



ROCKWOOL Foundation Berlin

Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES

262/26

Capital–Skill Complementarity Beyond Firm Boundaries

Gert Bijnens, Anders Humlum, Stijn Vanormelingen

Capital–Skill Complementarity Beyond Firm Boundaries

Authors

Gert Bijmens, Anders Humlum, Stijn Vanormelingen

Reference

JEL Codes: E22, J23, J24, J31, L24, O33

Keywords: Capital–skill complementarity; Firm boundaries; Outsourcing; Investment spikes; Production networks; Automation

Recommended Citation: Gert Bijmens, Anders Humlum, Stijn Vanormelingen (2026): Capital–Skill Complementarity Beyond Firm Boundaries. RFBerlin Discussion Paper No. 262/26

Access

Papers can be downloaded free of charge from the RFBerlin website: <https://www.rfberlin.com/discussion-papers>

Discussion Papers of RFBerlin are indexed on RePEc: <https://ideas.repec.org/s/crm/wpaper.html>

Disclaimer

Opinions and views expressed in this paper are those of the author(s) and not those of RFBerlin. Research disseminated in this discussion paper series may include views on policy, but RFBerlin takes no institutional policy positions. RFBerlin is an independent research institute.

RFBerlin Discussion Papers often represent preliminary or incomplete work and have not been peer-reviewed. Citation and use of research disseminated in this series should take into account the provisional nature of the work. Discussion papers are shared to encourage feedback and foster academic discussion.

All materials were provided by the authors, who are responsible for proper attribution and rights clearance. While every effort has been made to ensure proper attribution and accuracy, should any issues arise regarding authorship, citation, or rights, please contact RFBerlin to request a correction.

These materials may not be used for the development or training of artificial intelligence systems.

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com



Capital-Skill Complementarity Beyond Firm Boundaries*

Gert Bijmens[†]

Anders Humlum[‡]

Stijn Vanormelingen[§]

September 26, 2026

Abstract

New machinery has long been argued to raise the relative demand for skilled labor, yet firm-level evidence finds no such shift at investing firms. We resolve this puzzle by looking beyond firm boundaries. Linking Belgian data on capital investments, buyer–supplier transactions, and worker skills, we show that the skill bias of new machinery arises entirely at external suppliers. Following machinery investment spikes, purchases from suppliers raise firms’ total demand for high-STEM workers by 45%, three times the rise for other skill groups, while in-house hiring raises it by less than 5%, evenly across skill groups. This external demand comes first from machine integrators and later from complementary services such as software, and is stronger for smaller firms, lumpier investments, and rapidly advancing capital, consistent with fixed and adjustment costs of hiring high-STEM workers in-house. In a general-equilibrium model estimated on our event studies, the rise in machinery per worker from 2003 to 2021 raises relative demand for high-STEM workers by 8.8 log points and their wage premium by 3.1. Measured on payrolls alone, both would fall.

JEL codes: E22, J23, J24, J31, L24, O33.

*We thank Katharine Abraham, David Autor, Jim Bessen, Christian Dustmann, Maarten Goos, Federico Huneus, Erik Hurst, Xavier Jaravel, Patrick Kline, Pascual Restrepo, Chad Syverson, Felix Tintelnot, Joonas Tuhkuri, John Van Reenen, and seminar participants at the CESifo SI (2026), ESEM (2026), MEA (2024), NBER Productivity, Innovation, and Entrepreneurship (Spring 2026), NBER SI Labor Studies (2026), SOLE (2026), Bergamo, CBS, Chicago, Danmarks Nationalbank, KU Leuven, Masaryk University, the National Bank of Belgium, NYU Stern, RFBerlin, Rockwool, Southampton, UCLS, and Zurich for their helpful comments. We thank Chris Brijs (Crossroads Bank for Social Security) for his help with the data and Søren Peters (HowToRobot) for insightful conversations. Stijn Vanormelingen gratefully acknowledges financial support from F.W.O. Vlaanderen, grant number G0B7422N. The views expressed in this paper are those of the authors and do not necessarily reflect those of the National Bank of Belgium.

[†]National Bank of Belgium, gert.bijmens@nbb.be

[‡]University of Chicago, Booth School of Business, and NBER, anders.humlum@chicagobooth.edu

[§]KU Leuven, Department of Economics, stijn.vanormelingen@kuleuven.be

How do new machines affect the demand for worker skills? Since the capital–skill complementarity hypothesis of Griliches (1969), economists have argued that new machinery complements skilled labor. As Greenwood and Yorukoglu (1997) note, *“setting up and operating new technologies often involves acquiring and processing information. Skill facilitates this adoption process.”*

Yet despite its central role in many macro-labor models (Krusell et al., 2000), recent micro-level evidence of capital–skill complementarities has been elusive: armed with newly available firm-level data, a growing number of studies document that firms investing in new machinery do not alter their skill mixes (Aghion et al., 2025; Curtis et al., 2023; Hirvonen, Stenhammar and Tuhkuri, 2025).

In this paper, we argue that capital–skill complementarity operates *beyond* firm boundaries: investing firms mostly buy the skills needed to integrate and operate new machinery from external suppliers rather than hiring them in-house. Hence, while new machinery may appear skill-neutral on investing firms’ payrolls, it raises the demand for skills at their suppliers and thus in the labor market as a whole.

Robot integration illustrates this process. A rule of thumb in industrial robotics is that installation and integration account for two-thirds of the cost of a robot cell (IFR, 2022). Because most manufacturers lack the expertise to do this work in-house, they buy it from specialized integrators. The number of integrator firms in the U.S. rose by about 300% in the decade to 2020 (RIA, 2020).

To study this process systematically, we link Belgian firm-level data on capital investments, buyer–supplier transactions, and worker skills from 2003 to 2020, which lets us measure how firms adjust both in-house employment and external labor services around major capital investments. We measure a firm’s external employment using a key proportionality assumption: if a firm buys 10% of a supplier’s sales, it externally employs 10% of that supplier’s workforce. This assumption is standard in input–output analysis, and

we use our microdata to validate it in our setting.

To capture skill bias, we classify labor along two dimensions: the *level* of education (low refers to below-tertiary education; high to tertiary or higher) and the *field* of education (STEM vs. non-STEM). Under the capital–skill complementarity hypothesis, capital investment should increase the relative demand for high-STEM labor through two channels: installing and integrating new capital requires skilled workers (Goldin and Katz, 1998), as does operating more capital-intensive production (Griliches, 1969).

We adopt the “investment spikes” design from recent studies (e.g., Bessen et al., 2025; Aghion et al., 2025; Hirvonen, Stenhammar and Tuhkuri, 2025), analyzing how firms adjust their labor demand before and after large increases in machinery investment. Reflecting the lumpy nature of capital investment (Cooper and Haltiwanger, 2006), the machinery stock per worker more than triples (from about 19 to over 60 thousand euros) around a typical investment spike in our analysis sample. Our empirical strategy uses a difference-in-differences approach, comparing firms that experience an investment spike to similar firms in the same industry that have not yet done so.

While investment spikes are endogenous firm decisions, we conduct several validation checks of our identification strategy. First, the spikes design has been validated against quasi-experiments in prior work (Aghion et al., 2025; Hirvonen, Stenhammar and Tuhkuri, 2025). Second, spiking and non-spiking firms are on similar paths before the spikes, and controlling for their sales growth in the spike year leaves our conclusions intact. Finally, the skill bias of investment spikes is specific to capital embodying advanced technologies (Greenwood, Hercowitz and Krusell, 1997), arguing against generic confounders of large investments.

We establish three main findings on how firms adjust their skill demand around machinery investment spikes.

First, machinery spikes sharply raise the demand for high-STEM labor, and the skill

bias comes entirely from outside the firm. Total high-STEM employment (in-house plus external) rises by 50% in the spike year, against 11–20% for other skill groups, and remains 33% above baseline four years later. In-house employment grows evenly across skill groups, raising total employment by only 3–5% at the spike and 9–13% after four years, consistent with firm-level studies that observe only in-house payrolls.

Second, the external skill bias unfolds in two stages that match the two channels of capital–skill complementarity. Machine suppliers and integrators drive an initial burst, raising total high-STEM demand by about 15% in the spike year, as the investment–skill complementarity of Goldin and Katz (1998) predicts for the flow of new capital. The bias then persists but shifts toward skill-intensive services complementary to more capital-intensive production, such as software, as the capital–skill complementarity of Griliches (1969) predicts for the larger capital stock.

Third, the external skill bias varies systematically across investing firms and capital types. It is stronger for smaller firms and for services firms, which employ fewer high-STEM workers in-house before the spike. It is absent for intangible investments, which follow smoother investment profiles and rely far less on external suppliers. And it is specific to new machinery: building spikes generate a weaker bias that fades over time, while machinery from suppliers new to the firm, which is more likely to embody new technologies, generates the strongest bias of all, consistent with the capital-embodied technical change of Solow (1960) and Greenwood, Hercowitz and Krusell (1997).

Our main results are robust to the many ways one can construct staggered event studies: varying spike thresholds, single- versus repeat-spiking treated firms, matched versus cell-average comparison firms, clean-control versus later-spiking comparison firms, and balanced versus unbalanced panels. By comparing spiking firms only to not-yet-spiking firms, our stacked difference-in-differences design also circumvents econometric issues that may arise from staggered treatments under heterogeneous effects (Roth et al.,

2023).

To understand why firms do not hire machine-complementary skills in-house, we develop a conceptual framework that combines the canonical model of capital-skill and investment-skill complementarities (Goldin and Katz, 1998; Griliches, 1969) with two classical rationales for outsourcing: hiring capital-complementary skills entails both fixed and adjustment costs (Abraham and Taylor, 1996). Because of these costs, the high-STEM workers needed to deploy new machinery are often concentrated at specialized external suppliers, which can reap scale economies and smooth lumpy demand by serving multiple investing firms. The same costs also explain why these suppliers are especially important for smaller firms (for whom fixed costs are harder to absorb), for lumpy investment categories (which entail larger adjustment costs), and for technologically advanced capital types (which require more specialized skills to integrate and operate). This mechanism complements the wage-savings motive at the center of the literature on domestic outsourcing of lower-paid work (Goldschmidt and Schmieder, 2017): rather than shedding low-pay work to avoid firm wage premia, investing firms outsource their highest-skill work to specialized suppliers that pay a wage premium.

