Displacing and Reinstating Technologies in Production Networks: Workers and Firms

Stan de Ruijter*

November 15, 2026

Latest version

Abstract

Technological change can affect workers far beyond the firms that develop or adopt new technologies. I study how innovation at suppliers affects firms and workers downstream, using Danish matched employer–employee data and measures of labor-displacing and labor-reinstating technological change, and show that the effects are differentiated by technological type. A shift in suppliers' innovation toward displacing rather than reinstating technologies is associated with lower employment and labor share decline in buying firms, while value added per worker and operating margins rise, with little change in the pay of incumbent workers. Consequences for workers mirror those for firms: employment and wages for workers are lower when suppliers shift towards displacing rather than reinstating innovations. And worker effects extend throughout careers and beyond their original industries: those same shifts send movers to lower-paying employers and raise exits from employment, while wage losses persist after workers move elsewhere, and partly reflect employment at lower-paying firms. A task-based model rationalizes these findings through two channels: technological improvements embodied in purchased inputs and a relocation of economic activity away from buying firms. The findings suggest that understanding the labor-market consequences of technological change requires looking beyond the industries where innovation originates.

JEL Codes: J23, J31, O33, L14, C67

Keywords: Automation, technological change, production networks, input-output linkages, employer pay premia, matched employer-employee data, Denmark

*University College London, CReAM and RFBerlin. Email: uctpsld@ucl.ac.uk.

1 Introduction

Most industries encounter more new technology through the inputs they buy than through innovation of their own: in Denmark, 84 percent of the average industry's patent-weighted technology exposure originates with its suppliers.[1] Technological change at a supplier can alter the composition of the goods and services it sells, so that some intermediate inputs come to embody tasks previously performed by workers at buying firms, while others demand new types of work from those same workers. Consider a car manufacturer that shifts to supplying its components as sealed, pre-assembled modules. Repair shops no longer need to open the gearbox: they replace the unit, and the diagnostic and repair work their mechanics used to do is now built into the part itself. Other supplier innovations, such as sensors, driver-assistance equipment and electric drivetrains, may instead create demand for calibration, software maintenance and technical diagnostics at those same repair shops. In neither case does the repair shop develop a new technology itself. Yet the work performed by its employees changes.

These examples illustrate a broader feature of technological change in networked economies: innovation reallocates tasks not just within firms, but across them as well. Two literatures each address one part of this phenomenon. The task-based literature shows how technology displaces and reinstates labor where it is adopted (e.g. Acemoglu and Restrepo, 2019, 2022), and the production-network literature shows how sectoral shocks propagate to output and prices across industries (e.g. Acemoglu et al., 2012; Baqaee and Farhi, 2019; Carvalho et al., 2021). Neither has traced how technological change at one industry's suppliers reaches the wages and employment of individual workers in the industries that buy from them. This paper does so. It asks how technological change at a supplier industry affects firms and workers in the industries that purchase its products, and how these effects depend on the nature of the innovation.

To answer these questions, I use Danish matched employer–employee register data, which follow every worker in the country over four decades, and I measure the technology of the industries with which each worker's own industry trades. Each worker's exposure to technological change has three components: the technology of the industries her own industry buys from (the suppliers' channel), the technology of her own industry, and the technology of the industries it sells to (the customers' channel). The paper focuses on the suppliers' channel.[2] Several independent classifications distinguish labor-displacing innovations, which

[1] Exposure is each industry's stock of US breakthrough patents and the stocks of its suppliers, weighted by the predetermined total-requirements coefficients of Section 3.4. The share is computed on stock levels, averaged over 1995–2019.

[2] The own-industry channel is the setting of most existing work on automation and innovation (e.g.

allow purchased inputs to substitute for tasks previously performed within buying firms, from labor-reinstating innovations, which create new tasks for downstream workers. I estimate the effects of suppliers' displacing and reinstating innovations on buying firms' employment, pay, margins and labor shares, and on workers' wages, mobility and employment as they move across employers and industries.

I find that the direction of firm and worker effects is determined by whether their industries' suppliers develop displacing or reinstating technologies. In the firm register over 1999–2019, suppliers' displacing technologies reduce buying firms' employment and wage bills and lower the share of value added paid to labor, while suppliers' reinstating technologies raise all three. Buying firms retain the cost savings. Under displacing technologies value added does not fall; value added per worker rises by 0.201 (s.e. 0.092) and the operating margin by 0.047 (s.e. 0.017), while the pay of incumbent workers (-0.023, s.e. 0.016) and of new hires is minimally affected. Adjustment instead primarily takes place on the hiring margin: hiring falls, the workforce becomes older and more concentrated in routine occupations, and investment declines. There is no evidence of capital deepening at the buyer; suppliers' displacing technologies do not appear as equipment on the balance sheets of the firms that purchase from them.

The burden of adjustment falls on workers who started their careers in the most exposed industries, regardless of their later industry of employment. Over 1995–2019, suppliers' displacing technologies lower the pay of workers from the industries that buy from them, and reinstating technologies raise it by a similar amount. Under the process–product classification the two coefficients are -0.222 (s.e. 0.088) and 0.250 (s.e. 0.084) per log point of suppliers' patent stock, or -0.7 and 0.7 percent of hourly pay per standard deviation. Because the two patent stocks rise and fall together across industries, these estimates identify the effect of the composition of suppliers' technology: buying from suppliers whose technology shifts toward displacing lowers pay, and buying from suppliers whose technology shifts toward reinstating raises it. The negative effect of displacing technology runs against the standard logic of network propagation, under which productivity gains at suppliers lower input prices and raise labor demand in the industries they serve. That logic holds for one type of technology and is reversed for the other. Effects follow the worker rather than the employer: when the exposure of a worker's first industry and the exposure of her current employer's industry enter the same regression, the first industry accounts for the entire effect

Acemoglu and Restrepo, 2020; Webb, 2020; Kogan et al., 2017; Autor et al., 2024). I include it throughout and compare it with the suppliers' channel in Section 5.2. Estimates for the customers' channel are small and not substantially differentiated by type of technology. This is what standard production-network models predict: supply-side shocks such as technological change propagate downstream to buyers rather than upstream to suppliers (Acemoglu et al., 2016).

(−0.215 against 0.018 for displacing technology). Workers who had left their first industry lose as much as those who remained, and losses grow over time.

The changes in pay are accompanied by changes in employer, in employment and in exits from employment. Over 40 percent of the wage effect reflects the pay level of the employer at which a worker ends up. Among workers who change employer, the probability of moving to an employer in the bottom third of the pay-premium distribution rises with suppliers' displacing stock (0.335, s.e. 0.129) and falls with the reinstating stock (−0.413, s.e. 0.124). Across all workers, the probability of being employed falls with the displacing stock (−0.137, s.e. 0.043) and rises with the reinstating stock (0.106, s.e. 0.050). Exits from employment rise by 0.054 per year on a base rate of 0.060, concentrated among workers under 30 and consisting mostly of spells outside the labor force rather than unemployment, and annual earnings respond about twice as strongly as hourly wages. The losses are unevenly distributed. The hourly-wage effects are absent in the lowest fifth of the pay distribution and largest in the highest fifth; within a given job they fall on white-collar workers at large, high-premium employers and in occupations that coordinate production, whereas lower-paid workers adjust through hours and employment. Both types of technology reach buyers through the manufactured inputs they purchase. Splitting suppliers' exposure by type of supplier, I find that equipment makers and other manufacturers account for most of the effect, while the most patent-intensive suppliers that sell no production process, chemicals and pharmaceuticals, account for none of it. Weighting suppliers by buyers' equipment investment rather than by their purchases does not yield results.

The identifying variation I utilize is technological change that arrives through predetermined trade relationships. Exposure measures are constructed via combining fixed input–output shares with the subsequent technology of connected industries measured via US patent data, so that buyers cannot generate exposure by re-sourcing toward suppliers whose technology is changing, and workers retain their first observed industry, such that it cannot be generated via sorting either. The central econometric concern with network-transmitted shocks is that exposed and unexposed workers may differ systematically; the design addresses this by comparing workers with themselves over time. Three checks support this comparison. With municipality-by-year fixed effects the two coefficients are −0.214 and 0.247, almost unchanged from the baseline. Exposure that arrives three years in the future does not predict current wages. And the split appears separately within 1995–2003 and within 2004–2019, at 63 to 72 percent of its full-sample magnitude.

A model of the buying industry rationalizes these facts. The industry's tasks are performed by a purchased input below a task boundary and by its own workers above it; its wage moves with its employment, and cost changes pass through to its price only partially.

A shift of its suppliers' innovation toward process has two effects at once: the input performs more of the buyer's tasks and performs the tasks it already carried out more cheaply, and part of what the buyer used to sell is now produced by the supplier or by others. The buying firms' accounts determine the structure of the model. Purchases contract most, so the input's share of cost falls. Two explanations are consistent with this fall: either the input became cheaper per task and its share fell because buyers demand a given quantity of what it does, or the fall in demand lowered the wage until labor took tasks back. The margin rises, so pass-through is incomplete. Output per buyer falls while the number of buyers does not, a pattern that a cheaper input alone cannot generate and a demand shift can. The buyer's price, its output and its accounts identify the demand shift and the wage curve and bound the capability gain, and Section 6 estimates them from the reduced-form coefficients and reports a counterfactual within the support of the data.

The findings connect two literatures that have developed separately. The automation literature documents displacement where adoption takes place (Acemoglu and Restrepo, 2020; Graetz and Michaels, 2018; Dauth et al., 2021; Koch et al., 2021; Humlum, 2019), building on the task framework of Acemoglu and Restrepo (2019, 2022). The production-network literature shows that sectoral shocks propagate to output and prices (Long and Plosser, 1983; Acemoglu et al., 2012; Baqaee and Farhi, 2019; Carvalho et al., 2021; Barrot and Sauvagnat, 2016), but it has not traced them to individual workers' wages with worker-level identification. The closest antecedent is Acemoglu and Restrepo (2020), who estimate the total local effect of robot adoption on commuting zones. In their analysis production linkages enter as a control strategy rather than as network-propagated exposure, and the setting differs in an important way: the United States adopts frontier automation, whereas Denmark largely buys it embodied in inputs. Their focus is on the places that adopt; mine is on the industries that buy. The closest firm-level evidence is Bilgin et al. (2024), who find that a supplier's first robot purchase raises the productivity, employment and average wage of its customers in Turkish manufacturing. The two sets of results are more compatible than they first appear. On the firm side, suppliers' robot adoption in Denmark is likewise followed by higher capital per worker at buying firms, by somewhat higher employment, measured less precisely, and by no decline in firm-average wages. The papers differ in the object of study. Their wage outcome is the firm average, which moves with hiring and workforce composition, whereas mine is the pay and employment of workers followed from their first industry wherever they end up, the margin on which a move to a lower-paying employer or an exit from employment appears and a firm average does not. Employers are ranked using the pay premia of Abowd et al. (1999) and Card et al. (2013), used here as an outcome of a network-transmitted technology shock. The design draws on the shift-share literature

(Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022; Adao et al., 2019; Anderson and Rubin, 1949); the mechanism echoes the firm-boundary substitution of Feenstra and Hanson (1999), applied within the domestic production network; and the worker-flow channel relates to the literature on mobility spillovers (Balsvik, 2011; Stoyanov and Zubanov, 2012; Poole, 2013; Serafinelli, 2019).

The remainder of the paper proceeds as follows. Section 2 presents the conceptual framework. Section 3 describes the setting, the data and the measurement of technology and exposure, and Section 4 the estimating equations and the research design. Section 5 presents the reduced-form results, beginning with the buying firms and the comparison of own and suppliers' technology, followed by the worker design and the origin-industry effect, workers' pay by technology type together with its identification and scope, employers, employment and exits, the incidence of the losses, and the suppliers that transmit the effects. Section 6 estimates the framework on the reduced-form coefficients, and Section 7 concludes.

2 Conceptual Framework

Production in a firm consists of tasks, such as specifying a component, machining a part or processing an invoice. Each task can be performed by the firm's own workers or purchased already embodied in an intermediate input (Feenstra and Hanson, 1999). When suppliers' innovation shifts toward process technologies, the embodied version of a task becomes cheaper, and the buying firm purchases tasks that it used to perform in-house (Figure 1). The technology is installed at the supplier, but the labor it displaces is employed at the buyer. Innovation that creates new products shifts the boundary between purchased and in-house tasks in the opposite direction.

Two results of Section 5 shape the model. First, under suppliers' displacing technology the buying industry contracts on every nominal margin and purchases contract the most, so that the cost share of the purchased input falls (Section 5.1). A model in which tasks are reallocated to the input at a fixed intermediate share predicts the opposite. Second, output per buying firm falls while the number of buying firms does not, which a cheaper input cannot explain on its own. I therefore allow the input's cost share to respond, and I allow a shift of suppliers' innovation toward process to affect the buyer in two ways: the input performs more of the buyer's tasks at lower cost, and part of the output that the buyer used to sell is produced elsewhere.

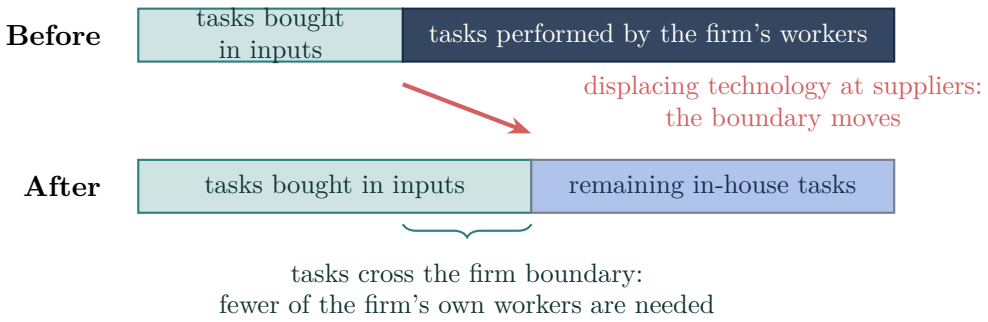


Figure 1: The task boundary inside a buying firm. Each bar represents the buying firm's production. The boundary separates the tasks purchased embodied in inputs from the tasks performed by the firm's own workers. When suppliers' technology shifts toward the displacing type, the embodied version of in-house tasks becomes cheaper and the boundary shifts, so that the firm needs fewer of its own workers to produce the same output. Reinstating technology shifts the boundary in the opposite direction. Both types reach buyers through the manufactured inputs they purchase (Section 5.8).

2.1 Setup

Consider a single buying industry that produces output y from a unit measure of tasks $z \in [N-1, N]$, with elasticity of substitution $\sigma > 0$ between tasks (Acemoglu and Restrepo, 2019, 2022). Each task can be performed either by a purchased manufactured input M, with productivity $\gamma_M(z)$ and price p_M, or by the industry's own labor L, with productivity $\gamma_L(z)$ and wage w. Labor's comparative advantage, $\gamma_L(z)/\gamma_M(z)$, is increasing in z. The input therefore performs all tasks below a boundary I and labor performs the rest, where I is the task at which the two are equally costly, $p_M/\gamma_M(I) = w/\gamma_L(I)$. Under constant returns this assignment determines the unit cost C. Let $c(z)$ denote the cost of task z under this assignment and $s(z) = (c(z)/C)^{1-\sigma}$ its cost share. The input's cost share s_M is the integral of $s(z)$ over the tasks below I, and labor's cost share is $s_L = 1 - s_M$. Two local properties of the task schedule appear in the results below: the cost share of the marginal task, $h = s(I)$, and the slope ρ of $\ln[\gamma_L(z)/\gamma_M(z)]$ at I.

The industry faces the demand curve $y = D\,p^{-\eta}$ and sets its price at a markup over unit cost, $p = \mu C$. Pass-through of cost into the price may be incomplete: $d\ln p = \varphi\, d\ln C$ with $\varphi \in (0, 1]$, so that $d\ln\mu = -(1-\varphi)\, d\ln C$. This is the local form of any pricing rule with variable markups, and $\varphi = 1$ corresponds to a constant markup. The industry's wage increases with its employment along a wage curve, $\ln w = \ln\bar{w} + \varepsilon^{-1}\ln L$, which is a reduced form for the wage differences between industries that the data identify. With turnover $R = py$, the wage bill is $wL = (s_L/\mu)R$ and purchases are $p_M M = (s_M/\mu)R$.

These two expressions imply the ratios that are observed in the firm accounts:

$$\frac{p_M M}{R} = \frac{s_M}{\mu}, \quad \frac{wL}{p_M M} = \frac{s_L}{s_M}, \quad \frac{wL}{\text{VA}} = \frac{s_L}{\mu - s_M}, \quad \frac{\text{VA} - wL}{R} = 1 - \frac{1}{\mu}, \quad \frac{\text{VA}}{L} = w\,\frac{\mu - s_M}{s_L}. \tag{1}$$

In the Danish accounts, purchases average 0.527 of turnover and the operating margin averages 0.124, which imply $\mu = 1.14$ and $s_M = 0.60$. The model abstracts from the industry's own capital. The buying firms' investment rate and capital per worker do not rise under suppliers' displacing technology, and the effect is transmitted through purchased inputs rather than through equipment (Section 5.8).

The shock. Suppose that suppliers' innovation shifts toward process by one unit (Section 2.3 defines the unit). The shift affects the buyer in two ways. First, the input's productivity in the tasks it performs rises by $g \geq 0$ at an unchanged price per unit, so that its cost per task falls by g and the boundary shifts. I refer to g as the capability gain. Second, the buyer's demand shifter falls by $\delta \geq 0$, that is, $d\ln D = -\delta$. This demand shift represents the part of the buyer's sales that, after the suppliers' process innovation, is produced by the supplier or by a third party. Setting $\delta = 0$ or $g = 0$ isolates each effect. A shift toward product innovation is the same shock with the opposite sign.

2.2 Comparative Statics

All responses are expressed per unit of the shock and evaluated at the initial equilibrium, with the wage adjusting along the wage curve; Appendix D.2 gives the derivations. Let $\omega = g + d\ln w$ denote the cost advantage that the input gains over labor at the marginal task, net of the wage response. The boundary shifts by ω/ρ. Because unit cost is continuous at the boundary, this reassignment of tasks has no first-order effect on cost, and $d\ln C = -s_M g + s_L\, d\ln w$. The input's cost share changes by

$$ds_M = \omega A s_L, \qquad A = \frac{h}{\rho\, s_L} - (1-\sigma)\, s_M. \tag{2}$$

The first term of A reflects the tasks that are reassigned to the input. The second reflects the tasks that the input already performs: they become cheaper, and when tasks are gross complements ($\sigma < 1$) their share of cost declines. Real output and turnover follow from the demand curve, $d\ln y = -\delta - \eta\varphi\, d\ln C$ and $d\ln R = -\delta + (1-\eta)\varphi\, d\ln C$. Combining labor

demand, $wL = (s_L/\mu)R$, with the wage curve gives

$$d\ln L = -\frac{[A + (1-\eta\varphi)\, s_M]\, g + \delta}{\kappa}, \qquad d\ln w = \frac{d\ln L}{\varepsilon}, \qquad \kappa = 1 + \frac{1 + A - (1-\eta\varphi)\, s_L}{\varepsilon}, \tag{3}$$

and the net cost advantage is

$$\omega = \frac{(\varepsilon + \eta\varphi)\, g - \delta}{\varepsilon + \Lambda}, \qquad \Lambda = A + s_M + \eta\varphi\, s_L, \tag{4}$$

where Λ is the own-wage elasticity of the industry's labor demand.

Three terms lower labor demand in equation (3). The first, A, is the reallocation of tasks to the input. The second, $(1-\eta\varphi)s_M$, is a scale effect: a cost reduction lowers the wage bill at given cost shares when demand is inelastic or when part of the reduction is retained in the markup. The third is the demand shift, which lowers all quantities in proportion. Because all responses are linear in (g, δ), the response to both shocks together is the sum of the responses to each. At a fixed wage, the demand shift alone lowers employment, output, turnover, purchases, the wage bill and value added by δ and leaves the cost shares, the price and the margin unchanged. With a wage curve it affects the cost structure only through the wage, which falls by $\delta/(\varepsilon + \Lambda)$. The input's net cost advantage is positive if and only if $(\varepsilon + \eta\varphi)\, g > \delta$. If the fall in demand is large enough, the wage falls by more than the input's cost per task, and labor takes tasks back from the input.

These expressions have five implications for the outcomes observed in the firm accounts and in the worker data. I state them for a shift of suppliers' innovation toward process; the signs are reversed for a shift toward product.

The input's cost share. The wage bill relative to purchases rises, and the input's cost share falls, if and only if $\omega A < 0$. This can happen in two ways. In the first, the input gains a net cost advantage ($\omega > 0$) and $A < 0$: the lower cost of the tasks that the input already performs reduces its cost share by more than the reassigned tasks raise it. This requires that tasks are gross complements and that comparative advantage changes steeply at the boundary, so that a capability gain reassigns few tasks while lowering the cost of many. In the second, the demand shift lowers the wage by more than the input's cost per task falls ($\omega < 0$), labor takes tasks back, and with $A > 0$ the input's cost share falls. I refer to these as the *input-gains case* and the *labor-retakes case*. The sign of ω distinguishes them, and the buyer's output price identifies it. In both cases purchases fall as a share of turnover.

Price and margin. Unit cost falls by $s_M g - s_L\, d\ln w$, which is positive whenever the wage falls. The buyer's output price falls by φ times this amount, and the operating margin rises by $(1-\varphi)/\mu$ times this amount. The margin therefore rises if and only if pass-through is incomplete.

Scale. Real output falls if and only if the demand shift exceeds the expansion that the lower price induces, $\delta > \eta\varphi\,|d\ln C|$. With the capability gain alone, real output rises. Nominal turnover falls if and only if $\delta > (\eta-1)\varphi\,|d\ln C|$, which always holds when demand is inelastic.

Employment and the wage. Employment and the wage fall if and only if $[A + (1 - \eta\varphi)s_M]\, g + \delta > 0$, and a demand shift lowers both for any value of A. The ratio of the wage response to the employment response is $1/\varepsilon$ whichever shock is at work. The two labor-market responses therefore cannot, on their own, attribute the loss to one shock or the other.

The labor share of value added. When the input's cost share falls, labor's cost share rises. The labor share of value added nevertheless falls if and only if the increase in the markup outweighs the increase in labor's cost share. With a constant markup, the labor share of value added rises whenever the wage bill relative to purchases rises.

Section 5 reports the estimate of each of these responses, and Section 6.1 collects what the estimates imply for the cases distinguished here.

2.3 Network Exposure and the Direction of Innovation

The input of industry j is a bundle of its suppliers' products, and the shock to industry j is the direction of its suppliers' innovation, weighted by its purchases. Let T denote a vector of technology measures over the $N = 121$ industries of the Danish input–output table. The exposures are

$$di_j \equiv T_j, \qquad us_j \equiv \sum_i \Big(L_{\text{total},ij} - \delta_{ij}\Big)\, T_i, \qquad ds_j \equiv \sum_i \Big(G_{\text{dom},ji} - \delta_{ij}\Big)\, T_i, \tag{5}$$

where L_{total} is the Leontief inverse including imported intermediates and G_{dom} is the Ghosh inverse of the domestic sales shares. Subtracting the identity removes the own-industry term, so that the direct and the network exposures are separately identified. The weights are predetermined, and workers are assigned to the first industry in which they are observed (Hummels et al., 2014).

The shock in the model corresponds to the direction $m_j = us_j^D - us_j^R$, where T^D and T^R are the log stocks of suppliers' process-oriented and product-oriented patents. I work with the difference because the two stocks have a correlation of 0.997 in levels. When the two exposures are rotated into a common level and a mix, the level has no effect on the wage (+0.028, s.e. 0.032) and the mix accounts for the whole effect (−0.472, s.e. 0.170). The same holds for the employer premium, employment and the firm outcomes. A common increase in both stocks reflects the overall pace of suppliers' innovation, which the year effects absorb. Industries differ in the direction of that innovation, and a shift toward product is the negative of a shift toward process. Patent stocks have no natural units, so the capability gain per unit of m_j has to be estimated; one residual standard deviation of the direction is 0.033. For any outcome Y, the coefficient on the mix is $\beta_Y = d\ln Y/dm$. It is identified from differences between sectors at decade frequency (Section 5.3).

Firm-level regressions with industry and year effects describe the average firm in an industry, including entrants and exiters. The corresponding within-firm estimates are smaller or zero, which indicates that the industry adjusts as much between firms as within them. The model is written for the industry and does not distinguish the two margins. The industry's own technology T_j affects the demand for its output rather than its tasks, and the estimates are consistent with this (Section 5.2).

2.4 Identification

Given the cost shares and the markup from the accounts, the parameters can be recovered from the reduced-form coefficients sequentially (Appendix D.4). The ratio of the wage coefficient to the employment coefficient equals $1/\varepsilon$. The coefficient for the wage bill relative to purchases gives ds_M and, together with the purchase share of turnover, the markup response, on which the margin provides an over-identifying restriction. The buyer's output price then gives the cost response, $d\ln C = \beta_p - d\ln\mu$. From it follow the capability gain, $g = (s_L\beta_w - d\ln C)/s_M$, the pass-through rate, the net cost advantage, $\omega = g + \beta_w$, and $A = ds_M/(\omega s_L)$. Finally, real output gives the scale response, $\beta_y = -\delta - \eta\beta_p$.

Three features of this mapping matter for the estimation. First, the demand elasticity and the demand shift are not separately identified. Real output, turnover and the wage bill all measure the same scale response, and η enters the model only together with δ. The data therefore identify the line $\delta = -\beta_y - \eta\beta_p$ but not a point on it. I set η to values in $[0.5, 2]$, the range of industry-level estimates in the trade and industrial organization literatures. Because the price response is small, this choice changes δ by about a tenth and leaves the other parameters unaffected. Second, the sign of ω, which distinguishes the input-gains case

from the labor-retakes case, is not determined at the precision of the price coefficient. The share moments identify the product $A\omega$ precisely but A only weakly when ω is close to zero. The estimation therefore restricts the parameters to the region in which the model is well defined ($A \geq -s_M$ and a labor demand curve that slopes down with a bounded elasticity) and reports the fit on each side of $A = 0$. Third, the task primitives σ, h and ρ enter only through A. They are not separately identified, and the analysis does not require them.

The model also has implications that do not depend on the parameter values. Several ratios of coefficients must take a known value: the wage curve across sample splits of the suppliers' channel, the markup implied by the margin relative to the markup implied by the shares, the labor share of value added, the share of value added in turnover, and value added per worker. In addition, the common level of the two patent stocks should have no effect, and the model's responses for the scale outcomes that are not used as targets can be compared with their estimates.

2.5 Limitations

Four limitations of the framework should be kept in mind. First, the coefficients that the model interprets are identified from differences between sectors at decade frequency, and the model adds no identifying variation of its own. The network is held fixed at predetermined weights, and the formation of supplier relationships is outside the model. Second, the model has one producer per industry, whereas the firm accounts describe the average firm, including entrants and exiters. Real output per firm falls, but the number of firms with employees does not (+0.71, s.e. 0.46, on 111 industries), and the estimate for the industry's aggregate real output is not statistically different from zero (+0.26, s.e. 1.23, on the 48 industries with a producer price index, a design that can detect effects larger than 2.4 per unit). The demand shift is therefore identified for the average buyer. Third, the model has a single wage per industry. The division of the wage response between the employer's pay premium and pay within the job is taken from the data (Section 5.6). Fourth, the model describes one buying industry and has no aggregate implications: the net loss across industries is bounded above by the sum of the industry responses and below by zero.

3 Setting, Data, and Measurement

3.1 The setting: a frontier-taking small open economy

Denmark is a buyer of technology rather than a producer of it. The country is small and exceptionally open. It accounts for a negligible share of global R&D and patenting, and the

goods in which new technology is embodied, such as machinery and components, industrial robots, ICT equipment and enterprise software, are for the most part developed at a technological frontier located elsewhere and then imported. Industrial robotics is a familiar example: the leading producers are Japanese and German equipment makers.[3] From the Danish perspective, the movement of the technology frontier is largely predetermined, whereas the timing with which particular technologies enter particular Danish industries reflects the decisions of individual firms. The empirical design rests on this distinction. Section 4 relies on frontier-driven shifters to identify the network coefficients and treats own-industry adoption as an endogenous margin that requires separate treatment.

Openness has a second implication, which concerns interpretation rather than identification. For goods traded at world prices, cost reductions among Danish suppliers need not translate into lower input prices for Danish buyers, because the world market sets the price. Transmission through the network may therefore differ between tradable and non-tradable inputs. The paper notes this heterogeneity but does not test it, because the price margin is measured too imprecisely to support such a test (Section 5.8).

The Danish labor market operates under the flexicurity system (Kreiner et al., 2016), which combines low dismissal costs with replacement rates of up to 90 percent for up to two years and with active retraining and job-search programs. A firm that needs fewer workers can dismiss them at little cost, and a worker who loses a job is insured against much of the income loss. The margins along which adjustment does and does not appear in the results (hiring relative to separations at the firm, employer changes and employment for the worker) are therefore chosen under few institutional constraints, which makes them informative about the underlying economics. Collective bargaining covers roughly 80 percent of workers and compresses wages more than in the United States, though less than in Norway or Sweden. Employment is about 2.9 million out of a population of 5.8 million, and manufacturing accounts for roughly 13 percent of employment.

3.2 Registers

The data are the administrative registers of Statistics Denmark, which are linked at the person, firm, and workplace level through unique identifiers and cover the universe of Danish residents and firms from 1980 to 2023. The linkage matters for this paper because the same person can be followed across firms, industries, and labor-market states for more than four decades, while the technology of the industries with which her industry trades is measured

[3]Denmark does host a robotics cluster around Odense, specializing in collaborative robots, but its output is a small fraction of the world market and is itself largely exported. The frontier that Danish *adopting* industries face is therefore set abroad.

alongside. Matched employer–employee data of this kind have been central to our understanding of how trade and technology shocks reach individual workers (Hummels et al., 2014), and the within-worker design of this paper depends on them. The data are accessed through Statistics Denmark's research scheme, and only aggregated results can be released.

The registers record hourly wages, annual earnings, employment status, industry, occupation, and education for every worker, and they extend the set of outcomes to margins that surveys cannot reach: hospital contacts by diagnosis, benefit receipt, pension contributions, retraining, and exit from the labor force. On the firm side, a survey of technology use records firms' actual adoption of industrial robots and AI, and the firm register reports each firm's value added, wage costs, operating result, fixed assets, and full-time-equivalent employment from 1999 onward. Appendix A.8 lists the registers, describes the construction of the outcome panels, and gives the coverage window of each outcome.

Because every worker is linked to her employer, a wage can be decomposed into a component that belongs to the worker and a component that belongs to the employer. I estimate the two-way decomposition of Abowd et al. (1999) on the matched panel and use the estimated employer effect as the *employer pay premium*: the amount an employer pays all of its workers relative to other employers, net of the characteristics of those workers (Card et al., 2013). The premium is available for 99.4 percent of the main sample. Since it is fixed for each employer, a worker's premium changes only when she changes employer, which makes it a natural measure of where workers end up.

The analysis sample consists of workers aged 18–64 who are observed in employment in at least one year between 1980 and 2023, excluding the self-employed and military personnel. It contains about 131 million worker-year observations on about 5.9 million workers. Each worker's *base industry* is the first industry in which she is observed with a valid code, and it is held fixed for all of her observations regardless of later moves. Section 4.1.2 explains why first-observed assignment is the appropriate intention-to-treat definition in a panel that begins in 1980.

3.3 Measuring technology

No single proxy for technology is beyond question, and the paper therefore does not rely on only one. The measures fall into three families: patent stocks that distinguish displacing from reinstating technology, indices constructed from industry accounts and occupations, and observed adoption. The distinction between displacing and reinstating technology follows the task framework of Acemoglu and Restrepo (2019). The three families are built from different sources and are subject to different kinds of error. Imputed measures err through

crosswalks and international averaging, surveys through sampling and self-reporting, and patents through classification. Agreement across families that share no common construction is informative, and I report the pattern of results across measures.

Technology type from patents. I draw the distinction between displacing and reinstating technology at the level of the individual patent. Starting from the breakthrough patents of Kelly et al. (2021), allocated to industries through patent-class concordances, two classifications divide each industry's patent stock in two. The first uses the automation classification of Mann and Püttmann (2023): automation patents form the displacing stock and all other breakthrough patents the reinstating stock. The second classifies breakthrough patents as process-dominant or product-dominant, with process patents forming the displacing stock and product patents the reinstating stock. The two classifications draw on the same universe of patents but sort it by different rules, and each stock enters the exposures in logs. A third typing, used only as a check on these two, does not rely on the claims of a patent at all. It compares the text of each breakthrough, as summarized by Liu et al. (2025), with the work tasks that existed in the year the patent was issued: the closer the patent is to an existing task, the more it counts toward the displacing stock, and the farther, the more it counts toward the reinstating stock (Appendix A.6). The concordances assign a patent to the industries that develop and manufacture its technology, not to the industries that use it. One assigns each technology class to sectors in proportion to the patents filed in that class by the sectors' firms (Neuhäusler et al., 2019); the other assigns each technology field to the sector that manufactures it (Schmoch, 2008). A patent on a textile machine therefore counts toward machinery and reaches the textile industry only through the textile industry's purchases from machinery. An industry's stock measures the technology developed in that industry rather than the technology the industry installs. A process patent claims a way of making or doing something, and a product patent claims a thing. The premise behind the exposures is that a supplier's technology reaches its customers inside the goods and services the supplier sells to them. Leave-Denmark-out versions of the stocks, built from patenting outside Denmark, provide the foreign shifter used in Section 4.1.1.

The automation index. A second displacing measure, and the one whose exposure has a natural scale, is the task-displacement component of the Acemoglu and Restrepo (2022) decomposition, constructed from the EU KLEMS industry accounts for Denmark (Stehrer et al., 2019). It measures the part of the change in an industry's labor share that is attributable to machines taking over tasks that workers previously performed, and it varies at the industry-year level. I refer to it as the automation index. The index is a ten-year trailing average of the annual task-content changes. Because the source data begin in 1995, its first value, for 2004, measures displacement over 1995–2004, and for 1995–2003 the index is set to

zero for every industry; the reinstatement index is constructed in the same way. Estimates based on the index therefore rely on a comparison of the years from 2004 onward with the years before, and the paper uses the index as a blunt check rather than as a main measure (Appendix B). The index records how much displacement an industry experienced, not which technologies produced it, so it does not separate the two types of technology in the way the patent stocks do. Its reinstating counterpart is the task-reinstatement component of the same decomposition, which captures the creation of new tasks in which labor has a comparative advantage. The ICT capital services share, also from EU KLEMS, provides a further time-varying measure, one that is directly observed rather than imputed. Occupation-based measures complement these: the automation probabilities of Frey and Osborne (2017), routine task intensity following Autor and Dorn (2013), the share of employment in occupations with emerging tasks from O*NET, and AI exposure scores (Webb, 2020; Felten et al., 2021; Hampole et al., 2025; Liu et al., 2025). These enter as robustness measures rather than as primary ones. Occupation-level scores are crosswalked from SOC through ISCO-08 to DISCO-08 and aggregated to industries using employment weights (Appendix A.1).

Observed adoption. The third family measures adoption directly. Indicators of firm-level use of industrial robots, AI, ERP, and cloud computing from the ITAV survey are aggregated to industries. These measures are used in the comparison of observed with imputed technology in Section 5.1 and Appendix B. The robot items begin around 2018 and the AI items in 2020, so the AI measure covers only a few survey years; I report it but place little weight on it.

Measurement windows. Each measure is available for the years covered by its source data, so the effective window differs across measures. This heterogeneity is by design. The EU KLEMS-based automation and innovation measures extend to 2021 in the source release, but the estimates in this paper use their coverage only through 2019, which keeps the pandemic years out of the identifying variation; the main worker regressions are therefore identified from variation over 1995–2019. The patent-based measures end in 2019 by construction: the underlying automation classification ends in 2014 and the breakthrough-patent universe in 2016, and because breakthrough status is defined by forward citations, recent patents cannot yet be classified. After the last year of patent flows the stocks only depreciate, so the variation in the final years of the sample comes from earlier patenting. The observed ITAV measures extend to 2023, so the paper's evidence for 2020–2023 comes from the observed adoption measures, while the estimates based on the imputed and patent-based measures are identified from the longer history.

3.4 The production network and exposure measures

The production network is taken from a hybrid input–output table. The domestic block comes from the Danish national supply–use tables, which distinguish 117 industries and are considerably finer than the OECD Inter-Country Input-Output tables; the international block combines the OECD ICIO tables (OECD, 2021) with bilateral trade data. The result is a table with 121 industries. This IO-121 classification also serves as the primary industry classification of the worker analysis, for an econometric reason: at the two-digit NACE level several IO industries collapse into a single code, and the direct, suppliers', and customers' exposures are correlated at 0.92–0.97, too collinear to identify separately. At the IO-121 level the correlations fall to workable levels. Appendix A.3 describes the construction.

The exposures implement the objects defined in Section 2.1. The input-coefficient and output-coefficient matrices are averaged over 2000–2004, a window that predates the acceleration of automation adoption in the sample, and are held fixed throughout. The reason is that the network is itself an outcome: if technological change rewires supply chains, weights taken from year t would embed the very endogenous response the design seeks to identify. Letting $\bar{L} = (I - \bar{A})^{-1}$ denote the predetermined Leontief inverse (Leontief, 1936), suppliers' exposure is

$$us_{jt} = \sum_i \tilde{\ell}_{ij}\, T_{it}, \qquad \tilde{L} \equiv \bar{L} - I, \tag{6}$$

and customers' exposure is constructed symmetrically from the predetermined domestic Ghosh inverse (Ghosh, 1958) under the same convention of subtracting the identity. Suppliers' exposure uses the total-requirements matrix, which includes imported intermediates, while customers' exposure uses the domestic matrix, since import rows have no Danish gross output. Subtracting the identity removes the accounting term that would otherwise impose a near-unit correlation between direct and suppliers' exposure, while retaining genuine roundabout linkages through the own industry; Appendix A.4 quantifies the difference. The roundabout term is not negligible. It averages 0.068 across base industries, so that an industry's purchases from its own industry account for about 7 percent of mean exposure, and the industry's own stock therefore enters through these purchases as well as through the separate own-industry regressor. Section 4.1.2 and Appendix A.5 report the estimates with this term removed. Two variants replace the predetermined weights, one using contemporaneous year-t tables and the other restricting attention to linkages of network distance two or more; both are reported alongside the main estimates (Table 9 for the patent stocks and Appendix B.6 for the automation index). A further split divides each suppliers' exposure according to the kind of supplier from which it originates, and a third weighting replaces the buyer's intermediate purchases with its equipment investment (Section 5.8).

The example from the introduction illustrates what the construction produces. Danish repair shops (repair and maintenance of motor vehicles, one of the 121 industries) buy about a fifth of their purchased inputs from motor-vehicle manufacturers on the 2000–2004 weights and about a third from car dealers and wholesalers, purchases that largely consist of trade margins (car dealers hold almost no patents, while wholesale trade holds some under the concordance). The remainder comes from tire makers, engine and pump makers, metal products, and services. Denmark builds no cars, so the motor-vehicle manufacturers are located abroad, and the exposure uses the total-requirements matrix precisely so that what they sell is counted. Each supplier industry contributes its log patent stock, by type, weighted by its share of the buyer's purchases. In 2005 the motor-vehicle industry held about 900 process-dominant and 1,800 product-dominant breakthrough patents, while the repair shops held almost none. On direct purchases alone, the repair shops rank tenth among the 117 industries with purchases in terms of their suppliers' displacing exposure, at roughly twice the value of the median industry. The Leontief inverse adds the suppliers' suppliers, and on the full weights of the regressor their exposure is about twice as large, 5.79 for displacing and 6.72 for reinstating technology in 2005. This places them in the top third of the 110 base industries in the main sample, though not in the top ten, which consists entirely of manufacturing industries whose suppliers' suppliers are also manufacturers. The units follow from the construction: a coefficient is expressed per log point of the suppliers' stock, and the magnitude statements in Section 5 use the standard deviation of the exposure net of the fixed effects, because the two stocks of a pair move together and never move independently by their total spread.

Table 1: Who is exposed: base industries by third of suppliers' displacing exposure, 2005

Third	Exposure (cut)	Industries	Workers (%)	Manufacturing (%)	White-collar (%)	Log wage (mean)	Premium (mean)	Equipment weight (%)
Bottom third	below 2.28	37	60	0.7	61	5.18	0.080	5
Middle third	2.28 to 4.58	36	19	4.6	46	5.17	0.095	6
Top third	above 4.58	37	21	71.4	38	5.18	0.101	11

Notes: The 110 base industries of the main wage sample in thirds of the suppliers' displacing exposure (process–product classification, $\bar{L} - I$ weights) in 2005; the cuts are in log patent units. Workers: share of the 2005 workers of the sample whose first industry is in the third. Manufacturing: share of those workers whose first industry is in manufacturing (NACE 10–33). White-collar, log hourly wage and employer pay premium: worker-weighted means in 2005. Equipment weight: the share of the third's supplier weights that falls on equipment makers (NACE 26–30), worker-weighted. The ten most exposed industries are ships and other transport equipment, computers and electronics, dairy, motor vehicles, engines, windmills and pumps, basic metals, electric motors, other food products, meat processing, paints and soap (exposure 8.60 to 6.83); the ten least exposed are employment activities, road freight, broadcasting, suburban trains and buses, social work, theatres and arts, water transport, oil and gas extraction, household employers, travel agents.

