73-77) for discussion of the IFR's treatment of missing data. Their analysis employs statistical yearbooks for labour-force-size estimates, which lack the level of detail needed here.

[7] This appears to be due to the full-time equivalent adjustment, which is absent in the LFS data. As the IFR data are proprietary, only approximate values can be reported here.

Fig. 2: Correspondence between robot installations and R&D tax credit data 2015-2021 (log-scale, centred by industry; cf. Tab. 2 for industry classification)

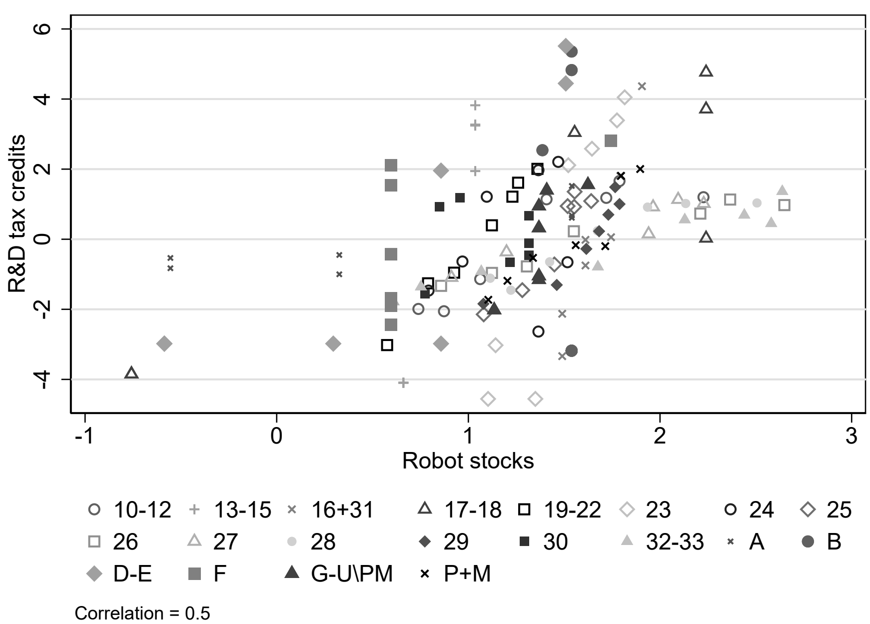

additional proxy for the exposure to automation technologies. Figure 2 shows that there is a fairly strong correlation between the R&D tax credits and the IFR measure of the operational robot stocks. The advantage of using the tax data is two-fold: (a) it resolves potential robot exposure at the *firm* level, which allows us to focus on the direct exposure as opposed to exposure from competitors within the same industry, and (b) it may capture automation technologies that do not fit into the definition of a robot discussed above, which may, nevertheless, have the potential to displace human labour.

4 Identification strategy

The main aim is to estimate the impact of exposure to robotic technologies on a person-specific labour market outcome Y (wages, job separations). Figure 3 is a schematic depiction of the general structure of the estimation problem.

Fig. 3: Structural model

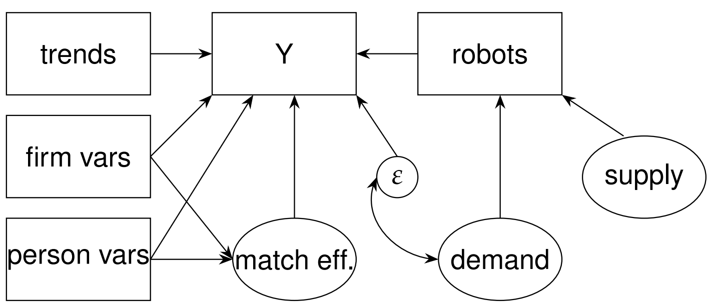

The outcome variable Y is expected to be influenced by national and global economic trends. As discussed in the previous section, business cycles are clearly reflected in the labour market aggregates and nominal wages tend to follow inflation trends. Furthermore, the overall trend in the rise of automation has to be accounted for to avoid spurious correlations.

Firm variables and personal variables determine the outcome either via a direct link, e.g. a worker's experience may directly increase her wage through a "tenure effect" or indirectly via a "match effect," which describes how well suited a particular worker might be for a particular job. In addition to the match effect, a provision must be made for unobserved effects ε, which may be connected with either firm-specific or worker-specific factors that are not recorded in the data, e.g. a firm's plans for re-organisation or worker's satisfaction on the job.

If the unobserved effects ε could be assumed to be uncorrelated with the robot installations, estimation could proceed immediately by constructing an

OLS model with fixed effects for employer-worker dyads and controls for firm and worker characteristics as appropriate. However, the hypothesis of labour displacement by technology as a cost-minimising measure implies that this is not the case. On this hypothesis, firms' wage bills are related to their (unobserved) demand for robotic solutions and hence their robot stocks will be endogenous to the outcome variable Y.

In order to sever the potential correlations between robot installations and the unobserved effects ε, we use only that part of the variation in robot installations that is related to the *supply* of robots via an instrumental variables (IV) approach. The supply of robots that are suitable for a particular industry in a given year is approximated by the robot installations by major robot-adopting economies outside Europe. The implicit identifying assumption here is that whatever demand shocks may be driving robot stocks outside Europe, these are uncorrelated with Slovak demand shocks after conditioning on global trends. If this identifying assumption is satisfied, the endogenous variation in robot installations (i.e. that related to the demand for robots in Slovakia) is purged in the first stage of the two-stage least squares (TSLS) estimation and the second stage thus provides consistent estimates of the effect of robots on labour market outcomes Y.

The argument for exogeneity of non-European robot installations rests on the interpretation of these variables as indicators of the availability of robotic technologies for particular uses and thus not correlated with the (endogenous) demand for these technologies among Slovak firms. A similar argument is advanced by Dauth et al. (2021) who instrument German robot stocks with those outside Germany. Here we take an additional precaution by using only non-European robot stocks, thus avoiding problems of potential Europe-wide labour-market shocks. Furthermore, to the extent that global robot stocks may be interpreted as responses to shifts in the world-wide labour market (e.g. adoption of robots may be a response to population ageing which makes it more difficult to find suitable workforce), these global trends that are already captured by time fixed effects and thus do not threaten identification.

4.1 IFR specification

The main specification is estimated by the following employer-worker panel regression:

$$\begin{aligned} y_{ijkt} = &\beta \text{Robot stock}_{kt} + \\ &+ \alpha_{ij} + \alpha_t + \mathbf{X}_{ijkt}\gamma + \varepsilon_{ijkt}, \end{aligned} \tag{16}$$

where y is the dependent variable (wages, job dismissals, job changes), i indexes individual workers, j indexes firms in industries k, t indexes time in years from 2004 to 2022. Parameter β measures the sensitivity of y with respect to exposure to robot installations. Equation (16) is estimated using data on robot installations as reported by the International Federation of Robotics (IFR), so that Robot stock$_{kt}$ refers to the installations present in firm j's industry code (see Table 2).

Robots stock$_{kt}$ indicates robots in Slovakia that are installed in firm j's industry (k) and hence firm index j is omitted.[8] Model (16) also includes fixed effects for each employer-worker dyad (α_{ij}) and for each year α_t.[9] Fixed effects for employer-worker dyads aim to eliminate concerns about selection of different types of workers to different firms. Specifically, firms that have the resources to adopt robotic technologies are likely to invest in their HR departments and thus might be able to find workers that are better suited to their positions. As a result, had we not controlled for α_{ij}, concerns about positive selection bias would be justifiably raised.