We organize our empirical analysis around a dynamic version of the framework with lumpy investment, as in Cooper and Haltiwanger (2006). This dynamic framework yields the outsourcing predictions we take to the data, and it formalizes the conditions under which investment spikes identify the skill bias of new machinery. In the model, a firm invests once its capital is far enough below target to justify the fixed cost of adjusting, so a spike is a discrete drop in its effective price of capital, the variation an ideal experiment would provide. Because output demand shocks raise all skill groups proportionally, spikes they trigger still identify the skill bias of new machinery. The dynamic framework also shows how our estimates map into the substitution elasticities of the nested-CES structure of Krusell et al. (2000).

Finally, we embed this dynamic firm problem in a quantitative general-equilibrium model in which firms differ in productivity, machinery intensity, and the hardware prices they face. We estimate the model to match our investment-spike event studies and other moments of the micro data, and take the elasticity of substitution between skill groups from the macro-labor literature, varying the labor supply elasticity. The estimated model implies that machinery substitutes about 1.15 times as easily for low non-STEM labor as for high-STEM labor in production, while integration makes an outsized contribution to the skill bias (36 percent of the rise in relative demand, despite employing only 4.3 percent of high-STEM workers). The model aggregates our micro-level estimates consistently and delivers a formula that expresses the equilibrium effects in terms of the aggregate rise in machinery, the shift in relative labor demand per unit of machinery, which our event studies measure, and the elasticities of relative labor supply and demand. The rise in machinery per worker between 2003 and 2021 raises the high-STEM wage premium by 3.1 log points and the relative employment of high-STEM workers by 4.6 log points, relative to low non-STEM workers. Measured on investing firms' payrolls alone, the same estimates would imply that new machinery lowered the skill premium.

By measuring skill bias beyond the boundaries of investing firms, our paper bridges two extensive but largely distinct literatures, on capital-skill complementarities and on firms' make-or-buy decisions. Empirically, we combine firm-level data on capital investments and employment (e.g., Aghion et al., 2025; Hirvonen, Stenhammar and Tuhkuri, 2025) with firm-to-firm transactions (e.g., Dhyne et al., 2025; Huneus, Kroft and Lim, 2025) to measure how the skill bias of new capital manifests within and outside investing firms. Approaches that do not observe firm-to-firm trade, whether aggregating to local labor markets (e.g., Acemoglu and Restrepo, 2020) or applying sectoral input-output tables (e.g., Humlum, 2022), miss most of this external skill bias, as we show in Section 5. Theoretically, we connect models of capital-skill and investment-skill comple-

mentarity (e.g., Goldin and Katz, 1998; Griliches, 1969) and make-or-buy decisions (e.g., Abraham and Taylor, 1996; Grossman and Helpman, 2002) to understand *where* capital-complementary skills are employed, and we embed the resulting firm-level skill bias in a general-equilibrium model to quantify *how much* it matters in the aggregate.

1 Conceptual Framework

To conceptualize our empirical analysis, we embed a canonical model of capital–skill and investment–skill complementarities (Goldin and Katz, 1998; Griliches, 1969) within firms’ decision to hire or outsource labor services. We use it to derive testable predictions about both complementarities and why firms outsource the required labor. The same framework guides the rest of the paper: Section 3.1.1 uses its dynamic version to motivate our investment-spikes design, and Section 6 solves a quantitative version in general equilibrium to quantify the aggregate implications.

1.1 Setup

Consider a firm j in year t . Production combines capital, skilled production labor, and unskilled production labor:

$$Y_{jt} = F(K_{jt}, S_{jt}^P, U_{jt}^P). \quad (1)$$

Deploying new capital requires hardware X_{jt} and skilled and unskilled integration labor, which together produce installed capital I_{jt} :

$$K_{jt+1} = (1 - \delta)K_{jt} + I_{jt}, \quad I_{jt} = G(X_{jt}, S_{jt}^I, U_{jt}^I), \quad (2)$$

where capital depreciates at a rate δ .

For each labor category $\mathbf{L} \in \{U^P, S^P, U^I, S^I\}$, firms can either hire in-house (H) or

source from external suppliers (O , for outsourced):

$$\mathbf{L}_{jt} = \mathbf{L}_{jt}^H + \mathbf{L}_{jt}^O. \quad (3)$$

Firms choose investment and labor (in-house and external) to maximize the present-discounted sum of profits. In-house labor may involve fixed and adjustment costs, and investment costs may be non-convex. Appendix A states the firm's full dynamic program.

1.2 Model Predictions

Our first prediction relates firms' capital investments to their demand for skills (tested in Section 3). Following Griliches (1969), operating a more capital-intensive production process may require additional skills, such as programmers and engineers running automated production lines. Following Goldin and Katz (1998), integrating new capital may require skilled workers such as design engineers and robot integrators.

Prediction 1 (Capital investment and the demand for skills). *Capital investment can raise the relative demand for skills through two channels:*

Prediction 1(a) (Capital–skill complementarity in production). *If capital complements skilled production labor more than unskilled production labor, a larger capital stock K raises the relative demand for skilled production workers S^P (Equation 1).*

Prediction 1(b) (Investment–skill complementarity in integration). *If hardware complements skilled integration labor more than unskilled integration labor, a larger flow of new hardware X raises the relative demand for skilled integration workers S^I (Equation 2).*

Because Prediction 1(a) works through the capital stock and Prediction 1(b) through the flow of new capital, an investment spike should raise skill demand persistently in production and in a burst around installation in integration. See Appendix A.1 for formal derivations.

The remaining predictions concern whether firms hire capital-complementary skills in-house or buy them from external suppliers (tested in Section 4), and Predictions 2 and 3 draw on the two outsourcing motives identified by Abraham and Taylor (1996).

First, hiring capital-complementary skills in-house, such as robot technicians or software engineers, may entail fixed costs: a firm must maintain a dedicated team, equip it, and build up expertise across projects. External suppliers can spread these costs across many investing firms (Abraham and Taylor, 1996, § II.C).¹ Second, capital investments often arrive in occasional spikes, generating volatile demand for integration labor. In-house adjustment costs (e.g., hiring or firing costs) may make this volatility difficult to absorb: an in-house integrator would have to be hired and let go around each investment event. Suppliers, by contrast, smooth their workload across clients with offsetting demand cycles (Abraham and Taylor, 1996, § II.B).

Prediction 2 (Fixed costs and outsourcing). *If hiring capital-complementary skills in-house entails a fixed cost, firms outsource these skills when their demand for them is too small to cover that cost. Smaller firms and firms in less capital-intensive sectors (e.g., services) should therefore rely more on external suppliers.*

Prediction 3 (Lumpy investments and outsourcing). *If adjusting in-house employment is costly (e.g., hiring and firing costs), firms outsource capital-complementary skills when their investment is lumpy enough that in-house staff would sit idle between investment episodes. Outsourcing should therefore be more prevalent for lumpier investment categories (e.g., machinery rather than intangibles).*

Prediction 4 (Advanced capital and outsourcing). *If these fixed and adjustment costs are*

¹Fixed costs are one source of scale economies in hiring capital-complementary skills. One such cost is acquiring and maintaining the expertise for an application, which a supplier pays once and uses across many clients: integrators and other machine suppliers compete on portfolios of past installations, signaling that they have this expertise and keep it current as the technology advances (cf. Prediction 4). The Belgian supplier of machine-tending automation RoboJob, for example, advertises its more than 2,000 installations (RoboJob, 2026); industry bodies collect similar case studies (IFR, 2026).

larger for more specialized skills, both outsourcing motives should be strongest for capital undergoing rapid technological change (Greenwood, Hercowitz and Krusell, 1997) (e.g., machinery rather than buildings, or machinery from new rather than existing suppliers).

See Appendix A.2 for formal derivations.

A final motive, and perhaps the most commonly considered, is to save directly on wages or benefit payments (Abraham and Taylor, 1996, § II.A), e.g. by avoiding firm wage premia (Bilal and Lhuillier, 2025; Goldschmidt and Schmieder, 2017). We examine it in Section 4.6 and Appendix D.6 by testing whether high-wage investing firms outsource capital-complementary skills more than low-wage ones, and by comparing the wages these workers earn at external suppliers and at investing firms.

2 Data and Measurement

2.1 Data

We use comprehensive administrative data from Belgium that link firms’ annual accounts on capital and investment with business-to-business transaction records and employee-level education data from 2003 to 2020. The datasets share firm identifiers, so we can merge them at the firm level; Appendix B describes them in detail.

Annual Accounts and VAT Declarations. We observe the annual accounts of all firms in Belgium, which disaggregate tangible fixed assets into buildings and land, machinery and equipment, and office and rolling stock, along with intangible fixed assets. We complement these with Value Added Tax (VAT) records on firms’ sales and purchases (by category and origin).

Business-to-Business Transactions. Our data on business-to-business transactions come from VAT customer listings: all VAT-liable Belgian firms report annual sales to each

customer firm above €250, yielding comprehensive domestic buyer–supplier linkages.²

Education Data. We obtain employee-level data from the social security data warehouse (KSZ), which records each worker’s highest education by *level* (e.g., below-tertiary vs. tertiary or higher) and *field* (e.g., STEM vs. non-STEM).³ Motivated by the technical skills required to deploy new machinery (Section 1), we classify workers into four categories along these two dimensions: high-STEM, high non-STEM, low STEM, and low non-STEM. Employees are matched to their employers via social security declarations, yielding firm-year counts in full-time equivalents for each category.⁴ We are the first to link these employee-level education data with firm-to-firm transactions.⁵

2.2 Measurement

2.2.1 Investment Spikes

We define a firm as experiencing a machinery investment spike if its change in machinery stock per worker reaches the 90th percentile of positive changes across all firm-years:

$$\text{Spike}_{jt} = \mathbf{1}[i_{jt} \geq i_{P90|+}] \quad \text{with} \quad i_{jt} = \frac{\Delta \text{Machinery}_{jt}}{\frac{1}{2}(L_{jt}^H + L_{jt-1}^H)}, \quad i_{P90|+} = \text{€}17,115 \quad (4)$$

²We observe only aggregate transaction values, without detail on transaction type or product code. Data construction follows Dhyne, Duprez and Komatsu (2023); the dataset has been used by Bernard et al. (2022), Baqaee et al. (2025), and Bijmens, Jäger and Schoefer (2025), among others. Our analysis restricts to pairs in which both firms file an annual account and the annual transaction value exceeds €1,000 (Appendix B).