Table 1 shows which workers are exposed. Dividing base industries into thirds by their suppliers' displacing exposure in 2005, the most exposed third holds 21 percent of the main sample's 2005 workers, of whom 71 percent are in manufacturing, and has a white-collar share of 38 percent. The least exposed third holds 60 percent of workers, with 0.7 percent in

manufacturing and a white-collar share of 61 percent. Mean log hourly wages are the same in every third (5.180 compared with 5.176), and the employer pay premium is somewhat higher in the most exposed third (0.101 compared with 0.080). Exposure thus sorts workers by sector rather than by pay level, which anticipates the sector-level scope statement of Section 5.5.

3.5 Summary statistics

Table 2 reports summary statistics for the index-based and observed exposure measures, together with the realized window of each measure. Suppliers' automation exposure has a mean of 0.0017 and a standard deviation of 0.0028, so that a one standard deviation increase corresponds to about 1.6 times the mean level. Direct exposure is more dispersed in absolute terms (its standard deviation is 0.0039) but is defined for fewer workers, because not every base industry has its own EU KLEMS measure, whereas nearly all are connected to suppliers with measured technology. Customers' exposure is roughly half as dispersed as suppliers' exposure. Each regression uses the years covered by its measures, and the table reports the window for each. The patent exposures have no natural scale, so their magnitudes are expressed per standard deviation of the variation the regressions use, net of the fixed effects (Section 5.4).

4 Empirical Design

The estimating equations relate a worker's wage to the technology of the industries from which her own industry buys, comparing each worker with herself over time. For a single technology measure, the specification is

$$y_{it} = \alpha_i + \delta_t + \beta_{di}\, di_{j(i),t} + \beta_{us}\, us_{j(i),t} + \beta_{ds}\, ds_{j(i),t} + X'_{it}\gamma + \varepsilon_{it}, \tag{7}$$

where y_{it} is the log hourly wage (or, in other specifications, log annual earnings, an employment indicator, an employer-change indicator, or log hours), $j(i)$ is worker i's base industry, defined as the first industry in which she is observed, and the exposures are those of Section 3.4. The parameter of interest is β_{us}, the response of a worker's wage to the technology of her industry's suppliers. Because the index does not distinguish between the two types of technology, its sign provides only a blunt check on the labor response derived in Section 2.2: the coefficient should be negative if the displacing component dominates what the index records. The worker fixed effects α_i absorb all time-invariant worker characteristics, including ability, preferences, and occupation-specific human capital, and the year fixed effects δ_t

Table 2: Summary statistics: technology exposure measures

	Mean	SD	Min	Max	Window
Direct automation	0.0013	0.0039	0.000	0.142	1995–2021
Suppliers' automation	0.0017	0.0028	0.000	0.101	1980–2023
Customers' automation	0.0008	0.0015	0.000	0.051	1980–2023
Direct innovation	0.0041	0.0076	0.000	0.350	1995–2021
Suppliers' innovation	0.0028	0.0044	0.000	0.151	1980–2023
Customers' innovation	0.0015	0.0029	0.000	0.197	1980–2023
Direct ICT capital	0.0376	0.0298	0.00099	0.235	1980–2019
Suppliers' ICT	0.0261	0.0154	0.000	0.154	1980–2023
Customers' ICT	0.0167	0.0157	0.000	0.163	1980–2023
Direct robot use (observed)	0.0887	0.1997	0.000	1.000	2001–2023
Suppliers' robot use	0.0626	0.1728	0.000	1.735	1980–2023
Customers' robot use	0.0335	0.1105	0.000	0.941	1980–2023

Notes: Worker-year panel moments of the exposure measures (computed on the assessment sample of the merged panel). Direct exposure is the worker's base industry's own technology measure; suppliers' (customers') exposure propagates the same measure through the off-diagonal Leontief (Ghosh) inverse with predetermined 2000–2004 weights. Automation, innovation, and ICT are imputed from EU KLEMS; robot use is observed firm adoption from the ITAV survey. Window is the span of years with observed values; estimation samples vary by measure accordingly.

absorb aggregate trends, including the aggregate wage adjustment implied by the framework. The control vector X_{it} includes firm tenure in the richest specification; education is absorbed by the worker fixed effects for the vast majority of workers, and age is collinear with the worker and year fixed effects. Standard errors are clustered at the base industry, the level at which treatment varies, which yields 121 clusters.

The main specification distinguishes technology by type. The framework of Section 2.2 implies that suppliers' displacing and reinstating technologies affect the buying industry's workers with opposite signs. To test this implication, I estimate

$$y_{it} = \alpha_i + \delta_t + \sum_{c \in \{di,\, us,\, ds\}} \left(\beta_c^D \, c^D_{j(i),t} + \beta_c^R \, c^R_{j(i),t} \right) + \varepsilon_{it}, \tag{8}$$

where c^D and c^R denote the exposures constructed from the displacing and the reinstating patent stock of a given classification. The two suppliers' coefficients, β^D_{us} and β^R_{us}, are the parameters of interest, and the framework implies that they have opposite signs. The paper's main results are the estimates of these two coefficients; equation (7) with the automation index serves as the blunt check reported in Appendix B. The remainder of the analysis uses three variants of equation (8). The first replaces the outcome with the pay of the worker's

employer. Writing $\psi_{J(i,t)}$ for the pay premium of worker i's employer in year t, the log wage can be written as the sum of the premium and a remainder, $y_{it} = \psi_{J(i,t)} + (y_{it} - \psi_{J(i,t)})$, and estimating the equation separately for each component decomposes the wage coefficient exactly into the part that reflects what the worker's employer pays all of its workers and the part that is specific to the worker. Because the premium is fixed for each employer, its coefficient measures movement between employers, up or down the pay ladder. A related outcome is the pay of the worker's coworkers in a given year, defined as the mean wage residual, net of worker and year effects, of the other workers at her employer. This outcome moves both when the worker changes employer and when the employer changes what it pays, and a firm-level regression of the same component on the firm's own exposure distinguishes the two. The second variant follows workers beyond the main sample: employment and annual earnings are estimated on every worker with a base industry, whether or not she is employed in a given year, and the destinations of job changers are estimated on the worker-years in which an employer change occurs. The third variant is estimated on the buying firms. The same exposures, assigned according to the firm's industry, enter firm-level regressions of employment, hiring and separation rates, the age and occupational composition of the workforce, the investment rate, wage costs, value added, and the labor share, with year effects and either firm or industry effects. These regressions test the implications of Section 2.2 for the buying firms' accounts. The worker estimates for the automation index are built up progressively, beginning with direct exposure alone and successively adding suppliers' exposure, customers' exposure, the three reinstatement exposures, and tenure, with a parallel set of specifications for the ICT measure. Where the text compares groups of workers, each group is estimated in a separate regression with its own year effects, because a pooled regression with group interactions and common year effects would attribute differences in the groups' wage trends to the exposure. The within-worker estimator addresses the first-order concern about selection: if high-ability workers sorted into exposed industries, a cross-sectional comparison would be contaminated, whereas the equations above compare the same worker across years in which the exposure of a fixed base industry is high or low. For the wage, the worker effects add precision rather than identification. Replacing them with base-industry and year effects yields the same estimate, so the comparison is effectively one of the same base industry over time; with year effects alone the signs reverse, because the workers of exposed industries earn more on account of their composition, a difference that controls for 2005 age, sex, and education remove (Section 5.3, Appendix Table 42). The employment result, in contrast, does require the within-worker comparison, since the pool of workers in a base industry changes over time. The within-worker comparison cannot, however, generate variation that the exposures do not contain. The patent stocks move

slowly within an industry, and Section 5.5 examines what this implies for the scope of the estimates.

4.1 Identification

The predetermined network weights and the first-observed industry assignment define the estimand and are not intended as approximations to a contemporaneous exposure. This estimand combines two conceptually distinct conditions. First, it holds each buyer's input mix at its pre-period values, so that $us_{j(i),t}$ moves only when the technology of the industries from which j bought in the base period moves. Second, it holds each worker at the industry in which she is first observed. The coefficient β_{us} is therefore an intention-to-treat effect: the response of incumbent workers' pay to supplier-side technology shifts, transmitted through the purchase structure that existed before those shifts arrived. By construction, the buyer's subsequent decisions (to change suppliers, to bring tasks in-house or contract them out, or to change industry) enter neither the weights nor the assignment. The endogenous re-sorting that a contemporaneous exposure would embed is thus excluded from the regressor rather than imperfectly controlled for. This is the object to which the comparative statics of Section 2.2 refer, namely the effect of the change in suppliers' technology on the workers who were in place when it arrived. The effect of *realized* exposure, meaning the pay response to the input mix a buyer actually ends up with, is a distinct object. The two can be compared for the automation index, on which the estimates in the remainder of this paragraph are based. For the network weights, the realized effect can be recovered by instrumenting realized, contemporaneous-weight exposure with the predetermined-network exposure. This yields -2.95 (s.e. 1.22, first-stage $F = 388$), close to the intention-to-treat estimate, which indicates that buyers' re-sourcing makes little difference. The patent estimates point in the same direction, with contemporaneous-weight coefficients close to the predetermined ones (Table 9). Realized assignment to the worker's current industry is a different matter. The recorded industry frequently changes without any change of job: in 6.2 percent of consecutive years in which a worker remains with the same employer, and in 23 percent across the 2008 classification change. As a result, 77 percent of stayers' worker-years carry an industry code other than the first-observed one, and exposure built on the current code inherits this noise. Least squares on the current-industry exposure gives -0.60 (s.e. 0.24); instrumenting it with the base-industry exposure gives -10.2 (s.e. 5.94, first-stage $F = 34.6$), about three times the intention-to-treat estimate but imprecise. The first-observed assignment is therefore the more stable of the two, and the intention-to-treat coefficient is the more conservative estimate.

The credibility of the design rests on two distinct elements. The first is a claim about the *source* of the identifying variation: where it comes from and why that source is unrelated to Danish wages. The second is defensibility, in the sense of surviving the standard battery of pre-trend, placebo, and concentration tests, which are covered in Section 5.5; a longer battery, run on the automation index, is reported in Appendix B.

4.1.1 A frontier-taking small open economy

I assume that Denmark takes the world technology frontier as given (Section 3.1), and take this as the source of the identifying variation. Global waves of invention and adoption cause Danish adoption but are not caused by conditions in Danish industries. The leave-Denmark-out patent shifters use this variation directly. The foreign adoption measures use it through other countries' realized choices, and the primary domestic exposures use it indirectly, because domestic adoption follows the timing and pricing of the frontier even though each industry chooses its own moment to move. The identification problem thus separates into a source claim, which is defended here, and a transmission claim, which is embodied in the predetermined network weights of Section 3.4.

The design that uses this source most directly regresses worker outcomes on the foreign wave itself. Let F_{kt} denote leave-Denmark-out foreign displacement patenting allocated to industry k's technology classes. Propagating it through the same predetermined weights gives

$$\text{ForeignUpstream}_{jt} = \sum_{k} \tilde{\ell}_{kj} \, F_{kt}, \tag{9}$$

the world displacement wave bearing on the industries that supply industry j, which enters the worker equation in place of the domestic suppliers' exposure. Since Denmark does not cause the wave, the source claim is as strong as the setting allows. The cost is a composite estimand. Greater foreign displacement innovation reaches Danish buyers through what their suppliers sell, but it also intensifies import competition in the same product classes. The reduced form is therefore presented as an incidence statement, the effect of the world technology wave on workers as transmitted through their industry's position in the network, rather than as a structural decomposition. The shift-share design provides the decomposition instead. The two designs are complementary, in that the first has the cleaner source of variation and the second the cleaner estimand.

4.1.2 Design-based shift-share

The shifters differ in how closely they track the frontier. The patent stocks are constructed from the universe of breakthrough patents allocated to industries by technology class. Danish

inventors account for a negligible share of that universe, so for the estimates by technology type the shifter is close to the frontier itself. The automation index, by contrast, is based on the task displacement realized in Danish industries, which is driven by the frontier but timed domestically. Three features of the design isolate the defensible part of the variation in both cases. First, the shares are elements of the predetermined Leontief inverse and are fixed before any outcome in the sample years can affect them. Second, the weights are off-diagonal: the shifters belong to industries other than the worker's own, and the remaining own-industry contribution runs only through genuine roundabout production chains, while the linear own-industry term is absorbed by $di_{j(i),t}$. This roundabout contribution, an industry's purchases from its own industry, averages 0.068 (at most 0.366) across base industries and accounts for about 7 percent of mean exposure. Removing it from the regressor leaves the firm results and the workers' employment and earnings results unchanged, and leaves the pay results with the same signs at about three-quarters of their size (Appendix A.5); part of the pay effect therefore runs through industries' purchases from their own industry. Third, any threat to the suppliers' coefficient requires outcomes in industry j to influence technological change in the industries that supply j, in a manner coordinated across many suppliers and weighted by their predetermined linkage strength. This is the shift-share design of Borusyak et al. (2022), in which as-good-as-random shifters are the identifying condition; the complementary route based on share exogeneity, due to Goldsmith-Pinkham et al. (2020), is assessed through the pre-trend tests. The concentration of the identifying variation in a few large supplier industries is a feature of the setting, and the paper reports the shock-level diagnostics rather than claiming many independent shocks (Appendix B.4, for the automation index).

Base-industry assignment completes the design. Fixing each worker at her first observed industry follows Hummels et al. (2014) and Autor et al. (2013): it is an intention-to-treat assignment under which later sorting attenuates rather than contaminates the estimates. Using the first observation, rather than a fixed calendar year such as 2005, avoids assigning pre-2005 observations to an industry the worker selected only later, possibly in response to the very network shocks under study, since the first-observed industry predates all of the worker's subsequent outcomes by construction.

4.2 Industry- and firm-level specifications

The industry-level counterpart aggregates outcomes to the industry-year level:

$$y_{jt} = \alpha_j + \delta_t + \beta_{di}\, di_{jt} + \beta_{us}\, us_{jt} + \beta_{ds}\, ds_{jt} + \varepsilon_{jt}, \tag{10}$$

where y_{jt} is log employment, the log average hourly wage, the labor share, or a measure of task composition, with industry and year fixed effects, the same progressive build-up of regressors, and standard errors clustered by industry. The firm-level counterpart is

$$y_{ft} = \alpha_f + \delta_t + \sum_{c \in \{di,\, us,\, ds\}} \left(\beta_c^D \, c_{j(f),t}^D + \beta_c^R \, c_{j(f),t}^R \right) + \varepsilon_{ft}, \tag{11}$$

where y_{ft} is the firm's log employment, log wage costs, log value added, operating result over value added, labor share, investment rate, or fixed assets per full-time equivalent, all measured in the firm register, or its hiring rate, separation rate, and the mean age and occupational composition of its workforce, constructed from the matched worker data; $j(f)$ is the firm's industry, and standard errors are clustered by industry. The firm effects α_f compare a buying firm with itself as its suppliers' technology changes. Replacing them with industry effects follows the firms of an industry over time while allowing the set of firms to change, so that entry and exit are counted. The same equation with a single exposure triple yields the firm-level estimates for the automation index and for observed robot adoption. These firm-level regressions test the implications of Section 2.2 for the buying firms' accounts. Long-difference variants over five- and ten-year windows complement the annual estimates, since the two designs weight short-run and cumulative variation differently. Heterogeneity by network position is examined by interacting the exposures with predetermined centrality, restricted to betweenness and eigenvector centrality: upstreamness and downstreamness are nearly collinear with the exposures themselves (variance-inflation factors of 24–51 and 22–78, compared with 2–4 for betweenness and 4–8 for eigenvector centrality), while Katz centrality lies in between (8–15) and is borderline.

5 Reduced-Form Results

This section presents the reduced-form evidence, first for the buying firm and then for its workers. Section 5.1 shows that firms whose suppliers develop displacing technologies see reductions in employment and in their labor share, while value added per worker and operating margins rise and the pay of incumbent workers is unchanged; under reinstating technologies each of these effects is reversed. Section 5.2 compares the technology of an industry's suppliers with the industry's own technology and finds that the former matters at least as much for its workers, with the opposite sign. Section 5.3 establishes that the worker-level effects attach to the worker rather than to the employer: a worker's pay responds to the technology of the suppliers of the industry in which she started, whether or not she still works there, and not to that of her current employer's suppliers. Given that design, Section 5.4

reports the main wage estimates, under which suppliers' displacing technology lowers pay and suppliers' reinstating technology raises it, and Section 5.5 shows that these estimates are stable within municipality and year, show no anticipatory response, appear in each half of the sample, and are identified mainly from differences between broad sectors. Section 5.6 traces the accompanying changes in employers, employment and exits. Section 5.7 shows that the losses are concentrated among higher-paid white-collar workers from exposed origin industries, employed at firms where pay is set at the firm level, whereas lower-paid workers adjust through hours and employment. Section 5.8 shows that both types of technology reach buyers through the manufactured inputs they purchase, from equipment makers and other manufacturers alike, rather than through the equipment they invest in. An index of suppliers' automation that does not distinguish the two types is used as a blunt check and is reported in Appendix B.

5.1 Buying Firms

Buying firms contract when their suppliers' technology shifts toward the displacing type and expand when it shifts toward the reinstating type, as the model of Section 2.2 implies for the buyer's scale and its employment. Table 3 reports estimates of equation (11) for employment, value added, wage costs, and the operating result as a share of value added, all taken from the firm register. With firm fixed effects and the automation–other classification, the coefficient on the displacing stock in the employment regression is −0.083 (s.e. 0.045) and the coefficient on the reinstating stock is 0.104 (s.e. 0.044) under predetermined weights; under contemporaneous weights the corresponding estimates are −0.103 (s.e. 0.032) and 0.121 (s.e. 0.032). The effects are larger still when firm effects are replaced by industry effects, which allow entry and exit to contribute. Under the process–product classification and predetermined weights, value added changes by −0.657 (s.e. 0.142) and 0.761 (s.e. 0.144), wage costs by −0.666 (s.e. 0.148) and 0.749 (s.e. 0.186), and employment by −0.397 (s.e. 0.130) and 0.535 (s.e. 0.150). In terms of the example of Section 3.4, as the car makers' process patents accumulate, Danish repair shops contract and their wage bill falls relative to their value added; as the product patents accumulate, the opposite occurs.

The distribution of value added moves together with firm size. The operating result as a share of value added rises under displacing technology and falls under reinstating technology: with firm effects and the process–product classification the coefficients are 0.056 (s.e. 0.028) and −0.058 (s.e. 0.027), and with industry effects they are 0.195 (s.e. 0.069) and −0.178 (s.e. 0.070). The share of value added not paid out as wages moves in the same way (Appendix Table 38). Since a higher operating-result share corresponds to a lower labor share, suppliers'

Table 3: Suppliers' displacing and reinstating technology: buying firms

	Log employment		Log value added		Log wage costs		Operating result / VA	
Network weights:	Predet.	Contemp.	Predet.	Contemp.	Predet.	Contemp.	Predet.	Contemp.
Panel A: firm and year fixed effects								
Process/product, displacing	-0.119	-0.094	0.019	0.086	-0.084	-0.011	0.056*	0.059*
	(0.070)	(0.074)	(0.144)	(0.144)	(0.118)	(0.116)	(0.028)	(0.026)
Process/product, reinstating	0.133*	0.106	0.008	-0.069	0.116	0.031	-0.058*	-0.059*
	(0.064)	(0.069)	(0.133)	(0.134)	(0.110)	(0.108)	(0.027)	(0.025)
Automation/other, displacing	-0.083	-0.103**	-0.043	-0.051	-0.039	-0.044	0.006	0.008
	(0.045)	(0.032)	(0.033)	(0.035)	(0.026)	(0.031)	(0.007)	(0.009)
Automation/other, reinstating	0.104*	0.121***	0.078*	0.071	0.082**	0.071*	-0.014	-0.014
	(0.044)	(0.032)	(0.034)	(0.038)	(0.029)	(0.035)	(0.008)	(0.009)
Panel B: industry and year fixed effects								
Process/product, displacing	-0.397**	-0.261*	-0.657***	-0.437**	-0.666***	-0.439***	0.195**	0.189**
	(0.130)	(0.112)	(0.142)	(0.164)	(0.148)	(0.126)	(0.069)	(0.060)
Process/product, reinstating	0.535***	0.262*	0.761***	0.391*	0.749***	0.420***	-0.178*	-0.179**
	(0.150)	(0.102)	(0.144)	(0.152)	(0.186)	(0.114)	(0.070)	(0.057)
Automation/other, displacing	-0.056	-0.105	-0.119	-0.122	-0.180	-0.070	0.061	0.036
	(0.067)	(0.060)	(0.136)	(0.068)	(0.116)	(0.078)	(0.038)	(0.022)
Automation/other, reinstating	0.157	0.151*	0.233**	0.134	0.207*	0.099	-0.081**	-0.057*
	(0.103)	(0.070)	(0.086)	(0.073)	(0.098)	(0.091)	(0.026)	(0.027)

Notes: Firm-level regressions of buying firms' outcomes on the suppliers' patent stocks of both types (own-industry and customers' stocks included), one regression per column and patent classification, through 2019. Employment is the number of employees; value added, wage costs and the operating result are amounts from the firm register; the operating result over value added is winsorized at the 1st and 99th percentiles. A higher operating-result share of value added is a lower labor share. Coefficients are per unit of log patent stock. Standard errors clustered by IO-121 industry; significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

displacing technology reduces the labor share of the firms that buy from them and suppliers' reinstating technology raises it. Not every estimate in Table 3 is individually significant, and the two classifications are precise in different places: the automation–other pair for employment with firm effects, and the process–product pair for the labor share and for firm size with industry effects. Taken together, however, the estimates in Table 3 and Figure 2 point consistently in one direction. The size and the labor share of buying firms decrease with suppliers' displacing technology and increase with suppliers' reinstating technology, and the sign of the employment coefficient, as well as that of the operating-result share, is the same in all eight specifications for each type.

Buying firms retain the resulting cost savings and do not reduce the pay of their incumbent workers. Table 4 examines where value added goes. Under suppliers' displacing technology, employment falls by -0.119 (s.e. 0.070) while value added does not (0.019, s.e. 0.144). Accordingly, value added per worker rises by 0.201 (s.e. 0.092), wage costs per worker by 0.102 (s.e. 0.058), and the operating margin by 0.047 (s.e. 0.017), while the labor share falls by -0.055 (s.e. 0.030). Reinstating technologies reverse each of these responses. The pattern differs from what a decline in demand would produce and resembles that of a firm that purchases a larger part of its production, employs fewer people, and preserves its margin.

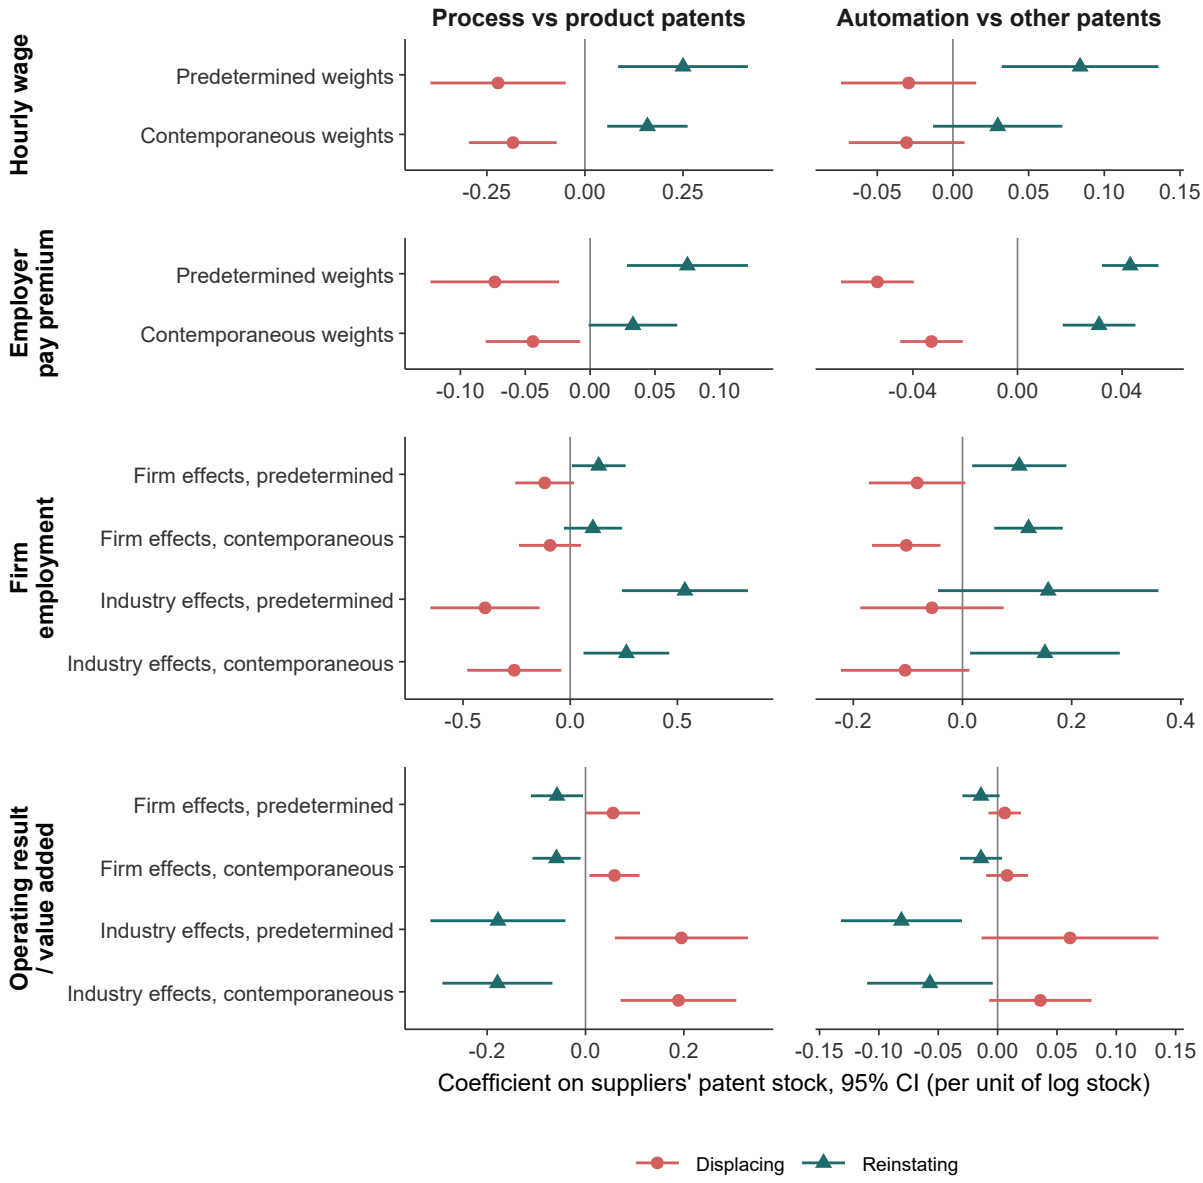


Figure 2: Suppliers' displacing and reinstating technology have opposite effects on workers and on buying firms. Each marker is the coefficient on the suppliers' displacing (circle) or reinstating (triangle) patent stock, estimated jointly, with its 95 percent confidence interval. Rows one and two are worker-level regressions with worker and year fixed effects on the main sample (Table 9); rows three and four are firm-level regressions on the firm register (Table 3). Columns correspond to the two patent classifications. Within each panel the specifications differ by network weights (predetermined or contemporaneous) and, for firms, by firm or industry fixed effects. A higher operating result over value added corresponds to a lower labor share. Coefficients are per unit of log patent stock; standard errors are clustered by IO-121 industry.

Table 4: Inside the buying firm: value added, margins and the pay of incumbents and new hires

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: firm outcomes, all firms in buying industries (firm and year effects)				
Log employment	-0.119	0.133*	-0.083	0.104*
	(0.070)	(0.064)	(0.045)	(0.044)
Log value added	0.019	0.008	-0.043	0.078*
	(0.144)	(0.133)	(0.033)	(0.034)
Log value added per worker	0.201*	-0.202*	0.068*	-0.085**
	(0.092)	(0.087)	(0.030)	(0.031)
Log wage costs per worker	0.102	-0.099	0.068**	-0.075***
	(0.058)	(0.055)	(0.022)	(0.022)
Operating margin	0.047**	-0.046**	0.002	-0.004
	(0.017)	(0.017)	(0.004)	(0.005)
Labor share	-0.055	0.057*	-0.008	0.015
	(0.030)	(0.029)	(0.007)	(0.008)
Panel B: the firm's pay component, worker-year pay net of worker effects (firm and year effects)				
All workers	-0.019	0.015	-0.031***	0.026**
	(0.015)	(0.014)	(0.009)	(0.008)
Incumbents	-0.023	0.018	-0.030**	0.023**
	(0.016)	(0.014)	(0.009)	(0.008)
New hires	-0.010	0.008	-0.032*	0.029*
	(0.020)	(0.018)	(0.014)	(0.013)
Incumbents minus new hires	-0.022	0.018	0.010	-0.014
	(0.014)	(0.012)	(0.009)	(0.008)

Notes: Each row is one pair of firm-level regressions (process–product and automation–other classification) with firm and year fixed effects on all firms in buying industries, weighted by employment; the coefficients are those on the suppliers' displacing and reinstating patent stocks. Panel A: log employment, log value added, log value added per full-time equivalent, log wage costs per full-time equivalent, the operating margin (operating result over turnover) and the labor share (wage costs over value added). Panel B: the firm-year mean of its workers' pay net of worker fixed effects (the firm's pay component), for all workers, for incumbents (employed at the firm the previous year), for new hires, and their difference. Firm-years in the process–product regression, in millions, by panel: Panel A: 5.0 (row 1), 1.6 (row 2), 1.5 (row 3), 1.4 (row 4), 1.7 (row 5), 1.6 (row 6); Panel B: 4.8 (row 1), 2.5 (row 2), 1.9 (row 3), 1.2 (row 4). Standard errors clustered by IO-121 industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Panel B asks whether the firm pays its workers less. Under the main classification it does not: the firm's pay component, defined as the mean pay of its workers net of worker fixed effects, changes by −0.019 (s.e. 0.015) for all workers, by −0.023 (s.e. 0.016) for incumbents, and by −0.010 (s.e. 0.020) for new hires, none of which is distinguishable from zero. Under the automation–other pair the pay components do fall, for incumbents and new hires alike, so that for this measure the two sources cannot be separated. The same conclusions hold both at high-premium firms and at the remaining firms, and neither entry nor concentration in the buying industry changes.

Table 5 shows how the contraction takes place. In the automation–other classification the hiring rate falls with the displacing stock and rises with the reinstating stock (−0.033, s.e. 0.013, and 0.026, s.e. 0.012); in the process–product classification the hiring coefficients have the same signs but are imprecisely estimated (−0.011, s.e. 0.029, and 0.008, s.e. 0.026). The firm-level separation rate does not change detectably in either classification (0.024, s.e. 0.026, and −0.021, s.e. 0.023, for the process–product pair). This estimate is imprecise, however, and Section 5.6 shows that exits from employment do rise at the worker level, so the two findings are not in conflict. Consistent with reduced hiring, the workforce ages under displacing technology and becomes younger under reinstating technology, by 1.34 (s.e. 0.67) and −1.31 (s.e. 0.63) years of mean age. The share of the firm's workers in routine occupations moves in the same direction (0.107, s.e. 0.047, and −0.086, s.e. 0.041), and the investment rate falls under displacing technology and rises under reinstating technology (−0.053, s.e. 0.022, and 0.052, s.e. 0.021). The labor share of value added, measured directly as wage costs over value added, behaves as the operating-result share implies: the coefficients are −0.055 (s.e. 0.030) and 0.057 (s.e. 0.029) with firm effects, at the margin of statistical significance, and −0.153 (s.e. 0.061) and 0.145 (s.e. 0.065) with industry effects. The automation–other pair yields the same signs for age, the routine share and the labor share, with smaller coefficients, and no response of the investment rate. Appendix Tables 39 and 40 report every outcome estimated, and Appendix Table 41 splits the worker flows by the firm's own pay premium; under the process–product pair the employment coefficients are of similar magnitude at low-, middle- and high-paying buying firms.

There is no evidence of capital deepening at the buyer. Fixed assets per full-time equivalent do not change significantly under displacing technology in either classification (0.092, s.e. 0.206, for the process–product pair and 0.070, s.e. 0.059, for the automation–other pair, both with firm effects), and the investment rate falls. Suppliers' displacing technology therefore does not appear as equipment on the balance sheets of the firms that buy from them; these firms invest less and retain an older workforce that performs more routine work. Under reinstating technology the pattern is reversed, and in the automation–other classification

Table 5: Inside the buying firm: hiring, workforce, investment and the labor share

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: firm and year fixed effects				
Log employment	-0.123	0.138*	-0.086	0.108*
	(0.071)	(0.066)	(0.046)	(0.045)
Hiring rate	-0.011	0.008	-0.033*	0.026*
	(0.029)	(0.026)	(0.013)	(0.012)
Separation rate	0.024	-0.021	-0.013	0.014
	(0.026)	(0.023)	(0.019)	(0.017)
Mean age of the workforce (years)	1.34*	-1.31*	0.63	-0.66*
	(0.67)	(0.63)	(0.34)	(0.32)
Share in routine occupations	0.107*	-0.086*	0.047	-0.030
	(0.047)	(0.041)	(0.029)	(0.023)
Investment rate	-0.053*	0.052*	-0.001	0.004
	(0.022)	(0.021)	(0.008)	(0.009)
Labor share of value added	-0.055	0.057	-0.008	0.015*
	(0.030)	(0.029)	(0.007)	(0.008)
Log fixed assets per FTE	0.092	-0.110	0.070	-0.106*
	(0.206)	(0.199)	(0.059)	(0.053)
Log turnover	-0.207	0.213	0.003	0.020
	(0.132)	(0.126)	(0.055)	(0.056)
Panel B: industry and year fixed effects				
Mean age of the workforce (years)	3.82*	-6.66**	-0.44	-1.64*
	(1.94)	(2.55)	(0.60)	(0.78)
Share in routine occupations	0.301*	-0.287	0.118**	-0.116*
	(0.129)	(0.181)	(0.037)	(0.054)
Investment rate	-0.139**	0.129**	-0.027	0.011
	(0.046)	(0.049)	(0.035)	(0.025)
Labor share of value added	-0.153*	0.145*	-0.057	0.080**
	(0.061)	(0.065)	(0.034)	(0.025)
Log fixed assets per FTE	-0.588	0.175	-0.259	-0.058
	(0.329)	(0.273)	(0.161)	(0.118)
Log turnover	-0.591	0.603	0.144	0.006
	(0.411)	(0.451)	(0.241)	(0.237)

Notes: Firm-level regressions of buying firms' outcomes on the suppliers' displacing and reinstating patent stocks of the firm's industry (own-industry and customers' stocks included), one regression per row and patent classification, through 2019. The hiring and separation rates are inflows and outflows of employees over employment. Mean age and the share in routine occupations (routine manual and routine cognitive) are computed over all workers of the firm. The investment rate is net investment over fixed assets and the labor share is wage costs over value added, from the firm register; rates and shares from the register are winsorized at the 1st and 99th percentiles. Firm-years in the process–product regression, in millions, by panel: Panel A: 4.5 (row 1), 4.4 (row 2), 4.5 (row 3), 5.0 (rows 4–5), 1.5 (row 6), 1.6 (row 7), 1.4 (row 8), 2.0 (row 9); Panel B: 5.0 (rows 1–2), 1.5 (row 3), 1.6 (row 4), 1.4 (row 5), 2.0 (row 6). Standard errors clustered by IO-121 industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

capital per worker falls as the firm adds workers (−0.106, s.e. 0.053).

These responses are the ones the model of Section 2.2 implies for firm size and the labor share. A buying industry whose suppliers turn toward process innovation loses employment, wage bill and value added, and its labor share of value added falls as the markup rises; when suppliers turn toward product innovation each of these responses is reversed. Size here refers to in-house size, that is, employment, wage costs and value added, which is the quantity the model describes. Turnover by technology type points in the same direction but is imprecisely estimated (−0.207, s.e. 0.132, and 0.213, s.e. 0.126). Taken together, these results indicate that firms adjust on the hiring margin and retain the rents. The question of who bears the adjustment therefore concerns the workers who are not hired, those who leave, and those whose skills were specific to the way the industry used to produce. The worker-level results address this question.

Observed robot adoption by suppliers, measured over a short and recent window, is followed by higher capital per worker at buying firms (0.266, s.e. 0.125) and somewhat higher wage costs (0.107, s.e. 0.056), with no significant change in value added or in the operating-result share. For the comparison with Bilgin et al. (2024), this implies that the Danish firm-level results do not contradict theirs when suppliers adopt robots: buying firms deepen capital, and their average wage does not fall. The two papers differ in the object studied, since a firm-level average does not follow the workers who leave or who are never hired.

5.2 Own against Suppliers' Technology

The technology of an industry's suppliers matters at least as much for its workers as the industry's own technology, and the two work in opposite directions. The industry's own patent stocks enter every regression in the paper. Table 6 places them beside the suppliers' and the customers' stocks from the same regression, estimated on the main sample of 47.3 million worker-years, and expresses all three in the one unit in which they can be compared: each coefficient multiplied by the residual standard deviation of its own regressor, net of the fixed effects and the other regressors. The own stock is the industry's log patent stock, whereas the suppliers' stock is a purchase-weighted sum of other industries' log stocks, so the two are not comparable per unit but are comparable per standard deviation. On that scale, an industry's own process patents are associated with hourly pay that is 0.66 percent higher and its own product patents with pay that is -0.42 percent lower, neither estimate being significant (per unit, 0.038, s.e. 0.024, and −0.028, s.e. 0.024). Its suppliers' process patents lower pay by -0.73 percent per standard deviation and their product patents raise

it by 0.74 percent (per unit, −0.223, s.e. 0.088, and 0.251, s.e. 0.085), both significant. Per log point, the suppliers' coefficient is 6 times the own coefficient; per standard deviation, the two are of roughly equal magnitude (1.1 times) and of opposite sign. For annual earnings and employment the own process stock does have a significant effect, of 2.01 percent and 0.41 percentage points per standard deviation, and there too its sign is opposite to that of the suppliers' term (-2.39 percent and -0.66 points). The suppliers' exposure is the only term whose two types have opposite signs for pay, earnings, employment and the employer pay premium under both classifications; the customers' stocks are significant only once, for pay under reinstating patents, and with the wrong sign for a demand channel. The firm panel gives the same result in the same units (Panels E and F): suppliers' process patents lower a buying firm's employment by -0.77 percent per standard deviation and product patents raise it by 0.94 percent, with the own stocks close to zero. In terms of the example, the car makers' own process patents are associated with higher pay and employment at the car makers themselves and with lower pay and employment at the repair shops.

When either channel is estimated alone it absorbs part of the other, and the omission biases both coefficients toward zero. Table 21 enters the own stocks alone, the suppliers' stocks alone, both together, and both together with the customers' stocks, on the same samples. The own and the suppliers' stock of a given type are positively correlated (0.21 across worker-years and 0.31 within worker and year on the main sample; 0.54 and 0.30 across industry-years), and their effects have opposite signs, so omitting one pulls the other's coefficient toward the sign of the omitted term. Estimated alone, the own process stock's pay coefficient is 0.032 (s.e. 0.028), compared with 0.042 when the suppliers' stocks are included, a reduction of 24 percent; for employment and earnings the own coefficients are 42 and 34 percent smaller when estimated alone and lose significance. The suppliers' coefficients estimated alone are between 7 and 38 percent smaller than when the own stocks are included, for the worker outcomes (for pay, −0.161, s.e. 0.086, against −0.214, s.e. 0.092), and unchanged in the firm panel, where the two stocks are less correlated within firm (0.16). The collinearity that matters is within each pair rather than across channels: the two own stocks correlate at 0.99 and the two suppliers' stocks at 1.00, and the other five terms explain 0.71 of the within-worker variance of the suppliers' process stock, almost entirely through the suppliers' product stock. Within manufacturing first industries and within the remaining industries, the suppliers' coefficients keep their signs but lose precision (Panels G and H), as the between-sector interpretation of the identification would suggest; the own process stock is negative within manufacturing and positive elsewhere. For adoption studies the implication is a bias toward the suppliers' sign: where suppliers' technology moves together with the industry's own and in the opposite direction, an own-industry coefficient estimated without

Table 6: Own-industry technology against suppliers' and customers'

	Coefficient per unit (s.e.)				Coefficient × SD, percent			
	Process vs product		Automation vs other		Process vs product		Automation vs other	
	Displ.	Reinst.	Displ.	Reinst.	Displ.	Reinst.	Displ.	Reinst.
Panel A: log hourly wage								
Own industry's stocks	0.038 (0.024)	-0.028 (0.024)	0.019 (0.021)	-0.007 (0.015)	0.66	-0.42	0.25	-0.14
Suppliers' stocks	-0.223* (0.088)	0.251** (0.085)	-0.029 (0.023)	0.084** (0.026)	-0.73	0.74	-0.23	0.81
Customers' stocks	0.046 (0.055)	-0.121** (0.046)	-0.040 (0.026)	-0.012 (0.025)	0.26	-0.52	-0.42	-0.13
Panel B: employed								
Own industry's stocks	0.019* (0.009)	-0.009 (0.013)	0.013 (0.008)	-0.001 (0.007)	0.41	-0.14	0.30	-0.03
Suppliers' stocks	-0.137** (0.043)	0.106* (0.050)	-0.052*** (0.012)	0.026 (0.018)	-0.66	0.41	-0.81	0.26
Customers' stocks	0.021 (0.032)	-0.028 (0.045)	0.003 (0.012)	-0.013 (0.019)	0.17	-0.17	0.06	-0.17
Panel C: log annual earnings								
Own industry's stocks	0.093* (0.038)	-0.073 (0.048)	0.051 (0.036)	-0.026 (0.027)	2.01	-1.14	1.14	-0.66
Suppliers' stocks	-0.493*** (0.137)	0.507** (0.167)	-0.154*** (0.044)	0.163** (0.058)	-2.39	1.97	-2.43	1.66
Customers' stocks	0.106 (0.091)	-0.179 (0.108)	-0.017 (0.047)	-0.044 (0.055)	0.84	-1.08	-0.33	-0.56
Panel D: employer pay premium								
Own industry's stocks	0.001 (0.006)	-0.011 (0.006)	0.001 (0.004)	-0.006* (0.003)	0.01	-0.17	0.01	-0.11
Suppliers' stocks	-0.074** (0.025)	0.075** (0.024)	-0.054*** (0.007)	0.043*** (0.005)	-0.24	0.22	-0.42	0.42
Customers' stocks	0.009 (0.021)	-0.024 (0.018)	-0.018** (0.006)	0.011 (0.007)	0.05	-0.10	-0.18	0.12
Panel E: buying firms' log employment (firm and year effects)								
Own industry's stocks	0.008 (0.021)	-0.008 (0.021)	0.020** (0.007)	-0.021** (0.007)	0.21	-0.19	0.85	-0.76
Suppliers' stocks	-0.119 (0.070)	0.133* (0.064)	-0.083 (0.045)	0.104* (0.044)	-0.77	0.94	-1.14	1.55
Customers' stocks	-0.021 (0.065)	0.006 (0.056)	0.029 (0.025)	-0.038 (0.023)	-0.23	0.07	0.52	-0.74
Panel F: buying firms' labor share (firm and year effects)								
Own industry's stocks	0.001 (0.009)	-0.001 (0.010)	0.009** (0.003)	-0.012*** (0.003)	0.04	-0.02	0.40	-0.44
Suppliers' stocks	-0.055 (0.030)	0.057* (0.029)	-0.008 (0.007)	0.015 (0.008)	-0.35	0.40	-0.10	0.23
Customers' stocks	0.017 (0.018)	-0.009 (0.016)	0.030** (0.011)	-0.026* (0.012)	0.18	-0.10	0.53	-0.52

Notes: Each panel reports, for each classification, one regression in which the industry's own displacing and reinstating patent stocks, its suppliers' stocks (purchase-weighted, predetermined 2000–2004 weights) and its customers' stocks (sales-weighted) enter jointly. Panels A and D: the main wage sample, worker and year effects; B and C: all workers with a first industry, worker and year effects; E and F: the firm panel, firm and year effects, exposure by the firm's industry. The first four columns report the coefficient per unit of log patent stock with its standard error in parentheses; the last four report the coefficient times the residual standard deviation of that regressor net of the fixed effects and the other regressors, in percent of the outcome (percentage points for employment and the labor share): the unit in which the three channels can be compared. Observations, in millions, by panel: Panel A, 47.3; Panel B, 117.9; Panel C, 100.3; Panel D, 47.1; Panel E, 5.0; Panel F, 1.6. Standard errors clustered by IO-121 industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

the suppliers' term understates the own effect, in this setting by a quarter to two fifths.