Eliminating the variation between individual employer-worker matches has the added benefit of addressing differences between individual employment contracts, which is important given our data limitations. Slovak Social security data do not contain information about the hours worked under a given contract, which introduces additional variation to the earnings variable. Part-time contracts with lower monthly earnings than the corresponding full-time earnings are not distinguishable in our data. However, this problem is largely remedied by using a fixed effect for each employment contract, thereby focusing only on changes in earnings *within* that contract.[10] In cases when y_{ijt} is a dismissal, we substitute α_{ij} with simple firm fixed effects α_j as single-failure survival models are unidentified when unit and time fixed effects are included simultaneously (e.g. Allison and Christakis, 2006).

[8] An earlier version of this paper included coefficients for robot installations upstream and downstream the supply chain similar to models used by Smarzynska Javorcik (2004). However, due to the availability of robot data only at the industry level, the interpretation of these coefficients was difficult. For this reason, only direct exposure is included here.

[9] In cases when a person works for a firm j on multiple occasions, these are subsumed into a single fixed effect if the hiatus was under three months, on the assumption that these are prolonged personal or administrative leaves. Hence, workers returning for seasonal work after longer periods are treated as a new match between a firm and a worker. Models with a single fixed effect for all work spells within the same firm-worker dyad on the assumption of perfect anticipation give similar results. Earnings from simultaneous contracts within a single employer-worker dyad are summed together.

[10] As an additional remedy, we study survey data containing information on hours worked, which allows us to model the labour demand directly. The results, presented in sec. 5.2, confirm the findings from administrative data.

Fixed effects for time periods (α_t) are included to eliminate secular trends in both robot stocks and in the general business cycle. In a time of crisis, standard economic arguments indicate that firms will be less likely to expend their reserves on major hiring efforts or pay increases. Time-fixed effects eliminate these overall trends, leaving identification of β only on differences between employer-worker matches in which a person faced a greater/lesser change in robot exposure. An additional advantage of using both employer-worker fixed effects alongside time-fixed effects is that time-fixed effects now effectively serve as controls for tenure effects. This is because in year t, a person i has worked $t - t_0$ years for firm j and thus, conditional on the start date t_0, which is captured by the employer-worker fixed effect, there is no additional tenure variation in yearly data.

Finally, there is a question of addressing time-varying heterogeneity across the employer-worker dyads. For this reason, we construct control variables $\mathbf{X}_{ijt}$ using the means of the dependent variables *excluding* individual j (suppressing industry index k for simplicity):

$$x_{ijt} = \frac{1}{N_{jt} - 1} \sum_{\iota \neq i} y_{\iota jt}, \tag{17}$$

where N_{jt} is the number of observations of variable y for firm j in year t. By computing the average earnings and separations we create a set of controls $\mathbf{X}_{ijt}$ that describes the personnel situation in firm j in year t. We also add the number of workers N_{jt} and proportion of male workers as additional controls. Thus, a rich, time-varying set of controls is available and its inclusion in (16) is unobjectionable unless the residual terms ε_{ijt} are correlated within firms. Intra-cluster correlations, combined with inclusion of the mean of the left-hand side variable as a regressor, would lead to endogeneity of $\mathbf{X}_{ijt}$ and to bias in β (or, $\mathbf{X}_{ijt}$ would be 'bad controls' in the terminology of Pischke and Angrist, 2008). However, results from estimating (16) under the exclusion of the time-varying controls yields similar results and hence the potential problem of endogeneity is not material to the discussion of robot exposure (see sec. 5.3.3).

As discussed above, we instrument Slovak robot installations using major robot installations outside Europe. Specifically, we use installations in China, India, Japan, Korea, Mexico, Thailand, Taiwan, and the USA, which collectively amounted to about 70% of robot installations worldwide in 2019. For convenience, let $\mathcal{C}$ denote this set of these major loci of robot installations outside

Europe. Thus, the first-stage equations take the form:

$$R_{ijkt} = \sum_{c \in \mathcal{C}} \pi_{1,c} \text{Robot stock}_{c,kt} + \\ + \omega_{ij} + \omega_t + \mathbf{X}_{ijt} \boldsymbol{\pi} + \xi_{ijkt}. \tag{18}$$

Hence, the robot exposure in Slovakia is projected on eight robot stocks, which are excluded from the second-stage regression. To avoid confusion with (16), fixed effects are denoted as ω_{ij} and ω_t.

4.2 R&D tax credit specification

Model (16) proxies the robot exposure within firm j by the industry-wide number of robot installations in j's industry. This plan relies on the assumption that firms within a given industry adopt robotic technologies at a comparable rate, or, at the very least, there are no dramatic outliers in adoption. Generally, one might assume that larger firms would be more active in procuring automation technologies than smaller ones due to economies of scale and potentially easier access to finance. While differences in the *levels* of robot stocks across firms are eliminated in model (16) by fixed effects for employer-worker matches, it is an open question how much of the differences in the *rate* of adoption is captured by the control variables (number of workers).

This consideration motivates our use of a different proxy for the robot installations, namely the tax credit for research and development (R&D). As R&D may well encompass other activities besides implementation of new technologies, we instrument this variable once again by robot installations in the same sector and year outside Europe. That is, we seek to separate the variation in R&D that is related to robot installations from other R&D activities. Projecting R&D on robot stocks outside Europe thus captures the R&D expenses related to availability of robotic technologies and allows us to discard other sources of variation in R&D (most notably other research and robot purchases driven by local labour market conditions).

Instrumenting a firm-level measure of R&D tax credits by industry-level IFR data, however, has an undesirable side effect of eliminating firm-level variation in the predicted values of robot exposure variables. To prevent this loss of information, we construct additional instruments following Lewbel (2012). Lewbel's Theorem 1 shows that in a model with an endogenous regressor and exogenous controls as contemplated here, one can construct properly excluded instruments as follows:

$$z_i = \tilde{x}_i \times \widehat{\xi}_i, \tag{19}$$

where $\tilde{x}_i$ is a centred exogenous variable (potentially included in the second stage) and $\widehat{\xi}_i$ is the estimated residual from projecting an endogenous regressor on all available exogenous regressors. Under heteroskedasticity of ξ, the newly-created instrument z_i will satisfy the condition of relevance. The condition of proper exclusion is satisfied if $\tilde{x}_i$ and the other controls used to obtain $\widehat{\xi}_i$ are exogenous (but *not* necessarily properly excluded from the second stage). To implement this plan, we isolate residuals ξ_{ijt} from (18) and interact them with centred robot stocks in countries $\mathcal{C}$ outside Europe creating eight Lewbel instruments

$$z_{jkt,c} = \left(\text{Robot stock in country } c_{kt} - \text{Avg. robot stock in country } c_k\right) \times \hat{\xi}_{jkt} \tag{20}$$

that vary not only across industries k but also across firms j thus enabling us to preserve inter-firm variation in data. Thus, the specification with R&D data follows the model above except the robot variables in (16) are replaced with R&D tax credits and the first-stage (18) is augmented by Lewbel-type instruments (20).

An additional bonus of using a firm-specific measure of robot exposure is the possibility for more flexible controlling for secular trends. Using industry-level variation to identify β in (16) effectively means that only a single common time trend can be absorbed by fixed effects α_t. By contrast, firm-level variation allows absorbtion of industry-specific time trends (α_{kt}) thereby allowing for the possibility that business cycles have different impacts on different sectors of the economy. Thus, in the firm-level version of the model, we use industry-specific time trends as opposed to a common trend.

The price to be paid for using tax data is a restriction on the window of time available. Our financial data extend to 2015 only and thus exclude a substantial portion of the time series available when using IFR data. However, a sizeable dataset is still available. Another limitation of the R&D data is that they record a yearly *flow* variable while IFR data estimate the *stock* of robot installations. Unfortunately, we cannot compute cumulative stocks of R&D spending due to the shorter time window for these data. Hence, the instrumented tax credits are best viewed as approximations of the stock values. Despite these limitations, the R&D data provide a very natural robustness check to the estimates obtained using the full sample with IFR data.