³STEM comprises ISCED fields 05 (Natural Sciences, Mathematics, Statistics), 06 (Information and Communication Technologies), and 07 (Engineering, Manufacturing, Construction). Appendix Tables B.1 and B.2 provide the full classification.


⁴We focus on employment rather than wages due to data limitations: worker-level earnings in the KSZ data are only reported in coarse, top-coded bins, and firm-level wage bills from the annual accounts are not broken down by skill group. That said, because machine suppliers pay higher wages than investing firms (Section 2.2.3), our focus on employment tends to understate the impact of new machinery on the skill bias in wage bills.

⁵Firm-to-firm transaction data have been linked to matched employer–employee records in Belgium (Dhyne et al., 2025), Chile (Huneus, Kroft and Lim, 2025), and Turkey (Demir et al., 2024). These studies examine how demand shocks affect earnings through the production network; we study the skill bias of new machinery, in-house and at suppliers.

where $L_{jt}^H = \mathbf{1}^T \mathbf{L}_{jt}^H$ denotes in-house employment of firm j in year t .⁶

Figure 1 illustrates the lumpiness of machinery investment: these spikes triple the machinery stock per worker (from 19 to 61 thousand euros). The investment flow that produces this step-up is a one-off event: total fixed investment peaks at nearly 20% of sales in the spike year before returning to baseline (Appendix Figure C.5).

Figure 1: Machinery Investment Spikes



Notes: This figure plots the machinery capital stock per worker around machinery investment spikes. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

2.2.2 Machine Supplier Industries

Machine suppliers comprise the builders, installers, distributors, and advisors of new machinery; robot integrators, for example, often combine distribution, installation, and advisory services. Appendix B.2.2 describes the industrial organization of machine supply, classifies the industry codes for these suppliers, and validates the classification against

⁶While the specific cutoff at the 90th percentile is arbitrary, Appendix C.3.1 shows that our results are robust to alternative thresholds. Because spikes are defined in machinery changes per worker, the cutoff also does not mechanically select firms by size; indeed, smaller and larger spiking firms experience similar-sized spikes (Appendix Figures D.1 and D.2).

external lists of robot suppliers in Belgium. Figure B.1 shows that firms undergoing an investment spike sharply increase their share of purchases from machine supplier industries, with total purchases averaging 1.7 times direct investment expenditures, a first sign that external suppliers may shape the employment effects of new machinery.⁷

2.2.3 External Employment

To measure external employment, we link firm-to-firm transactions to suppliers’ payrolls. We rely here on a proportionality assumption common in input–output analysis: for example, if firm j purchases 10% of firm k ’s sales, we assume that firm j externally employs 10% of firm k ’s workers in each skill group.⁸

Table 1: Firm Employment (Year of Investment Spike)

	Total (FTE)	Share of Total (%)			
		High STEM	High Non-STEM	Low STEM	Low Non-STEM
In-House Employment	39.7	5.0%	14.2%	25.5%	55.3%
External Employment	22.9	6.7%	15.1%	26.8%	51.4%
Machine Suppliers	3.6	13.9%	15.4%	35.9%	34.8%

Notes: This table shows the average in-house and external employment of spiking firms. *Sample:* Private sector firms in the year of their first machinery investment spike.

Firm j ’s first-order external employment, by skill group, is:

$$\mathbf{L}_{jt}^O = \sum_k \frac{T_{kjt}}{T_{kt}} \mathbf{L}_{kt}^H, \quad (5)$$

⁷Appendix B.2.2 further shows that the machine supplier industries account for the vast majority of industrial robot imports into Belgium; Békés et al. (2025) document a similar pattern in Hungary. While the physical machines are highly tradable (e.g., Belgium produces no industrial robots), integration is overwhelmingly done locally: domestic integrators import the machines and bundle them with installation and programming services for the end-user firm. The external market lists in the same appendix identify 606 robot supplier firms in Belgium in 2021. Our domestic-firm data therefore capture the integration channel but undercount the labor that goes into developing and producing the physical machines abroad.

⁸This assumption is more plausible for specialized suppliers whose workforce is dedicated to the same set of products or services, a good description of the machine integrators and other machine-complementary suppliers central to our findings. As Section 4.5 shows, these suppliers’ workforces are highly specialized, with stable skill mixes invariant to customer composition.

where T_{kjt} denotes the sales of firm k to firm j in year t , T_{kt} the total sales of firm k , and $\mathbf{L}_{kt}^H$ is the in-house employment of firm k . In effect, $\mathbf{L}_{jt}^O$ captures the skill content embodied in firm j 's purchases, a firm-level domestic analogue of the factor content of trade (e.g., Davis and Weinstein, 2001; Trefler, 1995). Section 3.3 shows that our findings are robust to including higher-order linkages (i.e., suppliers of suppliers).

Table 1 reports that the average spiking firm employs 40 workers in-house and sources external labor equivalent to 23 full-time workers, including around 4 from machine supplier firms. Appendix B.2.1 illustrates this measurement for a typical robot-integration project, detailing how the external labor for installation and servicing maps into the annual firm-to-firm transactions we observe.

Notably, machine suppliers' workforces are nearly three times as high-STEM-intensive as investing firms' in-house workforces (13.9% versus 5.0%), a first indication that machine suppliers are a source of skill bias from new machinery. Moreover, machine suppliers pay about 3.6% higher wages than investing firms do, holding skill composition fixed (Appendix Table D.6), consistent with the efficiency gains from outsourcing we explore in Section 4.

3 New Machinery and the Demand for Skills

In this section, we study how firms change their demands for skills around investment spikes in new machinery, analyzing both their in-house and external employment. Section 3.1 describes our empirical strategy and considers threats to identification, Section 3.2 presents our main results, Section 3.3 shows their robustness, and Section 3.4 connects them to theory.

3.1 Empirical Strategy

The ideal experiment to test our model’s predictions would exogenously vary investment costs to induce variation in investment I and capital K , and trace the resulting effects on labor demand for integration and production tasks, both within and outside the investing firm. We approximate this with the “investment spikes” design pioneered by Bessen et al. (2025) and applied in Aghion et al. (2025), Hirvonen, Stenhammar and Tuhkuri (2025), and others. The idea is to exploit the lumpiness of capital investment to isolate episodes with sharp increases in I and K . In our framework, as in Cooper and Haltiwanger (2006) and Bloom (2009), this lumpiness stems from non-convex investment costs, and Section 3.1.1 shows why the resulting spikes identify the skill bias of new machinery. We compare the labor adjustments of spiking firms (Section 2.2.1) to similar non-spiking firms in a stacked difference-in-differences framework.

We apply three sample restrictions. First, we focus on private sector firms with at least five FTE employees (in-house plus external) in the year before the spike. Second, we use firms’ first investment spike (28% of spiking firms experience follow-on spikes in the post-event window). Third, we study a balanced sample for four years before and after each spike. This yields an analysis sample of 3,863 spiking firms (Appendix Table C.1). Section 3.3 shows our results are robust to varying these sample restrictions.

For each spiking firm in year t , we randomly select one control firm from the same NACE 3-digit industry in year t that has at least five FTE employees in $t - 1$, has not spiked before t , and is observed for the full event-study window. Table 2 shows that the two groups have comparable in-house and external employment, and comparable skill compositions within each, two years before the spike, despite not being matched on these outcomes.⁹

⁹We choose a two-year pre-spike comparison because investing firms begin engaging machine integrators (especially for design and advisory services) before the physical investment: external high-STEM employment is already elevated in the year before the spike, while the machinery stock is still flat (Fig-

Neither group is restricted after the event year: treated firms may undertake follow-on spikes, and control firms may spike later in the window (Appendix C.3.5 shows that about 8% do so within the post-event window). Our estimates therefore identify the effect of investing at e relative to firms’ counterfactual investment policies, rather than relative to never investing. Still, our results are robust to restricting either group to firms that do not spike within the post-event window: the comparison group to “clean controls” (Appendix C.3.5, as in Cengiz et al., 2019) and the treatment group to “single-spiking” firms (Appendix C.3.2).

Our main specification is a stacked difference-in-differences in outcomes Y between the spiking ($\text{Spike} = 1$) and non-spiking ($\text{Spike} = 0$) firms around a spike-event year, with changes indexed relative to two years before the spike:

$$Y_{sjt} = \beta_1 \text{Spike}_{je} + \sum_k \beta_{0k} \mathbf{1}_{\{t=e+k\}} + \sum_{k \neq -2} \beta_{1k} \text{Spike}_{je} \mathbf{1}_{\{t=e+k\}} + \varepsilon_{sjt}, \quad (6)$$

where Spike_{je} equals one if firm j undergoes an investment spike (Equation 4) in event year e and zero for its comparison firms, $\mathbf{1}_{\{t=e+k\}}$ are event-time indicators that equal one if year t is k years from the event year e , and β_{1k} are the coefficients of interest, identifying the effect of investment spikes under parallel trends. We estimate Equation (6) by OLS separately for each skill group s , clustering the standard errors at the event (treated–control pair) level. Because our stacked difference-in-differences design compares each spiking firm only to not-yet-spiking firms in the same industry-year, it avoids the biases that can arise in two-way fixed-effects estimators when treatment is staggered and effects are heterogeneous (Callaway and Sant’Anna, 2021; Roth et al., 2023; Sun and Abraham, 2021).

Our outcomes of interest, Y_{sjt} , are firms’ total employment in each skill group s (high-STEM, high non-STEM, low STEM, and low non-STEM), measured in full-time

ure 2); see also Figure B.1 and the project timeline in Appendix B.2.1.