The own-industry coefficient should be interpreted as descriptive. An industry's own technology may respond to its own conditions in either direction, and the design that protects the suppliers' coefficient, namely predetermined shares, off-diagonal weights, and other industries' technology as the shifter, does not extend to it; Appendix B.5 reports the instruments tried for the own term and their failure. What the table does establish is that the technology that reaches an industry through its purchases affects its workers' pay and employment at least as much as the technology the industry patents itself, and in the opposite direction, and that it is the suppliers' term rather than the own term that affects every outcome in a consistent pattern.

5.3 From Firms to Workers

The worker-level analysis follows each worker from the industry in which she is first observed and includes a worker fixed effect. This section explains the reasons for that design and what it delivers. Exposure varies at the industry-by-year level. Workers in exposed industries are not a random draw from the workforce, and nearly all of the variation in exposure is between workers rather than within them: on the main sample, the standard deviation of the displacing exposure is 1.788 across workers and 0.156 within a worker over time (for the reinstating exposure, 1.911 and 0.103). A regression without worker effects would therefore compare workers from different industries and would largely measure who works where. The worker effect instead identifies the estimate from changes in the same worker's pay over time as the suppliers of her origin industry change their technology. Appendix Table 42 shows what this implies. With year effects alone the signs are reversed, with a coefficient of 0.214 (s.e. 0.104) on the displacing stock: workers in exposed industries earn more because of their composition, and controlling for 2005 age, sex and education removes this difference (0.058, s.e. 0.054). With first-industry and year effects the estimates are -0.217 (s.e. 0.069) and 0.237 (s.e. 0.074), the same as with worker effects (-0.223 and 0.251) and more precise. The wage result is thus a comparison of the same first industry over time, and the worker effects add precision rather than identification; the conclusion of Section 5.5 that the estimates are identified mainly between sectors is unaffected. The employment result is different. With first-industry effects it is -0.040 (s.e. 0.032), compared with -0.137 with worker effects, because the pool of workers attached to a first industry changes over time, so that for employment the within-worker comparison is what identifies the effect. Assigning the first observed industry rather than the current one gives the estimates an intention-to-treat interpretation. The current employer is itself a choice that the technology affects, and Section 5.1 has shown that firms respond on the hiring margin, so that the workers affected are those who are not hired and those who leave, whom an assignment by current employer would miss by construction.

Two findings show that the effects attach to the worker rather than to the employer. Table 7 assigns exposure either by the worker's first industry, as elsewhere in the paper, or by the industry of her current employer, using the same construction and weights on the same workers and years. When exposure is assigned by the first industry, white-collar workers' pay changes by -0.270 (s.e. 0.103) and, within the job they hold, by -0.136 (s.e. 0.053); when it is assigned by the current employer's industry, the corresponding coefficients are -0.025 (s.e. 0.022) and -0.007 (s.e. 0.011). When both exposures enter one regression (Panel B), the first industry's exposure accounts for the entire effect, -0.215 (s.e. 0.090) and 0.243 (s.e. 0.086), and the current employer's industry adds nothing, 0.018 (s.e. 0.021) and -0.007 (s.e.

Table 7: The wage effect follows the worker's first industry, not her current employer's

	Displacing	Reinstating
Panel A: exposure of the first industry or of the current employer's industry		
White-collar, wage, first industry	-0.270** (0.103)	0.305*** (0.092)
White-collar, wage, employer's industry	-0.025 (0.022)	0.031 (0.020)
White-collar, within the job, first industry	-0.136* (0.053)	0.142** (0.049)
White-collar, within the job, employer's industry	-0.007 (0.011)	0.006 (0.011)
Blue-collar, wage, first industry	-0.150 (0.079)	0.198* (0.085)
Blue-collar, wage, employer's industry	0.026 (0.028)	-0.010 (0.026)
Blue-collar, within the job, first industry	-0.051 (0.042)	0.071 (0.046)
Blue-collar, within the job, employer's industry	-0.029** (0.011)	0.031** (0.010)
Panel B: both exposures in one regression, all workers		
Wage, first industry	-0.215* (0.090)	0.243** (0.086)
Wage, current employer's industry	0.018 (0.021)	-0.007 (0.019)
Within the job, first industry	-0.097* (0.040)	0.103* (0.042)
Within the job, current employer's industry	-0.004 (0.008)	0.005 (0.007)
Panel C: workers in or out of the first industry by 2005 (log hourly wage)		
Still in the first industry	-0.221* (0.100)	0.201 (0.120)
Left the first industry	-0.215* (0.090)	0.259** (0.083)
Left, within the job	-0.102** (0.037)	0.117** (0.040)
Panel D: leavers, by years since leaving the first industry (log hourly wage)		
Before leaving	-0.068 (0.088)	0.093 (0.076)
0 to 2 years after	-0.081 (0.078)	0.107 (0.068)
3 to 5 years after	-0.169* (0.074)	0.185** (0.062)
6 to 10 years after	-0.197* (0.079)	0.207** (0.066)
11 or more years after	-0.169* (0.079)	0.178** (0.068)

Notes: Process–product classification; coefficients on the suppliers' displacing and reinstating patent stocks. Panel A assigns the exposure by the worker's first observed industry, as elsewhere in the paper, or by the industry of her current employer (same construction and weights), for white- and blue-collar workers; wage regressions have worker and year effects, within-the-job regressions worker-by-employer and year effects. Panel B enters both exposures in one regression. Panel C splits workers by whether they were still in their first industry in 2005. Panel D interacts the first industry's exposure with time relative to the year the worker left it. Worker-years, in millions, by panel: Panel A: 19.8 (rows 1–4), 15.5 (rows 5–8); Panel B, 47.3; Panel C: 10.9 (row 1), 33.8 (rows 2–3); Panel D, 42.3. Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

0.019); within the job the pattern is the same, apart from a small blue-collar exception noted in the table. Only 23 to 24 percent of these worker-years are spent in the first industry. The second finding is that leaving the first industry does not end the loss. 74 percent of the 2005 workers had already left their first industry by that year, and they lose as much as those who remained: −0.215 (s.e. 0.090) against −0.221 (s.e. 0.100), and −0.102 (s.e. 0.037) within their subsequent jobs. In event time (Panel D), the first industry's exposure has no effect on pay before the worker leaves (−0.068, s.e. 0.088) or in the first two years after leaving; the loss emerges three to five years after departure, −0.169 (s.e. 0.074), reaches −0.197 (s.e. 0.079) at six to ten years, and persists thereafter. Long-difference estimates imply the same horizon: the ten-year wage effect is −0.279 (s.e. 0.105) against 0.354, the twenty-year effect −0.260 against 0.300, and the five-year effect −0.118 against 0.151.

Table 8: Origin-industry effect or origin-industry trend: three checks

	Displacing	Reinstating
Panel A: ten-year differences, with the next decade's change in exposure		
Wage, same decade, baseline	-0.198** (0.065)	0.179 (0.092)
Wage, same decade, with the next decade	-0.189*** (0.053)	0.155* (0.079)
Wage, next decade's change (placebo)	0.009 (0.091)	0.052 (0.117)
Earnings, same decade, with the next decade	-0.407** (0.136)	0.490** (0.182)
Earnings, next decade's change (placebo)	-0.320 (0.182)	0.417 (0.257)
Panel B: the first industry's 1985–1994 wage and employment trends		
Wage, same sample, no trend controls	-0.301*** (0.065)	0.320*** (0.066)
Wage, with the trends	-0.288*** (0.055)	0.312*** (0.058)
Within the job, no trend controls	-0.119** (0.038)	0.122** (0.041)
Within the job, with the trends	-0.117** (0.038)	0.118** (0.041)
Panel C: relative to coworkers (employer-by-year effects)		
All workers, worker and year effects	-0.223* (0.088)	0.251** (0.085)
All workers, worker and employer-by-year effects	-0.093 (0.058)	0.118* (0.056)
White-collar, worker and employer-by-year effects	-0.144* (0.070)	0.175** (0.067)
White-collar, job and employer-by-year effects	-0.110* (0.044)	0.123** (0.045)
Blue-collar, job and employer-by-year effects	-0.017 (0.024)	0.035 (0.027)

Notes: Process–product classification; coefficients on the suppliers' displacing and reinstating patent stocks of the worker's first industry. Panel A: ten-year differences for stayers, start years 1985, 1990 and 1995 with start-year effects; the placebo rows are the coefficients on the change in the first industry's exposure over the following ten years. Panel B: the first industry's 1985–1994 wage and employment trends, projected forward, as controls, on the sample of first industries with 1985–1994 data. Panel C: employer-by-year fixed effects, so the comparison is among coworkers at the same employer in the same year. Worker-years, in millions, by panel: Panel A: 3.7 (rows 1–3), 4.3 (rows 4–5); Panel B, 46.8; Panel C: 47.3 (rows 1–2), 19.8 (rows 3–4), 15.6 (row 5). Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

These estimates are best interpreted as career effects of the first industry's exposure, and the open question is whether they reflect the technology or an origin-industry trend. Three checks in Table 8 bear on this question. First, the next decade's change in the first industry's exposure does not predict pay in the current decade: in ten-year differences for stayers, the placebo coefficient is 0.009 (s.e. 0.091), and the same-decade effect is unchanged at −0.189 (s.e. 0.053). Second, projecting the first industry's 1985–1994 wage and employment trends forward changes the estimate from −0.301 to −0.288 (s.e. 0.055). Third, among coworkers at the same employer in the same year, about half of the white-collar effect remains, −0.110 (s.e. 0.044) compared with −0.136 within the job; the remainder, and the whole of the blue-collar effect, reflects which employer a worker ends up at. What the design cannot do is distinguish a technology-driven origin effect from an origin-industry trend that is present neither in the pre-period nor in the next decade's exposure. The opposite signs on two kinds of patents from the same suppliers are the main argument against the latter interpretation, and the automation–other pair, which does not pass the next-decade placebo, is held to the weaker standard of Section 5.5. One interpretation consistent with all of these findings, though not tested here, is a vintage penalty: skills tied to the way an industry used to produce lose value once its suppliers' technology changes the inputs the industry buys, and the loss stays with workers after they leave.

5.4 Workers' Pay by Technology Type

Suppliers' displacing technology lowers the wages of workers from the industries that buy from them, and suppliers' reinstating technology raises them. Table 9 reports estimates of equation (8), which tests whether the two types enter with the opposite signs that Section 2.2 implies, on the main sample of workers followed from their first observed industry over 1995–2019, with worker and year fixed effects, so that the coefficients are identified from wage changes within a worker's career as the technology of her first industry's suppliers changes. Under the process–product classification and predetermined network weights, the coefficient on suppliers' displacing patent stock is −0.222 (s.e. 0.088) and the coefficient on suppliers' reinstating patent stock is 0.250 (s.e. 0.084). The estimates are very similar when the predetermined weights are replaced by contemporaneous ones, at −0.184 (s.e. 0.057) and 0.159 (s.e. 0.052). The coefficients under the automation-versus-other classification have the same signs in all four wage columns but are smaller and less precise; only the reinstating coefficient under predetermined weights (0.084, s.e. 0.026) is individually significant.

Two features of these estimates matter for their interpretation. First, the two coefficients in each pair are close to equal in magnitude and opposite in sign. The displacing and the

Table 9: Suppliers' displacing and reinstating technology: wages and employer pay premia

	Log hourly wage		Employer pay premium	
Network weights:	Predetermined	Contemporaneous	Predetermined	Contemporaneous
Panel A: process- versus product-dominant patents				
Suppliers' displacing	-0.222*	-0.184**	-0.074**	-0.044*
	(0.088)	(0.057)	(0.025)	(0.019)
Suppliers' reinstating	0.250**	0.159**	0.075**	0.033
	(0.084)	(0.052)	(0.024)	(0.017)
Worker-years	47,079,473			
Panel B: automation versus other breakthrough patents (Mann–Püttmann)				
Suppliers' displacing	-0.029	-0.031	-0.054***	-0.033***
	(0.023)	(0.020)	(0.007)	(0.006)
Suppliers' reinstating	0.084**	0.030	0.043***	0.031***
	(0.026)	(0.022)	(0.006)	(0.007)
Worker-years	45,362,287			

Notes: Each column within a panel is one worker-level regression with worker and year fixed effects on the main sample (1995–2019). The regressors are the own-industry, suppliers' and customers' log patent stocks of both types, entered jointly; the table reports the two suppliers' coefficients. The employer pay premium is the estimated employer effect from a two-way worker–employer wage decomposition; it is fixed for each employer, so a worker's premium changes only when she changes employer. Coefficients are per unit of log patent stock. Standard errors clustered by IO-121 base industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

reinstating stock are highly correlated across industries, at 0.99 for the process–product pair and 0.91 for the automation–other pair in the own-industry stocks and at 0.997 for the suppliers' exposures on the main sample (0.885 within a worker), because industries that patent heavily do so in both classes. What the pair identifies is therefore the composition of suppliers' patenting. Holding suppliers' total patenting constant, a shift toward the displacing class lowers the wages of workers in buying industries and a shift toward the reinstating class raises them, whereas additional patenting of both kinds, which the sum of the two coefficients captures, leaves wages roughly unchanged. Second, the coefficients are expressed per unit of log patent stock, and the relevant unit is the variation the regression actually uses. One standard deviation of the suppliers' exposure, net of the other regressors and the fixed effects, is 0.033 for the displacing stock and 0.029 for the reinstating stock, so a one standard deviation change alters hourly pay by about -0.7 and 0.7 percent, respectively; a one standard deviation shift in the composition of suppliers' patenting toward the displacing class, holding the total fixed, lowers hourly pay by -1.9 percent. Multiplying a coefficient by the total standard deviation of the exposure, 1.788, would be misleading, since the two stocks never move independently by that amount. These effects are modest for the average worker; Section 5.7 shows that they are concentrated at the top of the pay distribution.

The same opposite signs appear in the full worker panel: −0.179 against 0.148 under the process–product classification and −0.061 against 0.050 under the automation–other classification, all four significant (Appendix Table 35). Table 9 shows that they hold on the main sample and under both network weightings. The right-hand half of the table, discussed in Section 5.6, shows the same pattern for the pay premium of the worker's employer, and Figure 2 collects the worker-side and firm-side estimates by technology type.

5.5 Identification and Scope

The estimates by technology type are robust to comparing workers within the same municipality and year, show no significant response to future exposure, and are present in each half of the sample. They are identified mainly from differences between broad sectors, and industry-specific trends absorb them. Table 10 presents the evidence. The discussion focuses on the process–product pair, the main measure, and turns to the automation–other pair at the end.

Consider first local labor-market shocks. With municipality-by-year fixed effects, defined on the municipality in which the worker was first observed, the coefficients are −0.214 (s.e. 0.087) on the displacing stock and 0.247 (s.e. 0.084) on the reinstating stock, almost identical to the estimates without them. Workers from buying industries are therefore not simply located in places that fared worse or better.

Pre-existing trends are the second concern. If wages in a buying industry were already on a different path before its suppliers' technology changed, exposure arriving three years later should predict today's wage. I test this by adding the two suppliers' exposures three years ahead to the regression. The coefficients on future exposure are −0.067 (s.e. 0.071) and 0.062 (s.e. 0.069); neither is statistically significant, and the coefficients on current exposure remain at −0.197 (s.e. 0.069) and 0.244 (s.e. 0.072). The future coefficients do retain the signs of the main effect. Because the stocks are persistent, exposure three years ahead is a noisy measure of current exposure, and some loading on the lead is what such persistence would produce; it is small relative to the current coefficient. The next decade's change in exposure, in ten-year differences, does not predict pay either (Section 5.3).

The opposite signs are also present within each half of the sample. Within 2004–2019 the wage coefficients are −0.144 (s.e. 0.090) and 0.174 (s.e. 0.084); within 1995–2003 they are −0.141 (s.e. 0.068) and 0.182 (s.e. 0.093). Each is between 63 and 72 percent of its full-sample counterpart, and each is at the margin of conventional significance. Two conclusions follow. The estimate in Table 9 is not a comparison of the 1990s with the later years, since the same signs and similar magnitudes appear within each period; at the same time, part of

Table 10: Estimates by technology type: specifications, leads and sub-periods

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: log hourly wage, main sample (1995–2019)				
Worker and year effects	-0.223*	0.251**	-0.029	0.084**
	(0.088)	(0.085)	(0.023)	(0.027)
Municipality-by-year effects	-0.214*	0.247**	-0.025	0.084**
	(0.087)	(0.084)	(0.023)	(0.027)
Sector-by-year effects	-0.153	0.109	0.002	0.022
	(0.103)	(0.091)	(0.037)	(0.044)
Industry-specific linear trends	-0.027	0.027	-0.006	0.005
	(0.027)	(0.024)	(0.014)	(0.013)
Panel B: log hourly wage, with suppliers' exposure three years ahead				
Same sample, without the leads	-0.235**	0.273**	-0.012	0.103**
	(0.087)	(0.087)	(0.033)	(0.033)
Current exposure	-0.197**	0.244***	0.081	0.043*
	(0.069)	(0.072)	(0.071)	(0.018)
Exposure three years ahead	-0.067	0.062	-0.086	0.098*
	(0.071)	(0.069)	(0.049)	(0.045)
Panel C: sub-periods, log hourly wage				
2004–2019	-0.144	0.174*	-0.098**	0.117***
	(0.090)	(0.084)	(0.030)	(0.035)
1995–2003	-0.141*	0.182	-0.032	0.016
	(0.068)	(0.093)	(0.034)	(0.029)
Panel D: sub-periods, employer pay premium				
2004–2019	-0.035*	0.040*	-0.041***	0.038***
	(0.018)	(0.016)	(0.008)	(0.006)
1995–2003	-0.001	-0.062*	-0.045***	-0.010
	(0.022)	(0.028)	(0.011)	(0.011)

Notes: Each row is one worker-level regression per patent classification on the main sample, with worker fixed effects and the effects named in the row; the table reports the coefficients on the suppliers' displacing and reinstating patent stocks, entered jointly with the own-industry and customers' stocks. Municipality is the one in which the worker is first observed; sectors are eleven broad sectors of the worker's base industry; the trends are linear and specific to the base industry. Panel B adds the suppliers' two exposures of the base industry three years ahead and reports the current and the future coefficients of that regression. Panels C and D split the window. The first row is the specification of Table 9 on every worker-year of the main sample; that table uses the worker-years with an estimated employer premium, so the third decimal differs. Worker-years in the process–product regression, in millions, by panel: Panel A, 47.3; Panel B, 42.7; Panel C: 30.0 (row 1), 17.3 (row 2); Panel D: 29.8 (row 1), 17.2 (row 2). Standard errors clustered by IO-121 base industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

it does reflect such a comparison, since the full-sample coefficients exceed both sub-period estimates. Within six-year windows the residual variation in exposure is a third to a half of that in the full sample and the estimates are imprecise: −0.141 (s.e. 0.072) for 1995–2006, −0.058 (s.e. 0.079) for 2007–2019, and no effect for 2013–2019, after the patent flows have ended. The estimates by type thus identify decade-scale relationships. The employer pay premium behaves differently across the two periods (Panel D). In 2004–2019 it falls under displacing technology and rises under reinstating technology (−0.035, s.e. 0.018, and 0.040, s.e. 0.016), whereas in 1995–2003 it does neither (−0.001 and −0.062). The premium results of Section 5.6 should therefore be interpreted as results for the years from 2004 onward.

Two specifications remove most of the estimate, and they define its scope. With sector-by-year fixed effects for eleven broad sectors, the coefficients are −0.153 (s.e. 0.103) and 0.109 (s.e. 0.091): 69 percent of the displacing estimate and 43 percent of the reinstating estimate remain, and neither is distinguishable from zero. With industry-specific linear trends the coefficients are −0.027 (s.e. 0.027) and 0.027 (s.e. 0.024). The reason is that the exposures are stocks, and stocks move slowly. Within an industry, the suppliers' log patent stock is close to linear over the sample period, so an industry-specific trend removes the variation that identifies the coefficient along with any confounder, and because industries in the same sector buy from largely the same suppliers, sector-by-year effects remove much of it as well. What the design compares, then, is buying industries whose suppliers' technology mix evolved differently over 1995–2019, and a large part of that comparison is between broad sectors. These are limitations on the scope of the estimates rather than challenges to the credibility of the comparisons that survive: the opposite signs hold within municipality and year, within each sub-period, and without a sizeable lead. The paper does not claim identification from comparisons within sectors.

Sector-level wage trends unrelated to suppliers' technology could in principle account for the pattern. For this to be the case, such trends would have to be aligned with the composition of suppliers' patenting rather than its level, since more patenting of both kinds leaves wages unchanged; they would have to hold within municipalities and within each sub-period, without appearing three years in advance or in the next decade's exposure; and they would have to move employers' pay premia, the employment, margins and labor shares of buying firms, and the destinations of job changers in the same direction (Sections 5.1 and 5.6). This combination cannot be ruled out on the present evidence, and the causal interpretation rests on these tests together with the mechanism evidence that follows.

The automation–other pair is weaker on each of these counts and is subject to a pre-trend caveat. Its opposite wage coefficients appear only in 2004–2019 (−0.098, s.e. 0.030, and 0.117, s.e. 0.035); in 1995–2003 both coefficients are close to zero. When exposure

three years ahead is added to the regression, the coefficient on the future reinstating stock is significant (0.098, s.e. 0.045), the coefficient on the future displacing stock is −0.086 (s.e. 0.049), and the coefficient on the current displacing stock is no longer negative (0.081, s.e. 0.071). Sector-by-year effects and industry trends absorb the pair entirely. I therefore treat the automation–other classification as a secondary check on the signs and base the paper's conclusions on the process–product pair.

Both classifications infer a patent's type from the patent document, and the form of a first claim is partly a matter of drafting. A classification that does not use the claims gives the same worker estimates. For each breakthrough it measures how closely the text of the patent matches the work tasks that existed in the year of issue. Patents close to an existing task form the displacing stock, patents far from every existing task form the reinstating stock, and breakthroughs with no direct relation to work are left out (Appendix A.6). The coefficients on the suppliers' stocks of this pair are −0.277 (s.e. 0.093) and 0.275 (s.e. 0.102) for the hourly wage, −0.237 (s.e. 0.075) and 0.175 (s.e. 0.073) for employment, and −0.738 (s.e. 0.245) and 0.631 (s.e. 0.233) for annual earnings (Appendix Table 22). The two classifications are measured in different units, so I compare them per standard deviation of the exposure net of the fixed effects. On that scale the wage effects are -0.60 and 0.65 percent, against -0.73 and 0.74 percent for the process–product pair on the same worker-years; the employment effects are -0.58 and 0.43 percentage points against -0.66 and 0.41, and the effects on the employer pay premium are somewhat larger than under the process–product pair. The embedding model and every step of the construction were chosen by a rule written down before any of these regressions was run. Per standard deviation, the wage estimates are similar under two other embedding models, with a longer excerpt of the patent text, and when only the closest and the farthest third of patents are used. They are smaller, with the same signs, and no longer significant on one or both sides when patents are weighted by their importance, when closeness is measured against the task list of the preceding decade, and under the embedding model that agreed least with the validation sample.

The check has two limits. First, it does not extend to firms. The suppliers' stocks of the similarity pair affect neither the employment of buying firms (0.140, s.e. 0.187, and −0.116, s.e. 0.187) nor their labor share, so the firm results of Section 5.1 rest on the two claim-based classifications. Second, it does not establish which classification is preferable. When the suppliers' stocks of all three pairs enter one wage regression, the six terms are nearly collinear and the regression does not separate them. The displacing coefficient of the similarity pair keeps its size (−0.288, s.e. 0.152), the displacing coefficient of the process–product pair changes sign (0.104, s.e. 0.098), and neither pair's reinstating coefficient is distinguishable from zero (−0.110, s.e. 0.112, and 0.144, s.e. 0.126). The classifications

evidently draw on much the same variation in what suppliers patent, as one would expect if each measures the same distinction with error, and the regression cannot determine which measures it better.

5.6 Employers, Employment and Exits

The wage changes are accompanied by changes in where workers are employed, in whether they are employed, and in exits from employment. Table 11 reports the effects of the two exposures on the pay of the worker's employer, on employer changes, and on employment and annual earnings; Table 12 reports their effects on exits from employment and on the destinations of those exits.

Between 40 and 44 percent of the wage effect reflects the pay level of the employer at which a worker ends up. Panel A replaces the worker's own wage with the pay of her coworkers, defined as the mean wage residual, net of worker and year effects, of the other workers at her employer in that year. The coefficients are −0.099 (s.e. 0.032) on suppliers' displacing stock and 0.099 (s.e. 0.031) on suppliers' reinstating stock, compared with wage coefficients of −0.222 and 0.250 on the same sample. The employer's fixed pay premium, which changes only when the worker changes employer, accounts for 30 to 33 percent of the effect: −0.074 (s.e. 0.025) and 0.074 (s.e. 0.023). The premium falls under displacing and rises under reinstating technology in all eight cells of Table 9, and for the automation–other pair the premium estimates are considerably more precise than the wage estimates; as Section 5.5 showed, the premium results hold for the years from 2004 onward. The firm-level row of Panel A asks whether exposed firms themselves come to pay less. Under the process–product pair they do not: a firm's own pay component does not move with the exposure of its industry (0.002, s.e. 0.029, and −0.003, s.e. 0.027), consistent with what Table 4 showed for incumbents and new hires separately. The employer component of the wage effect therefore reflects which firms employ workers from exposed industries, rather than changes in what a given firm pays. Under the automation–other pair, firms' own pay does respond, by −0.037 (s.e. 0.014) and 0.035 (s.e. 0.015), so that for this measure the two sources cannot be separated.

Workers who change employer move to lower-paying employers under displacing technology and to higher-paying ones under reinstating technology. Among employer changes (Panel C), the probability that the new employer is in the bottom third of the premium distribution rises with the displacing stock (0.335, s.e. 0.129) and falls with the reinstating stock (−0.413, s.e. 0.124), while the probability that it is in the top third changes by −0.195 (s.e. 0.066) and 0.222 (s.e. 0.071). The automation–other pair gives the same four signs, each significant

Table 11: Suppliers' displacing and reinstating technology: employers, moves and employment

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: the wage and the pay of the worker's employer (worker and year effects)				
Log hourly wage	-0.222*	0.250**	-0.031	0.085**
	(0.088)	(0.085)	(0.023)	(0.026)
Pay of the employer's other workers that year	-0.099**	0.099**	-0.046***	0.052***
	(0.032)	(0.031)	(0.008)	(0.007)
Fixed employer pay premium	-0.074**	0.074**	-0.054***	0.043***
	(0.025)	(0.023)	(0.007)	(0.005)
Firm level: the firm's own pay component	0.002	-0.003	-0.037**	0.035*
	(0.029)	(0.027)	(0.014)	(0.015)
Panel B: employer changes (worker and year effects)				
Changes employer	0.097	-0.118*	-0.044*	-0.026
	(0.053)	(0.048)	(0.022)	(0.021)
Changes employer and industry	0.103*	-0.127**	-0.017	-0.023
	(0.049)	(0.045)	(0.014)	(0.012)
Panel C: at an employer change (base-industry and year effects)				
New employer in the bottom third by premium	0.335**	-0.413***	0.193***	-0.186***
	(0.129)	(0.124)	(0.048)	(0.031)
New employer in the top third by premium	-0.195**	0.222**	-0.088***	0.122***
	(0.066)	(0.071)	(0.020)	(0.025)
Change in the employer premium	-0.008	-0.005	0.009*	-0.012***
	(0.013)	(0.011)	(0.004)	(0.003)
Panel D: all workers (worker and year effects)				
Employed	-0.137**	0.106*	-0.052***	0.026
	(0.043)	(0.050)	(0.012)	(0.018)
Log annual earnings	-0.493***	0.507**	-0.154***	0.163**
	(0.137)	(0.167)	(0.044)	(0.058)

Notes: Coefficients on the suppliers' displacing and reinstating patent stocks, one regression per row and patent classification. Panel A, main sample: the pay of the employer's other workers is the mean wage residual, net of worker and year effects, of the worker's coworkers at her employer in that year; the fixed premium is the employer effect of the two-way wage decomposition. The firm-level row regresses the firm's mean wage residual on the exposures of the firm's industry with firm and year effects, weighted by workers. Panel B: indicators on consecutive worker-years with an employer in both. Panel C: worker-years with an employer change; thirds are those of the employer-premium distribution across workers. Panel D: every worker with a base industry, through 2019, whether or not she stays in the main sample. Observations in the process–product regression, in millions, by panel: Panel A: 46.3 (rows 1–3), 2.9 (row 4); Panel B, 42.6; Panel C, 9.6; Panel D: 117.9 (row 1), 100.3 (row 2). Standard errors clustered by IO-121 base industry (by the firm's industry in the firm-level row). Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

(0.193 and −0.186 for the bottom third; −0.088 and 0.122 for the top third). The change in the premium at the move itself does not respond under the process–product pair (−0.008, s.e. 0.013, and −0.005, s.e. 0.011), and under the automation–other pair it is small and of the opposite sign. This suggests that exposed workers' moves occur lower on the pay ladder at both origin and destination, rather than that each move represents a larger step down. The two classifications disagree on how often workers move. Under the process–product pair, employer changes that take the worker out of the industry become more frequent with the displacing stock (0.103, s.e. 0.049) and less frequent with the reinstating stock (−0.127, s.e. 0.045); under the automation–other pair neither coefficient behaves this way (Panel B).

Employment also responds. Among all workers with a base industry, whether or not they remain in the main sample (117.9 million worker-years through 2019), the probability of being employed falls with suppliers' displacing stock (−0.137, s.e. 0.043) and rises with the reinstating stock (0.106, s.e. 0.050); per standard deviation of exposure net of the fixed effects, these estimates correspond to -0.7 and 0.4 percentage points on an employment rate of 0.768. The automation–other pair agrees for displacing technology (−0.052, s.e. 0.012), while its reinstating coefficient is positive but not significant (0.026, s.e. 0.018). Annual earnings change by −0.493 (s.e. 0.137) and 0.507 (s.e. 0.167), about twice the hourly-wage coefficients (2.2 and 2.0 times), as one would expect if reduced time in work comes on top of lower pay per hour. Suppliers' technology therefore affects how many of a buying industry's workers are employed as well as how much the employed are paid.

The employment effect arises from exits that do not lead to unemployment. Among workers employed in the previous year, the probability of not being employed in the current year rises by 0.054 (s.e. 0.015) under displacing technology, relative to a base rate of 0.060, and falls by −0.051 (s.e. 0.019) under reinstating technology (Table 12, Panel A). The exits are largest for workers under 30, at 0.063 (s.e. 0.027) on a base rate of 0.094, and small at older ages. Across fifths of the worker's 2005 wage distribution they rise throughout, from 0.051 (s.e. 0.017) in the lowest fifth to 0.072 (s.e. 0.018) in the highest, against base rates that decline from 0.076 to 0.027 (Appendix Table 43); relative to the base rate, a standard deviation of exposure raises exits by about 3.2 percent in the lowest fifth and 12.9 percent in the highest, so that higher-paid workers respond most on this margin as well. Re-employment of workers who were not employed does not change detectably (−0.042, s.e. 0.042). Following employees under 60 into the next year (Panel C), the decline in employment, −0.045 (s.e. 0.013) on a base of 0.931, shows up as a status outside the labor force without benefits, 0.041 (s.e. 0.011), a category that includes students; registered unemployment, pension receipt and social assistance do not respond (−0.001, s.e. 0.004, for unemployment). Within a further three years most of these workers are back in employment, and exits followed by additional

Table 12: Exits from employment, re-employment and where the exits go

	Displacing	Reinstating	Mean
Panel A: not employed this year, workers employed last year			
All	0.054***	-0.051**	0.060
	(0.015)	(0.019)	
Under 30	0.063*	-0.045	0.094
	(0.027)	(0.027)	
30 to 49	0.016**	-0.003	0.036
	(0.006)	(0.007)	
50 to 59	0.034	-0.013	0.044
	(0.026)	(0.024)	
60 and over	-0.035	0.213	0.204
	(0.120)	(0.113)	
Panel B: employed this year, workers under 60 not employed last year			
All	-0.042	0.015	0.227
	(0.042)	(0.048)	
Panel C: status next year, employees under 60			
Employed	-0.045***	0.039*	0.931
	(0.013)	(0.016)	
Registered unemployed	-0.001	0.013*	0.025
	(0.004)	(0.006)	
Pension	0.001	0.000	0.005
	(0.003)	(0.004)	
Social assistance	0.000	0.001	0.003
	(0.001)	(0.001)	
Other status, no benefits	0.041***	-0.048***	0.022
	(0.011)	(0.014)	
Panel D: within three years of an exit, employees under 60			
Back in employment	0.036***	-0.036**	0.037
	(0.011)	(0.013)	
Not back in employment	0.010**	-0.003	0.032
	(0.004)	(0.005)	
Exit followed by more education	0.014***	-0.017***	0.009
	(0.004)	(0.005)	

Notes: Process–product classification; coefficients on the suppliers' displacing and reinstating patent stocks of the worker's first industry, in regressions with worker and year effects on all workers with a first industry. Panel A: the probability of not being employed this year among workers employed last year, by age. Panel B: the probability of being employed this year among workers under 60 not employed last year. Panel C: the worker's status the next year, employees under 60. Panel D: outcomes within three years of an exit. The last column is the mean of the outcome. Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

education account for about a third of the additional exits (Panel D). The employment effect thus consists mainly of young workers leaving employee jobs for a spell, rather than of layoffs into unemployment, which is consistent with the firm-level evidence that buying firms adjust by hiring less. Receipt of unemployment and sickness benefits falls under displacing technology and rises under reinstating technology over the years for which the benefit registers are available (through 2007).

5.7 Who Bears the Loss

The loss is concentrated among higher-paid, white-collar workers from exposed origin industries, employed at firms where pay is set at the firm level; lower-paid workers adjust through hours and employment. Figure 3 reports one regression per group of workers for the process–product pair, with groups defined by 2005 characteristics and each regression including its own year effects, so that differences in wage trends across groups are not attributed to the exposure.

The effects increase with pay. In the lowest fifth of the 2005 wage distribution neither coefficient differs from zero (-0.005, s.e. 0.081, and -0.004, s.e. 0.076). The coefficients then grow in absolute value with each successive fifth, reaching -0.302 (s.e. 0.094) and 0.299 (s.e. 0.083) in the highest. By occupation, the effects are largest for workers who held non-routine cognitive occupations in 2005 (-0.287, s.e. 0.106, and 0.327, s.e. 0.093) and absent for workers in routine manual occupations (-0.033, s.e. 0.104, and -0.009, s.e. 0.107), with routine cognitive and non-routine manual occupations in between. Sex and education do not separate workers in the same way: the two signs appear for men and for women and within every education group, with overlapping confidence intervals.

The gradient reflects pay within the job rather than the choice of employer. Table 13 decomposes each group's wage effect into the part operating through the employer pay premium and the part that is pay within the job the worker holds, the latter estimated with a fixed effect for each worker–employer pair. The premium coefficients are similar across the pay distribution, ranging from -0.042 in the lowest fifth to -0.084 in the highest (Appendix Table 36): workers in every pay group end up at lower-paying employers under displacing technology. The groups differ in pay within the same job. In the highest fifth it falls by -0.151 (s.e. 0.054), about 50 percent of that group's wage effect, and in the fourth fifth by -0.080 (s.e. 0.040), whereas in the lowest three fifths it does not change (0.052, s.e. 0.037, in the lowest fifth). Among white-collar workers, the within-job loss occurs at high-paying and at large employers: -0.266 (s.e. 0.074) within the job at employers in the top third of pay premia and nothing detectable in the lower two thirds (-0.054 and -0.017); -0.162 (s.e.