4.3 Task-based specification

We complement our baseline industry-level specification with a task-based measure of automation exposure that more closely aligns with the underlying theoretical framework. The task-based model emphasizes that automation operates at the level of tasks rather than industries, re-allocating specific activities from labour to capital depending on their technological substitutability. Accordingly, variation in exposure to robotisation could arise not only across industries, but also across occupations with different task content.

To capture this dimension, we combine the IFR data on the task content of industrial robot installations with concordance that maps these tasks into ISCO occupations (OECD, 2019). This allows us to construct a measure of job-specific exposure to robot-performed tasks. We then interact this occupational exposure with industry-level robot adoption to obtain a measure of effective exposure that reflects both the technological feasibility of automation and its realized adoption.

Formally, we define worker-level exposure as

$$\text{Exposure}_{ijkt} = \text{Task Exposure}_{it} \times \widetilde{\text{Robot Intensity}}_{kt} \tag{21}$$

where Task Exposure$_{it}$ captures the extent to which individual i performs tasks that are automatable by industrial robots, and $\widetilde{\text{RobotIntensity}}_{kt} \in [0, 1]$ is a normalized measure of robot stock in industry k at time t. The normalization ensures that exposure is bounded between zero and unity and thus by construction, occupations in industries with no robot adoption are assigned zero exposure, while workers in highly automated industries are fully exposed to the task-specific automation potential implied by the IFR data.

This approach has several advantages. First, it introduces within-industry heterogeneity in exposure, allowing us to distinguish between workers performing different tasks within the same production environment. This is particularly important given that automation technologies target specific activities rather than entire industries. Second, it provides a closer empirical counterpart to the task-based theory, in which automation reallocates a subset of tasks rather than uniformly affecting all workers in a sector. Third, by interacting technological feasibility (task exposure) with realized adoption (industry robot intensity), the measure captures both the supply of automation technologies and their economic implementation, mitigating concerns that task-based measures alone overstate effective exposure.

Overall, this task-based specification allows us to test more directly the core mechanism of the model, namely that automation affects labour-market outcomes through the selective reallocation of tasks, and to identify heterogeneous effects across workers depending on the task content of their occupations. However, this specification comes at a cost. Most importantly, it cannot be implemented using the administrative employer-employee data used in the baseline analysis, as these data do not contain information on occupations performed. As a result, we rely instead on survey data, specifically the Slovak Structure of Earnings dataset, which provides detailed information on workers' occupations and characteristics although the extract available covers a shorter time span (2017–2022) and a more limited set of firms.

A further limitation of the survey data is the absence of stable worker identifiers. Individual workers cannot be tracked across firms and, in many cases, not even within firms over time. This precludes the use of worker or matched employer-employee fixed effects, which are a central feature of the baseline specification. Consequently, unobserved heterogeneity at the worker level cannot be differenced out in the same way as in the administrative data. Furthermore, implementing the IV strategy described above would also be infeasible as we lack gender-differentiation in the data on labour force for several countries. Since gender has been found a potentially important source of heterogeneity (Ge and Zhou, 2020), maintaining its interaction with robot exposure should not yield to the sparse set of instruments.

We address these limitations by exploiting the richer set of observable worker characteristics available in the survey and adopting a flexible parametric specification that allows for heterogeneous effects of automation across demographic and occupational groups. In particular, we estimate models that fully interact robot exposure with occupation and gender:

$$\begin{aligned}Y_{ijkt} = {} & \beta_{gpr}\text{Exposure}_{ijkt} \times \text{Female}_i \times \text{ISCO}_{it} + \\ & + \beta_{gr}\text{Exposure}_{ijkt} \times \text{Female}_i + \\ & + \beta_{pr}\text{Exposure}_{ijkt} \times \text{ISCO}_{it} + \\ & + \beta_r\text{Exposure}_{ijkt} + \mathbf{X}_{it}\gamma + \alpha_j + \alpha_t. \end{aligned} \tag{22}$$

The set of estimates β_{gpr}, β_{gr}, β_{pr}, β_r allow for treatment effect heterogeneity of *r*obot exposure by *g*ender and *p*rofession while controlling for a comprehensive set of worker characteristics $\mathbf{X}_{it}$. These include tenure within the firm, constructed from the hiring date to either the end of the calendar year or the termination date, a part-time indicator, age and education categories. While this approach does not fully replicate the identification afforded by matched employer-employee data, it allows us to account for a rich set of observable differences across workers and firms and to capture systematic heterogeneity in the effects of automation across occupations and demographic groups. Moreover, as we show below, we find very little evidence of robot adoption responding to *local* labour market conditions (as has been found elsewhere in the literature), which means that the IV strategy, while useful to uncover the demand structure for robots does not yield meaningfully different estimates than

OLS and thus specification (22) is on a much more secure footing than first appearances might suggest.

Finally, we have to surmount an important limitation of the task-based specification and the baseline model. Both of these approaches capture only the realized employment relationships observed in the data. In particular, the analysis identifies the effects of automation through outcomes such as job separations or wage changes among workers who are already employed. By construction, it is silent about the set of workers who may not have been hired in the first place due to automation. This selection margin is potentially important, as firms may respond to robot adoption primarily by adjusting hiring policies and vacancy creation rather than by displacing incumbent workers.

To address this concern, we complement the worker-level analysis with a firm-level specification that captures the extensive margin of labour demand. Specifically, we construct measures of full-time equivalent (FTE) employment within individual occupations (by 1-digit ISCO codes) at the firm level and examine how these evolve with automation exposure. This allows us to observe whether robot adoption is associated with a decline in employment, in addition to shifts in the composition of employment across occupations.

In practice, we estimate models of firm-level employment in which the dependent variable is the number of FTE workers in a given firm and within occupation groups, and the key explanatory variable is the firm's exposure to automation, controlling for sector-specific time trends and firm fixed effects. This approach captures both the displacement of incumbent workers and changes in hiring, thereby supplementing the picture of labour demand responses gleaned from the worker-specific models. In particular, it allows us to detect whether occupations that are more exposed to robot-performed tasks experience a relative decline in employment within firms that adopt automation technologies.

Taken together, the worker-level and firm-level specifications provide complementary evidence on the effects of automation. While the former identifies impacts on wages and separations among employed workers, the latter captures broader adjustments in employment levels and occupational composition that may arise through changes in hiring and job creation.

5 Results

Throughout this section, we will present the estimation results as a unitless percentage of the last observed value of the dependent variable in the sample:

$$100 \times \widehat{\beta} \times \frac{\overline{R}_T - \overline{R}_0}{\overline{Y}_T} \tag{23}$$

where $\widehat{\beta}$ is the estimated effect of the change in robot exposure R on the dependent variable Y. Subscripts 0 and T denote the periods at the start and end of the sample period respectively and the overbar indicates the mean value. In other words (23) expresses the percentage of the the dependent variable at the end of the observed period that can be attributed to the change in exposure to robots. In this way, the differences in scale of both the dependent and predictor variables cancel out, which facilitates comparisons across specifications. For the purposes of standard error calculations the fraction in (23) is treated as a fixed constant because only the variance of $\widehat{\beta}$ is material to the determination of statistical significance.

5.1 Main IV estimates

The baseline results are obtained from TSLS regressions with robot exposure measured by the predicted number of installations or R&D expenditure. Table 3 reports the results for both specifications while Figure 4 presents the IFR results graphically.