Table 2: Firm Employment before Machinery Investment Spikes

	Spiking		Non-Spiking		Std. Difference	
	In-House (1)	External (2)	In-House (3)	External (4)	In-House (5)	External (6)
Total employment	38.6 (323.9)	19.2 (130.7)	42.9 (320.3)	18.5 (160.8)	-0.01	0.00
Share of employment (%)						
— High STEM	4.9 (10.5)	5.8 (5.8)	4.7 (10.1)	5.3 (5.4)	0.02	0.09
— High non-STEM	13.9 (18.9)	16.4 (10.4)	14.7 (19.0)	17.1 (11.2)	-0.04	-0.06
— Low STEM	26.1 (22.4)	24.8 (10.6)	24.3 (20.6)	23.8 (10.6)	0.08	0.09
— Low Non-STEM	55.1 (25.1)	53.0 (12.0)	56.4 (23.7)	53.8 (12.2)	-0.05	-0.07
Nr. Firms	3,863		3,863			

Notes: This table reports in-house and external employment in the treatment (Columns 1–2) and control (Columns 3–4) groups in year $t - 2$. The table reports total employment in full-time equivalents and the employment share of different skill groups. Columns 5–6 show the standardized difference between the treatment and control group means, i.e. the difference in means divided by $\sqrt{(s_T^2 + s_C^2)}/2$, where absolute values above 25% are a standard threshold for assessing imbalance (Stuart and Rubin, 2008). *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

equivalents and summing in-house and external labor. To obtain relative effects, we express these as percentages of pre-spike levels (averaged over years -3 to -1), and winsorize the resulting variables at the 5th and 95th percentiles to mitigate the influence of outliers. An advantage of studying total employment (i.e., in-house plus external) for each skill group is that it circumvents the challenge of studying relative effects for outcomes that often contain zeros (Chen and Roth, 2024): for example, 64% of spiking firms have zero in-house high-STEM employment prior to the spike, likely reflecting the fixed costs of hiring emphasized in Section 1.2, whereas all firms in our sample have positive total pre-spike employment in each skill group.

To understand the sources of the skill bias, we further decompose the skill-specific employment outcomes by supplier source: separating in-house from external labor, and within external labor, distinguishing machine suppliers from other industries. We express

each component as a share of the skill group’s total pre-spike employment:

$$Y_{sjt}^M = \frac{L_{sjt}^M}{L_{sjpre}^H + L_{sjpre}^O}, \quad M \in \{H, O\}, \quad (7)$$

where L^H and L^O denote in-house and external employment and the “pre” subscript averages over years -3 to -1 . Because both components share the same denominator, they are directly comparable and add up to the total: $Y_{sjt}^H + Y_{sjt}^O$ recovers the skill group’s total-employment response. This normalization also avoids dividing by the in-house zeros noted above. The larger external responses we estimate below are therefore not an artifact of a smaller base. If anything, the common denominator works against our findings: it scales external responses by a base dominated by in-house employment (Table 2), and machine-supplier responses by one dominated by other industries (Table 1).

3.1.1 Threats to Identification

Investment spikes are endogenous firm decisions. The threat to identification is that spikes coincide with changes in firm fundamentals — production (F), investment technology (G), or output demand — that themselves induce skill bias. We first assess this threat empirically through five validity checks. We then show that the dynamic version of our conceptual framework, a standard model of lumpy investment, implies the validity of our design.

Validity Checks. Our first validation of the investment-spikes design comes from the existing literature: Hirvonen, Stenhammar and Tuhkuri (2025) and Aghion et al. (2025) show that it yields conclusions similar to those from quasi-experiments that use technology subsidies and foreign machine-supply shocks to shift investment costs.

Second, in our setting, spiking and non-spiking firms are similar before the spike (Table 2), even though comparison firms are selected only on industry and year, and the event-study figures below show that their outcomes evolve in parallel until the year before

the spike, when investing firms begin engaging integrators.

Third, the spikes themselves are discontinuous jumps in investment and capital (Figure 1). Provided fundamentals evolve smoothly, the discrete onset of the spikes lends greater credibility to short-run responses around the spike year than to longer-run adaptations.

Fourth, Appendix C.1.1 directly considers confounding output demand or productivity shocks: sales evolve smoothly around the spikes, and controlling for their growth leaves our conclusions intact (Appendix Figure C.2).

Fifth, our heterogeneity analyses across capital types help argue against confounders correlated with undertaking large investments, which would affect all capital types similarly. Instead, building spikes generate weaker bias (Section 4.3), intangible spikes leave the skill mix essentially unchanged, and machinery from new suppliers generates the strongest bias of all (Section 4.4), consistent with capital-embodied technical change (Greenwood, Hercowitz and Krusell, 1997; Solow, 1960).

Identification in the Model. In the dynamic version of our framework (Appendix A.3), a firm invests once the gap between its target and actual capital is large enough to justify the fixed cost of adjusting, so a spike arrives as a discrete drop in the firm’s effective price of capital (our ideal experiment), whatever provided the final push. In particular, because output demand shocks raise all skill groups proportionally (Lindner et al., 2024), spikes they trigger still identify the skill bias of new machinery, our parameter of interest.

Our design is also consistent with dynamic production-function estimators that combine time-to-build with Markov dynamics of firm productivity (Doraszelski and Jaumandreu, 2013; Olley and Pakes, 1996). Under these assumptions, a firm’s pre-spike factor choices encode its expected future technology, so firms with similar pre-spike outcomes, as ours are (Table 2), share similar expected technologies after the spike (Appendix A.3).¹⁰

¹⁰The same logic explains why neither group is restricted after the event year: restricting firms to never

Together, these checks and arguments show that our identifying assumptions are both empirically plausible and consistent with standard models of capital investment. The skill bias around investment spikes thus reflects the investment episodes themselves rather than confounding changes in firm fundamentals.

3.2 Main Results

Figure 2 presents our first key result, plotting employment by skill group separately for in-house (Panel (a)) and external labor (Panel (b)). Both are measured relative to the group’s total pre-spike employment (Equation 7), so the two panels sum to the total effect.

Total employment (summing Panels (a) and (b)) shows a sharp rise in high-STEM demand, which jumps by 50% in the spike year and remains 33% above baseline four years later. The other skill groups, by contrast, grow by only 11–20% in the spike year and 19–27% four years later.

The key finding is that this skill bias is sourced entirely externally. In-house employment grows broadly evenly across skill groups, if anything less for high-STEM workers (Panel (a)), raising total employment by only 3–5% at the spike and 9–13% after four years, while the high-STEM surge comes from external suppliers (Panel (b)).¹¹

Which external suppliers drive this effect? Figure 3 decomposes external employment into machine suppliers (Panel (a)) and other industries (Panel (b)), revealing a two-stage pattern.


Machine suppliers drive much of the initial spike (Panel (a)), raising total high-STEM demand by 15% in the spike year. Their high-STEM contribution never fully disap-

spike within the event window implicitly conditions on future productivity innovations, selecting firms with systematically unfavorable realized shocks. Consistent with this, our estimated employment effects are larger if we restrict control firms to those that do not spike in the post-event window (Appendix C.3.5), and smaller if we restrict treatment firms to single-spiers (Appendix C.3.2).


¹¹Appendix Figure C.3 shows that the larger effects for external high-STEM labor reflect both that external employment increases more than in-house employment and that spiking firms shift their external employment toward high-STEM skills.

Figure 2: Employment of Skill Groups around Machinery Investment Spikes
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)

(a) In-House



(b) External




Notes: This figure shows the employment of high-STEM, high non-STEM, low STEM, and low non-STEM in percent of total pre-spike employment by skill group (Equation 7). Panel (a) considers in-house employment, while Panel (b) considers external employment at supplier firms. Total pre-spike employment is the sum of in-house and external employment. Pre-spike refers to the average over the years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.

pears, settling around 5% four years later, consistent with the recurring maintenance, support, and reprogramming services that follow an installation (see, for example, the robot-integration project in Appendix B.2.1). By contrast, employment of other skill groups through machine suppliers moves little in the investment year and fades entirely over time.

Other industries account for most of the longer-run gains (Panel (b)), raising total high-STEM demand by 23% at the spike and a persistent 16% four years later. The external skill bias of new machinery thus extends well beyond machine supply and integration. Appendix C.2.2 shows that this rise is widespread across supplier sectors but most pronounced in computer programming and software, skilled services complementary to a more machine-intensive production process. Nor are these gains confined to technical skills: high non-STEM employment through other industries also rises and remains elevated, reaching the second-largest response by year 4, which suggests that more capital-intensive production also draws on commercial and business services. The largest high non-STEM gains come from wholesale distribution, and management consulting and legal services add more high non-STEM than high-STEM workers (Appendix Table C.2).

Figure 3: External Employment by Supplier Industry
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)



Notes: This figure shows the external employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Panel (a) considers external employment of machine suppliers, while Panel (b) considers external employment from all other industries. Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. Because each outcome is winsorized separately at the 5th and 95th percentiles, the two components need not sum exactly to the external employment estimates in Figure 2. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.

3.3 Robustness

Our main result, that the high-STEM skill bias of machinery investments arises entirely at external suppliers, holds across the key choices in our design: how we define spikes, the sample, the comparison group, and how far up the supply chain we count external employment.

Investment Spikes. Our baseline defines an investment spike at the 90th percentile of positive changes in machinery stock per worker; our key patterns are highly similar at the 85th and 95th percentiles (Appendix C.3.1) and when we restrict the treatment group to single-spiking firms (Appendix C.3.2).

Unbalanced Panel. Our main results use a balanced panel of firms observed four years before and after the event year. Including firm exits as zero employment leaves our

findings intact (Appendix C.3.3).

Control Group. As Appendix C.3.5 shows, our results are robust to restricting the comparison group to “clean controls” that do not spike within the post-event window (as in Cengiz et al., 2019) or to “later spikers” that spike only after it (as in Bessen et al., 2025). They also remain virtually unchanged if we compare spiking firms with industry-year averages rather than a randomly selected not-yet-spiking firm (Appendix C.3.4). The reason for this robustness is simple: the sharp break at the spike comes entirely from the spiking firms, while the comparison firms evolve smoothly throughout. Appendix C.2.1 plots total employment separately for spiking and comparison firms.

Higher-Order Links. Our main analysis counts the external labor that firms employ through their first-order suppliers. Including higher-order links (suppliers of suppliers) only strengthens the importance of external relative to in-house employment (Appendix C.3.6).