Table 13: Who pays: pay within the same job by pay fifth, employer and collar

	Displacing	Reinstating
Panel A: by 2005 pay fifth (log hourly wage; pay within the same job)		
Lowest fifth, wage	-0.005 (0.081)	-0.004 (0.076)
Lowest fifth, within the job	0.052 (0.037)	-0.063 (0.042)
Second, wage	-0.133* (0.060)	0.135* (0.056)
Second, within the job	-0.031 (0.031)	0.022 (0.032)
Third, wage	-0.173* (0.072)	0.171** (0.065)
Third, within the job	-0.049 (0.033)	0.042 (0.033)
Fourth, wage	-0.220** (0.081)	0.216** (0.072)
Fourth, within the job	-0.080* (0.040)	0.082* (0.038)
Highest fifth, wage	-0.302** (0.094)	0.299*** (0.083)
Highest fifth, within the job	-0.151** (0.054)	0.144** (0.050)
Panel B: white-collar workers within the job, by the 2005 employer		
Pay premium, lowest third	-0.017 (0.031)	0.051 (0.030)
Pay premium, middle third	-0.054 (0.043)	0.070 (0.039)
Pay premium, highest third	-0.266*** (0.074)	0.243*** (0.069)
Under 50 employees	-0.056 (0.056)	0.077 (0.049)
50 to 249 employees	-0.150* (0.061)	0.151** (0.058)
250 or more employees	-0.162** (0.058)	0.162** (0.051)
Within the job and occupation	-0.105* (0.048)	0.111* (0.046)
Panel C: pay or hours, within the job		
Blue-collar, log hourly wage	-0.051 (0.042)	0.072 (0.046)
Blue-collar, log hours	-0.133*** (0.040)	0.147*** (0.038)
Blue-collar, log annual earnings	-0.183** (0.068)	0.219** (0.070)
White-collar, log hourly wage	-0.136* (0.053)	0.142** (0.049)
White-collar, log hours	-0.090* (0.038)	0.097** (0.036)
White-collar, log annual earnings	-0.226** (0.080)	0.239** (0.076)
Blue-collar, log hours, continuous jobs	-0.057** (0.020)	0.074** (0.024)
White-collar, log hours, continuous jobs	-0.025 (0.019)	0.032* (0.016)

Notes: Process–product classification; coefficients on the suppliers' displacing and reinstating patent stocks of the worker's first industry. Each row is one regression on the stated group of main-sample workers, with the group's own year effects. "Wage" rows have worker effects; "within the job" rows have worker-by-employer effects, so the coefficient is pay in the job the worker holds. Panel B groups white-collar workers by their 2005 employer's pay premium tercile and size; the last row adds worker-by-employer-by-occupation effects. Panel C: hours are annual earnings over the hourly wage; continuous jobs are worker-employer pairs observed in the year before and after. Worker-years, in millions, by panel: Panel A: 7.4 (rows 1–2), 8.9 (rows 3–4), 9.1 (rows 5–6), 9.7 (rows 7–10); Panel B: 6.0 (row 1), 6.8 (row 2), 6.9 (row 3), 3.5 (row 4), 3.0 (row 5), 13.3 (row 6), 19.8 (row 7); Panel C: 15.5 (rows 1–3), 19.8 (rows 4–6), 7.8 (row 7), 10.9 (row 8). Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

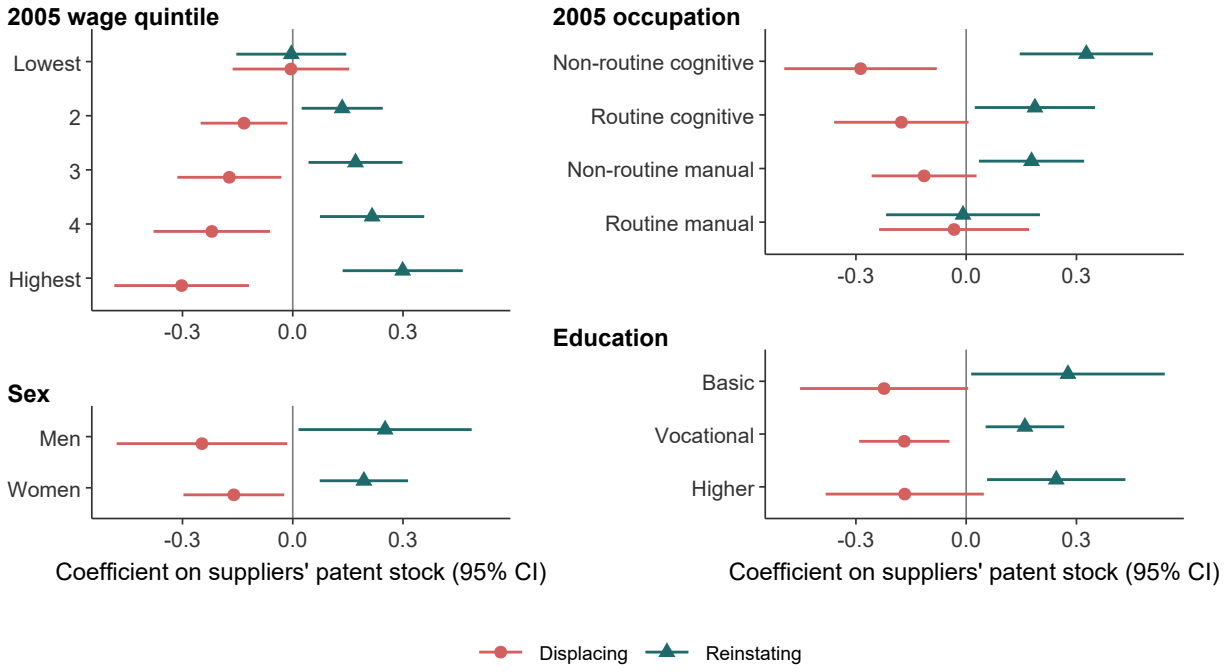


Figure 3: Who is affected, by technology type. Each pair of markers is the coefficient on the suppliers' displacing (circle) and reinstating (triangle) patent stock of the process–product classification in a regression of log hourly wage estimated separately for the workers in the stated group, with worker fixed effects and the group's own year effects; bars are 95 percent confidence intervals, clustered by IO-121 base industry. Groups are defined by 2005 characteristics, the occupation groups by the worker's 2005 occupation. Workers with unrecorded education and workers in agricultural and unclassified occupations are omitted from the figure. The estimates are reported in Appendix Table 36.

0.058) at employers with 250 or more employees against −0.056 (s.e. 0.056) at employers with fewer than 50. By occupation, the roles that coordinate production, namely production managers, engineers, technicians and logistics clerks, lose about twice as much within the job as other white-collar workers, −0.235 (s.e. 0.068) against −0.095 (s.e. 0.048); within the same job and occupation the loss is −0.105 (s.e. 0.048) out of −0.136, so that about three quarters of it is pay within the same role, and occupational upgrading within the employer does not decline. In terms of annual earnings both groups lose. Within the job, blue-collar workers lose through hours (−0.133, s.e. 0.040) rather than through the hourly wage (−0.051, s.e. 0.042), whereas white-collar workers lose through the hourly wage (−0.136, s.e. 0.053); in jobs observed in both the preceding and the following year, the white-collar hours effect disappears (−0.025, s.e. 0.019) while about half of the blue-collar effect remains (−0.057, s.e. 0.020). By major occupation group, the hourly wage falls for managers (−0.308,

s.e. 0.082), professionals (−0.335), technicians (−0.244) and clerks (−0.222), and does not change detectably for service and sales, craft, operator and elementary occupations; within the job, annual earnings fall in every group except service and sales, with the lower-paid groups adjusting through hours.

This result runs against the usual picture of who bears the cost of technological change, and the automation index yields the opposite gradient. Under the index, the loss is largest for workers with basic education, for men, for the lower-paid, and in manual occupations (Appendix B.10). The two measures do not describe the same workers. The index estimates rely on the comparison of pay from 2004 onward with the worker's own pay in 1995–2003 (Appendix B.1), whereas the patent estimates do not, which is one reason to place more weight on the patent gradient. The interpretation the evidence supports, although the paper does not test it, is that pay near the bottom of the Danish distribution is close to the rates set in collective agreements and adjusts through employment, whereas the pay of higher-paid, white-collar workers is set at the firm and moves with the firm; the roles that lose most are those that organized production around the inputs the displacing technology replaced, which is what a vintage penalty on production-management skills would look like (Section 5.3). One alternative can be ruled out: the affected workers are not those whose own tasks the patents describe. When workers are split by how exposed their own occupation is to patented technology, those in the least and the most exposed third show the same effects (Appendix Table 36).

A finer comparison, made within a buying industry and year and on the occupation in which a worker is first observed, does find that the pay effect is larger where the patents describe the occupation's tasks. Each breakthrough is divided among the occupations whose tasks its text describes, and for each buying industry and occupation I compute the share of the suppliers' 1994 patent stock that falls on that occupation (Appendix A.7). The suppliers' exposure is then interacted with this overlap, with first-industry-by-year effects absorbing the exposure itself, so that only differences between occupations within the same buying industry and year are used. One standard deviation more overlap changes the wage effect of suppliers' displacing technology by −0.056 (s.e. 0.010) and that of reinstating technology by 0.046 (s.e. 0.009), which is 25 and 18 percent of the two wage coefficients estimated on the same sample (Appendix Table 23). The estimates are similar among the half of workers whose first occupation code is mapped to the current classification with confidence (−0.065 and 0.058), under two other ways of assigning patents to occupations, and with the similarity pair; they are somewhat smaller for three-digit occupations (−0.030 and 0.039), and they remain significant when the standard errors are also clustered by occupation. The result has a narrower interpretation than this summary suggests, however. The gradient

reflects how much patented technology in general describes an occupation rather than how much the patents of the buyer's own suppliers do. When the overlap is measured relative to the occupation's share in all patents it has no effect (−0.015, s.e. 0.011, and −0.013, s.e. 0.020), and occupation-by-year effects remove the gradient (−0.026, s.e. 0.014, and 0.005, s.e. 0.012). The similarity pair is the exception, with a clear effect of the relative overlap (−0.084, s.e. 0.016, and 0.065, s.e. 0.012). The employment gradient has the opposite sign: employment falls less under displacing technology in the occupations the patents describe (0.024, s.e. 0.007, and −0.030, s.e. 0.008), and annual earnings show no gradient. Nor is there a gradient in the occupation held in 2005 (−0.008, s.e. 0.011, and −0.004, s.e. 0.014). The first occupation is determined before the exposure and the 2005 occupation is not, and only 45 percent of workers hold the same two-digit occupation at both dates. Within a buying industry, then, hourly pay responds more in the occupations that patents describe, while the network determines how exposed the industry is in the first place.

5.8 Mechanisms

Both types of technology reach buyers through the manufactured inputs they purchase, from equipment makers and other manufacturers alike, rather than through the equipment they invest in. If suppliers' technology merely made inputs cheaper, the type of supplier would be irrelevant; if instead the technology reaches the buyer inside what it buys, the type of supplier should matter, since equipment and components carry a production process into the buyer's plant whereas a chemical or a software licence does not. Table 14 decomposes each of the two suppliers' exposures into the parts originating from four kinds of supplier: equipment makers (electronics, electrical equipment, machinery, vehicles and other transport equipment), other manufacturing, chemicals and pharmaceuticals, and all remaining suppliers. The parts enter one regression, so that they sum to the headline exposure. Because the parts differ in dispersion, the comparable unit is each coefficient multiplied by the residual standard deviation of its own part, expressed in percent of hourly pay. On that scale, equipment makers' displacing stock lowers pay by -0.49 percent and their reinstating stock raises it by 0.32 percent (per unit, −0.560, s.e. 0.206, and 0.579, s.e. 0.259); other manufacturers' stocks contribute as much, at -0.48 and 0.59 percent; and the most patent-intensive suppliers of all, chemicals and pharmaceuticals, contribute nothing (0.13 and -0.03 percent, with the wrong signs and not significant), as is the case for service and other suppliers. The parts sum to -1.05 and 0.92 percent, of which manufacturing suppliers account for 92 and 99 percent. Panel B separates each buyer's purchases from its own industry, which the weights retain (Section 3.4) and which, for a manufacturing buyer, fall within its manufacturing part. These

Table 14: Which suppliers carry it: the parts of the suppliers' exposure by kind of supplier, in comparable units

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: four parts in one regression; each buyer's purchases from its own industry inside its part				
Equipment makers (NACE 26–30)	-0.560**	0.579*	-0.131	0.148*
	(0.206)	(0.259)	(0.083)	(0.071)
× SD, percent	-0.49	0.32	-0.22	0.31
Other manufacturing (NACE 10–33 without 20–21, 26–30)	-0.306*	0.300***	-0.016	0.099*
	(0.148)	(0.090)	(0.074)	(0.040)
× SD, percent	-0.48	0.59	-0.04	0.43
Chemicals and pharmaceuticals (NACE 20–21)	0.572	-0.113	-0.149	0.272
	(0.441)	(0.367)	(0.131)	(0.150)
× SD, percent	0.13	-0.03	-0.12	0.35
All other suppliers	-0.088	0.023	-0.012	-0.013
	(0.122)	(0.152)	(0.053)	(0.051)
× SD, percent	-0.21	0.04	-0.05	-0.06
Sum of the parts, × SD, percent	-1.05	0.92	-0.43	1.03
Panel B: the same with purchases from the own industry as a fifth part				
Purchases from the own industry	-0.575**	0.487*	-0.338**	0.264**
	(0.199)	(0.194)	(0.114)	(0.082)
× SD, percent	-0.67	0.48	-0.53	0.57
Equipment makers (NACE 26–30)	-0.742**	0.762*	-0.251	0.226
	(0.266)	(0.324)	(0.135)	(0.119)
× SD, percent	-0.46	0.32	-0.30	0.31
Other manufacturing (NACE 10–33 without 20–21, 26–30)	-0.191	0.236	0.125	0.041
	(0.224)	(0.124)	(0.090)	(0.054)
× SD, percent	-0.26	0.39	0.24	0.13
Chemicals and pharmaceuticals (NACE 20–21)	0.662	-0.117	-0.049	0.191
	(0.601)	(0.498)	(0.193)	(0.231)
× SD, percent	0.12	-0.02	-0.03	0.19
All other suppliers	0.017	-0.112	0.005	-0.066
	(0.131)	(0.175)	(0.064)	(0.072)
× SD, percent	0.03	-0.16	0.02	-0.26
Sum of the parts, × SD, percent	-1.24	1.02	-0.60	0.94

Notes: Log hourly wage on the main sample, worker and year effects, one regression per panel and classification. The suppliers' displacing and reinstating exposures are each split into the parts that come from the stated groups of supplier industries; the parts enter jointly with the own-industry and customers' stocks and sum to the suppliers' exposure of Table 9. In Panel A each buyer's purchases from its own industry are inside the part of the group it belongs to; in Panel B they are a separate part. For each part the first row is the coefficient per unit of log patent stock, the second its standard error, and the third the coefficient times the residual standard deviation of that part net of the fixed effects and the other regressors, in percent of hourly pay: the unit in which the parts can be compared and summed. Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

own-industry purchases account for about half of the total, 54 and 47 percent; equipment makers account for about a third, 37 and 31 percent, with their coefficients unchanged (-0.46 and 0.32 percent); other manufacturers account for the remainder and lose significance; and chemicals, pharmaceuticals and services still contribute nothing. The equipment makers' contribution is therefore not driven by the buyers' own patents. For the automation–other pair, the displacing side operates through equipment makers and own-industry purchases, and the reinstating side through manufacturing more broadly; chemicals and the remaining suppliers contribute nothing there either. The per-unit split into capital-good and other suppliers, in which the capital-good coefficients are several times the others, is reported in Appendix Table 44; because the capital-good part of the exposure varies about a third as much as the rest, per-unit coefficients overstate its share, and the comparable units do not support the statement that the effect runs "mainly through capital goods".

The second test asks whether the technology arrives through what buyers purchase or through what they invest in. Table 15 replaces the intermediate-purchase weights with investment-flow weights, which allocate each buying industry's equipment investment across the industries that produce the equipment; these weights are constructed from the capital-flow tables and predetermined to 2000–2004. On its own, the equipment-weighted exposure has no effect on pay (0.210, s.e. 0.960, and −0.416, s.e. 1.454). When both weightings enter one regression, the intermediate-purchase coefficients are unchanged, at −0.225 (s.e. 0.088) and 0.253 (s.e. 0.086), and the equipment-investment coefficients are −0.081 (s.e. 0.913) and −0.190 (s.e. 1.384). In comparable units the equipment terms are -0.0003 and -0.0003, against -0.0073 and 0.0074 for the intermediate terms, and their 95 percent confidence intervals ($[-0.0060, 0.0055]$ and $[-0.0043, 0.0037]$) just exclude an equipment effect as large as the intermediate one. The employer pay premium yields the same conclusion, as do investment matrices of a later vintage, suppliers' shares of equipment purchases, and transport-equipment and intellectual-property investment (Panels C to E). First-round purchases alone, excluding the own industry, give the same signs at a third to a half of the full-weight effect and are not significant (Panel D); the difference reflects the suppliers' suppliers and the own-industry purchases that the direct measure omits.

Two conclusions follow. The first is that the channel is not specific to displacing technology. A supplier's technology reaches its customers inside what it sells, and manufacturers sell the goods their customers produce with, such as the car parts of the example. If the technology is of the displacing kind, the product substitutes for work that the customer's own employees used to do; if it is of the reinstating kind, the product is something the customer's employees must work with. The split by type of supplier indicates that both arrive in this way. The second conclusion concerns where the technology is recorded. The

Table 15: Through what buyers purchase or through what they invest in: intermediate-purchase against investment-flow weights

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: log hourly wage, one weighting at a time				
Intermediate purchases, $\bar{L} - I$ (baseline)	-0.223*	0.252**	-0.029	0.084**
	(0.088)	(0.085)	(0.023)	(0.027)
× SD, percent	-0.72	0.74	-0.23	0.81
Equipment investment, 2000–2004 flows	0.210	-0.416	-0.021	0.056
	(0.960)	(1.454)	(0.279)	(0.417)
Equipment investment, 2014–2020 matrices	1.270	-1.540	0.389	-0.493
	(1.274)	(1.876)	(0.427)	(0.618)
Suppliers' shares of equipment purchases	0.130	-0.214	0.043	-0.092
	(0.081)	(0.112)	(0.083)	(0.057)
× SD, percent	0.43	-0.52	0.13	-0.45
Panel B: log hourly wage, both weightings in one regression				
Intermediate purchases	-0.225*	0.253**	-0.029	0.084**
	(0.088)	(0.086)	(0.024)	(0.027)
× SD, percent	-0.73	0.74	-0.23	0.81
Equipment investment	-0.081	-0.190	-0.080	0.055
	(0.913)	(1.384)	(0.295)	(0.419)
Panel C: employer pay premium, both weightings in one regression				
Intermediate purchases	-0.078**	0.078**	-0.054***	0.043***
	(0.026)	(0.024)	(0.007)	(0.005)
Equipment investment	-0.248	0.208	-0.073	0.051
	(0.257)	(0.389)	(0.083)	(0.083)
Panel D: log hourly wage, first-round purchases against equipment investment				
Direct requirements, own industry excluded	-0.137	0.223	0.028	0.151*
	(0.172)	(0.143)	(0.073)	(0.060)
× SD, percent	-0.25	0.41	0.08	0.64
Equipment investment	0.105	-0.321	-0.018	0.102
	(0.925)	(1.414)	(0.290)	(0.409)
Panel E: log hourly wage, investment by asset type in one regression				
Equipment	0.535	-0.807	0.123	-0.040
	(1.148)	(1.686)	(0.350)	(0.441)
Transport equipment	-0.399	-0.228	-0.640	0.487
	(0.840)	(0.746)	(0.486)	(0.381)
Intellectual property	-0.826	1.010	-0.282	0.207
	(0.733)	(0.958)	(0.264)	(0.332)

Notes: Main sample, worker and year effects, every regression with the own-industry and customers' stocks. The baseline weights suppliers by the buyer's intermediate purchases ($\bar{L} - I$, 2000–2004); the investment-flow weights allocate the buyer's gross fixed capital formation in equipment (or transport equipment, or intellectual property) across the industries that produce it, built from the capital-flow tables, predetermined to 2000–2004 unless stated. Where the run reported the residual standard deviation of a term net of the fixed effects and the other regressors, the third line of a row is the coefficient times that SD, in percent of hourly pay. Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

exposures weight suppliers by the buyer's purchases of intermediate inputs, so what comes from equipment makers in this analysis is what buyers purchase on current account, namely components, parts, small equipment and repair work, rather than the buyer's own investment in fixed assets, which the investment-flow weights show to have no effect. This is consistent with the absence of capital deepening documented in Section 5.1: the technology is embodied in what manufacturers deliver, and it does not appear on the buyer's balance sheet.

Direct evidence on the price margin is lacking. Producer-price and unit-value specifications are too imprecise to detect the declines in input prices that the substitution account implies, and no quantity response was detected in the short-window trade data. Further investigations of the channel were conducted with the automation index, and Appendix B.9 reports them with the caveats that apply to that measure: the loss under the index also concentrates in capital-good suppliers, it is as strong when exposure is reweighted toward relationship-specific inputs, and an exposure constructed from input similarity without a trade link bounds how much of it could reflect common shocks.

6 Structural Estimation

This section estimates the framework of Section 2 on the reduced-form coefficients. The model has five parameters and implies a response for each of the fourteen firm, worker and price outcomes used here. The estimation addresses two questions that the reduced-form coefficients cannot answer one at a time: what the fourteen coefficients identify when they are interpreted through a single mechanism, and how much of the loss in employment and pay is attributable to each of the two shocks.

Three objects are identified and one is not. The fall in demand for the buying industry's output is $\delta = 0.79$ per unit of the suppliers' shift toward process innovation (s.e. 0.15), and the wage-curve elasticity is $\varepsilon = 2.13$ (s.e. 0.68). The share response, $A\omega = -0.214$, implies that buyers shift spending from the purchased input toward their own labor. The capability gain g is identified only through the buyer's output price, and the data deliver a set for it rather than a point. The demand shift accounts for the entire fall in output per firm and for most or all of the fall in employment and pay. The data do not determine whether the improved input displaced the buyer's workers or complemented them.

6.1 The Reduced-Form Estimates and the Model

The reduced-form estimates of Section 5 determine which of the cases of Section 2.2 are relevant. All coefficients below are per unit of the suppliers' shift toward process innovation. The wage bill relative to purchases rises by $+0.52$ (s.e. 0.20) and the purchase share of turnover falls by -0.12 (s.e. 0.05), so that $\omega A < 0$; these two estimates do not determine whether the input-gains case or the labor-retakes case applies. The operating margin rises by $+0.11$ (s.e. 0.03) and value added per full-time worker does not fall ($+0.21$, s.e. 0.18), which requires incomplete pass-through, $\varphi < 1$. Real output per firm falls by -0.73 (s.e. 0.27), which requires a demand shift. Employment and the wage fall together, by -0.40 (s.e. 0.13) and -0.22 (s.e. 0.09), in the proportion given by the wage curve. The labor share of value added falls by -0.15 (s.e. 0.06), so that the increase in the markup outweighs the increase in labor's cost share. These estimates reject the model with full pass-through. They also reject the model with $\sigma \geq 1$ in the input-gains case, though not in the labor-retakes case.

Neither shock on its own is consistent with the estimates. The capability gain alone implies a rise in real output per firm, whereas output per firm falls. The relocation of a stage of production alone implies that buying firms exit, whereas the number of buying firms does not fall ($+0.71$, s.e. 0.46). The buyer's output price, which could in principle separate the two shocks, is too imprecisely estimated to do so (-0.13, s.e. 0.15). I specified these three tests before estimating the corresponding coefficients. The model with both

shocks is consistent with all of the estimates above, and it is the model estimated in this section.

6.2 Moments, Parameters and Estimator

The estimation uses 14 moments. Each is the coefficient on the displacing stock of the process–product classification per unit of the suppliers' exposure. Because the common level of the two stocks has no effect (Section 2.3), I interpret this coefficient as the response to the direction of the suppliers' innovation. Eleven moments are firm-account outcomes estimated with industry and year fixed effects (Table 3, Panel B, and the firm-register ratios of Section 5.1): employment, wage costs, value added, turnover, purchases, the purchase share of turnover, the wage bill over purchases, the labor share of value added, the operating result over value added, the margin over turnover, and value added per full-time worker. The remaining three are the hourly wage, estimated with worker and year fixed effects (Table 9), the buyer's producer price, and real output per firm. The last two are available for the 48 industries with a producer price index over 2005–2019. The firm moments describe the average firm in an industry-year, including entrants and exiters, which is the counterpart of the model's single producer. The joint covariance matrix of the moments is computed from industry-level score contributions on the 111 industries common to all of them.

I estimate five parameters: g, A, φ, ε and δ. The demand elasticity η is not identified separately from δ (Section 2.4) and is set to values in $[0.5, 2]$, with $\eta = 1.00$ as the reference value. The input's cost share and the markup are calibrated from the accounts using equation (1), which gives $s_M = 0.60$ and $\mu = 1.14$. The estimator is minimum distance with inverse-variance weights that are fixed before estimation (Appendix E.1). The joint covariance matrix enters the standard errors, the test statistic and the bootstrap.

The parameters are restricted to the region in which the model is well defined: $A \geq -s_M$, and a labor demand curve that slopes down with a bounded elasticity ($0 < \Lambda \leq 5$ in equation (4)). The search is repeated from starting values on both sides of every reflection of the objective function, so that no estimate depends on a single starting value. The upper bound on Λ is a choice, and Section 6.4 reports how the results change when it is tightened to the values found in the labor-demand literature. Because the share moments identify the product $A\omega$ precisely but A only weakly when the net cost advantage ω is close to zero, I report the fit on each side of $A = 0$. The two sides are the labor-retakes case ($A > 0$, $\omega < 0$) and the input-gains case ($A < 0$, $\omega > 0$) of Section 2.2.

Inference differs by parameter. I use delta-method standard errors for δ and ε, which lie in the interior of the admissible region, and confidence sets obtained by test inversion

Table 16: Structural estimates: the two shocks per unit of the suppliers' shift toward process innovation

	Case		
	Labor retakes tasks ($A > 0$, $\omega < 0$)	The input gains ($A < 0$, $\omega > 0$)	Intervals and notes
Demand shift δ	0.79 (0.15)	0.77 (0.15)	0.75 to 0.85 over η; bootstrap [0.52, 1.05]
Wage-curve elasticity ε	2.1 (0.7)	2.9 (1.0)	$\beta_w/\beta_L = 0.56$, Fieller [0.14, 1.51]
Share response $A\omega$	−0.214		the product the share moments identify; bootstrap [-0.42, 0.03]
Capability gain g	0.17 (0.23)	0.49 (0.23)	set over the retained share of cost changes [0.22, 1.38]; test inversion on the twelve mutually consistent moments [-0.1, 1.1]
Net cost advantage ω	−0.05 (0.26)	+0.31 (0.26)	not statistically different from zero in either case
Share response A	4.26	−0.60	at the bound on Λ; at the bound $-s_M$
Pass-through φ	0.34	0.65	differs between the two cases
Consistency of the moments, J (p)	19.6 (0.02)	16.5 (0.06)	9 accounting identities among the 14 moments; 4.2 on 7 ($p = 0.76$) without the two that fail
The demand shift's share of each response			
Employment	49 percent	100 percent	the input-gains share is 1.34; over the bootstrap draws [0.41, 1.58]
Hourly wage	49 percent	100 percent	over the draws [0.41, 1.58]
Real output per firm	1.18		Fieller [0.72, 2.02]; the capability gain alone would raise output

Notes: Minimum-distance estimates on the 14 moments at the reference demand elasticity $\eta = 1.00$, with $s_M = 0.60$ and $\mu = 1.14$ calibrated from the accounts. All quantities are per unit of the suppliers' shift toward process innovation; one residual standard deviation of the direction is 0.033. Delta-method standard errors in parentheses; bootstrap intervals from 199 score-bootstrap draws that repeat the full search; the test-inversion set for g collects the values not rejected at the 5 percent level. The two columns are the fits on either side of $A = 0$ within the region in which the model is well defined; the first fits better under the bound $\Lambda \leq 5$ and is the baseline, the second under a bound of 1.25. The J statistic is the distance between the fitted and the estimated moments; the nine restrictions it tests are accounting identities among the firm moments. The demand shift's shares are the demand component of each response divided by the total; a share of 100 percent means that the capability gain alone would have moved the outcome the other way.

for g, which does not. For quantities that depend on which constraint binds, I use a score bootstrap that repeats the full search in each of 199 draws. Ratios of coefficients receive Fieller confidence sets. Section 6.7 repeats the estimation with one input replaced at a time.

6.3 Estimates

Table 16 reports the estimates at the reference demand elasticity, with one column for each case. The labor-retakes case fits better under the bound $\Lambda \leq 5$, and I refer to it as the baseline. Appendix E compares the fitted value of every moment with its estimate (Table 47); the largest residual is 1.27 standard errors.

The demand shift is the most precisely estimated parameter. Per unit of the suppliers' shift toward process innovation, demand for the buying industry's output falls by $\delta = 0.79$ (s.e. 0.15). The estimate ranges from 0.75 to 0.85 as the demand elasticity varies over $[0.5, 2]$, its bootstrap interval is [0.52, 1.05], and it is the same in both cases. One residual standard deviation of the direction (0.033) corresponds to a fall in demand of 2.6 percent. The wage-curve elasticity is $\varepsilon = 2.13$ (s.e. 0.68). It is the inverse of the ratio of the wage response to the employment response, which is 0.56 with a Fieller set of [0.14, 1.51]. When a buying firm reduces employment by ten percent, the pay of the workers who remain falls by about five percent, so labor supply to the buying industry is far from perfectly elastic. The share response, $A\omega = -0.214$ with a bootstrap interval of [-0.42, 0.03], is the model's counterpart of the rise in the wage bill relative to purchases, +0.52 (s.e. 0.20): buyers reduce purchases by more than they reduce the wage bill. The margin over turnover rises by +0.11 (s.e. 0.03) because only a small part of the fall in cost is passed through to the price; the pass-through rate is $\varphi = 0.34$ in the baseline case.

The capability gain is not point-identified. In the baseline case $g = 0.17$ (s.e. 0.23), and in the input-gains case $g = 0.49$ (s.e. 0.23). As Section 2.4 explains, g is identified only as precisely as the buyer's output price is estimated. The price coefficient is -0.13 (s.e. 0.15), and without the price moment the standard error of g doubles, to 0.49. Expressed as an identified set over the share of cost changes retained in the markup, $1-\varphi$, the capability gain ranges from 0.22 to 1.38 as the retained share falls from 0.5 to 0.1. On the twelve moments that are mutually consistent (Section 6.5), the test-inversion set is [-0.1, 1.1]. The net cost advantage $\omega = g + \beta_w$ is -0.05 (s.e. 0.26) in the baseline case and +0.31 (s.e. 0.26) in the input-gains case, and it is not statistically different from zero in either.

The two cases differ in the pass-through rate and in the share response. In the labor-retakes case, $\varphi = 0.34$ and $A = 4.26$, which is at the upper bound on the elasticity of labor demand. In the input-gains case, $\varphi = 0.65$ and $A = -0.60$, which is at the lower bound $-s_M$. Which case fits better depends on the bound on Λ. Under $\Lambda \leq 5$ the labor-retakes case has $J = 19.6$ against 16.5 for the input-gains case. Under a bound of 1.25, the upper end of the range of industry-level estimates in the labor-demand literature, the better fit is the case in which the input gains, with $g = 0.49$ and $\omega = 0.31$.

The ratios of coefficients that must take a known value under the model are consistent

with it where they are informative. The ratio for the labor share of value added is 1.13 against a predicted value of one, with a Fieller set of [0.66, 8.96], and the ratio for value added per worker is 2.38, with a Fieller set of [-0.95, 17.46]. The ratios for the share of value added in turnover and for the two measures of the markup are uninformative at the precision of the moments that enter them.

6.4 The Contribution of Each Shock

Because every response is linear in (g, δ), each estimated response is the sum of a capability component (the model evaluated at $\delta = 0$) and a demand component (the model evaluated at $g = 0$). The bottom rows of Table 16 report the share of the demand component. For output per firm the result is the same in both cases: the demand shift accounts for 1.18 of the fall, with a Fieller set of [0.72, 2.02] and a bootstrap interval of [1.00, 1.85]. The entire fall in output per firm is therefore attributable to lower demand for the buyer's output. The capability gain alone would have raised output (Section 2.2), and the estimate of δ reflects how large the demand shift must be to more than offset it.

For employment and pay the decomposition is a set. In the baseline case the demand shift accounts for 49 percent of the fall in employment and 49 percent of the fall in the hourly wage, and the reallocation of the buyer's tasks to the input accounts for the remainder. In the input-gains case the demand component alone is 1.34 times the fall in employment. The capability gain is larger in that case and complements the tasks that remain with the buyer's workers, so that on its own it would have raised employment and pay. Across the admissible region the demand share ranges from 0.16 to 1.34. It is 0.49 at the bound $\Lambda \leq 5$ and 1.34 at a bound of 1.25, where the input-gains case fits better (Appendix Table 48). Over the bootstrap draws it lies in [0.41, 1.58] and exceeds one in 33 percent of them. The ratio of the two shocks, $g/\delta = 0.33$ with a Fieller set of [-0.57, 0.97], summarizes the same uncertainty.

The estimates therefore support one conclusion firmly and leave a second question open. When suppliers' innovation shifts toward process, most or all of the effect on buying firms and their workers comes from lower demand for the buyers' output. Part of what the buyer used to sell is produced elsewhere, and turnover, output per firm, employment and, through the wage curve, pay fall as a result, while the number of firms does not. The reallocation of tasks to a cheaper input is also visible in the accounts, in the rising wage bill relative to purchases and the falling purchase share. Its contribution to the loss in employment and pay lies between one half and zero, and may be negative. Which of these holds depends on the sign of the input's net cost advantage, and the price coefficient is not precise enough to determine it.

If demand falls, the fall should be visible among the buyers whose markets the suppliers' process innovation reaches. I split the 117 buying industries at the median share of output sold to households (0.08) and at the median share exported (0.10). Employment falls in both halves of both splits (by -0.33 and -0.25 above and below the household median, s.e. 0.21 and 0.21; by -0.42 and -0.22 above and below the export median, s.e. 0.19 and 0.17). The pay response is concentrated among buyers that sell domestically: it is -0.33 (s.e. 0.09) above the household median against -0.03 (s.e. 0.14) below it, and -0.29 (s.e. 0.09) below the export median against 0.09 (s.e. 0.13) above it, a difference of 2.4 standard errors. The employment response is thus not specific to one market, whereas the pay response is, and I cannot rule out that the fall in demand behind the pay response is specific to the domestic market.

6.5 What the Over-Identifying Restrictions Test

With fourteen moments and five parameters, the model appears to have nine over-identifying restrictions. These restrictions do not, however, test the mechanism. At a given demand elasticity, the model's fourteen responses lie in a five-dimensional linear subspace for any parameter values. The nine linear combinations orthogonal to this subspace are accounting definitions that relate the firm outcomes to six underlying objects (employment, the wage, the purchase share, the markup, the price and output), together with the cost identity $wL = s_L R/\mu$. For example, the operating result over value added changes by the negative of the change in the labor share, value added per worker changes by the change in value added less the change in employment, and wage costs change by the sum of the changes in the wage and in employment. The five parameters span the subspace, so the model is exactly identified given η, and the J statistic tests whether the moments are consistent with one another as measurements. A Wald test of the nine identities, which does not depend on the model, gives 19.9 on 9 degrees of freedom ($p = 0.019$). This is the value of the J statistic in Table 16, up to the first-order condition at the bound on Λ.

Two identities fail, and both failures reflect measurement. The sum of the operating result over value added and the labor share should not change, but it changes by 0.042 (s.e. 0.015). The operating result in the accounts is net of depreciation, so the share of value added that is neither labor cost nor operating result falls by this amount, and the model has no counterpart to it. The second identity relates value added per full-time worker to value added and employment. The difference should be zero and is 0.47 (s.e. 0.20), because full-time equivalents and headcount are measured on different samples of firms. On the 1,194,766 firm-years that report every outcome this discrepancy is 0.11, whereas the discrepancy in

the operating-result identity, at 0.052, does not shrink. The other seven identities hold jointly (4.2 on 7 degrees of freedom, $p = 0.76$), and the two moments that fail do not drive the parameter estimates. On the twelve moments that exclude them, $g = 0.20$ (s.e. 0.24), $\delta = 0.82$ (s.e. 0.14) and $\omega = -0.06$, with $J = 4.3$ on 7 degrees of freedom ($p = 0.75$). Dropping any one of the fourteen moments keeps δ between 0.69 and 0.82 and the better fit in the labor-retakes case.

The test of the mechanism is instead the comparison of Section 6.1. The capability gain alone implied that real output per firm would not fall, and the rejection of that implication is the reason the model includes a demand shift.

6.6 Incidence and Counterfactuals within the Support

The model delivers the average wage response through the wage curve. How that response is divided between the employer and the job is taken from the data. An incumbent's pay falls by -0.222 (s.e. 0.088) per unit of the direction. The employer's fixed pay premium accounts for -0.074 (s.e. 0.025) of this, about a third (33 percent, s.e. 0.17). The pay of the worker's coworkers falls by -0.099 (s.e. 0.032), while pay within the job the worker holds does not change (0.002, s.e. 0.029). The average incumbent therefore loses pay through the employer at which they work rather than through a pay cut in a continuing job. Table 17 repeats this decomposition by group. The employer component is common to all groups: the premium coefficient lies within two standard errors of the core set's in all 16 groups and ranges from -0.027 to -0.111. The groups differ in pay within the job. The within-job coefficient declines monotonically with the worker's pay rank, from 0.052 (s.e. 0.037) in the lowest fifth of 2005 pay to -0.151 (s.e. 0.054) in the highest, a difference of -0.203 (s.e. 0.065). It is also negative in the cognitive occupations (-0.133, s.e. 0.054, non-routine; -0.115, s.e. 0.042, routine) and absent in the manual ones. Workers at the bottom of the pay distribution thus lose only through the employer component, whereas workers at the top also lose within their job. The model's wage response in the baseline case, -0.224, is 0.02 standard errors from the reduced-form estimate.

The counterfactuals I consider are shifts in the direction of the innovation of a buying industry's suppliers within the support of the data. One residual standard deviation of the direction is 0.033, the interquartile range of the industry means is 0.135 and their range is 0.501. In the public exposure series, the suppliers of the median industry shifted by 0.167 toward process innovation between 1995 and 2019, which is 5.1 standard deviations and inside the range. Table 18 reports the model's responses in the baseline case. Moving a buying industry from the 25th to the 75th percentile of its suppliers' direction lowers its

Table 17: Incidence of the pay loss: the employer's premium and pay within the job, by worker group

Group	Pay (s.e.)	Employer premium (s.e.)	Pay within the job (s.e.)	Premium's share (s.e.)
All incumbents				
Core set	-0.222 (0.088)	-0.074 (0.025)	0.002 (0.029)	0.33 (0.17)
Gender				
Women	-0.160 (0.070)	-0.046 (0.019)	–	0.29 (0.17)
Men	-0.247 (0.118)	-0.076 (0.027)	–	0.31 (0.18)
Education				
Basic	-0.223 (0.117)	-0.068 (0.030)	–	–
Vocational	-0.168 (0.063)	-0.059 (0.020)	–	0.35 (0.18)
Higher	-0.167 (0.110)	-0.055 (0.017)	–	–
Pay quintile (2005)				
Lowest fifth	-0.005 (0.081)	-0.042 (0.025)	0.052 (0.037)	–
Second	-0.133 (0.060)	-0.066 (0.026)	-0.031 (0.031)	0.50 (0.30)
Third	-0.173 (0.072)	-0.077 (0.027)	-0.049 (0.033)	0.45 (0.24)
Fourth	-0.220 (0.081)	-0.080 (0.027)	-0.080 (0.040)	0.36 (0.18)
Highest fifth	-0.302 (0.094)	-0.084 (0.022)	-0.151 (0.054)	0.28 (0.11)
Occupation group				
Non-routine cognitive	-0.287 (0.106)	-0.068 (0.018)	-0.133 (0.054)	0.24 (0.11)
Routine cognitive	-0.176 (0.093)	-0.057 (0.018)	-0.115 (0.042)	–
Non-routine manual	-0.115 (0.073)	-0.047 (0.028)	-0.031 (0.034)	–
Routine manual	-0.033 (0.104)	-0.027 (0.038)	0.014 (0.052)	–
Agricultural	-0.217 (0.096)	-0.111 (0.026)	-0.330 (0.087)	0.51 (0.26)
Other	-0.153 (0.103)	-0.067 (0.029)	-0.047 (0.054)	–

Notes: Reduced-form coefficients on the suppliers' displacing exposure, process–product pair, per unit: the incumbent's log hourly wage, the employer's fixed pay premium, and pay within the worker–employer pair, each with worker and year fixed effects and standard errors clustered by first industry. The premium's share is the premium coefficient over the pay coefficient with a delta-method standard error that treats the two as independent; it is shown only where the pay loss is more than two standard errors from zero. At the core set the coworkers' pay component is -0.099 (0.032).

employment by -6.43 percent, the pay of its incumbents by -3.02 percent and its real output per firm by -9.73 percent, and raises its margin over turnover by 1.51 points. These total responses are the same in both cases and under every bound, because the fourteen moments determine them. The decomposition is not: across the two cases and the bounds on Λ, the demand component of the fall in employment per standard deviation ranges from -2.32 to -0.26 percent and the capability component from -1.33 to 0.59, which is the set of Section 6.4. A parametric bootstrap that draws the moments from their estimated covariance matrix and repeats the search in each of 199 draws agrees with the score bootstrap (Appendix Table 49). The interval for δ is [0.54, 1.03] against [0.52, 1.05], and the interval for the demand share of the fall in employment is [0.45, 1.72] against [0.41, 1.58].

The estimates do not deliver an aggregate effect. The model describes a single buying industry. Without reallocation, the loss across the buying industries is the sum of their responses, which is an upper bound on the net loss. With full reallocation the net loss is zero, because the demand that the buyers lose is met by other producers. The productivity

Table 18: Counterfactuals within the support: the response to a shift of the suppliers' direction

	One net s.d. (0.033)		Interquartile range (0.135)		Median 1995–2019 change (0.167)	
Outcome	Total	Demand part	Total	Demand part	Total	Demand part
Employment (percent)	-1.6	-0.8	-6.4	-3.2	-7.9	-3.9
Hourly wage (percent)	-0.7	-0.4	-3.0	-1.5	-3.7	-1.8
Wage costs (percent)	-2.3	-1.1	-9.5	-4.7	-11.7	-5.8
Real output per firm (percent)	-2.4	-2.5	-9.7	-10.4	-12.0	-12.9
Turnover (percent)	-2.6	-2.6	-10.6	-10.6	-13.1	-13.1
Purchases (percent)	-3.5	-3.7	-14.3	-15.2	-17.6	-18.8
Purchase share of turnover (points)	-0.5	-0.6	-1.9	-2.4	-2.4	-3.0
Margin over turnover (points)	0.4	0.1	1.5	0.3	1.9	0.4
Labor share of value added (points)	-0.5	0.2	-2.1	0.6	-2.6	0.8
Output price (percent)	-0.2	-0.0	-0.9	-0.2	-1.1	-0.2

Notes: The model's response in the baseline case to a shift of the direction of the suppliers' innovation toward process by one residual standard deviation of the exposure, by the interquartile range of the industry means, and by the median industry's change in the public exposure series between 1995 and 2019, all in the units of the exposure; the range of the industry means is 0.501 and no shift exceeds it. Log outcomes in percent, shares in percentage points. The demand part is the response at $g = 0$; the rest is the capability part, which varies across the two cases and the bounds on Λ (Table 48), whereas the totals do not.

term of a general-equilibrium calculation would require the capability gain in units that the moments do not provide, and I do not report it.

6.7 Robustness

Table 19 repeats the baseline estimation with one input replaced at a time. The demand shift is stable. It lies between 0.69 and 0.82 in every specification that keeps the baseline coding of the patent stocks, regardless of which wage coefficient is used, whether the exposure includes the own-industry term, whether the firm moments are estimated on a common sample or on each outcome's own sample, and whether the price or real output is excluded. The one exception is the specification that uses the wage coefficient estimated with industry trends, at 0.42. The capability gain varies as one would expect of a parameter that is not point-identified, from -0.02 with the industry-trend wage to 0.39 without the buyers affected by the floor on the patent stocks. In these two specifications the net cost advantage also turns positive, which places them in the input-gains case, with the demand shift again accounting for the entire loss in employment and pay. Excluding the price moment doubles the standard error of g and leaves its point estimate at 0.16; excluding real output has no material effect. Recoding the floored patent stocks as zero reduces every parameter, to $\delta = 0.16$ and $g = 0.25$, because that coding attenuates every reduced-form coefficient by about two thirds. The sign of the demand shift is unchanged. Leaving each supplying industry out of the exposure in turn changes δ by less than three hundredths and g by less than a tenth.