Tab. 3: Baseline IV estimates

Dep. var.	Data	Period	Effect (%)	SE	N	R^2
Switch	IFR	2004-2010	-0.816***	(0.256)	14,742,840	0.13
Switch	IFR	2011-2019	-0.0855	(0.1)	24,201,256	0.09
Switch	IFR	2020-2022	0.599**	(0.301)	8,169,688	0.07
Switch	R&D	2011-2019	0.261*	(0.151)	10,787,461	0.08
Dismissal	IFR	2004-2010	0.876***	(0.311)	15,027,993	0.23
Dismissal	IFR	2011-2019	-0.271	(0.212)	24,569,748	0.20
Dismissal	IFR	2020-2022	0.235	(0.226)	8,529,129	0.20
Dismissal	R&D	2011-2019	0.0157	(0.0919)	11,168,486	0.22
Wage	IFR	2004-2010	0.198	(0.219)	12,755,146	0.04
Wage	IFR	2011-2019	0.793***	(0.17)	19,354,196	0.17
Wage	IFR	2020-2022	0.131	(0.102)	6,223,045	0.28
Wage	R&D	2011-2019	0.0465**	(0.0204)	8,406,410	0.18

Notes: Estimates from TSLS fixed effects (FE) models on yearly employer-worker data. Standard errors (SE) clustered by industries. Significance codes: * 10%, ** 5%, *** 1%. All regressions control for number of workers, proportion of male workers, and yearly FEs. Earnings equations are estimated with FE for each employer-worker dyad, while the remaining models condition on employer FEs only. Dropping singleton observations leads to varying sample sizes for different models.

Fig. 4: Baseline IV estimates (IFR model)

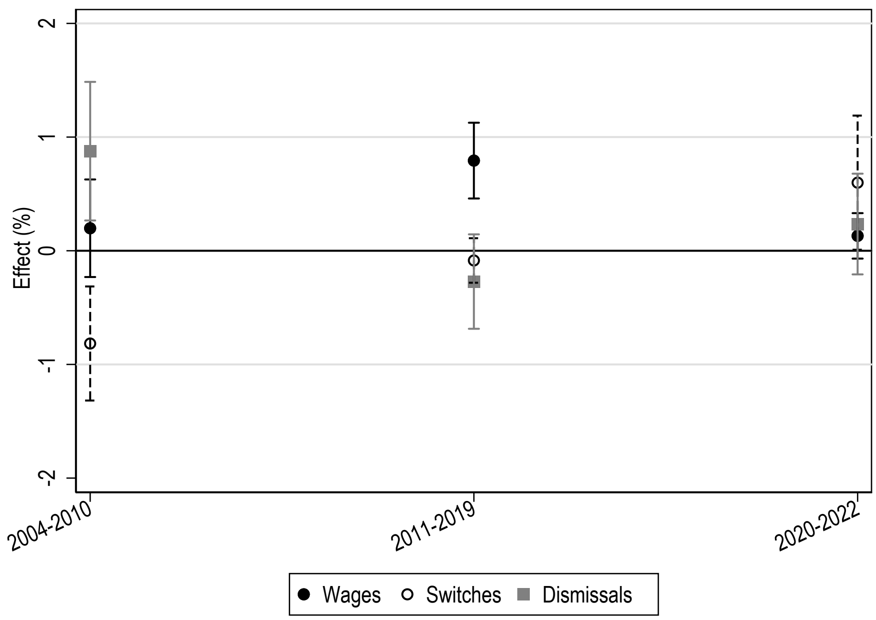

Across all specifications there is an agreement that workers exposed to robotic technologies benefited from higher wages at least in the period of recovery 2011-2019. In the IFR specification, we find that robotic installations accounted for about 0.8% of the wages in 2019. The R&D estimate is notably lower but still positive and statistically significant showing about 0.05% wage increase attributable to robots.

One possible interpretation of the difference between the IFR specification and its R&D counterpart could be related to the levels at which the key regressor is measured: while the IFR data show the robot adoption within the *industry*, the R&D data show the effect of installations within the workers' own firm. If this interpretation is correct, it would mean that the direct effect of robot exposure is smaller than the indirect effect, whereby robot installations among the firm's competitors lead to wage increases. This finding would be consistent with the previous literature (Graetz and Michaels, 2018; Acemoglu et al., 2023), which has identified positive effects on wages for workers that have been *indirectly* exposed to robotic technologies. A plausible economic intuition behind positive effect on wages for indirectly affected workers may be that the increased wages among robot-adopters induces wage increases among non-adopters in competition for workers.

As an alternative explanation, it is possible that the measurement error is somewhat heavier in the R&D specification, which makes it more difficult to isolate the effect of robot exposure. Nevertheless, it is encouraging that the different sources of information lead to the same substantive conclusion.

In terms of the fluctuations in the labour market, we find very small effects of robot exposure. In the 2004-2010 period, robot exposure was associated with fewer job switches, but the opposite effect has been found on dismissals. After that, a small positive effect on switches was identified in the IFR data in the crisis years 2020-2022, and perhaps also in 2011-2019 in the R&D data. Increased flux linked to robot adoption has been identified previously in the German data by Dauth et al. (2021) although their results pertained mostly to job re-assignments within the same workplace and diversion of younger workers away from jobs exposed to automation. Small and positive effects on job switches in response to robot exposure is consistent with the popular narrative of workers pro-actively seeking jobs with less potential for automation. Overall, however, the employment effects are very small and not statistically significant, which rules out the hypothesis of large displacement of labour by technology.

5.2 Task-based models

Figure 5 reports the estimated effects of robot exposure on wages and job separations across occupations and gender groups using the task-based specification. The plots display point estimates and 95% confidence intervals for ISCO major groups 2, 3, 7, 8, and 9. ISCO groups 1, 4, and 5 are excluded as they are not covered by the OECD (2019) task concordance, while ISCO group 6 is omitted due to its negligible representation in the data (less than 0.5 percent of observations), which precludes reliable estimation.

The results also show little evidence of economically meaningful effects of automation on either wages or separations. Across all occupations and demographic groups, point estimates are small in magnitude and tightly estimated around zero. In most cases, the confidence intervals are sufficiently narrow to exclude even modest effects. In this sense, the estimates can be characterized as precise zeros in that not only are they statistically insignificant, but they also rule out effects of a magnitude that would be economically relevant.

This pattern is evident in both outcomes. In the wage specification (Figure 5b), estimates fluctuate around zero with no systematic differences across occupations or between men and women. Similarly, in the separation specification (Figure 5a), the estimated effects are close to zero and precisely estimated, with no consistent pattern indicating increased displacement in more exposed

Fig. 5: Task-based results from survey data

(a) Separations

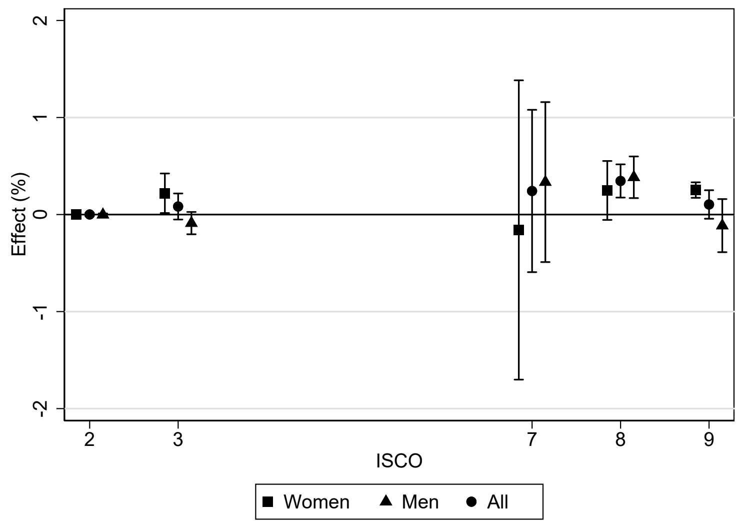

(b) Wages

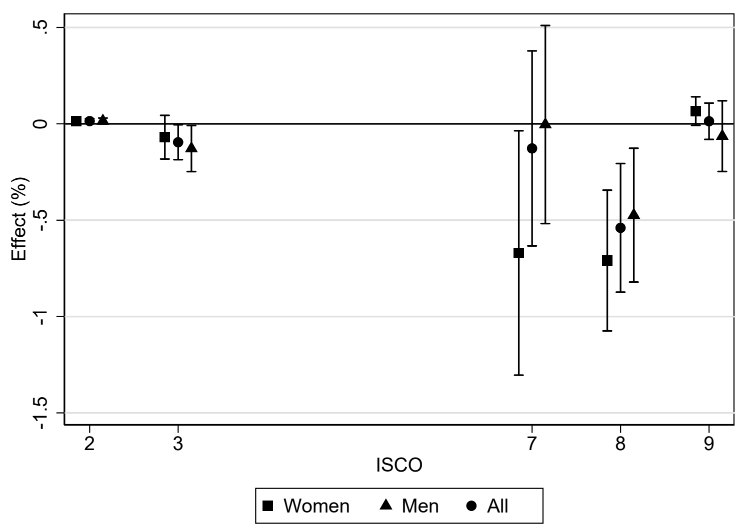

occupations. While a small number of coefficients display wider confidence intervals, particularly in ISCO group 7, these estimates remain centred near zero and do not suggest robust effects.