3.4 Interpreting the Results

Figures 2 and 3 provide clear empirical support for Prediction 1 of Section 1.2. Consistent with the investment–skill complementarity of Goldin and Katz (1998) (Prediction 1(b)), integrating new machinery sharply raises demand for skilled workers around the spike year, especially from machine suppliers. Consistent with the capital–skill complementarity of Griliches (1969) (Prediction 1(a)), operating more machine-intensive production further raises the demand for skilled labor in complementary services such as software, with effects that persist well beyond the investment event. These dynamics mirror the investment profile: the temporary integration response coincides with the one-off investment flow (Appendix Figure C.5), whereas the persistent production response tracks the machinery stock, which steps up and remains elevated (Figure 1).

The key surprise — and our main contribution — is that both complementarities operate entirely through external suppliers. The skill-neutral in-house effects in Figure 2(a) reproduce a growing set of firm-level studies showing that new machinery investments raise employment without raising the in-house skill share (Aghion et al., 2025; Curtis et al., 2023; Hirvonen, Stenhammar and Tuhkuri, 2025). Because these studies observe only in-house payrolls, they conclude that the evidence challenges the capital–skill complementarity hypothesis.¹²

Figure 2(b) resolves this apparent puzzle by showing that external suppliers are the missing channel linking capital accumulation to skill bias, a channel uniquely visible through our firm-to-firm transaction data.¹³ Indeed, the classical theories of capital–skill and investment–skill complementarities take no stand on where the required skills are employed: they are theories of production and investment, not theories of the firm. The importance of external suppliers for machinery investments corroborates the anecdotal evidence on robot integrators that initially motivated our study. Our comprehensive buyer–supplier data go further: the external high-STEM bias is pervasive across supplier types, and even the long-run increase in high-STEM labor, which reflects more capital-intensive production, is outsourced. Many of these are tradable services such as software, sourced mostly from outside the firm’s local labor market (Section 5.1), which helps explain why they too are not brought in-house.

Our estimates also speak to the substitution elasticities at the heart of the capital–skill

¹²Curtis et al. (2023) conclude that “*our results are not compatible with the capital–skill complementarity hypothesis.*” Motivated by the same framework, Hirvonen, Stenhammar and Tuhkuri (2022) write that “*The first part of the paper reports results in sharp contrast with the ideas that technologies necessarily reduce employment or are skill biased.*” Aghion et al. (2025) similarly note that their “*...results differ from the estimates of seminal structural models of capital–skill complementarity.*” Relatedly, Berlingieri et al. (2026) find modestly positive capital–skill complementarity at the firm level, though still an order of magnitude smaller than what would be required to explain the aggregate rise in skill demand attributed to capital by Krusell et al. (2000).

¹³Bilgin, Faia and Ottaviano (2026) use similar firm-to-firm transaction data to study how robot adoption affects the *performance* of *downstream* customer firms; we instead identify the *skill bias* of a broader set of capital investments operating through investing firms’ *upstream* supplier networks.

complementarity literature. Following Krusell et al. (2000), suppose that Equation (1) is a nested CES that combines machinery with high-STEM labor and that composite with low non-STEM labor. Appendix A.4 then maps the skill differential and the machinery deepening around a spike into the ratio of the outer-nest elasticity (σ_{HU}) to the inner-nest elasticity (σ_{KS}):

$$\frac{\sigma_{HU}}{\sigma_{KS}} = 1 + \frac{1}{s_K^H} \cdot \frac{\Delta \ln L_S - \Delta \ln L_U}{\Delta \ln K - \Delta \ln L_S}, \quad (8)$$

where $\Delta \ln L_S$, $\Delta \ln L_U$, and $\Delta \ln K$ are the four-year responses of high-STEM employment, low non-STEM employment, and the machinery stock, and s_K^H is machinery’s cost share within the inner nest. This mapping is invariant to the spike’s trigger, whether an investment-cost draw or growth in output demand (Appendix A.3). Evaluated on our estimates, however, it mixes two complementarities: four years after a spike, total employment still includes the integration labor that maintains the larger machinery stock. Section 6 therefore estimates the two stages jointly in a quantitative version of our framework, separating the reduced-form skill bias into capital–skill and investment–skill complementarity.

Finally, our results reinforce a key conclusion from existing firm-level estimates (e.g., Aghion et al., 2025; Hirvonen, Stenhammar and Tuhkuri, 2025): machinery investments are associated with increased employment, and the effect only grows once we count external employment.¹⁴

¹⁴This conclusion differs from the negative employment effects of industrial robots in local labor markets estimated by Acemoglu and Restrepo (2020). Several elements help explain the difference. First, because external labor complementary to new machinery is mostly sourced from suppliers outside the local labor market (see Section 5.1), this channel is not captured by a local-labor-market analysis. Second, local-labor-market analyses capture other spillovers not considered here, including competition in product markets (especially for non-tradables) and poaching in labor markets. Finally, as Benmelech and Zator (2025) show, industrial robots represent only a small fraction of the total firm-level machinery investments we study.

4 Why Firms Outsource Capital-Complementary Skills

When and why do firms outsource capital-complementary skills? Our framework (Section 1.2) predicts where outsourcing should be most pronounced: among smaller firms, whose demand is too small to cover the fixed costs of hiring in-house; for lumpier investments, where hiring and firing costs bite hardest; and for capital undergoing rapid technological change, where specialized skills matter most. We study heterogeneity across investing firms (Sections 4.1–4.2) and capital types (Sections 4.3–4.4), turn to direct evidence from machine suppliers (Section 4.5), and discuss what the patterns reveal in Section 4.6.

4.1 Firm Size

Figure 4 separates spiking firms into smaller and larger firms based on their pre-spike total employment, split at the median (13.1 FTE). Figure D.1 shows that both smaller and larger spiking firms experience a similar increase of roughly 40 thousand euros in machinery stock per worker, so differences in responses do not reflect differences in the size of the investment.¹⁵ Smaller firms rely much more heavily on external high-STEM workers to implement their investment spikes: external suppliers raise total high-STEM employment by 59% in smaller firms versus 37% in larger firms in the spike year, persisting to 35% vs. 19% four years later.¹⁶ By contrast, in-house employment expands similarly across firm sizes, raising total employment by less than 5% at the spike and roughly 10% after four years, with little difference across skill groups.


Larger firms respond less partly because they are more likely to employ high-STEM workers in-house before their spikes (Appendix Tables D.1 vs. D.2). The difference between smaller and larger firms in Figure 4 thus partly reflects where they already sourced

¹⁵Figure D.2 shows that their difference-in-difference estimates are similar, too.

¹⁶Consistent with smaller firms relying more on outsourcing to facilitate their investment spikes, Appendix Figure D.3 shows that the share of machine suppliers in their total purchases rises by 6 percentage points in the spike year, compared to 4 percentage points for larger firms.

these skills, rather than a weaker skill bias of new machinery at larger firms.¹⁷

Figure 4: Employment by Pre-Spike Firm Size
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)




Notes: This figure shows employment around machinery investment spikes, splitting by the size of the investing firm. Smaller firms have pre-spike total employment (year -1) below the median (13.1 FTE). Panels (a) and (b) consider in-house employment, while Panels (c) and (d) consider external employment at supplier firms. The figures show employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment. Pre-spike refers to the average over the years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

¹⁷This also suggests that the investment-spikes design understates the full skill bias of capital investment for larger firms: where a standing in-house capability already covers the skills a new machine requires, some of the corresponding hiring took place alongside earlier investments, outside the event window of any single spike.

4.2 Application Sector

Figure 5 separates spiking firms by sector, with manufacturing firms shown in Panels (a) and (c) and services firms in Panels (b) and (d).¹⁸

Figure 5: Employment around Machinery Investment Spikes: Sector Heterogeneity
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)



Notes: This figure shows employment around machinery investment spikes for manufacturing (Panels (a) and (c)) and services (Panels (b) and (d)). Panels (a) and (b) consider in-house employment, while Panels (c) and (d) consider external employment. The figures show employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment. Pre-spike refers to the average over the years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around first investment spike. Balanced sample for four years before and after the spike.

Although machinery investments are more prevalent in manufacturing,¹⁹ Panels (c)

¹⁸Figure D.4 shows that the remaining private-sector industries (agriculture, utilities, and construction) display similar patterns.

¹⁹Figure E.1 shows that the machinery stock per worker is 90,000 euros in manufacturing, compared to 60,000 in the broader economy.

and (d) show that the external skill bias is pervasive across sectors, extending well beyond robotics and manufacturing. If anything, the shift toward external high-STEM labor is larger in services, mirroring the pattern for smaller firms: services firms also employ fewer high-STEM workers to begin with (Appendix Tables D.3 vs. D.4). In-house employment again grows evenly across skill groups, and by similar amounts in manufacturing and services (Panels (a) and (b)).


4.3 Capital Types

Figure 6 replicates the investment-spikes design for buildings (Panels (a) and (c)) and intangibles (Panels (b) and (d)), restricting to manufacturing firms to hold the type of investing firm fixed.

Building spikes generate weaker and less persistent skill bias. External high-STEM employment grows only two to three times as much as low STEM and low non-STEM (compared to three to four times for machinery in Figure 5), and this differential largely dissipates by year two.²⁰ Intangible spikes generate weaker skill bias still: their external employment responses are essentially flat across skill groups.

²⁰Part of the skill bias from building spikes reflects simultaneous machinery installation (Appendix Figure D.6): isolating building spikes that do not coincide with machinery spikes further reduces the skill bias (Appendix Figure D.7), though it does not disappear entirely, likely reflecting the involvement of architects and building engineers.

Figure 6: Employment around Investment Spikes: Heterogeneity by Capital Type
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)




Notes: This figure shows in-house and external employment of manufacturing firms around investment spikes in buildings (Panels (a) and (c)) and intangibles (Panels (b) and (d)). Panels (a) and (b) consider in-house employment, while Panels (c) and (d) consider external. The figures show employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment. Pre-spike refers to the average over the years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. Figure 5, Panels (a) and (c) show the corresponding estimates for machinery spikes. Appendix Figure D.5 plots the capital stock per worker around each type of spike. *Sample:* Manufacturing firms around first investment spike. Balanced sample for four years before and after the spike.