Table 19: Robustness of the structural estimates

Fit	g	δ	ω	J (p), d.f.
Baseline	0.17	0.79	−0.05	19.6 (0.02), 9
Wage with sector-by-year fixed effects	0.11	0.84	−0.08	24.5 (0.00), 9
Wage with industry-specific trends	-0.02	0.42	−0.07	38.6 (0.00), 9
Sample ending in 2016	0.18	0.79	−0.05	19.9 (0.02), 9
Exposure rebuilt from the patent file	0.17	0.79	−0.05	19.6 (0.02), 9
Without the own-industry term	0.15	0.78	−0.05	18.5 (0.03), 9
Firm moments on a common sample of firm-years	0.11	0.69	−0.05	34.3 (0.00), 9
Without the price moment	0.16	0.78	−0.05	19.7 (0.01), 8
Without real output	0.17	0.79	−0.05	19.7 (0.01), 8
Twelve mutually consistent moments	0.20	0.82	−0.06	4.3 (0.75), 7
Floored stocks coded as zero	0.25	0.16	+0.22	18.7 (0.03), 9
Without the fifteen floor-affected buyers (13 moments)	0.39	0.76	+0.34	15.3 (0.05), 8

Notes: Each row replaces one input to the baseline fit of Table 16 and re-estimates at the reference demand elasticity; the last column gives the J statistic with its degrees of freedom. The two wage rows replace the worker-level wage coefficient by its estimate with sector-by-year fixed effects or with industry-specific linear trends. The rebuilt-exposure and own-term rows replace the wage coefficient; the variants that also recompute the firm moments give $g = 0.17$ and 0.16, $\delta = 0.79$ and 0.80, with $J = 19.6$ and 19.2. The one-sample row estimates the eleven firm moments on the firm-years that report every outcome. The twelve-moment row drops the operating result over value added and value added per full-time worker. The floor rows recode the floored patent stocks as zero, or drop the fifteen buying industries whose exposure the floor affects, with the wage excluded as a target in that sample.

Four limitations remain. The demand elasticity is imposed rather than estimated. The task primitives enter only through the share response A and are not separately identified. The framework has no aggregate implications beyond the two bounds. And the sign of the net cost advantage, which determines how the loss in employment and pay is divided between the two shocks, is not determined by the price coefficient at its current precision. A more precisely estimated price response would narrow the set.

7 Conclusion

This paper has asked how technological change at an industry's suppliers affects the workers and firms in the industries that buy from them. The answer depends on the type of technology. A supplier's technology reaches its customers embodied in what it sells. When the technology is displacing, the supplier's product substitutes for work that the customer's own employees used to perform, and those workers lose pay and employment. When it is reinstating, the product is something the customer's employees must work with, and the same trade linkages transmit gains. The conventional logic of network propagation, in which a supplier's progress benefits the industries it serves, holds for one type of technology and is

reversed for the other.

Four findings from Danish register data covering every worker and firm from 1980 to 2023 support this answer, beginning with firms. First, firms whose suppliers develop displacing technologies see reductions in employment and in their labor share, while value added per worker and operating margins rise and the pay of incumbents and new hires is unchanged; they hire less, their workforce ages, and they invest less, with no evidence of capital deepening. Firms adjust through size and hiring and retain the cost savings. Under reinstating technologies these effects are reversed. The technology of an industry's suppliers matters more for its workers than the industry's own technology, and with the opposite sign. Second, suppliers' displacing technologies lower the pay of workers from buying industries and reinstating technologies raise it by a similar amount, under both patent classifications, among workers in the same municipality and year, and within each half of the sample period. The effect follows the worker's first industry rather than her current employer's; workers who left that industry lose as much as those who stayed, and the loss emerges years after they leave. Third, the changes in pay are accompanied by changes in employer, employment and exits. Workers who change employer move to lower-paying employers, and between 40 and 44 percent of the wage effect reflects the pay level of the employer at which a worker ends up; fewer workers are employed, and exits from employment rise, mostly among young workers and mostly into spells outside the labor force rather than unemployment. Fourth, both types of technology reach buyers through the manufactured inputs they purchase, from equipment makers and other manufacturers alike and not through the equipment buyers invest in, and the pay losses fall on higher-paid white-collar workers at large, high-premium employers and in occupations that coordinate production, while lower-paid workers adjust through hours and employment. The design has limits, which define the scope of the estimates: the exposures are slow-moving stocks, the identifying variation lies largely between broad sectors and is absorbed by industry trends, and the origin-industry effect cannot be separated from an origin-industry trend that appears in neither the pre-period nor the next decade's exposure. An index of suppliers' automation that does not distinguish the two types agrees in sign with the displacing results; because it is zero by construction before 2004, it is reported only as a blunt check (Appendix B).

In the framework, a shift of suppliers' innovation toward process affects the buying industry in two ways: the purchased input performs more of the buyer's tasks and performs them more cheaply, and part of what the buyer used to sell is now produced by others. A capability gain alone predicts a rise in output per buyer, which the data reject, and a relocation of activity alone predicts exit, whereas the number of buyers does not fall. Only the two shocks together are consistent with both facts. The accounts alone do not determine

why purchases contract most. Either the input became cheaper per task and buyers demand a given quantity of what it does, or the fall in demand lowered the wage until labor took tasks back. The sign of the input's net cost advantage distinguishes the two explanations; the buyer's output price identifies this sign, though imprecisely. The structural estimates of Section 6 put the fall in the buying industry's demand at 0.79 per unit of the suppliers' shift (s.e. 0.15) and attribute to it the whole of the fall in output per firm and most or all of the fall in employment and pay. The capability gain of the input is visible in the accounts, but its contribution to the loss is not point-identified: it lies between one half and zero, and may be negative, and the price coefficient is too imprecise to narrow this range. Within the support of the data, moving a buying industry from the 25th to the 75th percentile of its suppliers' direction lowers its employment by -6.43 percent, the pay of its incumbents by -3.02 percent and its real output per firm by -9.73 percent.

Two implications follow. For measurement, the workers exposed to new technology are for the most part not the adopters' employees. Exposure arrives through purchases, and its sign depends on the type of technology, a distinction that a single automation index averages away and that a study confined to the adopting firm, industry or local labor market cannot observe. For policy, the losses from displacing technology do not take the form of layoffs at an adopting firm. They appear one step down the supply chain, as fewer hires and lower-paying subsequent jobs for the workers who started in the exposed industries, which is not where support aimed at workers displaced by automation is directed. The type of technology matters as well: the same network that transmits losses from displacing technology transmits gains from reinstating technology.

References

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333.

Acemoglu, Daron and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

_ **and** _ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **and** _ , "Tasks, Automation, and the Rise in US Wage Inequality," *Econometrica*, 2022, *90* (5), 1973–2016.

_ **, Ufuk Akcigit, and William R. Kerr**, "Networks and the Macroeconomy: An Empirical Exploration," *NBER Macroeconomics Annual*, 2016, *30* (1), 273–335.

_ **, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Anderson, Theodore W. and Herman Rubin, "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *Annals of Mathematical Statistics*, 1949, *20* (1), 46–63.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller, "New Frontiers: The Origins and Content of New Work, 1940–2018," *Quarterly Journal of Economics*, 2024, *139* (3), 1399–1465.

Autor, David H. and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 2013, *103* (5), 1553–1597.

_ **,** _ **, and Gordon H. Hanson**, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Balsvik, Ragnhild, "Is Labor Mobility a Channel for Spillovers from Multinationals? Evidence from Norwegian Manufacturing," *Review of Economics and Statistics*, 2011, *93* (1), 285–297.

Baqaee, David Rezza and Emmanuel Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Barrot, Jean-Noël and Julien Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Bilgin, Nuriye Melisa, Ester Faia, and Gianmarco I. P. Ottaviano, "Technology Spillovers, Diffusion and Rivalry in Firm Networks," CEPR Discussion Paper DP19804, Centre for Economic Policy Research 2024.

Borusyak, Kirill and Peter Hull, "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 2023, *91* (6), 2155–2185.

_ , _ , and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Card, David, Jörg Heining, and Patrick Kline, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 2013, *128* (3), 967–1015.

Carvalho, Vasco M., Makoto Nirei, Yukiko U. Saito, and Alireza Tahbaz-Salehi, "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321.

Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, and Nicole Wößner, "The Adjustment of Labor Markets to Robots," *Journal of the European Economic Association*, 2021, *19* (6), 3104–3153.

Feenstra, Robert C. and Gordon H. Hanson, "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979–1990," *Quarterly Journal of Economics*, 1999, *114* (3), 907–940.

Felten, Edward W., Manav Raj, and Robert Seamans, "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal*, 2021, *42* (12), 2195–2217.

Frey, Carl Benedikt and Michael A. Osborne, "The Future of Employment: How Susceptible Are Jobs to Computerisation?," *Technological Forecasting and Social Change*, 2017, *114*, 254–280.

Ghosh, Ambica, "Input-Output Approach in an Allocation System," *Economica*, 1958, *25* (97), 58–64.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Graetz, Georg and Guy Michaels, "Robots at Work," *Review of Economics and Statistics*, 2018, *100* (5), 753–768.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller, "Artificial Intelligence and the Labor Market," Working Paper 33509, NBER 2025.

Humlum, Anders, "Robot Adoption and Labor Market Dynamics," Working Paper, University of Chicago 2019.

Hummels, David, Rasmus Jørgensen, Jakob R. Munch, and Chong Xiang, "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, 2014, *104* (6), 1597–1629.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring Technological Innovation over the Long Run," *American Economic Review: Insights*, 2021, *3* (3), 303–320.

Koch, Michael, Ilya Manuylov, and Marcel Smolka, "Robots and Firms," *Economic Journal*, 2021, *131* (638), 2553–2584.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *Quarterly Journal of Economics*, 2017, *132* (2), 665–712.

Kreiner, Claus Thustrup, Søren Leth-Petersen, and Peer Ebbesen Skov, "Tax Reforms and Intertemporal Shifting of Wage Income: Evidence from Danish Monthly Payroll Records," *American Economic Journal: Economic Policy*, 2016, *8* (3), 233–257.

Leontief, Wassily W., "Quantitative Input and Output Relations in the Economic Systems of the United States," *Review of Economics and Statistics*, 1936, *18* (3), 105–125.

Liu, Huben, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller, "Technology and Labor Markets: Past, Present, and Future," Working Paper 34386, NBER 2025.

Long, John B. and Charles I. Plosser, "Real Business Cycles," *Journal of Political Economy*, 1983, *91* (1), 39–69.

Mann, Katja and Lukas Püttmann, "Benign Effects of Automation: New Evidence from Patent Texts," *Review of Economics and Statistics*, 2023, *105* (3), 562–579.

Neuhäusler, Peter, Rainer Frietsch, and Henning Kroll, "Probabilistic Concordance Schemes for the Re-Assignment of Patents to Economic Sectors and Scientific Publications to Technology Fields," Discussion Papers Innovation Systems and Policy Analysis 60, Fraunhofer ISI, Karlsruhe 2019.

OECD, "Inter-Country Input-Output (ICIO) Tables," 2021.

Poole, Jennifer P., "Knowledge Transfers from Multinational to Domestic Firms: Evidence from Worker Mobility," *Review of Economics and Statistics*, 2013, *95* (2), 393–406.

Schmoch, Ulrich, "Concept of a Technology Classification for Country Comparisons," Final Report to the World Intellectual Property Organisation (WIPO), Fraunhofer Institute for Systems and Innovation Research, Karlsruhe 2008.

Serafinelli, Michel, ""Good" Firms, Worker Flows, and Local Productivity," *Journal of Labor Economics*, 2019, *37* (3), 971–1004.

Stehrer, Robert, Anastasia Bykova, Kirsten Jäger, Oliver Reiter, and Monika Schwarzhappel, "Industry Level Growth and Productivity Data with Special Focus on Intangible Assets," EU KLEMS Release, The Vienna Institute for International Economic Studies (wiiw) 2019.

Stoyanov, Andrey and Nikolay Zubanov, "Productivity Spillovers Across Firms Through Worker Mobility," *American Economic Journal: Applied Economics*, 2012, *4* (2), 168–198.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," Working Paper, Stanford University 2020.

A Data Construction Details

A.1 Occupation Crosswalks

The automation probabilities of Frey and Osborne (2017) are defined for the 702 occupations of the US Standard Occupational Classification (SOC). Applying them to Danish workers requires two crosswalks. The first maps SOC codes to the International Standard Classification of Occupations (ISCO-08) using the crosswalk published by the Bureau of Labor Statistics. Because SOC is a national elaboration of ISCO-08, most SOC codes correspond to a unique ISCO-08 code; where several SOC codes map into a single ISCO-08 code, I take the employment-weighted average of their automation probabilities, with weights from the BLS Occupational Employment and Wage Statistics. The second step is straightforward: the Danish occupational classification, DISCO-08, is structurally identical to ISCO-08 at the four-digit level, so no further aggregation is needed. The routine task intensity index of Autor and Dorn (2013) is constructed in the same way, with O*NET task-content scores defined at the SOC level and carried through ISCO-08 to DISCO-08.

Occupation-level measures are aggregated to the industry level through the workers themselves. Each worker is assigned the value of her recorded DISCO-08 occupation (codes recorded before 2010 are converted with the crosswalk described below), and the industry measure is the average over the employed workers in each IO-121 industry and year, so that occupations employing a larger share of an industry's workforce receive proportionally more weight. The SOC-to-DISCO weights also incorporate Danish employment counts, so that mappings into occupations that are common in Denmark are weighted more heavily. The baseline cross-sectional measures fix each industry at its 2005 value, which prevents an industry's measured task content from drifting with the endogenous reallocation of workers across occupations; the time-varying versions retain the full industry-by-year panel.

The occupation codes change classification in 2010. Codes recorded from 1991 to 2009 follow the Danish version of ISCO-88 and codes from 2010 onward follow DISCO-08, and there is no complete official bridge between the two: the conversion table of Statistics Denmark lists only the 324 of the 570 pre-2010 codes of the wage statistics that translate into a single DISCO-08 code. Treating the remaining codes as if they were DISCO-08 codes misclassifies them wherever ISCO-08 created or renumbered groups, as it did for personal care workers (from group 51 to 53), clerks (41 to 43 and 44) and ICT occupations (21 and 31 to 25 and 35). Only 40 percent of the workers who stayed with their employer from 2009 to 2010 are recorded under the same two-digit number in both years. I therefore construct the crosswalk from these job stayers. Each pre-2010 four-digit code is mapped to the DISCO-08 code that is most common in 2010 among the workers who held the code in 2009 and remained with

the same employer, 1.57 million workers in all. A code with fewer than 30 stayers takes the official translation where one exists, and otherwise the DISCO-08 group most common among the stayers of its three-digit group. The crosswalk assigns a DISCO-08 code to the first recorded occupation of 96 percent of workers; the remaining first codes identify a major group only.

The crosswalk is validated out of sample. Employers are divided at random into two halves, the crosswalk is built from the job stayers of one half, and it is evaluated on the stayers of the other. For 79 percent of held-out stayers it returns the two-digit group in which the worker is recorded in 2010, compared with 53 percent for the official table with unmatched codes kept as recorded. The break in the occupational distribution at the change of classification, measured by the dissimilarity index between the two-digit employment shares of 2009 and 2010, falls from 0.31 to 0.06, against 0.02 between other adjacent years, and the mapped group explains more of the variation in 2009 hourly wages (an R^2 of 0.24 on the group's 2010 mean wage, against 0.21). A six-digit version of the crosswalk is slightly more accurate across 2009–2010, but the last two digits of the old codes were recoded in 2003, and a six-digit table would treat that recoding as changes of occupation, so I use the four-digit version. The break is not eliminated. About a fifth of job stayers still change two-digit group across 2010, so that changes of occupation measured across that year partly reflect the change of classification. The routine and non-routine occupation groups of Figure 3 and Table 36 are defined from the first digit of the recorded code and do not depend on the crosswalk.

A.2 Industry Classification Concordance

The Danish industry classification, DB07, coincides with NACE Revision 2 at the two-digit level. The national accounts of Statistics Denmark use a 117-industry classification that is finer than two-digit DB07 in some sectors and coarser in others. I construct a concordance between the two using the official crosswalk published by Statistics Denmark. When a national-accounts industry spans several DB07 codes, the input-output flows are allocated in proportion to each DB07 industry's share of gross output within the national-accounts grouping, with the shares computed from the supply table for 2000–2004.

EU KLEMS (Stehrer et al., 2019) reports its technology measures at the NACE Revision 2 level, which maps directly onto two-digit DB07. The OECD Inter-Country Input-Output tables (OECD, 2021) use a 45-industry classification, which I concord to DB07 through the ISIC Revision 4 correspondence tables published by the United Nations Statistics Division.

A.3 Hybrid Input-Output Table Construction

The baseline analysis relies on a hybrid input-output table that combines domestic and international linkages. The domestic block is taken from the national supply-use tables of Statistics Denmark, converted into symmetric industry-by-industry tables at the 117-industry level under the industry technology assumption. International linkages come from the OECD ICIO tables, which report bilateral intermediate-input flows among 76 countries at the 45-industry level.

The two sources are combined as follows. The 117-industry domestic table records intermediate flows between Danish industries. For each Danish industry j, the import column of the supply table gives the share of total intermediate inputs sourced from abroad. These imported inputs are then distributed across source countries and foreign industries using the bilateral trade shares in the OECD ICIO tables. Formally, let m_{ij}^c denote the value of intermediate inputs that Danish industry j imports from industry i in country c. I set $m_{ij}^c = M_{ij} \times s_{ic}$, where M_{ij} is the total imported intermediate input from industry i to Danish industry j (from the supply table) and s_{ic} is country c's share of Denmark's total imports of industry i's output (from the OECD ICIO). The resulting hybrid table combines a detailed domestic block with a coarser but complete international block and distinguishes 121 industries in total. The suppliers' exposures are built from the total-requirements input-coefficient matrix, so that technology embodied in imported intermediates counts toward a buyer's exposure; the customers' exposures use only the domestic output-coefficient block, since import rows have no Danish gross output and the forward chain is defined only for domestic producers. Annual tables are available for the full sample window and are used in the time-varying robustness variant.

A.4 Exposure Construction: the Subtract-I Convention

The exposures of Section 3.4 subtract the identity matrix from the predetermined Leontief and Ghosh inverses rather than setting their full diagonals to zero, and the distinction matters. The diagonal element $\bar{\ell}_{jj} = 1 + \bar{a}_{jj} + (\bar{a}^2)_{jj} + \cdots$ has two components: an accounting identity, since one unit of industry j's gross output is required to deliver one unit of final output of industry j, and genuine roundabout linkages within the industry, such as electricity used to generate electricity or steel used to make the machinery that makes steel. Subtracting I removes only the accounting identity, which would otherwise dominate us_{jt} mechanically and force a near-unit correlation with di_{jt}, while preserving the economically meaningful roundabout term $\bar{\ell}_{jj} - 1$, which reflects actual production chains. In the Danish hybrid table these roundabout coefficients are far from negligible: 27 percent of IO-121 industries

have $\bar{\ell}_{jj} - 1 > 0.10$, and forestry and food processing reach 0.40–0.48. Setting the entire diagonal to zero would discard these real linkages; in practice the two exposure variants are 98.6 percent correlated. Appendix A.5 examines how much of the results depends on the roundabout term.

A.5 The Own-Industry Term of the Suppliers' Exposure

Because the weights are $\bar{L} - I$, an industry's purchases from its own industry remain part of its suppliers' exposure, and its own patent stock therefore enters the regressor through these purchases as well as through the separate own-industry term. Among the base industries of the main sample, the roundabout coefficient $\bar{\ell}_{jj} - 1$ averages 0.068 and reaches 0.366. The own component, $(\bar{\ell}_{jj} - 1)\, T_{jt}$, accounts for about 7 percent of the mean exposure and has a correlation of 0.67 with it. In much of manufacturing the buying industry is its own largest supplier: basic metals (25 percent of the industry's weight sum), fabricated metals (20), basic chemicals (19), wood (18), textiles (17), computers (16), and engines (15). Table 20 removes this term exactly from the regressor and, in a second set of columns, enters the two components separately, for the main worker outcomes and for the firm outcomes. The firm results and the workers' employment and annual-earnings results do not depend on the own term. Firm log employment is -0.121 (s.e. 0.079) and 0.133 (s.e. 0.071) without it, compared with -0.119 and 0.133 with it; worker employment is -0.136 (s.e. 0.054) compared with -0.137; and annual earnings are -0.447 (s.e. 0.166) compared with -0.493. The pay results retain their signs at roughly three-quarters of their size in comparable units. The wage coefficients become -0.164 (s.e. 0.102) and 0.206 (s.e. 0.096), or -0.49 and 0.54 percent per standard deviation, compared with -0.73 and 0.74 percent with the own term included, and the displacing coefficient loses precision. The premium coefficients, -0.061 (s.e. 0.030) and 0.062 (s.e. 0.026), remain significant, and the pay results for the automation–other pair are unchanged. Part of the pay effect therefore operates through what industries buy from their own industry. I regard this as a statement about the scope of the estimates rather than as evidence against the channel. Entered on its own, the own term has opposite signs for the two types on every worker outcome (for the wage, -0.638, s.e. 0.210, and 0.597, s.e. 0.208). Since the weights are fixed and the own stock is already included as a regressor, this coefficient indicates that the effect of the own stock is more negative under displacing patents, and more positive under reinstating patents, in industries that buy more of their own output. That is what the supplier channel predicts for purchases from the own industry, and it is the opposite of the direct own-stock term in Table 6. Because the coefficient rests on differences in $\bar{\ell}_{jj}$ across a handful of manufacturing industries, I interpret it as consistent

with the channel and nothing more.

Table 21 enters the own stocks and the suppliers' stocks separately, jointly, and jointly with the customers' stocks, on the samples of Table 6, and repeats the pay regressions separately for manufacturing and for non-manufacturing first industries (Section 5.2).

A.6 Technology Type from the Patent Text

The process–product classification infers the type of a patent from its first claim, and the automation classification from the description of a device. The similarity pair of Section 5.5 classifies the same breakthrough patents by a different criterion, namely how close the invention is to work that people already did. For each breakthrough I take the plain-language summary of the patent prepared by Liu et al. (2025), which states what the invention does and what is new in it, and the list of work tasks that the same authors compile for each decade. A patent is compared with the most recent list dated at or before its issue year, so that it is measured only against tasks that existed when it was issued. Summaries and tasks are converted into text embeddings, and the closeness of a patent to existing work is the mean cosine similarity between its summary and its five nearest tasks. Patents are ranked by closeness within their issue year. The rank, scaled to lie between zero and one, is the patent's weight in the displacing stock, and one minus the rank is its weight in the reinstating stock, so that an invention that describes something workers already do counts mostly as displacing and an invention far from every existing task counts mostly as reinstating. The weighted patents are assigned to industries and accumulated into stocks in the same way as the other patent measures.

One restriction matters for the reinstating stock. Among all breakthroughs, the patents farthest from existing tasks are for the most part chemical compounds, components and materials, which bear no direct relation to work, and only a small minority are technologies that create new tasks. The main version of the measure therefore keeps the breakthroughs with a direct relation to work. These are identified by a classifier trained on 855 breakthroughs whose relation to work was rated by a large language model from the summary and the issue year alone; 71 percent of breakthroughs are kept, and the version on all breakthroughs is reported as a check. The same rated sample determined the embedding model: under a rule written down before any regression was run, the model chosen is the one whose closeness ranking agrees best with the ratings. The measure is related to an existing classification in the expected way. The share of patents that Mann and Püttmann (2023) classify as automation rises from 62 percent in the tenth of patents farthest from existing tasks to 83 percent in the closest tenth.

Table 20: The own-industry term of the suppliers' exposure

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: log hourly wage (wage sample, worker and year effects)				
Suppliers' exposure on $\bar{L} - I$ (baseline)	-0.223*	0.251**	-0.029	0.084**
	(0.088)	(0.085)	(0.023)	(0.026)
× SD, percent	-0.73	0.74	-0.23	0.81
Without the own term	-0.164	0.206*	-0.017	0.080**
	(0.102)	(0.096)	(0.027)	(0.030)
× SD, percent	-0.49	0.54	-0.12	0.67
Split: suppliers' exposure without the own term	-0.145	0.175	0.017	0.043
	(0.099)	(0.094)	(0.032)	(0.029)
× SD, percent	-0.43	0.45	0.11	0.32
Split: the own term, $(\bar{\ell}_{jj} - 1)\, T_{jt}$	-0.638**	0.597**	-0.307*	0.285**
	(0.210)	(0.208)	(0.132)	(0.092)
× SD, percent	-0.76	0.60	-0.52	0.66
Panel B: employer pay premium (wage sample, worker and year effects)				
Suppliers' exposure on $\bar{L} - I$ (baseline)	-0.074**	0.075**	-0.054***	0.043***
	(0.025)	(0.024)	(0.007)	(0.005)
× SD, percent	-0.24	0.22	-0.42	0.42
Without the own term	-0.061*	0.062*	-0.057***	0.044***
	(0.030)	(0.026)	(0.009)	(0.006)
× SD, percent	-0.18	0.16	-0.40	0.37
Split: suppliers' exposure without the own term	-0.056	0.053*	-0.046***	0.030***
	(0.029)	(0.025)	(0.010)	(0.006)
× SD, percent	-0.17	0.14	-0.29	0.23
Split: the own term, $(\bar{\ell}_{jj} - 1)\, T_{jt}$	-0.165*	0.184*	-0.100***	0.107***
	(0.070)	(0.077)	(0.025)	(0.016)
× SD, percent	-0.20	0.18	-0.17	0.25
Panel C: employed (all workers with a first industry, worker and year effects)				
Suppliers' exposure on $\bar{L} - I$ (baseline)	-0.137**	0.106*	-0.052***	0.026
	(0.043)	(0.050)	(0.012)	(0.018)
Without the own term	-0.136*	0.106	-0.053***	0.024
	(0.054)	(0.062)	(0.014)	(0.020)
Split: suppliers' exposure without the own term	-0.124*	0.101	-0.042**	0.019
	(0.054)	(0.060)	(0.015)	(0.020)
Split: the own term, $(\bar{\ell}_{jj} - 1)\, T_{jt}$	-0.197*	0.100	-0.112	0.057
	(0.077)	(0.088)	(0.067)	(0.062)
Panel D: log annual earnings (all workers, worker and year effects)				
Suppliers' exposure on $\bar{L} - I$ (baseline)	-0.493***	0.507**	-0.154***	0.163**
	(0.137)	(0.167)	(0.044)	(0.058)
Without the own term	-0.447**	0.465*	-0.152**	0.160*
	(0.166)	(0.200)	(0.051)	(0.064)
Split: suppliers' exposure without the own term	-0.388*	0.417*	-0.112*	0.128*
	(0.162)	(0.194)	(0.055)	(0.063)
Split: the own term, $(\bar{\ell}_{jj} - 1)\, T_{jt}$	-0.989***	0.820**	-0.405*	0.329
	(0.261)	(0.300)	(0.174)	(0.172)
Panel E: buying firms' log employment (firm and year effects)				
Suppliers' exposure on $\bar{L} - I$ (baseline)	-0.119	0.133*	-0.083	0.104*
	(0.070)	(0.064)	(0.045)	(0.044)
× SD, percent	-0.77	0.94	-1.14	1.55
Without the own term	-0.121	0.133	-0.104*	0.124**
	(0.079)	(0.071)	(0.042)	(0.040)
× SD, percent	-0.71	0.86	-1.32	1.68
Split: suppliers' exposure without the own term	-0.124	0.130	-0.097**	0.114**
	(0.079)	(0.070)	(0.037)	(0.035)
× SD, percent	-0.73	0.84	-1.19	1.47
Split: the own term, $(\bar{\ell}_{jj} - 1)\, T_{jt}$	-0.169	0.258	0.037	0.022
	(0.159)	(0.163)	(0.142)	(0.161)
× SD, percent	-0.44	0.68	0.16	0.09
Panel F: buying firms' labor share (firm and year effects)				
Suppliers' exposure on $\bar{L} - I$ (baseline)	-0.055	0.057*	-0.008	0.015
	(0.030)	(0.029)	(0.007)	(0.008)
× SD, percent	-0.35	0.40	-0.10	0.23
Without the own term	-0.044	0.048	-0.003	0.012
	(0.034)	(0.032)	(0.008)	(0.008)
× SD, percent	-0.26	0.31	-0.04	0.17
Split: suppliers' exposure without the own term	-0.045	0.051	-0.007	0.017
	(0.033)	(0.031)	(0.008)	(0.009)
× SD, percent	-0.27	0.33	-0.08	0.22
Split: the own term, $(\bar{\ell}_{jj} - 1)\, T_{jt}$	-0.077	0.039	-0.069*	0.048
	(0.042)	(0.041)	(0.033)	(0.039)
× SD, percent	-0.20	0.10	-0.30	0.20

Notes: The suppliers' exposure weights suppliers by $\bar{L} - I$, which removes the unit own term of the Leontief inverse and keeps each industry's purchases from its own industry, $\bar{\ell}_{jj} - 1$ (mean 0.068, maximum 0.366 across base industries; about 7 percent of the mean exposure). The baseline rows repeat the paper's estimates. The second row of each panel subtracts the own term, $(\bar{\ell}_{jj} - 1)\, T_{jt}$ with T_{jt} the industry's own stock, from the regressor; the third and fourth rows enter the two parts separately. Every regression includes the industry's own stocks and the customers' exposure. Where the run reported it, the third line of a row is the coefficient times the residual standard deviation of the regressor net of the fixed effects and the other regressors, in percent. Standard errors clustered by IO-121 industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 21: Own-industry and suppliers' technology, each alone and together

	Own stocks			Suppliers' stocks		
	alone	with suppliers'	and customers'	alone	with own	and customers'
Panel A: log hourly wage; 47.3 million observations						
Displacing	0.032	0.042	0.038	-0.161	-0.214*	-0.223*
	(0.028)	(0.024)	(0.024)	(0.086)	(0.092)	(0.088)
× SD, percent	0.59	0.75	0.66	-0.60	-0.77	-0.73
Reinstating	-0.023	-0.035	-0.028	0.188	0.218*	0.251**
	(0.026)	(0.025)	(0.024)	(0.096)	(0.090)	(0.085)
× SD, percent	-0.36	-0.53	-0.42	0.61	0.68	0.74
Panel B: employer pay premium; 47.1 million observations						
Displacing	-0.002	0.002	0.001	-0.067**	-0.072**	-0.074**
	(0.007)	(0.006)	(0.006)	(0.022)	(0.023)	(0.025)
× SD, percent	-0.03	0.03	0.01	-0.25	-0.26	-0.24
Reinstating	-0.009	-0.013**	-0.011	0.050**	0.069**	0.075**
	(0.006)	(0.005)	(0.006)	(0.019)	(0.022)	(0.024)
× SD, percent	-0.14	-0.19	-0.17	0.16	0.21	0.22
Panel C: employed; 117.9 million observations						
Displacing	0.011	0.019*	0.019*	-0.097*	-0.126**	-0.137**
	(0.012)	(0.009)	(0.009)	(0.040)	(0.043)	(0.043)
× SD, percent	0.27	0.42	0.41	-0.54	-0.66	-0.66
Reinstating	-0.006	-0.009	-0.009	0.058	0.089	0.106*
	(0.014)	(0.013)	(0.013)	(0.046)	(0.050)	(0.050)
× SD, percent	-0.09	-0.14	-0.14	0.26	0.38	0.41
Panel D: log annual earnings; 100.3 million observations						
Displacing	0.061	0.093*	0.093*	-0.320*	-0.455***	-0.493***
	(0.050)	(0.039)	(0.038)	(0.131)	(0.138)	(0.137)
× SD, percent	1.39	1.98	1.98	-1.72	-2.30	-2.33
Reinstating	-0.046	-0.073	-0.073	0.264	0.423**	0.507**
	(0.053)	(0.049)	(0.048)	(0.154)	(0.163)	(0.167)
× SD, percent	-0.76	-1.13	-1.13	1.16	1.74	1.93
Panel E: buying firms' log employment (firm and year effects); 5.0 million observations						
Displacing	-0.006	0.003	0.008	-0.132*	-0.134*	-0.119
	(0.019)	(0.020)	(0.021)	(0.053)	(0.057)	(0.070)
× SD, percent	-0.15	0.08	0.21	-0.91	-0.91	-0.77
Reinstating	0.007	-0.004	-0.008	0.141**	0.143**	0.133*
	(0.020)	(0.020)	(0.021)	(0.049)	(0.053)	(0.064)
× SD, percent	0.17	-0.08	-0.19	1.06	1.05	0.94
Panel F: buying firms' labor share (firm and year effects); 1.6 million observations						
Displacing	-0.001	0.003	0.001	-0.030	-0.035	-0.055
	(0.008)	(0.009)	(0.009)	(0.027)	(0.032)	(0.030)
× SD, percent	-0.03	0.05	0.03	-0.13	-0.15	-0.21
Reinstating	0.003	-0.002	-0.001	0.034	0.040	0.057*
	(0.009)	(0.010)	(0.010)	(0.026)	(0.031)	(0.029)
× SD, percent	0.05	-0.03	-0.02	0.16	0.17	0.23
Panel G: log hourly wage, manufacturing first industries; 7.9 million observations						
Displacing	-0.048	-0.035	-0.008	-0.185	-0.075	-0.105
	(0.025)	(0.029)	(0.026)	(0.105)	(0.115)	(0.118)
× SD, percent	-0.76	-0.38	-0.08	-0.65	-0.19	-0.26
Reinstating	0.035	0.023	0.015	0.161*	0.069	0.053
	(0.022)	(0.028)	(0.027)	(0.082)	(0.095)	(0.097)
× SD, percent	0.51	0.20	0.13	0.65	0.16	0.13
Panel H: log hourly wage, other first industries; 39.4 million observations						
Displacing	0.042	0.049*	0.045*	-0.100	-0.140	-0.182
	(0.024)	(0.021)	(0.022)	(0.103)	(0.104)	(0.105)
× SD, percent	0.79	0.89	0.79	-0.34	-0.48	-0.56
Reinstating	-0.033	-0.037	-0.033	0.093	0.075	0.151
	(0.027)	(0.025)	(0.024)	(0.118)	(0.102)	(0.107)
× SD, percent	-0.52	-0.58	-0.49	0.26	0.21	0.39

Notes: Process against product patents, predetermined 2000–2004 weights. Each panel reports four regressions on one sample: the industry's own displacing and reinstating log patent stocks alone; its suppliers' purchase-weighted stocks alone; both; and both with the customers' sales-weighted stocks (the specification of Table 6; the customers' coefficients are there). Panels A, B, G and H: the main wage sample, worker and year effects; C and D: all workers with a first industry, worker and year effects; E and F: the firm panel, firm and year effects, exposure by the firm's industry. For each term the first row is the coefficient per unit of log patent stock, the second its standard error, and the third the coefficient times the residual standard deviation of that regressor net of the fixed effects and the other regressors of that specification, in percent of the outcome (percentage points for employment and the labor share). Standard errors clustered by IO-121 industry. Significance from the ratio of the rounded estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 22 reports the estimates. Panel A gives the coefficients on the suppliers' stocks of the similarity pair for the four worker outcomes and Panel B those of the process–product pair on the same samples; the last two columns express both in comparable units. Panel C changes one element of the construction at a time, and Panel D reports the two firm outcomes, for which the similarity pair shows no effect through suppliers. The version on all breakthroughs does affect the labor share of buying firms (-0.079, s.e. 0.031, and 0.083, s.e. 0.031), but not their employment.

Table 22: Suppliers' technology typed from the patent text: the similarity pair

	Coefficient		× SD, ×100	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: typed from the patent text (similarity pair)				
Log hourly wage	-0.277** (0.093)	0.275** (0.102)	-0.60	0.65
Employer pay premium	-0.155*** (0.028)	0.138*** (0.028)	-0.34	0.33
Employed	-0.237** (0.075)	0.175* (0.073)	-0.58	0.43
Log annual earnings	-0.738** (0.245)	0.631** (0.233)	-1.82	1.54
Panel B: typed from the first claim (process–product pair), same samples				
Log hourly wage	-0.223* (0.088)	0.251** (0.085)	-0.73	0.74
Employer pay premium	-0.074** (0.025)	0.075** (0.024)	-0.24	0.22
Employed	-0.137** (0.043)	0.106* (0.050)	-0.66	0.41
Log annual earnings	-0.493*** (0.137)	0.507** (0.167)	-2.39	1.97
Panel C: similarity pair, log hourly wage, other constructions				
Longer excerpt of the patent summary	-0.254** (0.088)	0.254** (0.098)	-0.61	0.66
Embedding model Qwen3-Embedding-8B	-0.363* (0.160)	0.320* (0.145)	-0.56	0.58
Embedding model text-embedding-3-small	-0.261** (0.089)	0.264** (0.086)	-0.44	0.50
Closest and farthest third only	-0.129** (0.048)	0.137** (0.051)	-0.52	0.62
Stocks at the floor coded zero	-0.268** (0.085)	0.261** (0.091)	-0.60	0.65
All breakthroughs, no work-relation filter	-0.276** (0.106)	0.264** (0.088)	-0.66	0.75
Weighted by patent importance	-0.126 (0.067)	0.174* (0.077)	-0.26	0.44
Task list of the preceding decade	-0.168 (0.097)	0.179 (0.106)	-0.47	0.55
Embedding model bge-large	-0.146 (0.135)	0.107 (0.096)	-0.32	0.33
Panel D: similarity pair, buying firms				
Log employment	0.140 (0.187)	-0.116 (0.187)	0.39	-0.32
Labor share	-0.007 (0.058)	0.012 (0.058)	-0.02	0.03

Notes: Coefficients on the suppliers' displacing and reinstating patent stocks of the worker's first industry (Panel D: of the firm's industry), entered jointly with the own-industry and customers' stocks of the same pair; worker and year effects (Panel D: firm and year effects). In Panels A, C and D each breakthrough patent is typed by how closely the plain-language summary of its text matches the work tasks that existed when it was issued: its rank in closeness within the issue year is its weight in the displacing stock, and one minus the rank its weight in the reinstating stock; only breakthroughs with a direct relation to work enter (Appendix A.6). Panel B is the process–product pair of Table 9 on the same worker-years. The last two columns multiply each coefficient by the residual standard deviation of its term net of the fixed effects and the other regressors, times 100: percent for the log outcomes, percentage points for employment and the labor share. Panel C changes one element of the construction at a time; the list of constructions was fixed before any estimate was computed. Observations, in millions: hourly wage 47.3, employer premium 47.1, employed 117.9 and annual earnings 100.3 worker-years; firm employment 5.0 and labor share 1.55 firm-years. Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A.7 Suppliers' Exposure by Occupation

Section 5.7 asks whether the pay effect of suppliers' technology is larger in the occupations whose work the patents describe. The measure divides each breakthrough patent among occupations. With the embeddings of Appendix A.6, the unit weight of a patent is spread equally over its 20 nearest tasks on the 2010 task list of Liu et al. (2025), and each task is assigned to the occupation that performs it, through the occupation codes of the task list and a chain of crosswalks to DISCO-08. One task list is used for all patents because the lists differ between decades, and a split on the list in force at issue would move patents between occupations at the decade boundaries; that split is reported as a check, as is one based on the patent–task links published by Liu et al. (2025). Accumulating the divided patents as before gives, for every supplier industry and year, a stock by occupation that sums to the industry's stock.

Let m^T_{jo} denote the share of the 1994 stock of type T held by the suppliers of industry j, weighted by the predetermined network weights, that falls on occupation o, standardized over the estimation sample. With $o(i)$ the occupation in which worker i is first observed, the specification is

$$y_{it} = \alpha_i + \delta_{j(i),t} + \sum_{T \in \{D,R\}} \theta^T \, us^T_{j(i),t} \times m^T_{j(i),o(i)} + \varepsilon_{it}, \tag{12}$$

where the first-industry-by-year effects $\delta_{j(i),t}$ absorb the exposures of equation (8), so that θ^D and θ^R are identified from differences between occupations within a buying industry and year. A first occupation is recorded for 99 percent of the workers of the main sample at two digits and for 96 percent at three digits; almost all first occupations date from before 2010 and are converted with the crosswalk of Appendix A.1. On this sample the wage coefficients of equation (8) are -0.221 (s.e. 0.089) and 0.250 (s.e. 0.085).

Table 23 reports the estimates. Panel A varies the sample, the clustering, the assignment of patents to occupations and the patent pair. Panel B collects the results that limit the interpretation: the overlap measured relative to the occupation's share in all patents, occupation-by-year effects, the occupation held in 2005, and the outcomes other than the hourly wage. Panel C replaces the interaction with an exposure constructed for the occupation itself. Neither form improves on the main estimates: weighting each supplier's stock by the occupation's share in that supplier's 1994 patents gives a displacing coefficient of -0.026 (s.e. 0.012) and an imprecise reinstating one, and the stocks of the patents that describe the occupation, which also use the drift of patenting across occupations over time, have no detectable effect.