These findings strongly corroborate the baseline results obtained using the instrumental variables strategy, which also indicate no significant impact of robot adoption on wages or job separations. The task-based specification therefore provides an important robustness check: despite relying on a different identification strategy, a different dataset, and a more granular measure of exposure, it yields qualitatively identical conclusions.

Importantly, the absence of detectable heterogeneity across occupations suggests that the null effects are not driven by aggregation bias in the baseline specification. If automation had substantial effects concentrated in specific task groups, the task-based approach would be expected to uncover such patterns. The fact that it does not reinforces the conclusion that, in this setting, robot adoption has had limited impact on labour-market outcomes.

Finally, while the restriction to a subset of ISCO occupations reduces the scope of the analysis, it does not appear to drive the results. The occupations included span a wide range of task profiles, including both routine and non-routine activities, and there is no indication that the excluded categories would overturn the overall conclusion. Taken together, the evidence from the task-based specification provides strong support for the baseline finding of negligible effects of automation on wages and employment outcomes.

We complement the worker-level analysis with firm-level estimates of the effects of automation on wages and labour demand, measured in terms of full-time equivalent employment. Figure 6 reports the corresponding coefficients across ISCO groups.

Relative to the worker-level results, the firm-level estimates are noticeably noisier, with wider confidence intervals and greater dispersion across occupations. This is consistent with the more aggregated nature of the outcome variables and the reduced sample size at the firm level. Nonetheless, the substantive conclusions remain unchanged. Point estimates are generally small, if sometimes rather imprecisely estimated, and there is no systematic pattern indicating either declines in employment or wages in occupations more exposed to robot-performed tasks.

In particular, while some coefficients, especially for production-related occupations (ISCO 7 and 8), display larger variance, they remain centred close to zero and do not provide robust evidence of negative employment effects. Crucially, there is no indication of a consistent decline in firm-level labour demand in more exposed occupations, suggesting that automation has not led to measurable reductions in the number of workers employed within firms.

Fig. 6: Task-based models: firm-level employment on FTE-equivalent

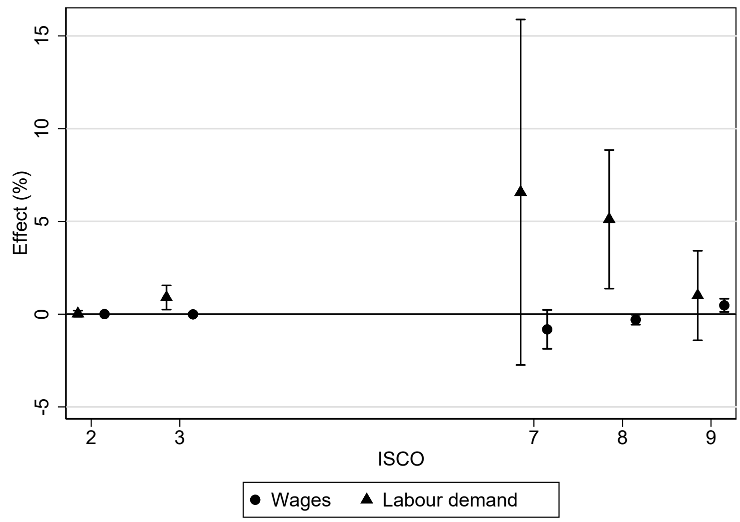

Taken together with the worker-level evidence, these results provide a coherent picture. The absence of effects on separations and wages among incumbent workers is not offset by hidden adjustments along the hiring margin. If firms were responding to automation primarily by reducing hiring rather than displacing existing workers, this would be reflected in declining employment at the firm level. The lack of such a pattern indicates that automation has had limited impact on both the intensive and extensive margins of labour demand.

Overall, the firm-level analysis reinforces the main conclusion of the paper. Despite the additional noise inherent in the aggregated specification, the estimates are consistent with the "precise zero" effects documented at the worker level. The convergence of evidence across datasets, levels of aggregation, and identification strategies strengthens the interpretation that robot adoption has had negligible effects on wages and employment outcomes in this setting.

5.3 Additional estimates

5.3.1 First stage results

Having discussed the baseline regressions, a brief look at the first-stage OLS models is in order. The object here is to separate the supply-side factors influencing the robot take-up with the demand-side ones. As argued above and illustrated in Figure 3, the demand-side reasons for installing robotic technologies could be correlated with the local labour market conditions and thus potentially endogenous in our regressions.

In particular, the decision *not* to adopt any robotic technologies at all may well be subject to firm-specific considerations, such as the availability of finance or availability of workforce, that are expected to influence wage determination and job dismissals. A firm experiencing a downturn in profits will have limited access to the credit necessary to implement a large automation programme. Such a firm may also find it more difficult to offer higher wages, which suggests a dynamic correlation between robot adoption and wages outside the causal channel.

In Figure 7, we show the relationship between the predicted intensity of robot adoption (or R&D activities) and the actual values. It can be immediately observed that, under both measures, a substantial sub-population of firms exists in which no adoption activity is taking place even though, based on their observable characteristics, there should be. This concentration of points with zero indication of actual robot adoption but positive predicted adoption indicates important unobserved factors influencing the adoption decision. In order to remove these unobserved factors from the estimation, only the *predicted* values are used in the two-stage least squares (TSLS) estimation to follow.

Notwithstanding the sub-population of firms whose decision not to adopt robotic technologies is insusceptible to explanation by factors observable to us, the first-stage regressions do show a reasonable fit for the data. The adjusted R-squared for the IFR specification (Figure 7a) is 0.966, which suggests that Slovak firms follows the global trends in robot adoption quite closely. This conclusion is supported by the trend overview in Figure 1. The corresponding Kleibergen-Paap F-statistics for the IFR specification with dynamic firm-level controls (17) for the entire observed period is 183,394.9 which indicates conventionally “strong" first stage. The first-stage under the R&D specification is appreciably weaker with adjusted R-squared of 0.504 but still high F-statistic of 193,256.9 including controls computed in (17). Excluding these controls actually leads to an increase in the F-statistics to 644,086.6 and 627,730.5 in IFR and R&D specifications respectively, which convincingly obviates the need to use weak-IV inference techniques.

5.3.2 OLS results

Table 4 shows results from models equivalent to the baseline IV regressions in Table 3 but dispensing with the IV correction for endogeneity. Here, we observe a result that has been broadly found in the literature, specifically, the en-

Fig. 7: First-stage results: firm-level observed stocks of robots vs. predicted stocks. Variables are plotted using the $\ln(x+1)$ transformation for clarity.