4.4 New versus Existing Machine Supply

To gauge whether machinery investments embody new technologies rather than merely replacing or deepening existing capital, Figure 7 separates machinery spikes by whether the dominant machine supplier is new to the investing firm. External high-STEM skill bias is substantially stronger for new-supplier spikes (Panel (a)) than for existing-supplier spikes (Panel (b)), especially around the event year. In-house employment responds similarly to new- and existing-supplier spikes (Appendix Figure D.9), and the stronger external high-STEM bias for new-supplier spikes appears at both machine suppliers and other industries (Appendix Figure D.10).

Figure 7: External Employment around Machinery Spikes: New versus Existing Suppliers (Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)



Notes: This figure shows external employment around machinery investment spikes, split by whether the dominant machine supplier is new to the investing firm (Panel (a)) or existing (Panel (b)). Spikes in which the increase in machine-supplier purchases comes predominantly from suppliers that are new to the firm (above the sample median of the new-supplier share) are classified as new-supplier spikes; the remainder are existing-supplier spikes (see Appendix D.4 for the exact definition). The figures show employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment. Pre-spike refers to the average over the years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

4.5 Supply-Side Evidence

So far we have studied new machinery from the side of *investing* firms. We now turn to the machine *suppliers*, where three facts support the fixed-cost (Prediction 2) and adjustment-cost (Prediction 3) rationales for outsourcing.

First, machine suppliers' workforces are highly specialized in high-STEM workers. As Section 2.2.3 shows, 13.9% of their employees are high-STEM, nearly three times the in-house share at investing firms (5.0%; Table 1). Appendix Table D.5 further shows that machine suppliers' skill mix is invariant to the share of revenue from spiking-firm clients, consistent with a structurally specialized workforce.

Second, each individual investing firm's demand for capital-complementary high-STEM labor is small in absolute terms. Summed across all its suppliers, the median spiking firm sources only 0.05 high-STEM full-time equivalents in the spike year, and the 90th percentile just 0.77 (Appendix Figure D.11), both well below the single full-time worker that would represent the minimum efficient scale of a dedicated in-house integration role.

Third, machine suppliers absorb the lumpy demand of individual investing firms by spreading their workforce across many uncorrelated clients. Appendix Figure D.12 shows that if each client had to staff its own demand, high-STEM employment would have to adjust 15.4 times as much as it does at the supplier, because one client's expansion offsets another's contraction.

4.6 Evidence on the Outsourcing Predictions

Each pattern maps to a prediction of our framework. The stronger external skill bias among smaller firms (Section 4.1) and firms in services (Section 4.2) is consistent with fixed costs of hiring high-STEM workers (Prediction 2). By contrast, larger firms, especially in manufacturing, can better amortize the in-house fixed cost, consistent with their larger in-house high-STEM workforces before the spikes. Békés et al. (2025) show simi-

larly that larger manufacturers are more likely to import robots directly, rather than rely on intermediation by domestic robot integrators, likely reflecting their greater internal know-how to identify and select appropriate robots.

Intangible spikes rely far less on external employment and generate no external skill bias, consistent with adjustment costs in hiring capital-complementary skills (Prediction 3): because intangibles follow a much smoother investment profile, they entail lower adjustment costs for in-house hiring. Skills complementary to intangible capital may also be harder to outsource due to their tacit and firm-specific nature.²¹

New machinery, especially from new suppliers, generates stronger skill bias than building spikes, consistent with capital-embodied technical change (Prediction 4): whereas a modern office or factory looks broadly similar to one from 20 years ago, modern machinery embodies very different technologies (Greenwood, Hercowitz and Krusell, 1997; Solow, 1960), a distinction Krusell et al. (2000) build explicitly into their model.

The supply-side evidence in Section 4.5 speaks to these mechanisms directly. Each investing firm’s demand is too small to sustain an in-house integrator team, as the fixed-cost rationale implies (Prediction 2). Machine suppliers smooth their workforce across many uncorrelated clients, absorbing adjustment costs that investing firms would otherwise bear (Prediction 3). And their skill mix does not depend on which clients they serve, consistent with a specialized workforce and with both rationales.

Two further patterns point to efficiency gains rather than direct wage savings as the primary motive. If firms outsourced these skills to pay less for them, suppliers would pay lower wages than the firms they serve, and the firms with the most to save, those paying the highest wages, would outsource the most (Abraham and Taylor, 1996, § II.A).

²¹Lindner et al. (2024) find skill bias for in-house employment in response to innovation shocks (the introduction of new products, services, processes, or organizational methods) and R&D tax reforms. One reconciliation with our evidence is that innovation-related skills are harder to outsource than skills complementary to new physical capital, consistent with our finding of no external skill bias for intangible investments. More broadly, this reconciliation highlights the distinction between “innovation-skill” and “capital-skill” complementarities.

Neither holds: machine suppliers pay a wage premium (Section 2.2.3), and if anything lower-wage firms outsource capital-complementary skills more (Appendix Figure D.13).²²

Taken together, our evidence fits theories of capital-embodied, skill-biased technical change once we recognize that the fixed and adjustment costs of hiring the required skills often lead investing firms to rely on external suppliers.

5 Measuring External Employment without Transaction Data

Most studies do not observe firm-to-firm transactions. We evaluate two approaches commonly used instead: aggregating the analysis to local labor markets (as in Acemoglu and Restrepo, 2020; Aghion et al., 2023; Dauth et al., 2021, among others), and applying fixed sectoral input-output matrices to firm-level changes in total input purchases (as in Humlum, 2022, among others). Both turn out to miss most of the skill bias from new machinery, because investing firms shift toward suppliers that are more skill-intensive than their sector’s average and mostly located outside their local labor market.

5.1 Local Labor Markets

Local labor market analyses capture external employment effects to the extent that the external labor is sourced locally. Our transaction data let us evaluate this directly.²³ Following Persyn and Torfs (2011), we split Belgium into 11 functional labor markets, as large as or larger than the commuting zones (CZs) used in the US and elsewhere

²²Abraham and Taylor (1996) provide evidence that wage-savings motives are more important for the outsourcing of janitorial services and other lower-pay occupations than for higher-pay occupations like engineering and machine maintenance. Hence, while wage-savings motives are a first-order driver of outsourcing for lower-pay jobs (Goldschmidt and Schmieder, 2017), our evidence shows that other drivers are at play for the high-pay jobs complementary to capital. A concern for internal wage structure may nonetheless operate in reverse here: just as high-wage firms shed low-pay work, investing firms may source high-pay integration work externally to keep those wages out of their own pay structure. The supplier wage premium and the greater outsourcing by lower-wage firms both fit this motive, which complements the fixed and adjustment costs we emphasize.

²³Local labor markets capture additional spillovers across firms not considered in this analysis, including competition in product markets (especially for non-tradables) and poaching in labor markets. Hence, this assessment should not be viewed as a comprehensive test of the validity of the local labor markets strategy.

(Acemoglu and Restrepo, 2020; Aghion et al., 2023; Dauth et al., 2021).²⁴

Figure 8 splits external employment by whether the suppliers are located inside or outside the local labor markets of investing firms. Most of the external employment effects from machinery investment spikes are generated outside the local labor market, especially for high-STEM skills and at machine supplier firms. For example, only about a third of machine suppliers’ effect on high-STEM employment in the spike year comes from suppliers inside the investing firm’s local labor market.

Why do local labor markets capture so little of this effect? First, machine investments are often lumpy and expensive events, making it worthwhile for an integrator to travel across commuting zones to meet these demands. Second, as Appendix C.2.2 shows, the increases in “other industries” often reflect tradable services, e.g., computer programming or software, where providers can meet demand without traveling to the client facilities. Dhyne and Duprez (2016) show that firm-to-firm flows in Belgium decrease with the distance between supplier and buyer, with heterogeneity across sectors that aligns with our findings.

5.2 Sectoral Input–Output


The second approach applies sectoral input–output matrices (e.g., from the National Accounts). This approach captures the external employment effects if these reflect a general shift toward the typical suppliers of the firm’s industry.²⁵ Figure 9 evaluates this method by scaling firms’ intermediate input purchases using sectoral input–output

²⁴Acemoglu and Restrepo (2020) use 722 CZs, with an average population around 0.5 million people. Dauth et al. (2021) use 402 CZs in Germany, with an average of 0.2 million people (the authors show robustness to making the boundaries broader, increasing the average local population to 0.5 million people). Aghion et al. (2023) use 287 CZs in France, with an average population of 0.24 million. By comparison, the 11 Belgian commuting zones have an average population of 1 million people. Larger zones make local sourcing more likely, so our conclusion that local labor markets matter little is conservative.

²⁵Both approaches measure the skill content embodied in a firm’s purchases (Section 2.2.3). They differ in whose skill mix they use: the sectoral tables use the industry’s typical suppliers, our transaction data the firm’s own.

tables.²⁶ In practice, it captures some of the external employment effects but still misses most of the high-STEM skill bias around the spike, because investing firms shift toward suppliers that are more high-STEM-intensive than their industry’s average.


Figure 8: External Employment by Local Labor Markets (LLM)
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)



Notes: This figure splits external employment by whether the suppliers are located inside (Panels (a) and (c)) or outside (Panels (b) and (d)) the local labor markets of the firm. Panels (a) and (b) consider the external employment of machine suppliers, while Panels (c) and (d) consider the external employment from all other industries. The figures show employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment. Pre-spike refers to the average over the years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. Because each outcome is winsorized separately at the 5th and 95th percentiles, the inside and outside components need not sum exactly to the estimates in Figure 3. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

²⁶To ensure consistency with our previous analysis, we construct the sectoral input-output tables directly from our firm-to-firm transaction data, ensuring that differences relative to the main analysis reflect only the level of aggregation of input-output relationships.

Figure 9: External Employment by Measurement Method
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)



Notes: This figure shows external employment of high-STEM, high non-STEM, low STEM, and low non-STEM in percent of total pre-spike employment by skill group (Equation 7). Panel (a) uses sectoral input-output tables to calculate external employment, while Panel (b) uses the business-to-business transactions (as in Figure 2). The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.