Table 23: The pay effect of suppliers' technology and the occupation's overlap with patented technology

	Displacing	Reinstating
Panel A: suppliers' exposure × the occupation's overlap, log hourly wage		
First-industry-by-year effects (main)	-0.056*** (0.010)	0.046*** (0.009)
Worker and year effects only	-0.081*** (0.014)	0.070*** (0.015)
Confidently mapped first occupations	-0.065*** (0.012)	0.058*** (0.013)
Clustered by first industry and by occupation	-0.056* (0.023)	0.046* (0.023)
Patents split on the task list at issue	-0.059*** (0.011)	0.058*** (0.012)
Patents split on published patent–task links	-0.044*** (0.012)	0.030* (0.012)
Three-digit occupations	-0.030*** (0.009)	0.039*** (0.010)
Similarity pair	-0.053*** (0.012)	0.035*** (0.010)
Panel B: scope		
Overlap relative to the occupation's economy-wide share	-0.015 (0.011)	-0.013 (0.020)
The same, similarity pair	-0.084*** (0.016)	0.065*** (0.012)
With first-occupation-by-year effects	-0.026 (0.014)	0.005 (0.012)
Occupation held in 2005	-0.008 (0.011)	-0.004 (0.014)
Outcome: employed	0.024*** (0.007)	-0.030*** (0.008)
Outcome: log annual earnings	-0.012 (0.013)	0.000 (0.020)
Panel C: occupation-specific exposure in place of the interaction, log hourly wage		
Suppliers' stocks weighted by the occupation's 1994 share	-0.026* (0.012)	0.010 (0.024)
Suppliers' stocks of patents describing the occupation	-0.020 (0.028)	-0.033 (0.035)

Notes: Process–product classification unless stated; worker fixed effects and, except in the second row, first-industry-by-year effects, which absorb the suppliers' exposure itself. Panels A and B report the coefficients on the suppliers' displacing and reinstating exposure of the worker's first industry, each interacted with the overlap of the worker's first occupation: the share of the suppliers' 1994 breakthrough stock of that type whose text describes the tasks of the occupation, standardized (Appendix A.7). A coefficient is the change in the effect of one log point of suppliers' stock per standard deviation of overlap. Occupations are two-digit DISCO-08 groups unless stated; confidently mapped first occupations are those whose pre-2010 code sends at least 80 percent of job stayers to one two-digit group (Appendix A.1). Panel C replaces the interaction with an exposure built for the occupation: the suppliers' log stocks, each weighted by the 1994 share of that supplier's patents that describe the occupation relative to the same share among all suppliers, and the suppliers' log stocks of the patents that describe the occupation. Worker-years, in millions: 47.1 in Panels A and C and the first three rows of Panel B, except confidently mapped occupations (24.7) and three-digit occupations (45.6); 2005 occupation 36.5; employed 110.2; annual earnings 96.1. Standard errors clustered by IO-121 first industry unless stated. Network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A.8 Register Sources, Sample, and Outcome Panels

The analysis window runs from 1980 to 2023. The demographic spine is the population register (BEF) from 1986 onward, with the earlier FAIN register covering 1980–1985. Employment status, industry, and occupation are taken from the labor-market registers (AKM, RAS, IDAS); the matched employer-employee register (IDAN) links workers to firms and records hourly wages; annual income comes from the income register (IND), unemployment histories from IDAP, and completed education from UDDA. From 2008 onward the job-spell register (BFL) provides exact start and end dates and hours for each employment spell. In-

dustry codes follow the Danish DB07 classification, which is equivalent to NACE Revision 2; older vintages are bridged through official concordances, so that a consistent industry series is available back to 1980.

A second group of registers supports the mechanism analysis. The ITAV survey records firms' use of information technology and, in its later waves, their use of industrial robots, artificial intelligence, ERP systems, and cloud computing. The firm register (FIRM, 1999–2022) reports each firm's turnover, value added, wage costs, operating result, fixed assets, net investment, and full-time-equivalent employment; the firm outcomes of Section 5.1 are built from these amounts, with the operating result over value added and the share of value added not paid out in wages winsorized at the 1st and 99th percentiles. The employer pay premium is the employer effect from a two-way decomposition of log hourly wages into worker and employer fixed effects (Abowd et al., 1999), estimated on the largest connected set of the matched panel. On the worker side, the national patient register (LPR) records hospital contacts by diagnosis, the benefit registers (KTHJ, SHSS, SYHB) record cash benefits, sickness pay, and related transfers, and the pension register (INPI) records contributions.

The outcome panels are constructed as follows. The labor-market outcomes are log hourly wage (from IDAN), log annual earnings (from IND), an employment indicator, an industry-change indicator, and log hours. Hospital contacts from LPR are grouped by diagnosis, with musculoskeletal and mental-health contacts of particular interest; because the register runs only through 2019 and contact counts are heavily skewed, these outcomes enter as indicators of any contact, and the corresponding estimates apply to that window only. Benefit outcomes enter in two forms, log amounts per recipient and receipt indicators, so that the intensive and extensive margins can be interpreted separately rather than being combined in raw kroner. Pension contributions come from INPI. Labor-force exit, retraining (re-enrollment in education), and geographic moves are constructed from the core registers. Each outcome panel inherits the coverage window of its source register, and the tables state the estimation window for each outcome. A *stayer* indicator, which marks workers observed with a valid industry code in 2005, is used for heterogeneity splits by labor-market attachment.

A.9 The Investment Window

In May 2012 the Danish Parliament passed *Lov nr 1354*, which allowed firms to deduct 115 percent of the cost of new machinery and equipment in the first year. The provision remained in force through December 2013, and its stated purpose was to stimulate business investment during the European debt crisis. Two features make it useful for this paper. First, its timing and motivation were unrelated to industry-specific labor-market conditions: the legislation

applied uniformly across industries and was designed as macroeconomic stimulus rather than as sectoral support. Second, its effective strength varied across industries, since industries that relied more heavily on machinery before the policy stood to gain more from accelerated depreciation.

The policy is nonetheless of limited use for identification. As an instrument for own-industry technology adoption it fails: with only two treatment years and an incentive that covered all machinery rather than automation equipment specifically, the induced first stage is far too weak to be useful (Appendix B.5). An own-industry difference-in-differences design around the window is not clean either, because machinery-intensive industries were on different trajectories through the debt crisis, and its pre-trends are reported as a caveat rather than relied upon. What the policy could support is a customer-side network difference-in-differences: the windfall to machinery-intensive *customers*, propagated through predetermined supply linkages to their suppliers, generates variation in the adoption incentives facing a supplier's customer base that is independent of the supplier's own conditions. With two treatment years and an incentive that covered all machinery, that design has little power, and no result in the paper relies on it.

B The Automation Index: Estimates and Diagnostics

This appendix reports the estimates based on the automation index: the magnitudes, the identification checks, the adjustment margins and the employer pay ladder, the response of buying firms, the results by supplier type, and the incidence across workers. The main text uses the index only as a blunt check. The first subsection explains why, and the explanation qualifies every estimate that follows.

B.1 What the Index Estimates Compare

The index is zero for every industry in 1995–2003. It is a ten-year trailing average of the annual task-content changes of Acemoglu and Restrepo (2022) in the EU KLEMS industry accounts, which begin in 1995; the first complete window ends in 2004, and the years before it are coded zero rather than missing, so they enter every regression. The reinstatement index is built the same way. In later years the index is positive only where the task content of an industry's labor share fell, so many industry-years remain at zero.

Three consequences follow for the estimates below. First, a within-worker regression on 1995–2019 relies mainly on the comparison of a worker's pay from 2004 on with her own pay in 1995–2003, according to the displacement her industry's suppliers underwent over that decade. This is visible when the window is split: the coefficient is −2.24 (s.e. 1.12) on 1995–2004, which contains the first measured year and nine years of zeros, and −0.99 (s.e. 1.12) on 2005–2019. On 2004–2019, the years on which the index is defined, it is −0.74 (s.e. 1.17), and 0.13 (s.e. 0.26) for the employer pay premium; neither is distinguishable from zero. Second, specifications that remove the 2004 comparison (sector-by-year effects, industry-specific trends, or a later start) remove most of the estimate, and I report them as defining the scope of the measure. Third, the index cannot show that suppliers' automation grew faster in one decade than in another; where the text speaks of 1995–2004 it means the decade over which the index is first measured. The patent stocks of the main text are annual from 1980 and are not affected.

B.2 Magnitudes and the Other Two Channels

A one standard deviation increase in suppliers' automation, with standard deviations computed on the estimation sample, lowers hourly wages by 0.91 percent at the column-(3) estimate of Table 24, about three times the same-standard-deviation own-industry association (0.31 percent). The same shift lowers annual earnings by 2.37 percent (Appendix Table 32, column (4)), with an employment effect of −1.73 (s.e. 1.42, $p = 0.23$) that is

Table 24: Technology exposure and log hourly wages: worker-level estimates

	(1)	(2)	(3)	(4)	(5)
Direct automation	-0.830**	-0.547*	-0.495	-0.433	-0.449
	(0.304)	(0.266)	(0.282)	(0.235)	(0.233)
Suppliers' automation		-3.502*	-3.308*	-2.992*	-2.902*
		(1.523)	(1.392)	(1.311)	(1.299)
Customers' automation			-0.891	-0.960	-1.154
			(1.850)	(1.810)	(1.819)
Direct innovation				0.436	0.445
				(0.310)	(0.307)
Suppliers' innovation				-1.550*	-1.598*
				(0.630)	(0.618)
Customers' innovation				-2.310**	-2.460**
				(0.820)	(0.814)
Worker and year FE	Yes	Yes	Yes	Yes	Yes
Tenure control	No	No	No	No	Yes
Observations	47,412,865	47,346,210	47,346,210	47,346,210	44,549,339
Within R^2	0.0003	0.0007	0.0007	0.0013	0.0018

Notes: Worker-year panel; the estimation sample spans 1995–2019, the years the imputed automation measures cover. Dependent variable is log hourly wage. Exposures use predetermined (2000–2004 average) input–output weights; suppliers' (customers') exposure weights the automation of industries the worker's base industry buys from (sells to), computed from the off-diagonal Leontief (Ghosh) inverse. Standard errors in parentheses, clustered by IO-121 base industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

negative but imprecise. Because exposure varies at the industry-year level, these standard deviations describe differences across industries: a worker whose industry's suppliers automate at the 84th percentile loses about one percent of hourly pay relative to a worker at the mean.

The customers' exposure has no detectable effect. This is consistent with the framework: adoption among an industry's customers combines two offsetting forces, and Section 2.3 does not sign their sum. I also draw no causal conclusion from the customers' coefficient, because it fails the placebo test below. The direct coefficient is stable across specifications but fails its own placebo test, and I treat it as descriptive throughout.

An accounting exercise illustrates what a design that considers only the adopting industry leaves out. I compute each industry's predicted wage effect from its base-year exposure twice: once with the own-industry coefficient alone, as an adopter-focused design would, and once with the column-(3) estimates. Of the twenty most exposed industries under the adopter-focused calculation, 10 are not among the twenty most exposed under the network calculation, and the predicted wage effect for the median worker's industry is about 7.2

times as large under the network calculation. The ratio is largest where own adoption is smallest (its covariance with own-industry exposure is negative), so the network calculation is not a rescaling of the adopter-focused one. It assigns predicted losses to industries, air transport among them, that an adopter-focused design would not examine. The exercise applies within-worker coefficients across industries and does not show that either calculation predicts better. It shows which industries the adopter-focused calculation leaves out, given the index estimates.

B.3 Pre-Trends and Dynamics

Table 25 collects the diagnostics that separate the suppliers' channel from the exposures the paper does not interpret causally. The within-worker placebo interacts each worker's *future* exposure with a pre-period indicator. For suppliers' automation the interaction is 0.566 ($p = 0.297$), so workers who would later be exposed were not on different wage paths beforehand. Direct and customers' exposure both fail this test clearly, which is why I treat them as descriptive. Their failure also indicates that the placebo has power to detect pre-existing differences where they exist. At the industry-year cell level, the two leads of suppliers' automation are jointly indistinguishable from zero ($p = 0.379$). The accounts-based reinstatement index fails the same test ($p = 0.030$ and $p = 0.002$), so results for that index are subject to a pre-trend caveat wherever they appear.

Two further results on dynamics are relevant. First, in the distributed lead–lag specification at the worker level the individual lead terms are negative and significant even though the joint cell-level test and the placebo pass; with a persistent, slowly-diffusing regressor the lead terms absorb part of the serially correlated exposure path, and with an index that is zero until 2004 they may also pick up displacement that accumulated before the index first measures it. Figure 4 reports the full profile of leads and lags. The effect builds over three to four years after exposure (from -1.40 one year after exposure to -4.78 at four years), with a contemporaneous coefficient near zero. Second, the five-year cell-level long difference is near zero and insignificant (Panel C). The wage losses thus appear as within-worker changes that accumulate over several years rather than as a permanent five-year change in the ranking of industry-average wages. Cell-level differencing also absorbs the persistent component of exposure that identifies the panel estimate.

B.4 Inference When Shocks Are Few

Shift-share variation of this kind concentrates in a small number of large supplier industries, and I treat that concentration as a limit on the scope of the estimates. Collapsing to

Table 25: Pre-trend diagnostics: placebo, joint leads, and long difference

	Estimate	(s.e.)		
Panel A: within-worker placebo (future exposure × pre-period)				
Direct exposure	5.941***	(1.178)		
Suppliers' exposure	0.566	(0.540)		
Customers' exposure	-52.229***	(11.259)		
Panel B: cell-level joint leads (industry × year, shock-level clustering)				
	Current	Lead 1	Lead 2	Joint p
Suppliers' auto	-1.049	-0.135	-0.788	0.379
Customers' auto	-2.023	-0.942	-1.685	0.185
Suppliers' inno	-0.844	-0.218	-0.818	0.030
Customers' inno	-2.457	-0.550	0.861	0.002
Panel C: five-year long difference (cell level)				
Suppliers' auto	0.192	($p = 0.632$)		
Customers' auto	-0.386	($p = 0.262$)		

Notes: Panel A interacts each worker's future (post-window) exposure with a pre-period indicator in the within-worker specification; a nonzero coefficient indicates differential pre-trends. Panel B regresses industry-year cell wages on current exposure and two leads with shock-level clustering; the joint p tests both leads equal to zero. Panel C is the five-year cell-level long difference. Suppliers' automation passes both lead diagnostics; the long difference is imprecise and near zero (see text). Direct and customers' exposures fail the placebo, and the innovation family fails the joint leads test; these channels carry a pre-trend caveat throughout. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

shock level following Borusyak et al. (2022) leaves roughly 7.1 effective shocks for suppliers' automation, with 70 percent of the identifying leverage in the five largest, led by extraction and refining; at that level of aggregation the automation coefficient is not distinguishable from zero, so I do not rely on many-shock asymptotics for the index estimate. Appendix Table 33 reports the shock-level regressions and Appendix Figure 7 the Rotemberg leverage weights.

Two results support the estimate instead. First, it does not depend on any single shock: re-estimating the cell-level specification with each supplier industry's shock removed in turn keeps the suppliers' automation coefficient strictly negative in all 89 re-estimations (range -2.16 to -1.30 around the full-sample -1.66). Removing extraction, which has the largest Rotemberg weight, strengthens the estimate, and removing refining attenuates it by about a fifth without changing its sign. Second, measures that do not share this concentration give the same result: the automation index (-3.308), observed robot adoption from firm surveys (-0.073, $p = 0.040$; -0.065, $p = 0.115$ with customers' exposure included), and the foreign patent wave below are all associated with lower wages in the buying industry. The

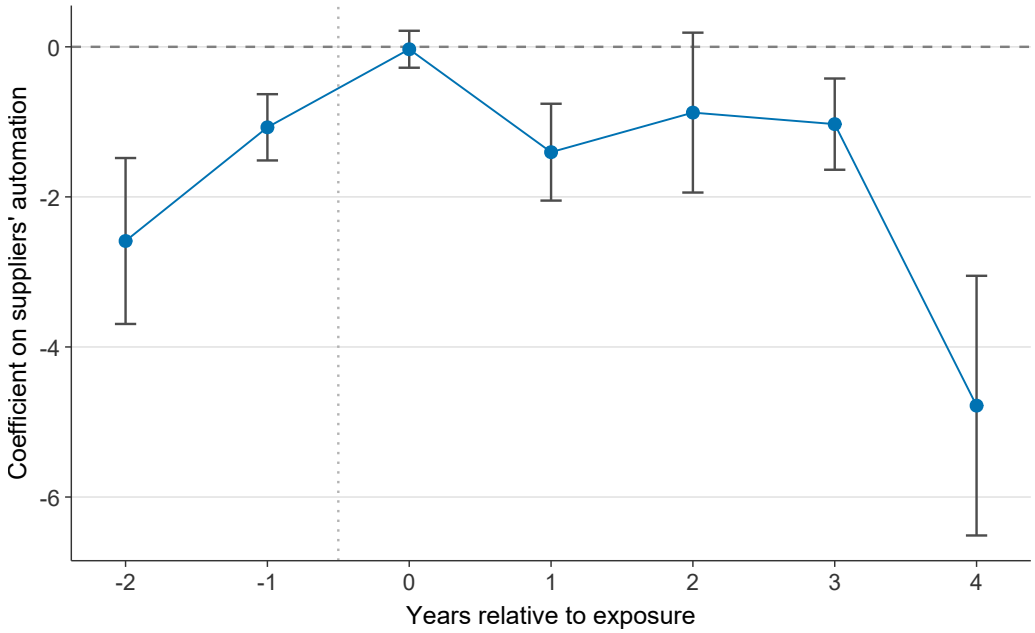


Figure 4: Dynamics of the suppliers'-automation wage effect. Coefficients on leads and lags of suppliers' automation exposure (predetermined weights) in the within-worker specification; whiskers are 95 percent confidence intervals, clustered by IO-121 base industry. The figure includes the lead terms. The cell-level joint test of the leads does not reject zero (Table 25), and the text discusses the interpretation of the individual lead estimates.

concentration reflects the structure of the Danish supplier network, which has hubs, and it is likely to hold for the patent stocks as well.

Two further checks address the same concern from different directions. The first keeps the industry as the cluster and asks whether 109 clusters, a few of which account for most of the leverage, are sufficient for inference. They are: the CV3 jackknife standard error is 1.45 against 1.39, the restricted wild cluster bootstrap gives $p = 0.012$ with a 95 percent interval of $[-6.06, -0.55]$, and dropping one base industry at a time leaves the coefficient between -3.80 (public administration out) and -2.26 (retail out). The effective number of clusters is only 11.5–28.5, with 57 percent of the leverage in five industries, which is why I rely on the bootstrap rather than on the conventional standard error. The second check is design-based. It asks how negative the coefficient would be if the observed adoption paths had landed on different suppliers, with every buyer keeping its purchase weights. Permuting supplier identities 100 times gives coefficients that average -1.84 with a standard deviation of 2.05, and the actual estimate lies well inside that distribution (permutation $p = 0.495$). Part of the estimate therefore comes from how much a buying industry purchases from other industries at all, multiplied by the economy-wide adoption path, rather than from which of its suppliers automated. The purchase share on its own is not the source of the problem: controlling for

each industry's total requirement from other industries changes the contemporaneous-weight estimate only from −3.16 to −2.98. Its product with the common adoption path, however, is not absorbed by year effects. Recentering exposure on its permutation mean, following Borusyak and Hull (2023), removes that component and leaves −2.20 (s.e. 1.32); controlling for the permutation mean instead gives −2.43 (s.e. 1.35). The two checks answer different questions. With the industry clusters as the unit of inference the null of no effect is rejected at the five percent level; with about seven effective shocks, the data cannot distinguish the realized assignment of adoption to suppliers from a random one. With an index whose adoption path is a step from zero in 2004, the common path is a large part of the variation, and this result is what one would expect. It concerns the scope of the estimate: the index identifies the effect of buying from an automating supplier base, not of buying from one automating supplier rather than another.

B.5 The Foreign Patent Wave

Automation patenting outside Denmark is a source of variation that Danish industries cannot influence. I propagate this foreign wave of displacing technology through the fixed domestic network. The predetermined weights rule out re-sorting toward exposed suppliers, and the threat to exclusion is trade exposure correlated with the same wave, which I discuss below. Table 26 reports the results. The reduced form is −0.0127 (s.e. 0.0045): buying industries whose suppliers are exposed to the foreign wave have lower wages. The reduced form does not involve the index, and it is evidence for the displacing side of the patent results. The first stage on the index is strong ($F = 24.3$), though part of that strength is mechanical, because the instrument and the exposure are propagated through the same input–output matrix. The design therefore rests on the exogeneity of patenting outside Denmark rather than on the size of F. Instrumenting suppliers' automation with the foreign wave gives −14.8 with a weak-instrument-robust Anderson–Rubin interval of $[-30.8, -4.4]$. This is the only bounded interval among the paper's instruments, and it is entirely negative. The instruments for own-industry adoption (the accelerated-depreciation window, demographic aging, exchange rates, and leave-out EU adoption) would bear on a causal claim about direct exposure, but they are too weak to use (first-stage F between 0.04 and 2.5). I report them as such and do not instrument the direct coefficient.

The IV point estimate is about 4.5 times the OLS coefficient. I interpret the gap as a local effect: the complier variation is foreign-wave-driven automation among traded-input suppliers, where transmission is plausibly strongest, and classical measurement error in imputed exposure attenuates OLS. The exclusion restriction may fail if the foreign wave also

reaches Danish workers through import competition rather than through suppliers' technology. For this reason I use these estimates to corroborate the main design and not to replace it.

Table 26: Foreign-wave variation in suppliers' automation

	Log hourly wage
Reduced form: foreign displacement wave	-0.013**
	(0.004)
IV: suppliers' automation (instrumented)	-14.818*
	(6.596)
OLS reference (Table 24, col. 3)	-3.308*
	(1.392)
First-stage F	24.349
Anderson–Rubin 95% CI	[−30.786, −4.400]
Observations	46,653,580

Notes: The instrument propagates automation patenting measured outside Denmark (leave-Denmark-out) through the predetermined domestic IO network to the worker's base industry. All rows control for direct and customers' exposure with worker and year fixed effects; standard errors clustered by IO-121 base industry. The Anderson–Rubin interval is weak-instrument robust and is the only bounded interval across the instrument battery; own-adoption instruments (accelerated depreciation window, demographic aging, exchange rates, leave-out EU adoption) have first-stage F below 3 and are not used. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

B.6 Stability and Scope

The suppliers' automation coefficient is stable across measurement, weighting, local-labor-market, trade, and inference choices. The checks that do change it show where the identifying variation comes from, and they define the scope of the measure. Table 27 collects these checks, re-estimating the specification of Table 24, column (3), one variation at a time.

I begin with measurement. Beyond the observed robot adoption of Appendix B.4, the same negative supplier-side transmission appears for a patent-based automation exposure (-0.155, $p = 0.001$) and routine task intensity (-0.289, $p < 0.001$), while ICT capital, the measure most likely to capture complementarity, has no effect and ICT growth enters positively. Turning to the weights, the main specification fixes the input–output structure at its 2000–2004 average; letting it vary contemporaneously gives -3.28, and restricting transmission to linkages two or more steps removed -4.07, all significant at the five percent level (Appendix Table 32). On the predetermined weights themselves, direct suppliers alone give

Table 27: Robustness of the suppliers' automation coefficient

	Suppliers' automation	(s.e.)	N
Baseline (worker FE, cluster IO-121)	-3.752**	(1.234)	50,298,238
Year-specific IO weights	-3.801**	(1.244)	50,231,583
Base year 1990 assignment	-1.782	(2.103)	49,896,246
Winsorize 0.5/99.5	-3.782**	(1.245)	50,298,238
Winsorize 5/95	-2.890**	(0.948)	50,298,238
Two-way clustering (industry, year)	-3.752**	(1.332)	50,298,238
Geographic exposure control	-3.609**	(1.201)	42,541,731
Manufacturing subsample	0.633	(1.078)	8,369,109
Services subsample	-1.939	(1.906)	36,313,885
Post-2000 window	-2.808*	(1.117)	41,326,932

Notes: Each row re-estimates the log hourly wage specification of Table 24, column (3), under the stated variation; the reported coefficient is suppliers' automation exposure. These checks run with contemporaneous weights on the full exposure window (through 2021), which is why its baseline exceeds the main predetermined 1995–2019 estimate of Table 24; the weight-timing comparison on a common window is Table 32. Standard errors clustered by IO-121 base industry except where stated. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

−4.58 (s.e. 2.10) and suppliers' suppliers alone −5.70 (s.e. 2.88); entered together both stay negative (−2.74 and −3.91) but neither is individually significant, since the two exposures are highly correlated and account for about equal shares of the variation. Extending the exposure window through 2021, the last years the source measures cover, gives −3.974 (s.e. 1.450). The within-industry standard deviation of suppliers' automation exposure is 56 percent of its total (82 percent for direct and 62 percent for customers' exposure), so industry effects do not absorb the exposures; given the construction, part of that within-industry variation is the step in 2004. Winsorizing wages at 0.5/99.5 through 5/95 gives coefficients between −3.78 and −2.89, all significant. Two-way clustering by industry and year widens the standard error modestly ($p = 0.009$). Controlling for distance-weighted geographic exposure, so that local labor market spillovers are not attributed to network transmission, leaves −3.61. Comparing only workers who live in the same municipality in the same year changes little: with municipality-by-year fixed effects on the municipality where the worker was first observed the coefficient is −3.17 (s.e. 1.39), and on the current municipality −2.54 (s.e. 1.24). Controls for the buying industry's own capital do not change it either: own ICT capital growth leaves it at −2.83 against −2.83 on the same sample, and industry capital intensity and its growth leave it at −1.79 against −1.82 on the smaller sample covered by the firm register. Nor does trade exposure account for it: with import penetration and offshoring from the Danish input–output tables the coefficient is −3.10 against −3.09 on the same 1996–2019 sample, and on 2007–2019, the years covered by the bilateral trade data,

adding China and low-wage import competition and exports changes the same-window coefficient from −1.54 (s.e. 1.07) to between −1.06 and −1.34 (−1.15 with all controls together).

Assignment and window choices matter more. Assigning workers their 1990 industry rather than their first observed industry attenuates the coefficient to −1.78 (insignificant): an assignment fixed fifteen to thirty years before the identifying variation misclassifies the exposure of every worker who changed industry in between, and the attenuation is what classification error predicts. Restricting to the post-2000 window gives −2.81 ($p = 0.013$). Splitting the window in two, the coefficient is −2.24 (s.e. 1.12) in 1995–2004 and −0.99 (s.e. 1.12) in 2005–2019, as Appendix B.1 explained. The later years also contribute. Starting in 2010 or 2012 and assigning workers the industry they were in at that date gives −1.45 (s.e. 0.69) and −1.22 (s.e. 0.62), about a third of the full-window coefficient. The pandemic years cannot contaminate the coefficient: the index ends in 2019, so no 2020–2021 observation enters the estimation. One further diagnostic is uninformative, and I report it without using it. A coarse cell-level pre-period placebo (regressing 2000–2004 outcomes on later exposure) produces large and unstable positive coefficients. It has the same cell-level aggregation problem as the long difference in Table 25, and the index is zero in those years.

Most of the identifying variation lies between broad sectors. With sector-by-year fixed effects for eleven broad sectors the suppliers' coefficient is −0.54 (s.e. 1.21, $p = 0.655$), and −0.58 when the direct coefficient may also differ by sector; industry-specific linear trends give 0.00 (s.e. 0.30). The standard error barely changes with the sector effects, so this is not a loss of precision. Suppliers' automation as the index measures it is to a large extent a sector-level event: industries in the same sector buy from much the same suppliers, and their measured exposure steps up in 2004 and accumulates almost linearly afterwards, so sector-by-year effects and industry trends absorb most of the treatment along with any confounder. The estimates by technology type keep more of their size under the same sector-by-year effects (Table 10).

The sample splits show where the index effect arises. Within manufacturing the supplier-side coefficient is +0.63 and insignificant on 8.4 million worker-years; the loss is in services and construction (−1.94 and, for the non-tradable definition, −2.57, $p = 0.064$). The pattern reflects where automated inputs are bought rather than made: the workers exposed to suppliers' automation are disproportionately in logistics, distribution, construction, and services that purchase automated components and business services. Within the manufacturing core, where firms both buy and sell automated inputs, the two directions offset each other.

Table 28: Worker adjustment margins

	Direct	Suppliers'	Customers'
Log hourly wage	-0.495	-3.308*	-0.891
	(0.282)	(1.392)	(1.850)
Log annual earnings	-1.689	-8.614*	4.346
	(0.911)	(4.240)	(5.532)
Employed	-0.610	-1.727	1.929
	(0.360)	(1.425)	(1.733)
Labor-force exit	0.125*	-0.030	-0.015
	(0.053)	(0.289)	(0.396)
Retraining	0.040	0.285	0.154
	(0.033)	(0.194)	(0.239)
Unemp. benefit (any)	-0.126	-1.023	0.490
	(0.145)	(0.631)	(0.797)
Log unemp. benefit	-1.346	-7.721	5.103
	(1.457)	(6.438)	(7.499)
Hospital contact (any)	0.242	-0.941	-1.675
	(0.148)	(1.023)	(1.502)
Mental-health contact	0.002	-0.004	0.00011
	(0.002)	(0.012)	(0.011)
Geographic move	0.088	0.194	-0.154
	(0.059)	(0.332)	(0.366)

Notes: Each row is one worker-level regression of the stated outcome on the three automation exposures (specification (3), predetermined weights, worker and year fixed effects). Benefit outcomes are receipt indicators and log amounts; hospital outcomes come from the patient register (LPR, through 2019); unemployment-benefit registers end 2007. Standard errors clustered by IO-121 base industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

B.7 Adjustment Margins and the Employer Pay Ladder

Under the automation index no employment or separation margin responds. Table 28 runs the specification of Table 24, column (3), across the adjustment margins the registers observe. Wages and earnings fall. Labor-force exit, unemployment-benefit receipt (indicator and amount), geographic moves, and hospital contacts, including mental-health contacts, show no significant response to suppliers' automation, and retraining rises only insignificantly. Workers do not change employer more often (-0.18, s.e. 1.24, on an annual rate of 22.7 percent) or move to another industry more often (0.48, s.e. 0.60). Their employers close less often, not more (-3.35, s.e. 1.61), and exit among buying firms does not respond (-0.64, s.e. 0.75). Direct exposure behaves differently: direct automation predicts labor-

force exit, whereas suppliers' automation does not. The patent stocks give a different answer on employment, which falls with suppliers' displacing technology and rises with suppliers' reinstating technology (Section 5.6); the nulls of this table are results for the index and should not be read as results for suppliers' technology in general.

Table 29: Suppliers' automation and the employer pay ladder

	Log hourly wage	Employer pay premium
Panel A: decomposition		
Main sample with an employer premium	-3.31*	-2.06***
	(1.39)	(0.46)
Wage net of the employer premium	-1.25	
	(1.28)	
Panel B: where the loss sits		
Worker × employer effects (within the job)	-0.84	
	(0.55)	
Worker × industry effects	-1.65*	
	(0.71)	
Employer × year effects (against coworkers)	-0.44	
	(0.85)	
Panel C: timing and starting position		
1995–2004	-2.24*	-1.90**
	(1.12)	(0.59)
2005–2019	-0.99	-0.04
	(1.12)	(0.20)
Year effects by first premium decile and first wage quintile	-2.45**	-0.85**
	(0.90)	(0.28)

Notes: Coefficients on the suppliers' automation index in the specification of Table 24, column (3), with worker and year fixed effects unless stated. Panel A: the wage coefficient on the 47.1 million worker-years with an estimated employer pay premium equals the sum of the coefficient on the premium and the coefficient on the wage net of it. Panel B replaces the worker effects with worker-by-employer or worker-by-industry effects, or adds employer-by-year effects. Panel C splits the window and adds year effects by decile of the worker's first employer premium and quintile of her first wage. Standard errors clustered by IO-121 base industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The index does affect the employers at which workers are employed. Table 29 divides the wage coefficient into the pay premium of the worker's employer and the remainder. Of the −3.31 on the sample with an estimated premium, −2.06 (s.e. 0.46) is a lower employer pay premium, and −1.25 (s.e. 1.28), which is not distinguishable from zero, is pay net of the premium: about 62 percent of the loss arises because exposed workers come to work at employers that pay all of their workers less. Panel B examines the margins on which a pay cut in a continuing job would appear and finds no loss there. Within a worker's spell at one employer the coefficient is −0.84 (s.e. 0.55); within a worker's spell in one industry it is −1.65 (s.e. 0.71); and against coworkers at the same employer in the same year it is −0.44 (s.e. 0.85).

The premium is lost when workers change employer. Exposed workers move at ordinary rates, but among all employer changes in a given year the premium falls by more at the move the more exposed the worker's first industry: −1.00 (s.e. 0.25) per unit of suppliers' automation with year effects, and −0.56 (s.e. 0.27) with sector-by-year effects. The comparison is one between industries. With first-industry effects, so that only changes in an industry's exposure over time are used, neither the premium change at a move (0.17, s.e. 0.17) nor the wage change (0.06, s.e. 0.25) responds, and neither does wage growth within the job (−0.05, s.e. 0.09). Figure 5 shows the pattern in levels for the workers observed in both 1995 and 2019. Workers whose first industry was in the quartile with the highest 2019 level of suppliers' automation, 44.1 percent of them in manufacturing, started at high-premium employers: a mean premium of 0.121, with 50.9 percent at employers in the top third. By 2019 their mean premium was 0.103 and 38.5 percent were at top-third employers. Workers in the lowest quartile moved the other way, from 0.071 to 0.100.

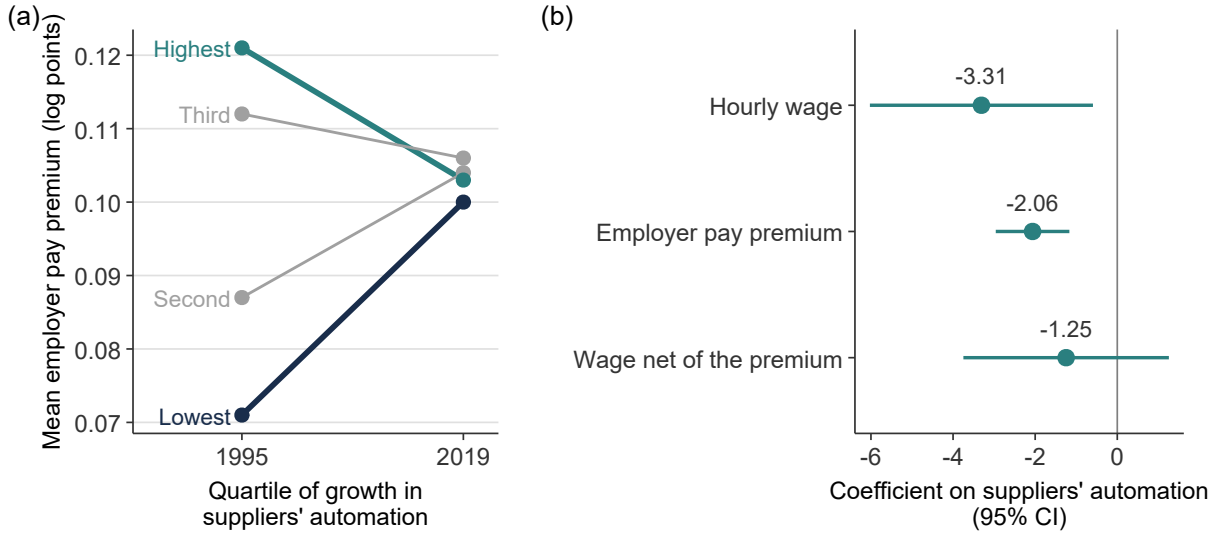


Figure 5: The employer pay ladder under suppliers' automation. Panel (a): mean employer pay premium in 1995 and 2019 for workers observed in both years, by quartile of the 1995–2019 change in suppliers' automation of the worker's first industry; because the index is zero in 1995, the change is the 2019 level. Workers in the highest quartile start at high-premium employers and end near the average; workers in the lowest quartile move up. Panel (b): the coefficient on the suppliers' automation index for the log hourly wage and its two parts, the employer pay premium and the wage net of the premium, with 95 percent confidence intervals (Table 29, Panel A).

At first glance, panel (a) might suggest mechanical convergence: workers who start at high-premium employers can only move down. However, the loss survives a comparison

among workers who started in the same place. With year effects that differ by decile of the worker's first employer premium and by quintile of her first wage, the wage coefficient is −2.45 (s.e. 0.90) and the premium coefficient −0.85 (s.e. 0.28) (Table 29, Panel C). The timing is that of the index: the premium coefficient is −1.90 (s.e. 0.59) on 1995–2004 and −0.04 (s.e. 0.20) after 2005, with a smaller loss reappearing in 2010–2019 (−0.25, s.e. 0.11). The premium loss is not specific to the automation index. Per standard deviation of suppliers' exposure it is 0.57 percent for the automation index, 0.73 percent for the accounts-based reinstatement index and 0.78 percent for ICT capital, and it is negative and significant for routine-task intensity, observed robot adoption, and foreign displacement patenting. The decomposition by technology type is in Section 5.6.

B.8 Buying Firms

The automation index shows no firm response on the firm register. With firm fixed effects and the predetermined exposure, the coefficient is 1.79 (s.e. 3.85) for employment, −1.20 (s.e. 2.34) for value added and −0.77 (s.e. 0.92) for the operating-result share. Neither the separation rate (−1.33, s.e. 1.63) nor the hiring rate (−1.77, s.e. 1.37) responds, and the firm's mean wage does not respond either (0.06, s.e. 1.08; 0.99, s.e. 0.74, among its incumbents). This is what one would expect from a measure that does not separate the two types when the types affect buying firms in opposite directions (Section 5.1). Over the full window the index is associated with a lower labor share at buying firms (−6.31, $p = 0.049$; Table 30, Panel B), but that estimate depends on extreme values: with firm fixed effects it is −12.08 (s.e. 6.15) without winsorizing and 0.27 ($p = 0.756$) once the outcome is winsorized at the 1st and 99th percentiles. I therefore do not rely on it.

B.9 Supplier Type, Timing and Channel Probes

Under the index, too, the loss is concentrated in exposure to capital-good suppliers. Panel A of Table 30 separates suppliers' automation by supplier type: the coefficient for capital-good suppliers is −9.797 (s.e. 3.04) against −2.554 (s.e. 1.18) for other suppliers. For the accounts-based reinstatement index the coefficient for capital-good suppliers is zero and the (smaller) effect comes from the other suppliers. The patent stocks do not share that asymmetry: under them the effects of both types are concentrated in capital-good suppliers (Table 44). Buying firms' capital intensity rises with suppliers' automation in the industry-level specification (38.8, $p = 0.013$; Panel B), and with firm fixed effects and the predetermined exposure the estimate has the same sign but is less precise (7.76, s.e. 4.80, $p = 0.109$). The patent stocks do not show capital deepening at the buyer (Section 5.1). Transmission under the index

Table 30: Mechanism: capital-good suppliers and buying firms' capital

	(1)	(2)
Panel A: worker wages by supplier type (embodied-capital split)		
	Automation	Innovation
Direct	-0.416	0.384
	(0.227)	(0.265)
Suppliers': capital-good	-9.797**	-1.394
	(3.041)	(2.160)
Suppliers': material	-2.554*	-1.957**
	(1.178)	(0.624)
Panel B: buying firms, firm-register outcomes (suppliers' automation)		
	Estimate	(s.e.)
Log fixed assets per full-time equivalent	38.841*	(15.302)
Labor share of value added	-6.313*	(3.174)

Notes: Panel A splits suppliers' exposure by supplier type: capital-good suppliers (NACE 26–30, 41–43) versus all other (material and service) suppliers; each column is one worker-level wage regression with worker and year fixed effects. Panel B: firm-level regressions of buying firms' outcomes, measured in the firm register, on the suppliers' automation index over the full window; the labor-share estimate is sensitive to extreme values (see text). Standard errors clustered by IO-121 industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

also comes mainly from the suppliers with the largest measured displacement. Separating suppliers' automation into the part that comes from the 26 industries in the top fifth of the automation change (45 percent of the exposure) and the part that comes from all other suppliers gives −6.08 (s.e. 2.23) for the first and −0.47 (s.e. 0.91) for the second.

Across industries the pattern of Appendix B.7 is also visible in long differences. Between 1995 and 2019 the mean employer premium of the workers who started in an industry fell with the growth of its suppliers' automation (−4.24, s.e. 0.67) and so did their mean wage (−7.49, s.e. 1.71); as with the worker-level estimates, the relationship is one between sectors and is absent within them (−0.38, s.e. 0.71).

The channel probes reweight the suppliers' automation exposure (Appendix Table 34). Reweighting toward suppliers with differentiated, relationship-specific output, for which switching supplier or buying at arm's length is hardest, leaves the effect as large as in the baseline. Transmission therefore does not depend on the buyer's ability to switch suppliers, which is consistent with technology that is embodied in what an established supplier delivers. One result limits this interpretation. An exposure built from input-profile similarity with no trade linkage (structural equivalence) is also negative and significant, so part of the raw co-movement reflects common shocks to industries with similar input structures,

and the single-channel designs bound the trade-specific component without isolating it. A technological-proximity weighting enters positively, consistent with knowledge spillovers that raise wages, so a general decline of similar industries does not explain the pattern. Worker flows are a second and smaller channel. When input–output linkages and labor-flow linkages are both used to propagate firm investment shocks, both enter (+0.577 and +0.065, each $p < 0.001$), with the input–output linkage the larger of the two, and movers from high-automation firms earn a premium at destination (+0.112, s.e. 0.020) that does not grow with task similarity between origin and destination occupations (interaction −0.001, s.e. 0.013). Network position moderates transmission in the direction that amplification implies, although the evidence is imprecise: every network-exposure interaction that is significant at conventional levels indicates that central industries are affected more, and most of the other estimates are imprecise (Appendix Table 46).