(a) IFR data

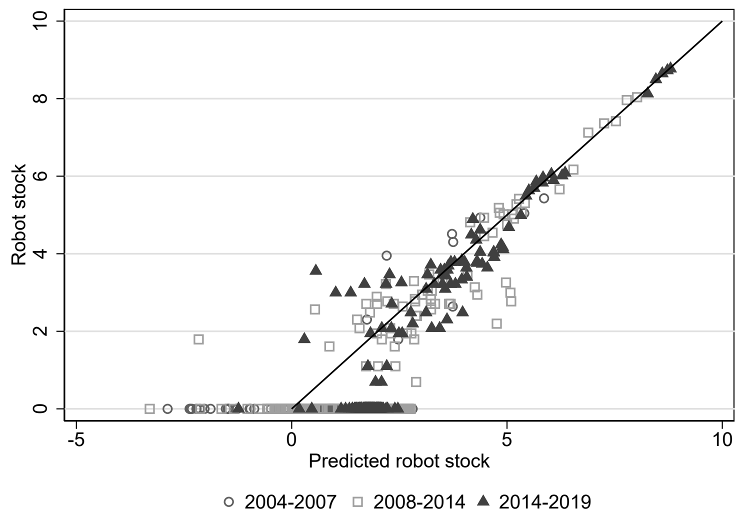

(b) R&D tax credit data

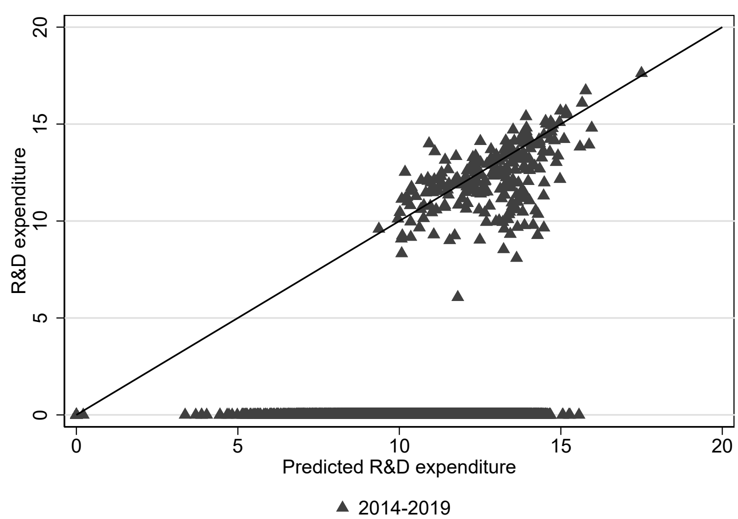

dogeneity correction yields only a minor change in the estimated coefficients, cf. Graetz and Michaels (2018, Tables 1-4) or Dauth et al. (2021, esp. Table 3). These results would support the proposition that robot adoption is generally driven by supply as opposed to demand and thus the robot stocks are (nearly) exogenous with respect to labour market outcomes.

Tab. 4: OLS results with baseline specifications

Dep. var.	Data	Period	Effect (%)	SE	N	R^2
Switch	IFR	2004-2010	-0.84***	(0.232)	14,742,840	0.13
Switch	IFR	2011-2019	-0.0201	(0.0843)	24,201,256	0.09
Switch	IFR	2020-2022	0.374	(0.249)	8,169,688	0.07
Switch	R&D	2011-2019	-0.0132	(0.0372)	10,787,461	0.08
Dismissal	IFR	2004-2010	0.572***	(0.199)	15,027,993	0.23
Dismissal	IFR	2011-2019	-0.346*	(0.193)	24,569,748	0.20
Dismissal	IFR	2020-2022	0.312*	(0.183)	8,529,129	0.20
Dismissal	R&D	2011-2019	-0.04**	(0.0179)	11,168,486	0.22
Wage	IFR	2004-2010	0.32**	(0.124)	12,755,146	0.04
Wage	IFR	2011-2019	0.624***	(0.133)	19,354,196	0.17
Wage	IFR	2020-2022	0.138*	(0.0727)	6,223,045	0.28
Wage	R&D	2011-2019	0.0785***	(0.00804)	8,406,410	0.18

Notes: Estimates from OLS fixed effects (FE) models on yearly employer-worker data. Standard errors (SE) clustered by industries. Significance codes: * 10%, ** 5%, *** 1%.

5.3.3 Robustness checks

In the above specifications, we estimated the impact of robot exposure expressed in relative terms, i.e. number of robots per 1000 workers. This specification is certainly preferable to the absolute number of robots within a firm or an industry as it describes the density of robots within a particular industry. Nevertheless, this plan is open to the objection that it intensifies measurement error by using estimated sizes of the workforce engaged by particular industries. This is especially problematic for China and Taiwan where these data are very sparse (see the Data section).

As Table 5 shows, using absolute numbers of robots does not change the substantive conclusions but the results are somewhat more erratic. In particular, a very large negative coefficient on job switches is found for the 2004-2010 period. As this is a time of generally low robot exposure, the absence of normalisation by the number of workers is likely causing sensitivity to outlying

Tab. 5: Alternative IV regressions

DV	Data	Period	Effect (%)	SE	Bad cn.	Spec.	Ind. tr.
Swi.	IFR	2004-2010	-10.99***	(0.534)	—	Abs	—
Swi.	IFR	2004-2010	-0.995***	(0.375)	Yes	Abs	—
Swi.	IFR	2011-2019	-0.00276	(0.163)	—	Abs	—
Swi.	IFR	2011-2019	-0.141**	(0.0598)	Yes	Abs	—
Swi.	IFR	2020-2022	0.0412	(0.205)	Yes	Abs	—
Swi.	IFR	2020-2022	0.118	(0.268)	—	Abs	—
Swi.	R&D	2011-2019	0.129	(0.286)	Yes	Rel	Yes
Swi.	R&D	2011-2019	0.0474*	(0.0259)	Yes	Abs	—
Swi.	R&D	2011-2019	0.122*	(0.0675)	Yes	Abs	Yes
Dis.	IFR	2004-2010	1.26**	(0.52)	Yes	Abs	—
Dis.	IFR	2004-2010	1.72**	(0.865)	—	Abs	—
Dis.	IFR	2011-2019	-0.0916	(0.193)	—	Abs	—
Dis.	IFR	2011-2019	-0.158	(0.152)	Yes	Abs	—
Dis.	IFR	2020-2022	0.399*	(0.241)	Yes	Abs	—
Dis.	IFR	2020-2022	0.351	(0.277)	—	Abs	—
Dis.	R&D	2011-2019	0.0571	(0.0611)	Yes	Abs	Yes
Dis.	R&D	2011-2019	-0.00719	(0.0158)	Yes	Abs	—
Dis.	R&D	2011-2019	-0.0709	(0.165)	Yes	Rel	Yes
Wage	IFR	2004-2010	0.143	(0.359)	Yes	Abs	—
Wage	IFR	2004-2010	0.217	(0.357)	—	Abs	—
Wage	IFR	2011-2019	1.02***	(0.153)	—	Abs	—
Wage	IFR	2011-2019	0.654***	(0.122)	Yes	Abs	—
Wage	IFR	2020-2022	0.187***	(0.0576)	Yes	Abs	—
Wage	IFR	2020-2022	0.357***	(0.0943)	—	Abs	—
Wage	R&D	2011-2019	-0.00212	(0.0023)	Yes	Abs	—
Wage	R&D	2011-2019	0.0182	(0.016)	Yes	Rel	Yes
Wage	R&D	2011-2019	0.000626	(0.0024)	Yes	Abs	Yes

Notes: Estimates from TSLS fixed effects (FE) models on yearly employer-worker data. Standard errors (SE) clustered by industries. Significance codes: * 10%, ** 5%, *** 1%. Bad cn. = inclusion of “bad controls" from eq. (17); DV = Dependent variable; Spec. = specification of the endogenous regressor with “Abs" standing for absolute number of robot installations and “Rel" for installations per 1000 workers; Ind. tr. = industry-specific time trends.

observations.