6 Aggregate Implications

A key finding of this study is that new machinery shifts labor demand toward high-STEM workers through external suppliers. We now circle back to the macro-labor question that motivated our analysis (Krusell et al., 2000) and ask what these effects imply for aggregate labor market inequality. Between 2003 and 2021, machinery per worker in the Belgian private sector almost tripled (Figure E.1). How much did this increase raise the wages of high-STEM workers relative to low non-STEM workers (the skill premium) and their relative employment? To answer this question, we build a quantitative general-equilibrium version of our conceptual framework (Section 1) that contains the four ingredients our evidence emphasizes: capital-skill and investment-skill complementarities, firm heterogeneity, lumpy investment, and employment beyond firm boundaries.

The model serves three purposes. First, it lets us interpret our reduced-form evidence

in terms of structural parameters (Section 6.1). Second, it lets us evaluate the aggregate implications of our firm-level estimates in general equilibrium (Section 6.2). Finally, it delivers a formula for how these equilibrium impacts depend on a few sufficient statistics, including the demand shift that our event study measures (Section 6.3).

6.1 Model and Calibration

This section describes our model and calibration in brief, with Appendix E.1 providing the formal details.

Model Setup. Production follows the nested CES of Krusell et al. (2000), in which machinery combines with high-STEM labor before low non-STEM labor, and integrating new machinery combines hardware with the labor of both skill groups in a second nested CES that generalizes Goldin and Katz (1998).²⁷ Firms differ in size, machinery intensity, productivity, and the hardware prices they face, and fixed adjustment costs make them invest in lumpy episodes (Cooper and Haltiwanger, 2006; Khan and Thomas, 2008; Winberry, 2021). Our baseline model counts firms’ in-house and external workers alike, as our event studies do, and Section 6.4 shows that our counterfactual results are robust to endogenizing this make-or-buy choice through the fixed costs of Prediction 2. Firms sell differentiated varieties under monopolistic competition, so each faces a downward-sloping demand curve and chooses its scale. They make forward-looking investment decisions taking prices as given, and wages adjust so that demand from all firms meets an upward-sloping supply of each skill group.

²⁷Our conclusions carry over if production also uses other inputs, such as other capital, materials, high non-STEM labor, and low STEM labor, as long as it is weakly separable in them (Appendix E.1). Krusell et al. (2000) similarly assume a production function that is Cobb–Douglas in structures. Consistent with this assumption, building spikes generate a much weaker skill bias that largely dissipates within two years (Section 4.3).

Estimation. We estimate 12 parameters so that the model reproduces our event studies: we apply the same investment-spike design to simulated firms and match its estimates to those in the data (Figure E.2). We set the remaining parameters outside the model, including the substitution elasticities between skill groups and within integration as well as the labor supply elasticity (Appendix E.3).²⁸ In brief, we find that machinery substitutes about 1.15 times as easily for low non-STEM labor as for high-STEM labor in production, and integration makes an outsized contribution to the skill bias (36 percent of the rise in relative demand, despite employing only 4.3 percent of high-STEM workers).

The ratio in production confirms capital–skill complementarity in firms’ production technology, although less strongly than the aggregate ratio of about 2.5 implied by Krusell et al. (2000).²⁹

Counterfactual. The counterfactual compares two stationary equilibria that differ only in the price of hardware, which Belgian firms take as given as a small open economy. The size of the decline is set to match the observed rise in machinery per worker.

6.2 Aggregate Skill Bias from New Machinery

Table 3 reports the results. Each entry is a change for high-STEM workers relative to low non-STEM workers, in log points (100 times the log change). In the model, the rise in machinery shifts the relative demand for high-STEM workers by 8.8 log points at unchanged wages (Column 3). In general equilibrium, the shift raises the skill premium by 3.1 log points (Column 4) and relative employment by 4.6 (Column 5). Integration accounts for 3.2 log points of the shift and production for 5.7. Within both stages the skill

²⁸The integration elasticities imply that installation and integration remain a roughly stable share of what a machine costs as hardware gets cheaper, in line with the rule of thumb of Appendix B.2.1.

²⁹Their aggregate production function combines the production and integration stages that we model separately. When we fit that same production function to the economy of our model, the ratio rises to 1.21 (Appendix E.3), still well below 2.5. Re-estimating the aggregate model on a sample extended to 2019, Ohanian, Orak and Shen (2023) report preferred estimates implying a ratio of about 2.3, and Castex, Cho and Dechter (2022) document a decline in capital–skill complementarity over time.

bias arises at external suppliers (Column 2), while the in-house part is slightly negative (Column 1). How the shift divides between wages and employment depends on the elasticity of relative supply: with perfectly inelastic supply the whole shift raises the skill premium, and as supply becomes more elastic more of it raises employment (Appendix Table E.3).³⁰

Table 3: Aggregate Skill Bias from New Machinery
(Log Points, Relative to Low Non-STEM)

	Shift in relative demand			General equilibrium	
	In-house (1)	External (2)	Total (3)	Skill premium (4)	Relative employment (5)
High STEM	-1.5	10.3	8.8	3.1	4.6
Integration	0.0	3.2	3.2	1.1	1.6
Production	-1.5	7.2	5.7	2.0	3.0

Notes: This table reports the model’s counterfactual for the rise in machinery per worker between 2003 and 2021, in log points (100 times log changes), for high-STEM workers relative to low non-STEM workers. The hardware price falls until machinery per worker has risen by the observed amount, and wages clear both skill markets. Column (3) is the shift in relative labor demand at the initial wages; Columns (4) and (5) are the general-equilibrium changes in the skill premium and relative employment. The integration and production rows split the model’s shift by stage: integration reflects investment–skill complementarity (Prediction 1(b)), production capital–skill complementarity (Prediction 1(a)). Columns (1) and (2) split each stage between in-house and external employment in the proportions of our event-study estimates (years 0–2 for integration, year 4 for production), which add up to the total because our outcomes scale both by the same pre-spike employment. Rows may not add up because of rounding.

6.3 Sufficient Statistics

The equilibrium impacts in Table 3 depend on three sufficient statistics: the shift in relative labor demand per unit of machinery, which our event study measures at the firm level, and the elasticities of relative labor supply and demand. Let S denote high-STEM and U low non-STEM workers. Relative supply rises with the skill premium at a constant elasticity ε^S , and relative demand depends on the skill premium and on

³⁰The two polar cases nest the approach of Krusell et al. (2000), with perfectly inelastic supply ($\varepsilon^S = 0$), and that of Autor, Dorn and Hanson (2013) and Mian and Sufi (2014), with perfectly elastic supply ($\varepsilon^S = \infty$). Solving for equilibrium wages also supplies the “missing intercept” of Moll (2021): the wage responses common to all firms, which a comparison of investing and non-investing firms differences out.

aggregate machinery per worker M , which rose by $\Delta \ln M = 108.0$ log points between 2003 and 2021. Between the two equilibria, the change in relative demand splits exactly into a shift at the initial wages, $\Psi_M \cdot \Delta \ln M$ (Table 3, Column 3), and a movement along the demand curve at the new level of machinery, $-\varepsilon_M^D \cdot \Delta \ln(w_S/w_U)$. Equating the change in demand with the change in supply gives

$$\Delta \ln \frac{w_S}{w_U} = \frac{\Psi_M \cdot \Delta \ln M}{\varepsilon^S + \varepsilon_M^D}, \quad \Delta \ln \frac{L_S}{L_U} = \varepsilon^S \cdot \Delta \ln \frac{w_S}{w_U}. \quad (9)$$

Appendix E.4 derives the formula and verifies it against the full nonlinear model.

The Demand Shift. Ψ_M is the rise in the relative demand for high-STEM workers per log point of machinery per worker, at unchanged wages. In the model, $\Psi_M = 0.0817$, a shift of 8.8 log points over the observed rise in machinery, which the Integration and Production rows of Table 3 split into the two stages of our framework. The dynamics of the event study discipline this split: the skill differential in the spike year identifies how intensively integration uses high-STEM labor, and the differential four years later the capital–skill complementarity in production. How does Ψ_M compare with our event-study estimate? Taking the event study at face value, dividing the year-four skill differential by the rise in machinery per worker, gives $\Psi_M = 0.108$, which over the same rise in machinery implies a shift of 11.7 log points against 8.8 in the model. Appendix E.5 shows that the gap reflects two features of the event study: it follows selected firms through one long jump, from below the economy’s machinery per worker to well above it, where production labor responds more per log point of deepening; and it holds installation labor per machine fixed, whereas cheaper hardware lowers it.

The Demand Elasticity. As the skill premium rises, firms substitute away from high-STEM workers, and ε_M^D measures how much relative demand falls, holding machinery

per worker fixed.³¹ In the estimated model, $\varepsilon_M^D = 1.399$. Substitution within production accounts for nearly all of it (1.357), and substitution within integration adds 0.038. The remaining terms are negligible in the model: the adjustment in the price of installed capital that holds machinery per worker fixed, reallocation across stages and, as in Oberfield and Raval (2021), across firms with different cost shares, and a dynamic residual from lumpy investment and uncertainty (Appendix E.4.2, Table E.2). Appendix E.6 shows how the results change with the elasticities that we set outside the model and with a larger dispersion of skill shares across firms.

6.4 In-House versus External Employment

Counting External Employment. External employment is critical for our conclusions about the aggregate skill bias of new machinery. With the measured shift $\Psi_M = 0.108$, Equation (9) implies that the rise in machinery raises the skill premium by 4.0 log points when we count all employment. Counting only investing firms’ in-house employment, as payroll data would, the skill premium instead *falls* by 1.4 log points: four years after a spike, in-house employment is up 14 percent for high-STEM workers but 18 percent for low non-STEM workers, when converted with aggregate in-house shares, as an aggregate of payrolls requires (Appendix E.5).

Endogenizing the Make-or-Buy Decision. In Appendix E.7, we extend the model to endogenize the make-or-buy decision: each year, a firm either hires its high-STEM production workers in-house, paying a fixed cost (as in Prediction 2), or buys them from suppliers at a markup equal to the 3.6 percent wage premium of machine suppliers (Table D.6). With the fixed cost set so that the in-house share of high-STEM employment matches the data, our conclusions about the skill bias of new machinery are robust to

³¹Because $\Delta \ln M$ is the observed rise in machinery, it already contains any response of machinery to wages, which the elasticity must not count again (Appendix E.4).

this extension, which leaves the rise in the skill premium nearly unchanged (Appendix Table E.7). Hence, while counting external employment is critical for quantifying the aggregate implications of new machinery, endogenizing the make-or-buy decision is not.