B.10 Who Loses under the Index

Under the automation index the loss falls on workers with basic education, on men, and on the lower-paid. This is the reverse of the gradient under the patent stocks (Section 5.7). Figure 6 reports one regression per group of workers, with groups defined by 2005 characteristics and each regression including its own year effects. The coefficient is −5.09 (s.e. 2.09) for workers with basic education and −1.93 (s.e. 1.01) for workers with vocational education, while for workers with higher education it is positive and not distinguishable from zero (2.06, s.e. 1.93). Men lose (−4.07, s.e. 1.74) and women do not (−0.90, s.e. 1.24). The loss declines up the wage distribution, from −3.65 (s.e. 1.32) in the lowest 2005 wage quintile to −2.11 (s.e. 1.26) in the highest. It also varies with career moves: workers who never change employer lose little (−1.27, s.e. 0.80), workers who change employer within their industry lose more (−2.16, s.e. 0.62), and workers who change industry lose most (−3.78, s.e. 1.58). The loss does not vary systematically with age, tenure, or the size of the 2005 employer.

The part of the loss that is due to the employer pay premium is common to all groups (Appendix Table 37). It lies between −0.98 and −2.25 across education groups, is −1.51 for women and −1.94 for men, and is about the same for workers who change employer within their industry and workers who change industry (−2.25 and −2.03). Estimated separately by type of current occupation, the loss is largest in non-routine manual occupations (−2.79, s.e. 0.67), followed by routine manual occupations (−1.71, s.e. 1.00), and there is no effect in cognitive occupations (−0.43 for routine and −0.40 for non-routine cognitive occupations). The occupation is the current one, so this split, unlike those of Figure 6, conditions on a characteristic that can itself respond.

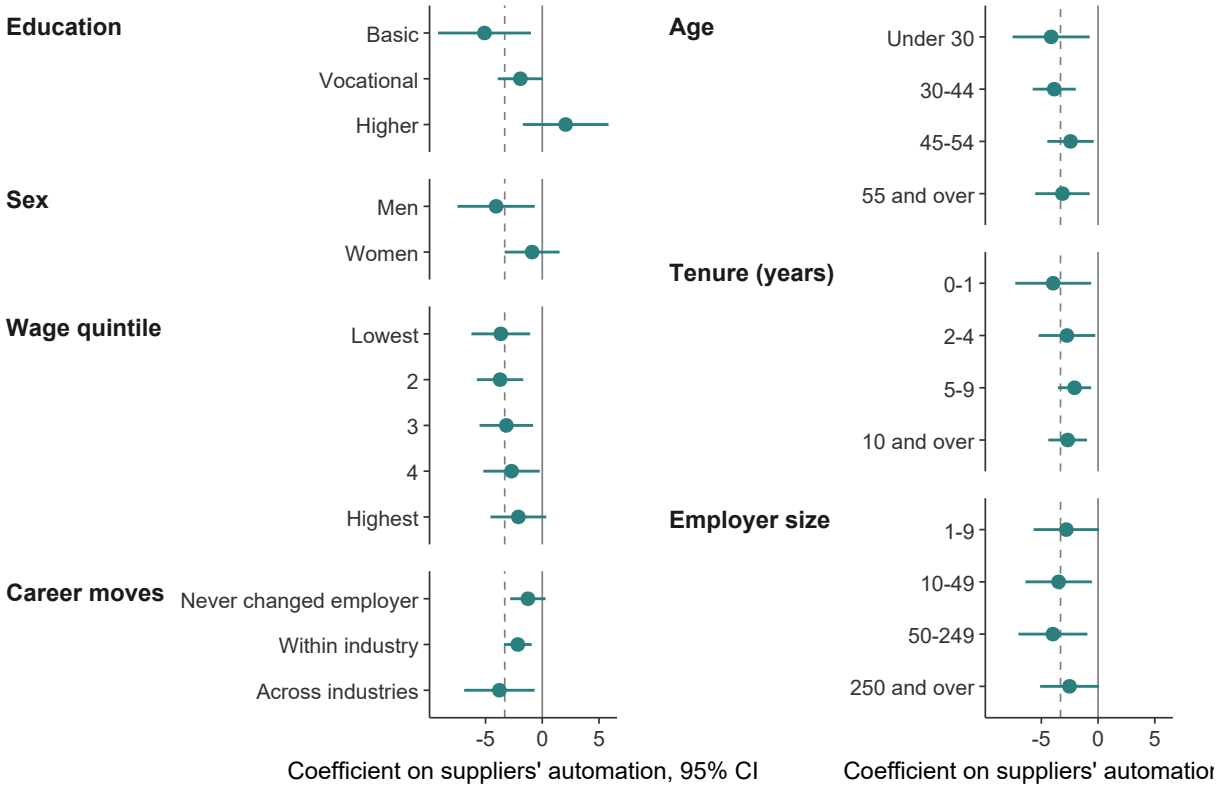


Figure 6: Who loses under the automation index. Each marker is the coefficient on the suppliers' automation index in a regression of log hourly wage on the three automation exposures, estimated separately for the workers in the stated group with worker fixed effects and the group's own year effects; bars are 95 percent confidence intervals, clustered by IO-121 base industry. Groups are defined by 2005 characteristics; career moves by whether the worker ever changes employer or industry. The dashed line is the estimate for all workers. Workers with unrecorded education are omitted. The estimates are in Appendix Table 37. The gradient under the patent stocks is the reverse (Figure 3).

The two measures disagree. A likely reason lies in what the index compares: its estimates rely on how pay changed between 1995–2003 and the years from 2004, and the groups that lose most under it are those whose pay fell behind over that break in the buying industries with the largest measured displacement. The patent stocks vary every year and do not single out that break, and under them the lowest wage fifth shows no effect of either type.

C Additional Tables and Figures

Table 31: Technology Measures: Summary

Measure	Type	Source	Construction	Time-varying	Coverage
Panel A: Automation / Labor-Displacing Measures					
Task displacement	Automation	EU KLEMS	Labor share decomposition into displacement component	Yes (annual)	NACE Rev 2, 1995–2021
Automation probability	Automation	O*NET / Frey and Osborne (2017)	Employment-weighted avg of occupation auto. probabilities	No (cross-section)	702 SOC codes → DISCO → industry
Routine task intensity (RTI)	Automation	O*NET	Routine minus non-routine task content, by occupation → ind.	No (cross-section)	O*NET task measures → DISCO → industry
ICT capital intensity	Automation	EU KLEMS	ICT capital services as share of value added	Yes (annual)	NACE Rev 2, Denmark
ICT capital growth	Automation	EU KLEMS	Annual log change in ICT capital services	Yes (annual)	NACE Rev 2, Denmark
Panel B: Innovation / Labor-Reinstating Measures					
Task reinstatement	Innovation	EU KLEMS	Labor share decomposition, reinstatement component	Yes (annual)	NACE Rev 2, 1995–2021
Emerging tasks	Innovation	O*NET	Count of tasks classified as "emerging" per occ. → ind.	No (cross-section)	O*NET task definitions → DISCO → industry
Occupational diversity	Innovation	Admin data	Shannon entropy of occupation distribution within industry	Yes (annual)	Danish registers, 2000–2019

Notes: This table summarizes the technology measures used in the analysis. "Automation" measures capture the displacement of labor from existing tasks by capital or algorithms. "Innovation" measures capture the creation of new tasks or work activities in which labor has a comparative advantage. Cross-sectional measures are assigned at the base year (2005); time-varying measures exploit within-industry variation over the sample period. The task displacement and reinstatement measures follow the decomposition in Acemoglu and Restrepo (2022). Occupation-based measures are crosswalked from SOC to ISCO-08 to DISCO-08 using official concordances and aggregated to the industry level using employment-weighted averages. See Section 3.3 for details.

Table 32: Robustness to IO-weight timing, and outcome margins

	Log hourly wage			Log earnings	Employed
	Predet.	Contemp.	Dist.≥2	Predet.	Predet.
Direct automation	-0.495	-0.442	-0.687*	-1.689	-0.610
	(0.282)	(0.241)	(0.301)	(0.911)	(0.360)
Suppliers' automation	-3.308*	-3.276**	-4.072*	-8.614*	-1.727
	(1.392)	(1.206)	(1.928)	(4.240)	(1.425)
Customers' automation	-0.891	-1.942	-8.391	4.346	1.929
	(1.850)	(1.809)	(9.548)	(5.532)	(1.733)
Observations	47,346,210	47,412,865	47,346,210	51,558,390	55,768,070

Notes: Columns (1)–(3): log hourly wage with predetermined (2000–2004), contemporaneous, and distance-≥2 (off-diagonal minus first-order) IO weights; column (1) reproduces specification (3) of Table 24. Columns (4)–(5): log annual earnings and an employment indicator under predetermined weights (specification (3)). Estimation window 1995–2019 throughout; worker and year fixed effects; standard errors clustered by IO-121 base industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 33: Shock-level inference for the shift-share exposures

	β (cell)	t (cell)	t (shock)	Effective shocks	Top-5 share
Auto, Suppliers'	-1.657	-1.384	-0.178	7.071	0.695
Auto, Customers'	-3.588	-2.015	-0.284	5.493	0.724
Inno, Suppliers'	-1.653	-2.807	-3.115	2.542	0.891
Inno, Customers'	-1.859	-2.490	-2.838	3.057	0.849

Notes: Cell = IO-121 industry × year; outcome is the cell-mean residualized log wage. Cell columns cluster by industry; the shock column aggregates to shock (supplier-industry × year) level following Borusyak, Hull, and Jaravel (2022) and clusters by shock industry. Effective shocks is the inverse Herfindahl of Rotemberg weights; top-5 share is the leverage of the five largest shocks. The largest weights for suppliers' automation sit in extraction (IO 060000, weight 0.30) and refining (190000, 0.15): identifying variation is concentrated, and the paper treats many-shock asymptotics as unavailable for the automation family (see text).

Table 34: What the network transmits: reweighted supplier exposures

	Network term	(s.e.)	Direct term	(s.e.)
Trade linkage (baseline)	-3.617**	(1.318)	-0.524*	(0.236)
Structural equivalence	-8.352*	(3.884)	-0.680*	(0.309)
Technological proximity	11.887**	(3.716)	-0.870**	(0.322)
Specificity-weighted linkage	-4.556*	(1.889)	-0.569*	(0.272)

Notes: Each row is one worker-level wage regression in which the suppliers' automation exposure is reweighted along one candidate transmission channel: input–output trade linkage (baseline), structural equivalence of input profiles (composition falsification, no trade link), Jaffe patent-class proximity, and Rauch differentiation of the supplier's output (specificity). Worker and year fixed effects; standard errors clustered by IO-121 base industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 35: Displacing versus reinstating technology on the full worker panel: the structural wage targets

	Mann–Püttmann		Process/product	
	Displacing	Reinstating	Displacing	Reinstating
Direct	0.024	-0.010	0.034**	-0.024
	(0.015)	(0.012)	(0.013)	(0.014)
Suppliers'	-0.061***	0.050**	-0.179***	0.148***
	(0.014)	(0.016)	(0.038)	(0.036)
Customers'	-0.013	0.016	-0.004	0.010
	(0.017)	(0.019)	(0.032)	(0.030)

Notes: Worker-level log hourly wage regressions with worker and year fixed effects on the full worker panel over the patent window (85.8 million worker-years); each column pair is one regression using one patent decomposition: Mann–Püttmann automation versus non-automation breakthroughs, and process-dominant versus product-dominant patents. Exposures are log breakthrough-stock measures propagated through the year-t (contemporaneous) IO network; coefficients are per unit of log stock and not comparable in magnitude to the exposure-share measures in Table 24. The main-sample estimates under both network weightings are in Table 9. Standard errors clustered by IO-121 base industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 36: Suppliers' displacing and reinstating technology by worker group

	Log hourly wage		Employer pay premium	
	Displacing	Reinstating	Displacing	Reinstating
2005 wage quintile				
Lowest	-0.005	-0.004	-0.042	0.039
	(0.081)	(0.076)	(0.025)	(0.022)
Second	-0.133*	0.135*	-0.066*	0.064**
	(0.060)	(0.056)	(0.026)	(0.024)
Third	-0.173*	0.171**	-0.077**	0.078**
	(0.072)	(0.065)	(0.027)	(0.025)
Fourth	-0.220**	0.216**	-0.080**	0.080**
	(0.081)	(0.072)	(0.027)	(0.026)
Highest	-0.302**	0.299***	-0.084***	0.081***
	(0.094)	(0.083)	(0.022)	(0.023)
2005 occupation				
Non-routine cognitive	-0.287**	0.327***	-0.068***	0.066***
	(0.106)	(0.093)	(0.018)	(0.017)
Routine cognitive	-0.176	0.187*	-0.057**	0.047**
	(0.093)	(0.084)	(0.018)	(0.015)
Non-routine manual	-0.115	0.178*	-0.047	0.051*
	(0.073)	(0.073)	(0.028)	(0.026)
Routine manual	-0.033	-0.009	-0.027	0.002
	(0.104)	(0.107)	(0.038)	(0.034)
Agricultural	-0.217*	0.265*	-0.111***	0.152***
	(0.096)	(0.113)	(0.026)	(0.029)
Other	-0.153	0.202	-0.067*	0.069*
	(0.103)	(0.110)	(0.029)	(0.033)
Education				
Basic	-0.223	0.277*	-0.068*	0.074*
	(0.117)	(0.134)	(0.030)	(0.032)
Vocational	-0.168**	0.160**	-0.059**	0.052**
	(0.063)	(0.055)	(0.020)	(0.018)
Higher	-0.167	0.245*	-0.055***	0.054***
	(0.110)	(0.096)	(0.017)	(0.016)
Not recorded	-0.059	0.083	-0.066*	0.073*
	(0.070)	(0.087)	(0.029)	(0.033)
Sex				
Men	-0.247*	0.251*	-0.076**	0.074**
	(0.118)	(0.120)	(0.027)	(0.027)
Women	-0.160*	0.194**	-0.046*	0.037*
	(0.070)	(0.061)	(0.019)	(0.018)

Notes: Each row is one pair of regressions (wage, premium) on the main-sample workers in the stated group, with worker fixed effects and the group's own year effects; the reported coefficients are those on the suppliers' displacing and reinstating patent stocks of the process–product classification. Groups are defined by 2005 characteristics, the occupation groups by the worker's 2005 occupation. Worker-years in the wage regression, in millions, by panel: 2005 wage quintile: 7.4 (row 1), 8.9 (row 2), 9.1 (row 3), 9.7 (rows 4–5); 2005 occupation: 17.2 (row 1), 4.7 (row 2), 8.9 (row 3), 7.2 (row 4), 0.2 (row 5), 6.6 (row 6); Education: 10.4 (row 1), 19.3 (row 2), 14.8 (row 3), 0.2 (row 4); Sex: 23.0 (row 1), 21.8 (row 2). Standard errors clustered by IO-121 base industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 37: Suppliers' automation by worker group: wages and employer pay premia

	Log hourly wage		Employer pay premium	
	Estimate	(s.e.)	Estimate	(s.e.)
Education				
Basic	-5.09*	(2.09)	-2.25***	(0.53)
Vocational	-1.93	(1.01)	-1.52***	(0.39)
Higher	2.06	(1.93)	-0.98***	(0.26)
Sex				
Men	-4.07*	(1.74)	-1.94***	(0.49)
Women	-0.90	(1.24)	-1.51***	(0.31)
Wage quintile				
Lowest	-3.65**	(1.32)	-1.77***	(0.37)
Second	-3.72**	(1.04)	-2.13***	(0.44)
Third	-3.17*	(1.20)	-2.22***	(0.49)
Fourth	-2.71*	(1.27)	-2.18***	(0.54)
Highest	-2.11	(1.26)	-1.71***	(0.41)
Career moves				
Never changed employer	-1.27	(0.80)		
Changed employer within the industry	-2.16**	(0.62)	-2.25***	(0.64)
Changed industry	-3.78*	(1.58)	-2.03***	(0.45)
Age				
Under 30	-4.14*	(1.73)		
30–44	-3.86***	(0.97)		
45–54	-2.43*	(1.04)		
55 and over	-3.14*	(1.22)		
Tenure (years)				
0–1	-3.96*	(1.70)		
2–4	-2.75*	(1.27)		
5–9	-2.07**	(0.75)		
10 and over	-2.68**	(0.87)		
Employer size				
1–9	-2.80	(1.47)		
10–49	-3.47*	(1.49)		
50–249	-3.98*	(1.55)		
250 and over	-2.52	(1.32)		

Notes: Each row is one regression on the workers in the stated group, with worker fixed effects and the group's own year effects; the reported coefficient is the suppliers' automation index. Groups are defined by 2005 characteristics, and career moves by whether the worker ever changes employer or industry over the sample. The premium columns exist for the groups where the split was run; a worker who never changes employer has no within-worker variation in the premium. Standard errors clustered by IO-121 base industry; significance in the premium columns from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 38: Suppliers' displacing and reinstating technology and buying firms: all outcomes

	Employees	FTE	Value added	Wage costs	Op. result	Op. result / VA	(VA − wages) / VA	Mean wage	Assets / FTE
Panel A: firm and year fixed effects, predetermined weights									
Process/product, displacing	-0.119	-0.115	0.019	-0.084	0.179	0.056*	0.055	0.014	0.092
	(0.070)	(0.166)	(0.144)	(0.118)	(0.169)	(0.028)	(0.030)	(0.015)	(0.205)
Process/product, reinstating	0.133*	0.140	0.008	0.116	-0.168	-0.058*	-0.057*	-0.015	-0.110
	(0.064)	(0.157)	(0.133)	(0.110)	(0.160)	(0.027)	(0.029)	(0.014)	(0.199)
Automation/other, displacing	-0.083	-0.083	-0.043	-0.039	0.007	0.006	0.008	0.011	0.070
	(0.045)	(0.052)	(0.033)	(0.026)	(0.048)	(0.007)	(0.007)	(0.008)	(0.059)
Automation/other, reinstating	0.104*	0.123*	0.078*	0.082**	-0.002	-0.014	-0.015	-0.014	-0.106*
	(0.044)	(0.057)	(0.034)	(0.029)	(0.050)	(0.008)	(0.008)	(0.008)	(0.053)
Panel B: firm and year fixed effects, contemporaneous weights									
Process/product, displacing	-0.094	-0.116	0.086	-0.011	0.288	0.059*	0.053	0.023	0.091
	(0.074)	(0.147)	(0.144)	(0.116)	(0.167)	(0.026)	(0.028)	(0.016)	(0.142)
Process/product, reinstating	0.106	0.128	-0.069	0.031	-0.274	-0.059*	-0.053*	-0.024	-0.098
	(0.069)	(0.135)	(0.134)	(0.108)	(0.159)	(0.025)	(0.027)	(0.015)	(0.137)
Automation/other, displacing	-0.103**	-0.127**	-0.051	-0.044	0.002	0.008	0.011	0.009	0.071
	(0.032)	(0.042)	(0.035)	(0.031)	(0.049)	(0.009)	(0.008)	(0.010)	(0.042)
Automation/other, reinstating	0.121***	0.159***	0.071	0.071*	-0.003	-0.014	-0.016	-0.013	-0.096*
	(0.032)	(0.045)	(0.038)	(0.035)	(0.053)	(0.009)	(0.009)	(0.009)	(0.039)
Panel C: industry and year fixed effects, predetermined weights									
Process/product, displacing	-0.397**	-0.453	-0.657***	-0.666***	-0.146	0.195**	0.153*	0.100	-0.588
	(0.130)	(0.355)	(0.142)	(0.148)	(0.280)	(0.069)	(0.061)	(0.072)	(0.329)
Process/product, reinstating	0.535***	0.624	0.761***	0.749***	0.075	-0.178*	-0.145*	-0.189	0.175
	(0.150)	(0.400)	(0.144)	(0.186)	(0.274)	(0.070)	(0.065)	(0.104)	(0.273)
Automation/other, displacing	-0.056	-0.052	-0.119	-0.180	-0.137	0.061	0.057	-0.016	-0.259
	(0.067)	(0.126)	(0.136)	(0.116)	(0.137)	(0.038)	(0.034)	(0.021)	(0.161)
Automation/other, reinstating	0.157	0.099	0.233**	0.207*	-0.016	-0.081**	-0.080**	-0.015	-0.058
	(0.103)	(0.188)	(0.086)	(0.098)	(0.099)	(0.026)	(0.025)	(0.035)	(0.118)
Panel D: industry and year fixed effects, contemporaneous weights									
Process/product, displacing	-0.261*	-0.482	-0.437**	-0.439***	0.143	0.189**	0.165**	0.074	-0.168
	(0.112)	(0.309)	(0.164)	(0.126)	(0.228)	(0.060)	(0.056)	(0.064)	(0.250)
Process/product, reinstating	0.262*	0.481	0.391*	0.420***	-0.170	-0.179**	-0.152**	-0.083	0.180
	(0.102)	(0.262)	(0.152)	(0.114)	(0.221)	(0.057)	(0.053)	(0.061)	(0.236)
Automation/other, displacing	-0.105	-0.167*	-0.122	-0.070	-0.030	0.036	0.041	-0.020	-0.006
	(0.060)	(0.085)	(0.068)	(0.078)	(0.069)	(0.022)	(0.021)	(0.025)	(0.099)
Automation/other, reinstating	0.151*	0.246**	0.134	0.099	-0.026	-0.057*	-0.056*	0.019	0.031
	(0.070)	(0.089)	(0.073)	(0.091)	(0.075)	(0.027)	(0.026)	(0.025)	(0.109)

Notes: The specifications of Table 3 for every firm outcome. Employees, full-time equivalents (FTE), value added, wage costs, the operating result and fixed assets per FTE are in logs; the mean wage is the firm's mean log hourly wage; the two ratios are winsorized at the 1st and 99th percentiles. Standard errors clustered by IO-121 industry; significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 39: Inside the buying firm by technology type: all outcomes, firm and year fixed effects

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Mean age (years)	1.34*	-1.31*	0.63	-0.66*
	(0.67)	(0.63)	(0.34)	(0.32)
Share routine occupations	0.107*	-0.086*	0.047	-0.030
	(0.047)	(0.041)	(0.029)	(0.023)
Share manual occupations	0.035	-0.021	0.011	0.005
	(0.044)	(0.040)	(0.019)	(0.016)
Share basic education	-0.023	0.023	-0.037	0.036
	(0.023)	(0.020)	(0.022)	(0.019)
Share higher education	0.019	-0.022	0.003	-0.009
	(0.013)	(0.012)	(0.007)	(0.006)
Share women	0.023	-0.028	-0.002	-0.005
	(0.021)	(0.019)	(0.008)	(0.008)
Investment rate	-0.053*	0.052*	-0.001	0.004
	(0.022)	(0.021)	(0.008)	(0.009)
Labor share of value added	-0.055	0.057	-0.008	0.015*
	(0.030)	(0.029)	(0.007)	(0.008)
Purchases over turnover	-0.005	0.005	-0.006	0.008
	(0.028)	(0.028)	(0.008)	(0.009)
Log fixed assets per FTE	0.092	-0.110	0.070	-0.106*
	(0.206)	(0.199)	(0.059)	(0.053)
Log turnover	-0.207	0.213	0.003	0.020
	(0.132)	(0.126)	(0.055)	(0.056)
Log gross profit	-0.111	0.129	-0.037	0.070
	(0.127)	(0.116)	(0.034)	(0.037)

Notes: The specifications of Table 5 for every outcome estimated, with firm and year fixed effects. Workforce shares are computed over all workers of the firm; manual occupations are routine and non-routine manual. Purchases over turnover and gross profit are from the firm register. Firm-years in the process–product regression, in millions: 5.0 (rows 1–6), 1.5 (row 7), 1.6 (row 8), 1.7 (row 9), 1.4 (row 10), 2.0 (row 11), 1.6 (row 12). Standard errors clustered by IO-121 industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 40: Inside the buying firm by technology type: all outcomes, industry and year fixed effects

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Mean age (years)	3.82*	-6.66**	-0.44	-1.64*
	(1.94)	(2.55)	(0.60)	(0.78)
Share routine occupations	0.301*	-0.287	0.118**	-0.116*
	(0.129)	(0.181)	(0.037)	(0.054)
Share manual occupations	-0.118	0.368***	0.064	0.101**
	(0.089)	(0.101)	(0.037)	(0.031)
Share basic education	0.059	-0.072	-0.030	0.007
	(0.037)	(0.051)	(0.025)	(0.025)
Share higher education	0.029	-0.078	-0.043*	0.019
	(0.044)	(0.055)	(0.022)	(0.024)
Share women	0.060	-0.086	0.030	-0.034
	(0.068)	(0.091)	(0.020)	(0.026)
Investment rate	-0.139**	0.129**	-0.027	0.011
	(0.046)	(0.049)	(0.035)	(0.025)
Labor share of value added	-0.153*	0.145*	-0.057	0.080**
	(0.061)	(0.065)	(0.034)	(0.025)
Purchases over turnover	-0.120*	0.051	-0.057	0.023
	(0.053)	(0.045)	(0.033)	(0.029)
Log fixed assets per FTE	-0.588	0.175	-0.259	-0.058
	(0.329)	(0.273)	(0.161)	(0.118)
Log turnover	-0.591	0.603	0.144	0.006
	(0.411)	(0.451)	(0.241)	(0.237)
Log gross profit	-0.991***	0.984***	-0.191	0.246
	(0.198)	(0.222)	(0.151)	(0.131)

Notes: The specifications of Table 5 for every outcome estimated, with industry and year fixed effects. Workforce shares are computed over all workers of the firm; manual occupations are routine and non-routine manual. Purchases over turnover and gross profit are from the firm register. Firm-years in the process–product regression, in millions: 5.0 (rows 1–6), 1.5 (row 7), 1.6 (row 8), 1.7 (row 9), 1.4 (row 10), 2.0 (row 11), 1.6 (row 12). Standard errors clustered by IO-121 industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 41: Buying firms' employment, hiring and separations by the firm's own pay premium

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Lowest third of the pay premium				
Log employment	-0.119*	0.114*	-0.031	0.035
	(0.053)	(0.050)	(0.042)	(0.040)
Hiring rate	-0.001	-0.002	-0.036*	0.028*
	(0.036)	(0.032)	(0.015)	(0.014)
Separation rate	0.017	-0.016	-0.024	0.020
	(0.032)	(0.029)	(0.021)	(0.019)
Middle third				
Log employment	-0.147	0.161*	-0.083	0.106*
	(0.079)	(0.071)	(0.055)	(0.052)
Hiring rate	-0.035	0.029	-0.043**	0.035**
	(0.028)	(0.025)	(0.014)	(0.013)
Separation rate	0.024	-0.022	-0.013	0.013
	(0.025)	(0.022)	(0.020)	(0.018)
Highest third				
Log employment	-0.110	0.142	-0.123**	0.158***
	(0.083)	(0.078)	(0.044)	(0.045)
Hiring rate	0.008	-0.007	-0.019	0.015
	(0.031)	(0.027)	(0.012)	(0.011)
Separation rate	0.020	-0.016	-0.008	0.011
	(0.024)	(0.021)	(0.016)	(0.014)

Notes: The first three rows of Table 5, estimated separately for buying firms in each third of the distribution of estimated employer pay premia, with firm and year fixed effects. Firm-years in the process–product regression, in millions, by panel: Lowest third of the pay premium: 1.4 (row 1), 1.3 (row 2), 1.4 (row 3); Middle third, 1.7; Highest third, 1.4. Standard errors clustered by IO-121 industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 42: Where the identifying variation lives: the same regression under four sets of fixed effects

	Process vs product		Automation vs other	
	Displacing	Reinstating	Displacing	Reinstating
Panel A: log hourly wage, suppliers' stocks				
Year effects	0.214*	-0.207*	0.067*	-0.073*
	(0.104)	(0.095)	(0.031)	(0.030)
Year effects and 2005 age group, sex and education	0.058	-0.061	0.015	-0.022
	(0.054)	(0.051)	(0.018)	(0.018)
First-industry and year effects	-0.217**	0.237**	-0.028	0.066**
	(0.069)	(0.074)	(0.019)	(0.023)
Worker and year effects	-0.223*	0.251**	-0.029	0.084**
	(0.088)	(0.085)	(0.023)	(0.026)
Panel B: log hourly wage, customers' stocks				
Year effects	-0.150***	0.141***	0.006	0.008
	(0.040)	(0.037)	(0.021)	(0.022)
Year effects and 2005 age group, sex and education	-0.075*	0.073*	0.018	-0.010
	(0.037)	(0.033)	(0.016)	(0.016)
First-industry and year effects	0.020	-0.098*	-0.054*	0.005
	(0.048)	(0.040)	(0.021)	(0.022)
Worker and year effects	0.046	-0.121**	-0.040	-0.012
	(0.055)	(0.046)	(0.026)	(0.025)
Panel C: employed (all workers), suppliers' stocks				
Year effects	-0.047	0.045	0.013	-0.009
	(0.034)	(0.032)	(0.015)	(0.013)
Year effects and 2005 age group, sex and education	-0.078**	0.073**	0.000	0.002
	(0.029)	(0.028)	(0.012)	(0.011)
First-industry and year effects	-0.040	0.032	0.007	-0.024
	(0.032)	(0.039)	(0.012)	(0.013)
Worker and year effects	-0.137**	0.106*	-0.052***	0.026
	(0.043)	(0.050)	(0.012)	(0.018)

Notes: Each row is one regression per classification of the stated outcome on the industry's own, its suppliers' and its customers' displacing and reinstating stocks, with the stated fixed effects; the own stocks are not shown. Panels A and B: the main wage sample (47.3 million worker-years); Panel C: all workers with a first industry (117.9 million). The 2005 demographics are the worker's age group, sex and education in 2005. Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 43: Exits from employment by the worker's 2005 wage fifth

	Process vs product		
	Displacing	Reinstating	Mean
Lowest fifth	0.051**	-0.051*	0.076
	(0.017)	(0.021)	
Second	0.062***	-0.071***	0.048
	(0.016)	(0.021)	
Third	0.068***	-0.080***	0.036
	(0.019)	(0.024)	
Fourth	0.066**	-0.080**	0.030
	(0.020)	(0.026)	
Highest fifth	0.072***	-0.094***	0.027
	(0.018)	(0.023)	

Notes: The probability of not being employed this year among workers employed last year, by the fifth of the 2005 hourly-wage distribution the worker was in, one regression per fifth with worker and year effects and the industry's own and customers' stocks; coefficients on the suppliers' displacing and reinstating stocks of the worker's first industry, process–product classification. The mean is the exit rate of the group. Worker-years, in millions: 8.7 (row 1), 11.9 (row 2), 12.5 (row 3), 13.5 (row 4), 13.6 (row 5). Standard errors clustered by IO-121 first industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

D The Model: Derivations and Identification

This appendix presents the model of Section 2 in full: the environment, the derivation of each closed-form expression, the fourteen moments and the accounting identities that link them, and the identification argument of Section 2.4.

D.1 Environment

The model describes a single buying industry, and I omit the industry subscript. The industry uses two inputs: labor L and a manufactured input M purchased from supplier industries. The industry's own capital is omitted because the buying firms' investment rate and capital per worker do not respond to the direction of the suppliers' innovation, and because the effect is transmitted through purchased inputs rather than through equipment on the buyers' balance sheets (Section 5.8).

Table 44: Suppliers' technology by type of supplier: capital-good and other suppliers

	Log hourly wage		Employer pay premium	
	Process vs product	Automation vs other	Process vs product	Automation vs other
Capital-good suppliers				
Displacing	-0.638**	-0.208*	-0.260***	-0.034
	(0.207)	(0.090)	(0.067)	(0.024)
Reinstating	0.517	0.178**	0.285***	0.038**
	(0.266)	(0.065)	(0.084)	(0.014)
Other suppliers				
Displacing	-0.143	-0.004	-0.046	-0.056***
	(0.095)	(0.031)	(0.027)	(0.008)
Reinstating	0.190*	0.062*	0.050*	0.045***
	(0.086)	(0.029)	(0.025)	(0.007)
Worker-years	47,346,210	45,620,372	47,079,295	45,362,110

Notes: Each column is one worker-level regression with worker and year fixed effects on the main sample. The suppliers' displacing and reinstating exposures are each split into the part that comes from capital-good suppliers (industries producing electronics, electrical equipment, machinery, vehicles and other transport equipment, and construction; NACE 26–30 and 41–43) and the part that comes from all other suppliers; the four parts enter jointly with the own-industry and customers' stocks. The two parts sum to the suppliers' exposure of Table 9 (correlation 0.9999 with the exposure used there). Standard errors clustered by IO-121 base industry. Coefficients are per unit of log patent stock; network weights are predetermined (2000–2004 average). Significance from the ratio of the estimate to its standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 45: Industry-level employment estimates (structural targets)

	(1) Direct	(2) + Suppliers'	(3) + Customers'	(4) + Innovation
Direct automation	-4.463	-4.480	-5.795**	-5.836*
	(2.666)	(2.455)	(2.180)	(2.452)
Suppliers' automation		0.494	0.079	-1.247
		(6.859)	(7.073)	(10.227)
Customers' automation			12.303	17.368*
			(8.165)	(7.780)
Direct innovation				-1.492
				(1.962)
Suppliers' innovation				-2.649
				(6.430)
Customers' innovation				12.190*
				(6.023)

Notes: Industry-year panel of IO-121 industries; dependent variable is log employment. Column headers follow the progressive design of Table 24 at the industry level; industry and year fixed effects throughout, standard errors clustered by industry. These six automation coefficients (and the innovation coefficients of column (4)) are the industry-employment moments of the structural estimation in Section 6. Industry-level estimates cannot separate within-worker wage changes from composition, which is why they serve as structural targets rather than primary evidence; see the compositional reading in Section 6. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 46: Network amplification: exposure × centrality interactions

	Betweenness	Eigenvector	Katz
Panel A: Log employment			
Direct × centrality	-13.382	22.867	1.551
	(68.225)	(17.215)	(1.019)
Suppliers' × centrality	407.434	-152.348**	-11.646**
	(320.915)	(49.676)	(3.920)
Customers' × centrality	-346.324	-22.817	-1.201
	(322.210)	(119.746)	(8.121)
Panel B: Mean log wage			
Direct × centrality	-5.054	9.785**	0.651**
	(21.858)	(3.450)	(0.214)
Suppliers' × centrality	-55.276*	2.840	0.284
	(26.680)	(5.720)	(0.463)
Customers' × centrality	-8.776	-87.232	-6.256*
	(35.848)	(44.757)	(2.912)

Notes: Industry-year panel of IO-121 industries (N = 2,727). Each column is one regression per panel of the outcome on the three automation exposures and their interactions with the stated centrality measure, with industry and year fixed effects; the centrality main effect is time-invariant and absorbed. Centrality is computed from the predetermined (2000–2004 average) IO table. Katz centrality is a mathematical cousin of the Leontief inverse and carries borderline variance inflation; upstreamness and downstreamness interactions are excluded as near-collinear with the exposures (Section 4.2). Standard errors in parentheses, clustered by industry. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Tasks. Output is a task aggregate over a unit measure of tasks $z \in [N-1, N]$ with elasticity of substitution $\sigma > 0$,

$$y = \Big[\int_{N-1}^{N} q(z)^{\frac{\sigma-1}{\sigma}} \, dz \Big]^{\frac{\sigma}{\sigma-1}},$$

where $q(z) = \gamma_M(z)\, m(z)$ if the input performs task z and $q(z) = \gamma_L(z)\, l(z)$ if labor does. Labor has a comparative advantage in higher tasks: $\gamma_L(z)/\gamma_M(z)$ is increasing in z, with slope $\rho(z) = d \ln[\gamma_L(z)/\gamma_M(z)]/dz > 0$.

Prices and assignment. The input costs p_M per unit and labor w. The unit cost of task z is $c_M(z) = p_M/\gamma_M(z)$ if the input performs it and $c_L(z) = w/\gamma_L(z)$ if labor does. A task is performed by whichever is cheaper, so by comparative advantage the input performs $z \leq I$ and labor $z > I$, where the boundary $I \in (N-1, N)$ solves $p_M/\gamma_M(I) = w/\gamma_L(I)$.

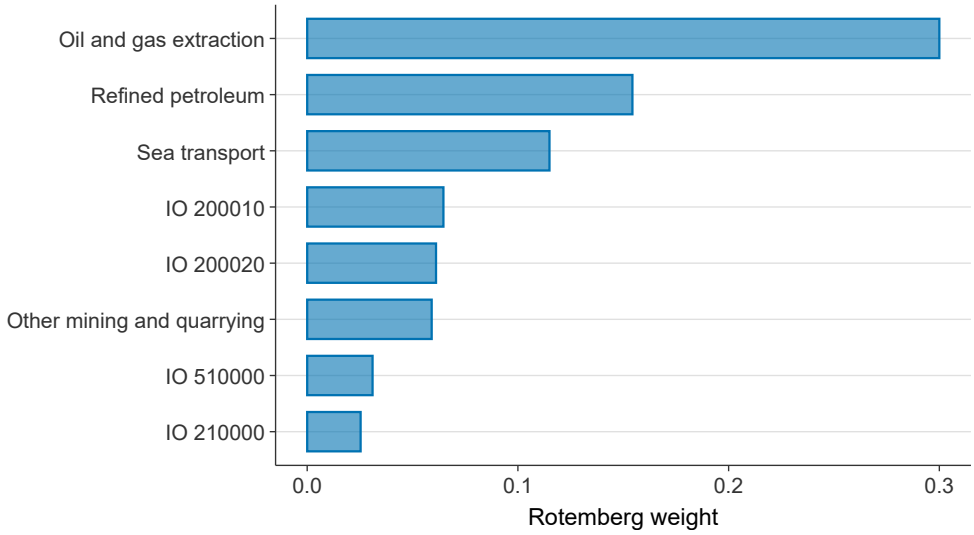


Figure 7: Rotemberg weights for suppliers' automation exposure: the eight largest shock industries. The weights measure each supplier industry's leverage over the shift-share estimate; extraction and refining together account for just under half of the total. The leave-one-shock-out re-estimation discussed in Appendix B.4 keeps the suppliers' coefficient negative in all 89 reruns.

Unit cost and shares. With constant returns,

$$C = \left[\int_{N-1}^{I} c_M(z)^{1-\sigma}\, dz + \int_{I}^{N} c_L(z)^{1-\sigma}\, dz \right]^{\frac{1}{1-\sigma}},$$

the cost-share density of task z is $s(z) = (c(z)/C)^{1-\sigma}$, the input's share of cost is $s_M = \int_{N-1}^{I} s(z)\, dz$, labor's is $s_L = 1 - s_M$, and $h = s(I)$ is the density at the boundary.

Demand, price and markup. The industry faces $y = D\, p^{-\eta}$ with $\eta > 0$ and sells at a markup over unit cost, $p = \mu C$. Pass-through of cost into the price may be incomplete,

$$d \ln p = \varphi\, d \ln C, \qquad d \ln \mu = -(1-\varphi)\, d \ln C, \qquad \varphi \in (0, 1], \tag{13}$$

which is the local form of any pricing rule with variable markups; $\varphi = 1$ corresponds to a constant markup. Turnover is $R = py$.

Labor supply. The industry's wage increases with its employment along a wage curve, $\ln w = \ln \bar{w} + \varepsilon^{-1} \ln L$. The wage curve is a reduced form for the wage differences between industries that the data identify, since the wage coefficient is estimated from variation across workers' first industries and years.

Accounting. The wage bill, purchases, value added and the operating surplus are

$$wL = \frac{s_L}{\mu}R, \qquad p_M M = \frac{s_M}{\mu}R, \qquad \text{VA} = R\Big(1 - \frac{s_M}{\mu}\Big), \qquad \text{VA} - wL = R\Big(1 - \frac{1}{\mu}\Big),$$

which imply the ratios of equation (1). In the firm accounts, purchases average 0.527 of turnover and the margin averages 0.124 of turnover, which imply $\mu = 1.14$ and $s_M = 0.60$. These are the calibrated values used in the estimation.

The shock. The direction of the suppliers' innovation shifts toward process by one unit (Section 2.3 defines the unit). The shift has two effects on the buyer. First, the input's productivity in the tasks it performs rises by $g \geq 0$, that is, $d\ln\gamma_M(z) = g$ for z in $[N-1, I]$ and in a neighbourhood of I, at an unchanged p_M. The input's cost per task therefore falls by g and the boundary shifts. Second, the buyer's demand shifter falls by $\delta \geq 0$, that is, $d\ln D = -\delta$. The demand shift represents the part of the buyer's sales that, after the suppliers' process innovation, is produced by the supplier or by a third party. The capability gain alone ($\delta = 0$) and the demand shift alone ($g = 0$) are the limiting cases. A shift toward product innovation is the same shock with the opposite sign.

D.2 Derivations

All derivatives are per unit of the shock and are evaluated at the initial equilibrium, with the wage adjusting along the wage curve. Let $x = d\ln w$, and let $\omega = g + x$ denote the cost advantage that the input gains over labor at the marginal task, net of the wage response.