Another potential issue may be from the inclusion of “bad controls" from Equation (17), which could cause endogeneity problems due to within-cluster correlations of the error terms. Once again, results in Table 5 show that the impact of these variables is modest. If anything, results including the “bad controls" tend to be even closer to zero; for instance, the wage difference for 2020-2022 from IFR data was estimated at 0.357% without the “bad controls" and their inclusion reduces the estimate to 0.187%, but both of these estimates are significant at the 1% level.

Next, we address the concern about potential trend heterogeneity across different industries. Assuming that the industries are moving asynchronously along the business cycle, the baseline specification with only a single set of time fixed effects could have insufficient controls for the temporal variation between firms. This criticism can be addressed only within the R&D specification in which the robot exposure varies on the firm level. However, inclusion of industry-specific time trends also has a negligible impact on the estimates.

Finally, we consider potential concerns related to the choice of countries used to construct the instrumental variables approximating the supply of robotic technologies. If the exogeneity of the instruments were compromised, or if there were unobserved heterogeneity in treatment effects, alternative choices of instruments could lead to substantially different estimates.

Table 6 reports estimates based on robot installations in Asia (i.e., using installations in China, India, Japan, Korea, Thailand, and Taiwan as instruments) and compares them to a specification that relies exclusively on installations in North America (i.e., the United States and Mexico). In both cases, the results are closely aligned with the baseline specification, including the more volatile estimates for employer transitions in the 2004–2010 period.

These findings suggest that potential linkages between European markets and the United States do not pose a threat to identification. More broadly, there is no evidence that alternative instrument sets uncover systematically different effects.

6 Conclusion

This paper studies the effects of robot adoption on wages and employment in one of the most automation-exposed economies in Europe and the world leader in production of cars per capita. Despite this unusually strong exposure, the Slovak experience over the period we study closely mirrors that of larger European economies undergoing rapid technological change. Across a

Tab. 6: Restricted instrument sets

Y	Data	Period	Effect (%)	SE	IV	Spec.
Sw.	IFR	2004-2010	-0.953***	(0.362)	Asia	Abs
Sw.	IFR	2004-2010	-10.5***	(0.505)	America	Abs
Sw.	IFR	2011-2019	-0.0596	(0.0644.)	Asia	Abs
Sw.	IFR	2011-2019	-0.0659	(0.059)	America	Abs
Sw.	IFR	2020-2022	-0.0482	(0.14)	America	Abs
Sw.	IFR	2020-2022	-0.119	(0.282)	Asia	Abs
Dis.	IFR	2004-2010	2.15***	(0.606)	America	Abs
Dis.	IFR	2004-2010	1.26**	(0.528)	Asia	Abs
Dis.	IFR	2011-2019	-0.164	(0.173)	America	Abs
Dis.	IFR	2011-2019	-0.175	(0.152)	Asia	Abs
Dis.	IFR	2020-2022	0.494	(0.301)	Asia	Abs
Dis.	IFR	2020-2022	0.214*	(0.12)	America	Abs
Wage	IFR	2004-2010	0.382	(0.38)	America	Abs
Wage	IFR	2004-2010	0.126	(0.367)	Asia	Abs
Wage	IFR	2011-2019	0.696***	(0.14)	America	Abs
Wage	IFR	2011-2019	0.652***	(0.126)	Asia	Abs
Wage	IFR	2020-2022	0.151***	(0.0319)	America	Abs
Wage	IFR	2020-2022	0.181***	(0.0665)	Asia	Abs

Note.: Estimates from TSLS fixed effects (FE) models on yearly employer-worker data. Standard errors (SE) clustered by industries. Significance codes: * 10%, ** 5%, *** 1%. Y = Dependent variable, while Sw. = Switches; Dis. = dismissals; Wage = wages. IV = subset of countries used for constructing instrumental variables; Spec. = specification of the endogenous regressor: “Abs" = Absolute number of installations.

range of specifications, datasets, and identification strategies, we find no evidence of economically meaningful effects of robotisation on labour-market outcomes. Our estimates are generally not only statistically insignificant but tightly concentrated around zero, allowing us to reject effects of plausible economic magnitude. In this sense, the results can be interpreted as *precise zeros*.

The absence of substantial employment effects is consistent with evidence from Czech and Dutch firms and stands in contrast to findings of displacement in the United States and China. At the same time, the modest positive wage effects we detect are in line with evidence from Germany, the Netherlands, and cross-country studies (Graetz and Michaels, 2018; Dauth et al., 2021; Acemoglu et al., 2023). We also find suggestive indications that wage gains are not confined to robot-adopting firms but extend to their competitors. This pattern is suggestive of general equilibrium spillovers, whereby increased demand for skilled labour among adopters raises wages more broadly, forcing non-adopting firms to adjust compensation in order to retain workers.

Taken together, these findings are consistent with the predictions of task-based models of automation. In these frameworks, automation simultaneously displaces labour from a subset of tasks and increases productivity in remaining and potentially new tasks, generating offsetting forces on wages and employment. Our empirical results suggest that, in the Slovak context, these forces largely cancel out. Importantly, this conclusion is not driven by a particular specification or dataset. Rather, it emerges consistently across administrative and survey data, worker-level and firm-level analyses, and alternative measures of automation exposure. The convergence of evidence across these approaches strengthens the interpretation that the net effects of robot adoption on labour-market outcomes are limited.

At an intuitive level, a plausible interpretation of these results is that firms primarily use robotic technologies to augment production processes rather than to substitute directly for labour. In this case, automation raises productivity without substantially reducing the demand for workers, consistent with models in which robot adoption is largely neutral to labour market outcomes.

The policy implications are as straightforward as they are forceful. First, the results provide little support for protectionist responses to automation. Restricting the adoption of robotic technologies would likely reduce firm competitiveness without delivering measurable labour-market benefits. Second, while aggregate employment effects appear limited, the literature suggests that adverse impacts may arise in specific contexts or for particular groups of workers. Policies that facilitate worker reallocation, including training and re-skilling, remain important for ensuring that workers can transition into tasks where labour retains a comparative advantage.

Several limitations point to directions for future research. Our empirical analysis relies on proxy measures of robot adoption rather than direct measures at the workplace level. In addition, the available data do not allow for a detailed decomposition of effects across narrowly defined occupations or task categories. Further work exploiting richer data on task content and workplace-level technology adoption would help to refine the mechanisms underlying the aggregate patterns documented here.

Overall, the evidence suggests that, even in a highly exposed economy, the labour-market effects of robotisation are modest and well captured by a framework in which opposing forces of displacement and productivity gains largely offset each other.

References

D. Acemoglu and P. Restrepo. Robots and jobs: Evidence from US labor markets. *National Bureau of Economic Research Working Paper Series*, No. 23285, 2017. doi: 10.3386/w23285. URL http://www.nber.org/papers/w23285.

D. Acemoglu and P. Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018. URL https://ideas.repec.org/a/aea/aecrev/v108y2018i6p1488-1542.html.

D. Acemoglu and P. Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2): 3–30, 2019.

D. Acemoglu, C. LeLarge, and P. Restrepo. Competing with robots: Firm-level evidence from France. *National Bureau of Economic Research Working Paper Series*, No. 26738, 2020.

D. Acemoglu, H. R. A. Koster, and C. Ozgen. Robots and workers: Evidence from the Netherlands. *National Bureau of Economic Research Working Paper*, No. 31009, 2023.

P. D. Allison and N. A. Christakis. Fixed-effects methods for the analysis of nonrepeated events. *Sociological Methodology*, 36:155–172, 2006. ISSN 00811750, 14679531. URL http://www.jstor.org/stable/25046694.

L. Aymerich-Franch and I. Ferrer. Liaison, safeguard, and well-being: Analyzing the role of social robots during the COVID-19 pandemic. *Technology in Society*, 70:101993, 2022. ISSN 0160-791X. doi: https://doi.org/10.1016/j.techsoc.2022.101993. URL https://www.sciencedirect.com/science/article/pii/S0160791X22001348.

D. Boddin and T. Kroeger. Servitization, inequality, and wages. *Labour Economics*, 77:102011, 2022. ISSN 0927-5371. doi: https://doi.org/10.1016/j.labeco.2021.102011. URL https://www.sciencedirect.com/science/article/pii/S0927537121000464.

M. Bělín and M. Veselková. The job retention scheme in Slovakia: Impact on dismissals and firm survival in the COVID-19 pandemic. *Eastern European Economics*, 62(3):361–383, 2023. ISSN 0012-8775. doi: 10.1080/00128775.2023.2254744. URL https://www.tandfonline.com/doi/abs/10.1080/00128775.2023.2254744.

F. Caselli and I. Coleman, Wilbur John. The world technology frontier. *American Economic Review*, 96(3):499–522, 2006. doi: 10.1257/aer.96.3.499. URL https://www.aeaweb.org/articles?id=10.1257/aer.96.3.499.

W. Dauth, S. Findeisen, J. Suedekum, and N. Woessner. The adjustment of labor markets to robots. *Journal of the European Economic Association*, 2021. ISSN 1542-4766. doi: 10.1093/jeea/jvab012. URL https://doi.org/10.1093/jeea/jvab012.

R. C. Feenstra, R. Inklaar, and M. P. Timmer. The next generation of the Penn World Table. *American Economic Review*, 105(10):3150–82, 2015. doi: 10.1257/aer.20130954. URL https://www.aeaweb.org/articles?id=10.1257/aer.20130954.

S. Ge and Y. Zhou. Robots, computers, and the gender wage gap. *Journal of Economic Behavior & Organization*, 178:194–222, 2020. ISSN 0167-2681. doi: https://doi.org/10.1016/j.jebo.2020.07.014. URL https://www.sciencedirect.com/science/article/pii/S0167268120302365.

O. Giuntella and T. Wang. Is an army of robots marching on Chinese jobs? *IZA Discussion Paper 12281*, 2019.

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

F. Jaccoud, F. Petit, T. Ciarli, and M. Savona. Employment over the technology life cycle: Evidence from European regions. *CESifo Working Paper*, No. 10987, 2024.

Š. Jurajda and M. Bělín. *Polarizace českých podniků v důsledku Průmyslu 4.0*. Svaz průmyslu a dopravy České Republiky, Praha, 2020. URL https://www.spcr.cz/images/Studie_320_polarizace_firem_prumysl40.pdf.

A. Jurkat, R. Klump, and F. Schneider. Tracking the rise of robots: The IFR database. *Jahrbücher für Nationalökonomie und Statistik*, 242(5-6):669–689, 2022. doi: doi:10.1515/jbnst-2021-0059. URL https://doi.org/10.1515/jbnst-2021-0059.

J. Laurenson. How Slovakia became the world's number one carmaker, 2026. https://www.bbc.com/news/articles/cm2yvz838deo, Accessed: 2026/04/17.

A. Lewbel. Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. *Journal of Business & Economic Statistics*, 30(1):67–80, 2012. ISSN 0735-0015. doi: 10.1080/07350015.2012.643126. URL https://doi.org/10.1080/07350015.2012.643126.

Å. Madesäter. *World Robotics 2005*. IFR Statistical Department, VDMA Services GmbH, Frankfurt am Main, Germany, 2005.

C. Müller. *World Robotics 2023 – Industrial Robots*. IFR Statistical Department, VDMA Services GmbH, Frankfurt am Main, Germany, 2023.

C. Müller and N. Kutzbach. *World Robotics 2020 – Industrial Robots*. IFR Statistical Department, VDMA Services GmbH, Frankfurt am Main, 2020.

H. Nakamura and M. Nakamura. Production function through automation in a task-based model. *Available at SSRN*, https://ssrn.com/abstract=3461175, 2019. doi: http://dx.doi.org/10.2139/ssrn.3461175.

H. Nakamura, M. Nakamura, and S. Moriwaki. The macroeconomics of automation. *Journal of Economic Growth*, 2026. ISSN 1573-7020. doi: 10.1007/s10887-026-09265-x. URL https://doi.org/10.1007/s10887-026-09265-x.

L. Nedelkoska and G. Quintini. Automation, skills use and training. *OECD Social, Employment and Migration Working Papers*, 202, 2018. doi: doi:https://doi.org/10.1787/2e2f4eea-en. URL https://www.oecd-ilibrary.org/content/paper/2e2f4eea-en.

OECD. Determinants and impact of automation: An analysis of robots' adoption in oecd countries. *OECD Digital Economy Papers*, 277, 2019. doi: https://doi.org/10.1787/ef425cb0-en.

J.-S. Pischke and J. D. Angrist. *Mostly Harmless Econometrics*. Princeton University Press, 2008.

H. K. Pyo, K. Fukao, and T. Miyagawa, 2011. URL http://www.asiaklems.net.

Y. Shen, D. Guo, F. Long, L. A. Mateos, H. Ding, Z. Xiu, R. B. Hellman, A. King, S. Chen, C. Zhang, and H. Tan. Robots under COVID-19 pandemic: A comprehensive survey. *IEEE Access*, 9:1590–1615, 2021. ISSN 2169-3536 (Print) 2169-3536. doi: 10.1109/access.2020.3045792.

B. Smarzynska Javorcik. Does foreign direct investment increase the productivity of domestic firms? In search of spillovers through backward linkages. *American Economic Review*, 94(3):605–627, 2004.

W. Sun, Z. Zhang, Y. Chen, and F. Luan. Heterogeneous effects of robots on employment in agriculture, industry, and services sectors. *Technology in Society*, 75:102371, 2023. ISSN 0160-791X. doi: https://doi.org/10.1016/j.techsoc.2023.102371. URL https://www.sciencedirect.com/science/article/pii/S0160791X23001768.

H. X. Wu and K. Ito. Reconstructing China's supply-use and input-output tables in time series. *RIETI Discussion Paper Series*, 15-E-004, 2015.

Abstrakt

Tato práce odhaduje dopady zavádění robotických technologií na mzdy a zaměstnanost na Slovensku, jedné z ekonomik nejvíce vystavených automatizaci v Evropě. S využitím párovaných administrativních údajů o zaměstnavatelích a zaměstnancích v kombinaci s průzkumovými ukazateli míry vystavení pracovních úkolů automatizaci identifikujeme dopad nových technologií pomocí strategie instrumentálních proměnných, která využívá exogenní variaci zavádění robotizace v mimoevropských ekonomikách. V rámci všech modelových specifikací jsme nezjistili žádné důkazy o ekonomicky významných dopadech zavádění robotů na vývoj na trhu práce. Odhadované dopady na mzdy a ukončení pracovního poměru jsou malé a přesně odhadované kolem nuly. Slovenské odhady poukazují na zvýšenou produktivitu u robotizovaných firem, což vytváří nové pracovní příležitosti, místo aby nahrazovalo lidskou práci.

Working Paper Series
ISSN 2788-0443

Individual researchers and the on-line versions of CERGE-EI Working Papers (including their dissemination) are supported by RVO 67985998 from the Economics Institute of the CAS

Specific research support and/or other grants are acknowledged at the beginning of the paper.

Published by
Charles University, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute of the Czech Academy of Sciences (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, Czech Republic
Phone: + 420 224 005 153
Email: office@cerge-ei.cz
Web: https://www.cerge-ei.cz/

Editor: Byeongju Jeong

The paper is available online at https://www.cerge-ei.cz/working-papers/.

Electronically published June 1, 2026

ISBN 978-80-7343-628-5 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)