Boundary. Differentiating the assignment condition $\ln p_M - \ln\gamma_M(I) = \ln w - \ln\gamma_L(I)$, with γ_M shifted by g at and around I, gives $-g - (\ln\gamma_M)'(I)\,dI = x - (\ln\gamma_L)'(I)\,dI$, so that $\rho\,dI = g + x$ and

$$dI = \frac{\omega}{\rho}, \qquad \rho = \rho(I). \tag{14}$$

Cost. In the differential of $C^{1-\sigma} = \int_{N-1}^{I} c_M^{1-\sigma} + \int_I^N c_L^{1-\sigma}$, the boundary terms cancel because $c_M(I) = c_L(I)$, which is the envelope property of the cost-minimizing assignment. Dividing by $(1-\sigma)\,C^{1-\sigma}$ gives $d\ln C = \int_{N-1}^{I} s(z)\,d\ln c_M\,dz + \int_I^N s(z)\,d\ln c_L\,dz$, and with $d\ln c_M = -g$ and $d\ln c_L = x$,

$$d\ln C = -s_M\,g + s_L\,x. \tag{15}$$

Shares. For $z \leq I$, $d\ln s(z) = (1-\sigma)[d\ln c_M - d\ln C] = (1-\sigma)[-g(1-s_M) - s_L x] = -(1-\sigma)\,s_L\,\omega$. Integrating against $s(z)$ over $[N-1, I]$ and adding the Leibniz term $h\,dI$ for

the tasks that cross the boundary,

$$ds_M = -(1-\sigma)\, s_M s_L\, \omega + \frac{h}{\rho}\,\omega = \omega A s_L, \qquad A = \frac{h}{\rho\, s_L} - (1-\sigma)\, s_M, \tag{16}$$

which is equation (2) of the text, and $ds_L = -ds_M$ (for labor's tasks $d\ln s = (1-\sigma)\, s_M\, \omega$, which integrates to $(1-\sigma)\, s_M s_L\, \omega - h\,\omega/\rho$). The first term of A reflects the tasks that are reassigned to the input. The second reflects the tasks that the input already performs, which become cheaper and, when tasks are gross complements, account for a smaller share of cost. $A < 0$ when the second term dominates.

Price, markup and scale. From (13) and the demand curve with $d\ln D = -\delta$,

$$\begin{aligned} d\ln p &= \varphi\, d\ln C, \qquad & d\ln \mu &= -(1-\varphi)\, d\ln C, \\ d\ln y &= -\delta - \eta\varphi\, d\ln C, \qquad & d\ln R &= -\delta + (1-\eta)\,\varphi\, d\ln C. \end{aligned} \tag{17}$$

Labor and the wage. Log-differentiating $wL = (s_L/\mu)R$ gives $x + d\ln L = d\ln s_L - d\ln\mu + d\ln R = -A\omega + (1-\eta\varphi)\, d\ln C - \delta$. Substituting (15) and collecting the terms in x,

$$x + d\ln L = -(\delta + Bg) + (1-E)\, x, \qquad B = A + (1-\eta\varphi)\, s_M, \qquad E = 1 + A - (1-\eta\varphi)\, s_L.$$

The wage curve implies $d\ln L = \varepsilon x$, so that $(\varepsilon + E)\, x = -(\delta + Bg)$ and

$$\begin{aligned} d\ln w &= -\frac{\delta + [A + (1-\eta\varphi)\, s_M]\, g}{\varepsilon + E}, \\ d\ln L = \varepsilon\, d\ln w &= -\frac{[A + (1-\eta\varphi)\, s_M]\, g + \delta}{\kappa}, \qquad \kappa = 1 + \frac{E}{\varepsilon}, \end{aligned} \tag{18}$$

which is equation (3) of the text. Using $s_M + s_L = 1$, $E = \sigma s_M + \eta\varphi s_L + h/(\rho s_L)$, and $E - B = \eta\varphi$. Three terms in the numerator lower labor demand: the reallocation of tasks to the input, A; a scale effect, $(1-\eta\varphi)s_M$, which arises because a cost reduction lowers nominal turnover when demand is inelastic and raises the markup when pass-through is incomplete, both of which reduce the wage bill at given cost shares; and the demand shift δ. The coefficient B can be verified independently. At a fixed wage, $L = y\int_I^N (c_L/C)^{-\sigma}/\gamma_L\, dz$ gives $d\ln L = d\ln y + \sigma\, d\ln C - (l(I)/L)\, dI = -\delta - (\sigma - \eta\varphi)\, s_M g - (h/s_L)(g/\rho)$, using $w\, l(I) = h\, Cy$ and $wL = s_L Cy$, and the coefficient on g is again $-B$.

The net cost advantage. Substituting x from (18) into $\omega = g + x$ and using $E - B = \eta\varphi$ gives

$$\omega = \frac{(\varepsilon + \eta\varphi)\, g - \delta}{\varepsilon + \Lambda}, \qquad \Lambda = E = A + s_M + \eta\varphi\, s_L, \tag{19}$$

which is equation (4) of the text. Here Λ is the own-wage elasticity of the industry's labor demand, which combines the boundary, substitution and scale margins, and $\varepsilon + \Lambda > 0$ is the condition for a determinate equilibrium on the wage curve. With the capability gain alone, $\omega > 0$: the wage response dampens the boundary and share responses but never reverses them, and ds_M has the sign of A. With the demand shift, $\omega > 0$ if and only if $(\varepsilon + \eta\varphi)\, g > \delta$. If the fall in demand is large enough, the wage falls by more than the input's cost per task, labor takes tasks back from the input, and ds_M has the sign of $-A$. The closed-form expressions in terms of primitives follow by substitution. In particular, $d\ln C = -[g(\varepsilon s_M + \sigma s_M + h/(\rho s_L)) + s_L\delta]/(\varepsilon + \Lambda)$, and (17) then gives the price, the markup, real output and turnover.

Verification of the labor equation. The left-hand side of the labor equation is $d\ln(wL) = (1 + \varepsilon)\, x$, and the right-hand side is $d\ln s_L - d\ln\mu + d\ln R = -(\delta + Bg) + (1 - E)\, x$. At $x = -(\delta + Bg)/(\varepsilon + E)$ both sides equal $-(1 + \varepsilon)(\delta + Bg)/(\varepsilon + E)$.

Each shock separately. Every response is linear in (g, δ), so the response to both shocks together is the sum of the responses to each. With the demand shift alone and a fixed wage ($\varepsilon \to \infty$), employment, real output, turnover, purchases, the wage bill and value added fall by δ, and the boundary, the cost shares, unit cost, the price and the margin are unchanged. With a wage curve the wage falls by $\delta/(\varepsilon+\Lambda)$, and the demand shift affects the cost structure only through this channel: unit cost falls by s_L times the fall in the wage, the price falls by φ times the fall in cost, the markup rises by $(1 - \varphi)$ times the fall in cost, and the input's cost share changes by $-As_L$ times the fall in the wage. With the capability gain alone and a fixed wage, $\omega = g$, $dI = g/\rho$, $d\ln C = -s_M g$, $ds_M = As_L g$, $d\ln p = -\varphi s_M g$, $d\ln y = \eta\varphi s_M g > 0$ and $d\ln L = -Bg$.

The accounting ratios. From equation (1),

$$d\ln\frac{p_M M}{R} = \frac{ds_M}{s_M} - d\ln\mu, \qquad d\ln\frac{wL}{p_M M} = -\frac{ds_M}{s_L s_M}, \tag{20}$$

$$d\ln\frac{wL}{\mathrm{VA}} = -\,ds_M\left[\frac{1}{s_L} - \frac{1}{\mu - s_M}\right] - \frac{\mu\, d\ln\mu}{\mu - s_M}, \quad d\left(1 - \frac{1}{\mu}\right) = \frac{d\ln\mu}{\mu}, \tag{21}$$

$$d\ln\frac{\mathrm{VA}}{L} = d\ln w + \frac{\mu\, d\ln\mu - ds_M}{\mu - s_M} + \frac{ds_M}{s_L}, \qquad d\ln\mathrm{VA} = d\ln R + \frac{-\,d(s_M/\mu)}{1 - s_M/\mu}. \tag{22}$$

Numerical verification. I solve the equilibrium numerically (the boundary from the assignment condition, the cost shares by quadrature, the wage from the wage curve) at 7 parameter points, for the capability gain alone, the demand shift alone and both shocks together, including the limiting cases and the additivity of the two shocks. The numerical solution reproduces every closed-form expression of this appendix: 186 of 186 comparisons agree, with a largest relative error of 1.3×10^{-9}. Two independent derivations from the environment, one for the capability gain and one for the general case, give the same expressions.

D.3 The Moments and the Accounting Identities

The fourteen moments. Each moment is the coefficient on the suppliers' displacing stock of the process–product classification per unit of the exposure. Because the common level of the two stocks has no effect (Section 2.3), I interpret it as the response to the direction of the suppliers' innovation. Eleven moments are firm outcomes estimated with industry and year effects, which describe the average firm in an industry-year: employment ($d \ln L$), wage costs ($d \ln wL$), value added, turnover ($d \ln R$), purchases ($d \ln p_M M$), the purchase share of turnover ($d(s_M/\mu)$), the wage bill over purchases ($-ds_M/(s_L s_M)$), the labor share of value added ($d[s_L/(\mu - s_M)]$), the operating result over value added (minus the labor share if the operating result is value added less the wage bill), the margin over turnover ($d \ln \mu/\mu$) and value added per full-time worker ($d \ln(\mathrm{VA}/L)$). The other three are the hourly wage of stayers, estimated with worker and year effects ($d \ln w$), the buyer's output price ($d \ln p$) and real output per firm ($d \ln y$, turnover deflated by the industry's producer price). The last two are available for the 48 industries with a price index over 2005–2019. The list of moments was chosen before the model was estimated, and their values are reported in Table 47.

Parameters and calibration. The mapping from parameters to moments has six parameters, $(g, A, \varphi, \eta, \varepsilon, \delta)$, with s_M and μ calibrated from the accounts using equation (1). The estimation holds η fixed at each value on a grid and estimates the other five.

The nine identities. At a fixed η the fourteen responses are linear in five objects, so they lie in a five-dimensional subspace for any parameter values. Numerically, the singular values of the centred matrix of responses at random admissible parameter points fall to machine zero after the fifth (the dimension of the image is 5, and the nine orthogonal combinations are constant to 1.6×10^{-15} across points). Six objects describe the firm: employment $d \ln L$, the wage $d \ln w$, the purchase share $d(s_M/\mu)$, the markup $d \ln \mu$, the price $d \ln p$ and output $d \ln y$. The cost identity $wL = s_L R/\mu$ links them, which leaves five degrees of freedom. The nine combinations are the definitions of the remaining moments in terms of these six:

1. turnover is price times quantity, $d\ln R = d\ln p + d\ln y$;
2. wage costs are the wage plus employment, $d\ln wL = d\ln w + d\ln L$;
3. purchases are the purchase share plus turnover, $d\ln p_M M = d\ln(s_M/\mu) + d\ln R$;
4. value added is turnover less purchases, $d\ln \mathrm{VA} = d\ln R - d(s_M/\mu)/(1 - s_M/\mu)$;
5. the wage bill over purchases is wage costs less purchases;
6. the labor share of value added is a function of the purchase share and the markup, (21);
7. the operating result over value added is minus the labor share;
8. value added per worker is value added less employment;
9. the cost identity, $d\ln wL = ds_L/s_L - d\ln\mu + d\ln R$, with ds_M from the purchase share and the markup.

The five free parameters span the subspace (the recovery exercise of Appendix E.2), so the model is exactly identified given η, and the minimum-distance statistic on nine degrees of freedom tests whether the moments are consistent with one another as measurements. A Wald test of the nine identities with the joint covariance matrix, which does not depend on the model, gives 19.9 on 9 degrees of freedom ($p = 0.019$). Two identities fail. Identity 7 fails because the operating result in the accounts is net of depreciation (the discrepancy is 0.042, s.e. 0.015, $t = 2.8$). Identity 8 fails because full-time equivalents and headcount are measured on different samples of firms (the discrepancy is 0.47, s.e. 0.20, $t = 2.4$; on the 1,194,766 firm-years that report every outcome it is 0.11). The other seven give 4.2 on 7 degrees of freedom ($p = 0.76$). On the twelve moments that exclude these two, the image is again 5-dimensional and the seven remaining identities hold.

D.4 Identification

Let β_Y denote the mix coefficient for outcome Y, that is, the coefficient on the displacing stock of the process–product classification per unit of the suppliers' exposure. Taking the cost shares and the markup from the accounts (equation (1)), the equations of the model

can be solved sequentially:

$$
\begin{aligned}
\beta_w/\beta_L &= 1/\varepsilon,\\
ds_M &= -\,s_L s_M\,\beta_{wL/p_M M},\\
d\ln\mu &= ds_M/s_M - \beta_{p_M M/R} = -\,s_L\,\beta_{wL/p_M M} - \beta_{p_M M/R},\\
\beta_{\text{margin}} &= d\ln\mu/\mu \qquad \text{(over-identifying)},\\
d\ln C &= \beta_p - d\ln\mu \qquad \text{(the buyer's output price)},\\
g &= (s_L\,\beta_w - d\ln C)/s_M, \qquad \varphi = \beta_p/d\ln C, \qquad \omega = g + \beta_w,\\
A &= -\,s_M\,\beta_{wL/p_M M}/\omega,\\
\beta_y &= -\,\delta - \eta\,\beta_p \qquad \text{(real output)},\\
\beta_R &= \beta_p + \beta_y \qquad \text{(an identity)},\\
\beta_{wL} &= -\,ds_M/s_L + (1-\eta\varphi)\,d\ln C - \delta = -\,ds_M/s_L + d\ln C + \beta_y \quad \text{(the same again)}.
\end{aligned}
$$

The last three lines show why the demand elasticity and the demand shift are not separately identified. Real output, nominal turnover, and the wage bill all measure the same scale response, $d\ln y = -\delta - \eta\varphi\, d\ln C$, and nothing else in them involves η or δ. The labor equation $wL = s_L C y$ is an identity of the model, and once g, A, φ, and ε are known, the wage equation (3) reduces to the same identity. The data therefore identify g, A, φ, ε, and $d\ln y$, but of the pair (η, δ) they identify only the line $\delta = -\beta_y - \eta\,\beta_p$, whose slope is the price response; every moment of the model is unchanged along that line. Because real output and nominal turnover together determine the price, $\beta_R = \beta_p + \beta_y$, the capability gain is identified at the precision of the price index whether or not the price moment itself is included, and the role of the price moment is to provide an over-identifying check on that identity. Without the price and without real output, $d\ln C$ is not observed. In that case the purchase share and the wage bill over purchases deliver $d\ln\mu = -(1-\varphi)\,d\ln C$, and turnover and the wage bill deliver the scale response, so the data identify ε, $A\omega$, $(1-\varphi)\,d\ln C$, and $(1-\eta\varphi)\,d\ln C - \delta$, four combinations of six parameters; the capability gain is then presented as a function of the pass-through over a range of values.

The decomposition. Because the mapping is linear in (g, δ), every moment is the sum of a capability component (the model evaluated at $\delta = 0$) and a demand component (the model evaluated at $g = 0$). The share of each component in the employment, wage, turnover, and real-output responses is a function of the parameters, for which a delta-method interval follows from the structural covariance matrix and a percentile interval from the score bootstrap. The demand component of the labor responses operates through $\delta/(\varepsilon + \Lambda)$ and the

capability component through $[A + (1 - \eta\varphi)s_M]\, g/(\varepsilon + \Lambda)$. The two labor moments alone cannot separate the two components (Section 2.2), but the price and real output can.

The sign of ω. The wage bill over purchases rises if and only if $\omega A < 0$ (Section 2.2). The price moment identifies $\omega = (\beta_w - \beta_p + d\ln\mu)/s_M$ directly, so in principle the data determine which case applies: either $\omega > 0$ and $A < 0$, the input-gains case, in which the input gains a net cost advantage, or $\omega < 0$ and $A > 0$, the labor-retakes case, in which labor takes tasks back as the demand shift lowers its wage. The condition $\omega > 0$ is equivalent to $(\varepsilon + \eta\varphi)\, g > \delta$. The standard error of the price moment enters g one for one, scaled by s_M, and enters δ through $\eta\,\beta_p$. As a result, g is bounded by the price interval and may be zero, whereas δ is determined to roughly the precision of the real-output moment. The demand shift is thus the better-measured of the two shocks.

Stability and the bounds on A. An equilibrium on the wage curve requires $\varepsilon + \Lambda > 0$, where $\Lambda = A + s_M + \eta\varphi s_L$ is the own-wage elasticity of labor demand. The share moments measure the product $A\omega$, and when ω is close to zero A is poorly determined: the objective function is flat along a direction in which A grows and ω shrinks, passing through $\Lambda \leq 0$ on the negative side and $\Lambda \to \infty$ on the positive side. The estimation therefore restricts the search to the region in which the model is well defined, $A \geq -s_M$, which the task model implies ($h/\rho \geq 0$ and $\sigma > 0$), and $0 < \Lambda \leq \bar{\Lambda}$ with $\bar{\Lambda} = 5$, so that labor demand slopes downward with a bounded elasticity (long-run own-wage elasticities in the literature are below one). A solution at one of these bounds is reported as such. It indicates that the data identify $A\omega$ and δ but not the position between the two cases, and the fit is then reported on each side of $A = 0$. Table 48 shows how the results depend on the bound.

Unit-free ratios. Each of the following ratios must take a known value under the model and therefore provides a test:

$$\frac{\beta_w}{\beta_L} = \frac{1}{\varepsilon} \quad \text{(wage curve; each split of the suppliers' channel)}$$

$$\frac{\mu\,\beta_{\text{margin}}}{-\,s_L\,\beta_{wL/p_M M} - \beta_{p_M M/R}} = 1 \quad \text{(markup: margin against shares)}$$

$$\frac{\beta_{wL/\text{VA}}}{-\,ds_M\,[1/s_L - 1/(\mu - s_M)] - \mu\, d\ln\mu/(\mu - s_M)} = 1 \quad \text{(labor share of value added)}$$

$$\frac{\beta_{\text{VA}} - \beta_R}{-\,d(s_M/\mu)/(1 - s_M/\mu)} = 1 \quad \text{(value added's share of turnover)}$$

$$\frac{\beta_{\text{VA}/L} - \beta_w}{(\mu\, d\ln\mu - ds_M)/(\mu - s_M) + ds_M/s_L} = 1 \quad \text{(value added per worker).}$$

Given η, two further ratios involve δ: the capability gain relative to the demand shift, and the demand shift's share of the fall in real output,

$$\frac{g}{\delta} = \frac{(s_L \beta_w - \beta_p + \mu\, \beta_{\text{margin}})/s_M}{-\beta_y - \eta\, \beta_p}, \qquad \frac{\delta}{-\beta_y} = 1 + \eta\, \frac{\beta_p}{\beta_y},$$

These are ratios of linear combinations of the moments. I report Fieller confidence sets for them, which remain valid when the denominator is imprecisely estimated (the price enters both). Using the joint covariance matrix of the coefficients, each ratio receives a delta-method standard error or a Fieller set. The most precisely measured firm moments are the wage bill, value added, and employment with industry fixed effects, together with real output per firm; turnover is imprecise, so the fourth ratio is the least informative.

Normalization. Counterfactuals within the support of the data, such as a shift in the direction of the suppliers' innovation by one residual standard deviation (0.033), require no unit of measurement, because they are the coefficients themselves. The model adds the restrictions that connect the coefficients (the ratios that must agree and the signs that must hold) and the decomposition of each response into the two shocks. When a unit is needed, for instance to report g or to compare the cost reduction with the fall in the price, it comes from the price moment through $d\ln C = \beta_p - d\ln\mu$, which is estimated with a wide interval. The demand shift requires no unit beyond η: it is a log change in demand per unit of the direction.

E Structural Estimation: Details and Further Results

This appendix describes the estimator of Section 6 and its inference, reports the tests of the estimator on simulated data, and presents the fitted value of every moment, the sensitivity of the results to the bound on the elasticity of labor demand, and the two bootstraps.

E.1 Estimator and Inference

Objective. The estimator is minimum distance. Let $\hat{\beta}$ denote the vector of the 14 moments and $m(\theta)$ the model's mapping at the calibrated shares. The estimate $\hat{\theta}$ minimises $(\hat{\beta} - m(\theta))' W (\hat{\beta} - m(\theta))$, where W is the diagonal matrix of inverse squared standard errors and is fixed before estimation. The joint covariance matrix of the moments is used for inference and for the tests. The search starts from many points, half on each side of $A = 0$, and each solution is refined by a local search within the admissible region $A \geq -s_M$, $0 < \Lambda \leq \bar{\Lambda}$. I

report the best solution on each side of $A = 0$, which gives the two cases of the text, and I indicate when a parameter is at a bound.

Covariance of the moments. The covariance matrix is computed from the per-industry score contributions of the regressions, on the 111 industries common to all moments, including the price and real output. The largest absolute correlation is between the operating-result share and the labor share, which are near-complements by construction.

Standard errors and tests. The covariance matrix of the parameters is the delta-method (sandwich) formula with the fixed weights and the joint covariance matrix of the moments. The minimum-distance statistic J is computed with the joint covariance matrix and has $14 - 5 = 9$ degrees of freedom; Appendix D.3 explains what it tests. I also report profiles of the objective over g, the identified set for g over the retained share $1 - \varphi$ without the price moment, and confidence sets for g obtained by test inversion (the values at which the model with g fixed is not rejected at the five percent level), on the fourteen moments and on the twelve mutually consistent ones. A score bootstrap resamples the industries' score contributions and repeats the full search in each of 199 draws, and ratios receive Fieller sets. I use the delta method for δ and ε, which lie in the interior of the admissible region, test-inversion sets for g, and bootstrap and test-inversion intervals for A, φ and the decomposition into the capability and demand components.

Robustness specifications. Each specification replaces one input to the baseline (Table 19). The specifications are: the wage coefficient estimated with sector-by-year effects and with industry trends; a sample ending in 2016; the exposure rebuilt directly from the patent file; the moments without the own-industry term of the exposure; the eleven firm moments estimated on the common sample of firm-years that report every outcome; the moments without the price; the moments without real output; the twelve mutually consistent moments; the floored patent stocks recoded as zero; and the sample without the fifteen buying industries whose exposure the floor affects. In the last specification the wage is not used as a target, a choice made before the placebo for that sample was estimated. In addition, each supplying industry is left out of the exposure in turn.

E.2 Tests of the Estimator

I verified the implementation in three steps before estimating the model on the actual moments. First, the model's mapping reproduces the closed-form expressions of Appendix D.2 and the numerical equilibrium at the test points. The labor equation holds as an identity

at random parameter points, every moment equals the sum of its capability and demand components, and moving (η, δ) along the line of Section 2.4 changes no moment.

Second, I estimate the model on noise-free moments generated at 8 random parameter points. With η fixed at its true value, the five free parameters are recovered to 6.6×10^{-13}, and the objective has a single minimum at zero. With η free, the estimate can lie anywhere on the (η, δ) line, while $(g, A, \varphi, \varepsilon)$ and $d \ln y$ are recovered. Without the price and real output, the four identified combinations are recovered and g is not.

Third, I simulate score contributions with correlated moments, holding η fixed, and examine the coverage of each inference procedure. Of 48 tests, 48 passed. Delta-method intervals cover $(g, A, \varphi, \varepsilon, \delta)$ at rates of 0.97, 0.94, 0.95, 0.94 and 0.96. The statistic on nine degrees of freedom rejects at a rate of 0.033 at a nominal level of 0.05. Test-inversion sets for g cover at 0.92 (mean width 1.21). Score-bootstrap intervals cover at 0.95, 0.95, 1.00, 0.95 and 0.97, and Fieller sets cover the seven ratios at 0.949 to 0.986. I then repeat the exercise with the estimated covariance matrix in place of the simulated one (300 replications for the delta method, 60 for the test inversion and 40 for the bootstrap). Of 48 tests, 47 pass. Delta-method coverage is 0.94, 0.96, 0.94, 0.94 and 0.94; the statistic rejects at 0.057; test-inversion sets for g cover at 0.98 (width 1.09); the score-bootstrap intervals cover at 0.80 for g and at 0.92, 0.92, 0.92 and 0.92 for the other parameters; and Fieller sets cover at 0.949 to 0.975. The one test that fails is the coverage of the bootstrap interval for g. For this reason I report test-inversion sets for g and no bootstrap interval.

E.3 The Fitted Moments

Table 47 and Figure 8 compare the fitted values with the estimates. No residual exceeds two standard errors in the baseline case. The two moments whose identities fail have the largest residuals among the firm-account moments after value added. Excluding them gives $g = 0.20$ (0.24), $\delta = 0.82$ (0.14), $\omega = -0.06$ and $J = 4.3$ on 7 degrees of freedom ($p = 0.75$), so these two moments do not drive the parameter estimates.

The unit-free ratios of Appendix D.4, each with its Fieller set, are as follows. The ratio for the labor share of value added is 1.13 against a predicted value of one [0.66, 8.96]. The ratio for the share of value added in turnover is -0.26 [-15.12, 2.28], and the ratio for value added per worker is 2.38 [-0.95, 17.46]. The markup ratio is 6.03 with an unbounded set. The wage curve is 0.56 [0.14, 1.51], the ratio g/δ is 0.33 [-0.57, 0.97], and the demand shift's share of the fall in real output is 1.18 [0.72, 2.02].

The test-inversion set for g on the fourteen moments is [0.5, 0.5]. This narrow set is a consequence of the rejection of the accounting identities: the unrestricted statistic already

Table 47: The fourteen moments: estimates, fitted values and the two components

		Labor retakes tasks (baseline)				The input gains
Moment	Estimate (s.e.)	Fitted	Residual (s.e.)	Demand part	Capability part	Fitted
Employment	-0.397 (0.130)	-0.477	0.61	-0.235	-0.242	-0.526
Wage costs	-0.666 (0.148)	-0.700	0.23	-0.345	-0.355	-0.706
Value added	-0.657 (0.142)	-0.486	-1.20	-0.408	-0.078	-0.483
Turnover	-0.591 (0.411)	-0.787	0.47	-0.787	0.000	-0.766
Purchases	-0.809 (0.464)	-1.056	0.53	-1.126	0.070	-1.019
Purchase share of turnover	-0.120 (0.052)	-0.142	0.42	-0.179	0.037	-0.133
Wage bill over purchases	0.520 (0.197)	0.356	0.84	0.781	-0.425	0.313
Labor share of value added	-0.153 (0.061)	-0.158	0.08	0.046	-0.204	-0.164
Operating result over value added	0.195 (0.069)	0.158	0.54	-0.046	0.204	0.164
Margin over turnover	0.110 (0.033)	0.112	-0.07	0.025	0.087	0.113
Value added per full-time worker	0.213 (0.176)	-0.010	1.27	-0.173	0.163	0.042
Hourly wage (workers)	-0.222 (0.088)	-0.224	0.02	-0.110	-0.113	-0.180
Output price	-0.132 (0.152)	-0.066	-0.44	-0.015	-0.051	-0.240
Real output per firm	-0.730 (0.269)	-0.721	-0.03	-0.772	0.051	-0.526

Notes: Coefficients per unit of the suppliers' shift toward process innovation, process–product classification. The estimate is the reduced-form coefficient with its standard error clustered by industry. Fitted values are the model in the baseline case (the fit under $\Lambda \leq 5$ at $\eta = 1$) and in the input-gains case ($A = -s_M$); the residual is the estimate less the baseline's fitted value in units of the standard error. The demand part is the model at $g = 0$ and the capability part the model at $\delta = 0$; the two add to the fitted value.

Table 48: The two cases under alternative bounds on the elasticity of labor demand

Case and bound	Objective	J (p)	g	ω	δ	Demand share, employment	Best fit under the bound
The input gains ($A = -s_M$, any bound)	6.58	16.5 (0.06)	0.49	+0.31	0.77	1.34	
Labor retakes tasks, $\Lambda \leq 1.25$	7.05	19.3 (0.02)	0.03	-0.23	0.74	0.95	the input gains
Labor retakes tasks, $\Lambda \leq 2$	5.84	18.7 (0.03)	0.09	-0.15	0.77	0.83	labor retakes
Labor retakes tasks, $\Lambda \leq 5$	5.34	19.6 (0.02)	0.17	-0.05	0.79	0.49	labor retakes
Labor retakes tasks, $\Lambda \leq 20$	5.24	19.8 (0.02)	0.21	-0.01	0.79	0.16	labor retakes

Notes: Fits of the fourteen moments at $\eta = 1$ with the parameters restricted to $A \geq -s_M$ and $0 < \Lambda \leq \bar{\Lambda}$, on each side of $A = 0$. In the input-gains case the solution is at $A = -s_M$ for every bound; in the labor-retakes case it is at the bound. The objective is the diagonal-weight distance, J the statistic with the joint covariance. The demand share is the demand part of the employment response over the total. The baseline is the row at $\bar{\Lambda} = 5$; the bound of 1.25 is the upper end of the range the labor-demand literature estimates for an industry.

exceeds the critical value, so the only value of g that is not rejected is the one at which the best fit switches from one case to the other. On the twelve mutually consistent moments the set is [-0.1, 1.1], which is the set reported in the text.

E.4 The Bound on the Elasticity of Labor Demand

Table 48 and Figure 9a show how the results depend on the choice of $\bar{\Lambda}$. In the input-gains case the solution is at $A = -s_M$ for every bound, and the demand component is 1.34 times the fall in employment. In the labor-retakes case the solution is at the bound, and the demand shift's share of the fall in employment declines as the bound is relaxed, to 0.16 at



Figure 8: The fourteen moments. Each estimate is shown with its 95 percent confidence interval (black, grey), together with the fitted value in the baseline, labor-retakes case (filled) and in the input-gains case (hollow). The largest residual in the baseline case is 1.27 standard errors (value added per full-time worker); among the twelve mutually consistent moments it is 1.20 (value added). Source: Table 47.

a bound of 20. Under a bound of 1.25 the better fit is the case in which the input gains; under bounds of 2, 5 and 20 the labor-retakes case fits better. The demand shift and the total responses do not depend on the bound. The decomposition does, and across the two cases and the bounds the demand share lies in [0.16, 1.34].

E.5 The Two Bootstraps

Table 49 compares the score bootstrap with the parametric bootstrap of Section 6.6. The score bootstrap gives δ in [0.52, 1.05], ε in [1.17, 5.27], $A\omega$ in [-0.42, 0.03], A in [-0.60, 4.40] and φ in [0.00, 0.78]. The two bootstraps give similar intervals for the objects that are identified (δ, $A\omega$ and the total responses). For the demand shift's share of the fall in employment, the interval includes one in both, and the parametric interval is wider.

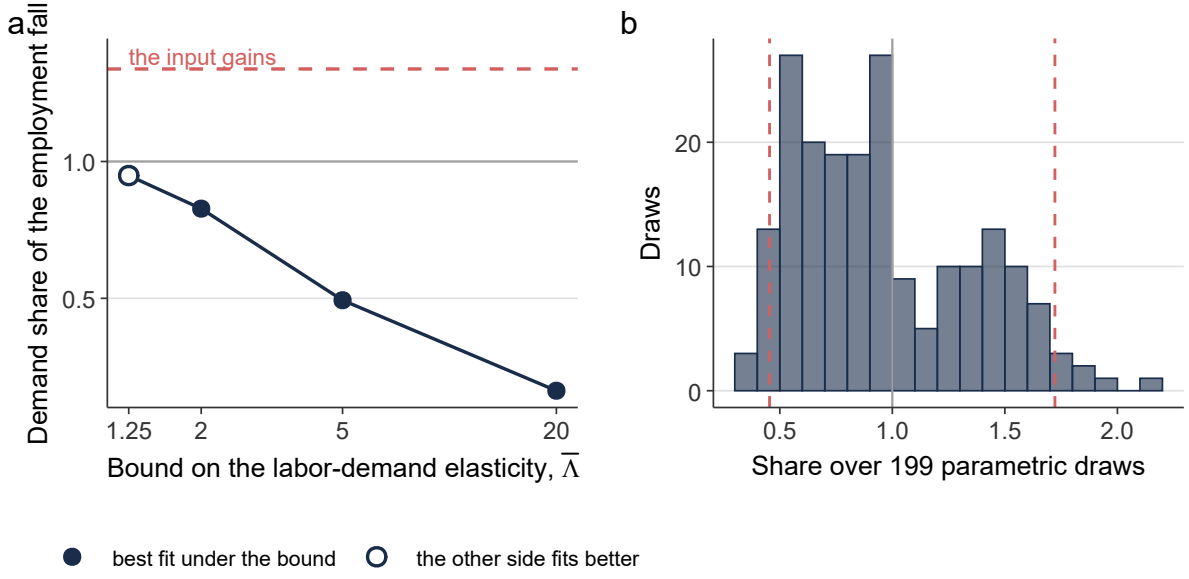


Figure 9: The demand shift's share of the fall in employment. Panel a shows the share in the labor-retakes case under each bound on the elasticity of labor demand (filled where this case is the better fit under that bound) and in the input-gains case, which does not depend on the bound. Panel b shows the distribution of the share over the 199 parametric bootstrap draws, with the 95 percent interval dashed; the share exceeds one in 36 percent of the draws. Sources: Tables 48 and 49.

E.6 Figures for the Counterfactual and the Incidence

Figure 10 shows the interquartile-range counterfactual of Table 18 with its demand and capability components. Figure 11 shows the incidence decomposition of Table 17 by group: the employer's premium component is common across groups, and pay within the job falls only at the top of the pay distribution and in the cognitive occupations.

Table 49: Two bootstraps: the score bootstrap and the parametric bootstrap

	Parametric bootstrap		Score bootstrap
Object	95 percent interval	Median	95 percent interval
Demand shift δ	[0.54, 1.03]	0.79	[0.52, 1.05]
Wage-curve elasticity ε	[1.28, 7.38]	2.32	[1.17, 5.27]
Share response $A\omega$	[-0.44, 0.03]	-0.23	[-0.42, 0.03]
Net cost advantage ω	[-0.30, 0.72]	-0.06	–
Demand share of the employment fall	[0.45, 1.72]	0.90	[0.41, 1.58]
Demand share of the output fall	[1.00, 2.03]	1.06	[1.00, 1.85]
Employment per net s.d. (percent)	[-2.4, -1.0]	-1.6	–
Wage per net s.d. (percent)	[-1.1, -0.2]	-0.7	–
Real output per net s.d. (percent)	[-3.3, -1.0]	-2.3	–

Notes: The score bootstrap resamples the industries' score contributions and repeats the full search in each of 199 draws; it is the interval reported in the text. The parametric bootstrap draws the fourteen moments from the normal distribution with their estimated covariance and repeats the full search in each of 199 draws; it serves as a check on the score bootstrap. Percentile intervals. Neither is used for g, whose sets come from test inversion. The demand share of the employment fall exceeds one in 36 percent of the parametric draws and 33 percent of the score draws; 36 percent of the parametric draws have their best fit in the input-gains case.

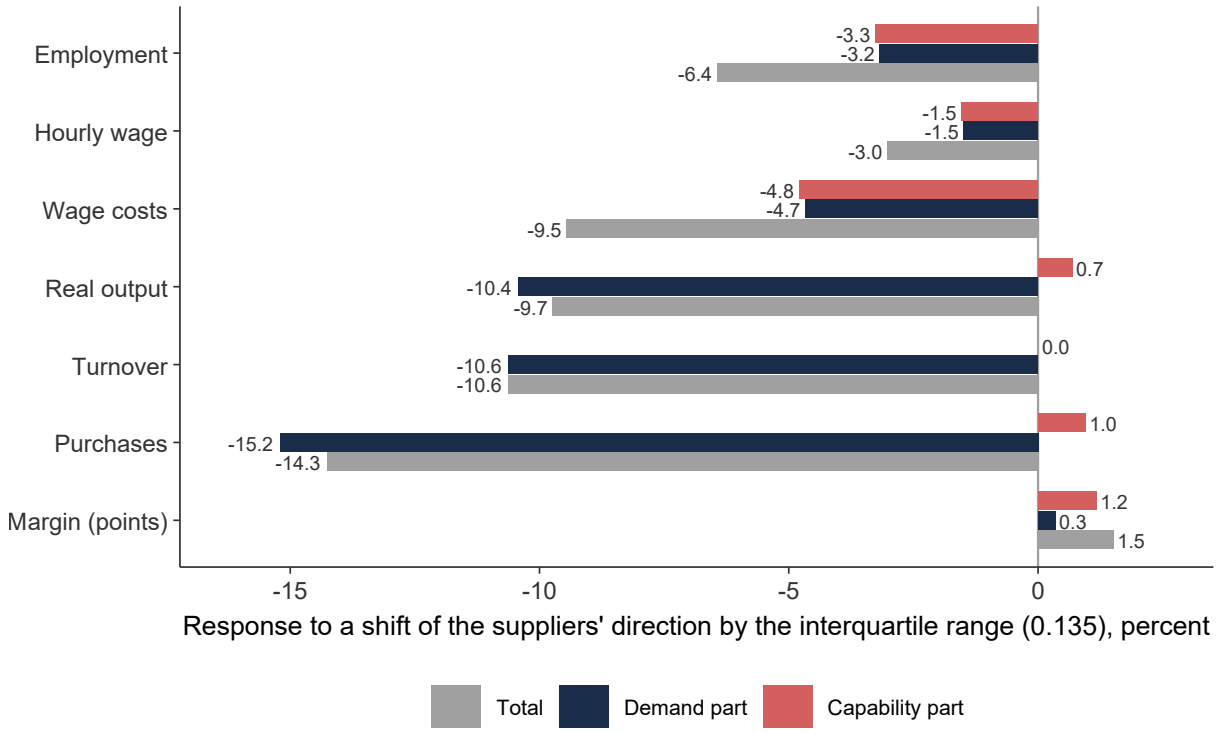


Figure 10: The response to a shift of a buying industry's suppliers from the 25th to the 75th percentile of the direction of their innovation, in the baseline case, divided into the demand component (the model at $g = 0$) and the capability component (the model at $\delta = 0$). The total responses are the same in both cases and under every bound; the decomposition varies over the set of Table 48. Source: Table 18.

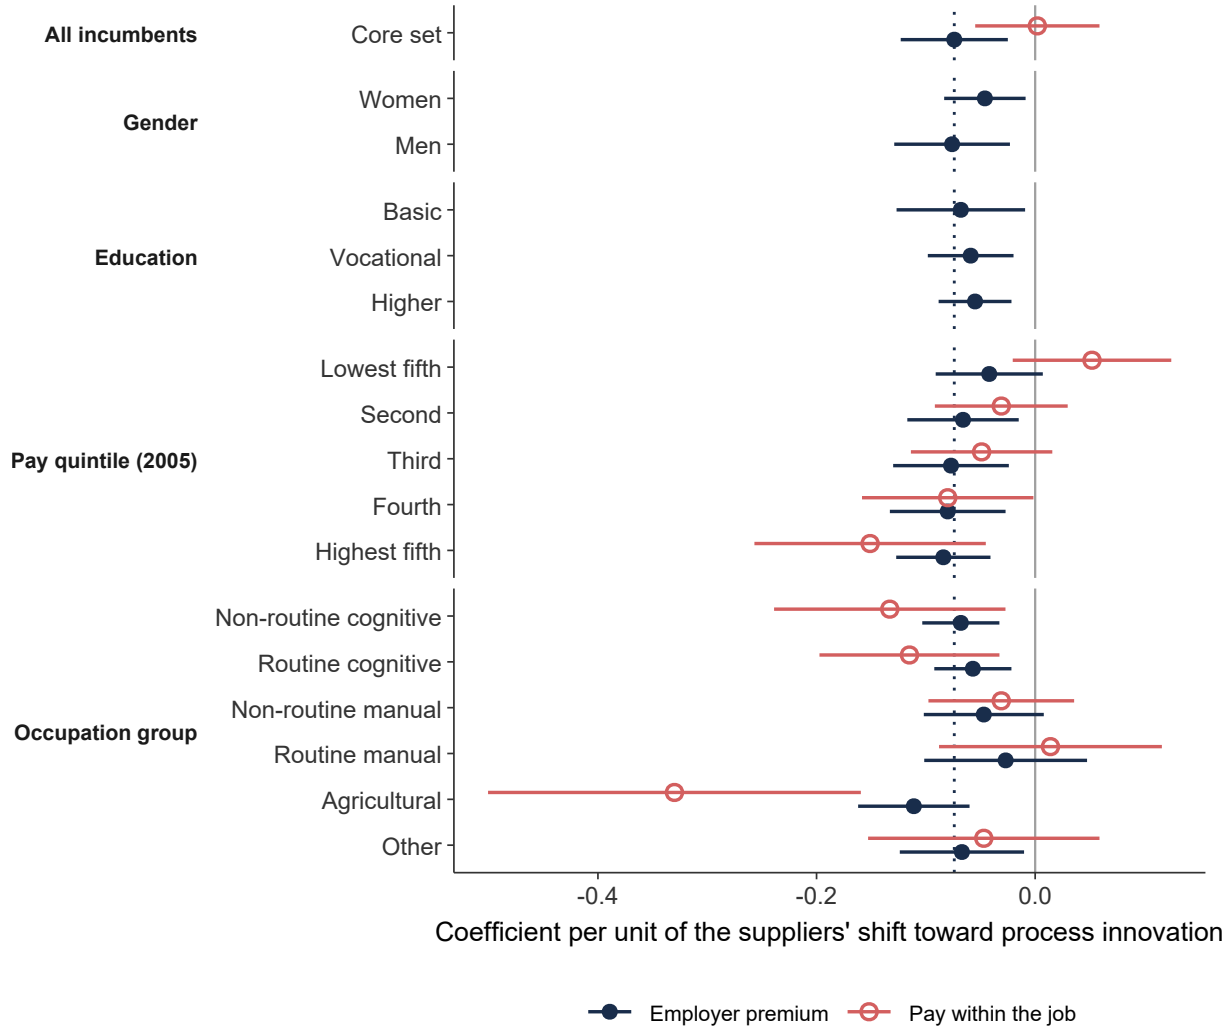


Figure 11: Incidence by worker group. The figure shows the coefficients per unit for the employer's fixed premium (filled) and for pay within the job (hollow, where estimated), with 95 percent confidence intervals. The dotted line is the premium coefficient for the core set. Source: Table 17.