7 Conclusion

The capital–skill complementarity hypothesis is a cornerstone of a large literature on technological change and labor market inequality (Krusell et al., 2000). Recent firm-level studies, however, find that firms investing in new machinery leave their skill mixes unchanged, casting doubt on the hypothesis. In this paper, we use new data to revive it: new machinery is skill-biased, but the bias arises at the external suppliers of investing firms. This pattern is widely understood by industry practitioners, as the rise of robot integrators and IT consultants illustrates, yet largely overlooked by the academic literature, in part for lack of suitable data.

Our findings have implications for both measurement and theory. On the measurement side, they show that capturing the skill bias of new machinery requires looking *beyond* the payrolls of investing firms. Aggregating our estimates in a general-equilibrium model, we find that the rise in machinery per worker between 2003 and 2021 raises the high-STEM skill premium by 3.1 log points, whereas payroll data alone would imply a fall. More generally, the model shows what an aggregation of this kind requires: the labor demand shift that event studies measure, and a few elasticities.

On the theory side, our paper emphasizes the central role of external suppliers in facilitating technology adoption. We propose a conceptual framework that embeds capital–skill complementarities into firms’ joint decisions to invest in capital and to hire or outsource the necessary skills. By spreading the fixed and adjustment costs of these skills across many clients, suppliers lower the cost of adoption, so the supply of integration services can shape who adopts and how fast. We welcome further research on the causal

role of external suppliers in shaping the speed and incidence of technology diffusion.³²

Finally, our findings raise a broader question: which other phenomena operate primarily beyond firm boundaries? The answer depends on whether the market can supply the affected labor more cheaply than firms can employ it themselves. In the case of artificial intelligence (AI), external suppliers may employ both the complementary skills (AI consultants) and the labor at risk of displacement (freelancers), so firm-level data would miss both sides of the adjustment. The data infrastructure, conceptual framework, and quantitative model we develop here are well-suited to tackle these questions.

³²The central role of external suppliers in technology adoption echoes historical evidence on the emergence of millwrights on the eve of the Industrial Revolution (Mokyr, Sarid and van der Beek, 2022) and recent work documenting the role of robot integrators (Békés et al., 2025; Brynjolfsson et al., 2023).

References

- Abraham, Katharine G, and Susan K Taylor.** 1996. “Firms’ use of outside contractors: Theory and evidence.” *Journal of Labor Economics*, 14(3): 394–424.
- Acemoglu, Daron, and Pascual Restrepo.** 2020. “Robots and Jobs: Evidence From US Labor Markets.” *Journal of Political Economy*, 128(6): 2188–2244.
- Aghion, Philippe, Celine Antonin, Simon Bunel, and Xavier Jaravel.** 2023. “The Local Labor Market Effects of Modern Manufacturing Capital: Evidence From France.” Vol. 113, 219–223, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel.** 2025. “Modern Manufacturing Capital, Labor Demand, and Product Market Dynamics: Evidence from France.” *American Economic Journal: Macroeconomics*. conditionally accepted.
- Autor, David H, David Dorn, and Gordon H Hanson.** 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review*, 103(6): 2121–2168.
- Baqae, David, Ariel Burstein, Cédric Duprez, and Emmanuel Farhi.** 2025. “Consumer Surplus from Suppliers: How Big is it and Does it Matter for Growth?” *Econometrica*, 93(6): 2043–2081.
- Békés, Gábor, Alessandra Bonfiglioli, Rosario Crinò, and Gino Gancia.** 2025. “Integrators and Robot Adoption: Facts from Hungary.” Centre for Economic Policy Research Discussion Paper 20026.
- Benmelech, Efraim, and Michał Zator.** 2025. “Robots and firm investment.” *Journal of Financial Economics*, 174: 104183.
- Berlingieri, Giuseppe, Filippo Boeri, Danial Lashkari, and Jonathan Vogel.** 2026. “Capital-skill complementarity in firms and in the aggregate economy.” National Bureau of Economic Research 33000. Revised June 2026.
- Bernard, Andrew B., Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes.** 2022. “The Origins of Firm Heterogeneity: A Production Network Approach.” *Journal of Political Economy*, 130(7): 1765–1804.
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge.** 2025. “What Happens to Workers at Firms that Automate?” *The Review of Economics and Statistics*, 107(1): 125–141.
- Bijnens, Gert, Simon Jäger, and Benjamin Schoefer.** 2025. “What Does Consulting Do?” National Bureau of Economic Research 34072.
- Bilal, Adrien, and Hugo Lhuillier.** 2025. “Outsourcing, Inequality and Aggregate Output.” National Bureau of Economic Research NBER Working Paper 29348.
- Bilgin, Nuriye Melisa, Ester Faia, and Gianmarco Ottaviano.** 2026. “Technology Spillovers, Diffusion and Rivalry in Firm Networks.” Working paper. March 2026.
- Bloom, Nicholas.** 2009. “The Impact of Uncertainty Shocks.” *Econometrica*, 77(3): 623–685.
- Brynjolfsson, Erik, Catherine Buffington, Nathan Goldschlag, J Frank Li, Javier Miranda, and Robert Seamans.** 2023. “The Characteristics and Geographic Distribution of Robot Hubs in U.S. Manufacturing Establishments.” National Bureau of Economic Research 31062.
- Callaway, Brantly, and Pedro HC Sant’Anna.** 2021. “Difference-in-differences with multiple time periods.” *Journal of Econometrics*, 225(2): 200–230.
- Castex, Gonzalo, Sang-Wook Cho, and Evgenia Dechter.** 2022. “The Decline in Capital-Skill Complementarity.” *Journal of Economic Dynamics and Control*, 138: 104363.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer.** 2019. “The Effect of Minimum Wages on Low-Wage Jobs.” *The Quarterly Journal of Economics*, 134(3): 1405–

- Chen, Jiafeng, and Jonathan Roth.** 2024. “Logs with zeros? Some problems and solutions.” *The Quarterly Journal of Economics*, 139(2): 891–936.
- Cooper, Russell W, and John C Haltiwanger.** 2006. “On the Nature of Capital Adjustment Costs.” *The Review of Economic Studies*, 73(3): 611–633.
- Curtis, E Mark, Daniel G Garrett, Eric C Ohrn, Kevin A Roberts, and Juan Carlos Suárez Serrato.** 2023. “Capital Investment and Labor Demand.” Working Paper.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner.** 2021. “The Adjustment of Labor Markets to Robots.” *Journal of the European Economic Association*, 19(6): 3104–3153.
- Davis, Donald R, and David E Weinstein.** 2001. “An Account of Global Factor Trade.” *American Economic Review*, 91(5): 1423–1453.
- Demir, Banu, Ana Cecília Fieler, Daniel Yi Xu, and Kelly Kaili Yang.** 2024. “O-Ring Production Networks.” *Journal of Political Economy*, 132(1): 200–247.
- Dhyne, Emmanuel, and Cédric Duprez.** 2016. “Three Regions, Three Economies?” *Economic Review*, , (iii): 59–73.
- Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot.** 2025. “Firm Responses and Wage Effects of Foreign Demand Shocks with Fixed Labor Costs and Monopsony.” *American Economic Review*, 115(12): 4328–4368.
- Dhyne, Emmanuel, Cédric Duprez, and Toshiaki Komatsu.** 2023. “The Belgian business-to-business transactions dataset 2002-2021.” National Bank of Belgium 444.
- Doraszelski, Ulrich, and Jordi Jaumandreu.** 2013. “R&D and Productivity: Estimating Endogenous Productivity.” *Review of Economic Studies*, 80(4): 1338–1383.
- Goldin, Claudia, and Lawrence F Katz.** 1998. “The Origins of Technology-Skill Complementarity.” *The Quarterly Journal of Economics*, 113(3): 693–732.
- Goldschmidt, Deborah, and Johannes F Schmieder.** 2017. “The rise of domestic outsourcing and the evolution of the German wage structure.” *The Quarterly Journal of Economics*, 132(3): 1165–1217.
- Greenwood, Jeremy, and Mehmet Yorukoglu.** 1997. “1974.” *Carnegie-Rochester Conference Series on Public Policy*, 46: 49–95.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell.** 1997. “Long-run implications of investment-specific technological change.” *The American Economic Review*, 87(3): 342–362.
- Griliches, Zvi.** 1969. “Capital-Skill Complementarity.” *The Review of Economics and Statistics*, 51(4): 465–468.
- Grossman, Gene M., and Elhanan Helpman.** 2002. “Integration versus Outsourcing in Industry Equilibrium.” *Quarterly Journal of Economics*, 117(1): 85–120.
- Hirvonen, Johannes, Aapo Stenhammar, and Joonas Tuhkuri.** 2022. “New Evidence on the Effect of Technology on Employment and Skill Demand.” The Research Institute of the Finnish Economy.
- Hirvonen, Johannes, Aapo Stenhammar, and Joonas Tuhkuri.** 2025. “What Do Technology Grants Do?” *Quarterly Journal of Economics*. conditionally accepted.
- Humlum, Anders.** 2022. “Robot Adoption and Labor Market Dynamics.” ROCKWOOL Foundation Research Unit Study Paper 175.
- Huneus, Federico, Kory Kroft, and Kevin Lim.** 2025. “Earnings Inequality in Production Networks.” Working paper.
- IFR.** 2022. “IFR Annual Report.” International Federation of Robotics.
- IFR.** 2026. “Case Studies.” International Federation of Robotics. Available at <https://ifr.org/case-studies>.
- Khan, Aubhik, and Julia K Thomas.** 2008. “Idiosyncratic Shocks and the Role of Noncon-

- vexities in Plant and Aggregate Investment Dynamics.” *Econometrica*, 76(2): 395–436.
- Krusell, Per, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante.** 2000. “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis.” *Econometrica*, 68(5): 1029–1053.
- Lindner, Attila, Balázs Muraközy, Balázs Reizer, and Ragnhild Schreiner.** 2024. “Firm-level Technological Change and Skill Demand.” *American Economic Review*. conditionally accepted.
- Mian, Atif, and Amir Sufi.** 2014. “What Explains the 2007–2009 Drop in Employment?” *Econometrica*, 82(6): 2197–2223.
- Mokyr, Joel, Assaf Sarid, and Karine van der Beek.** 2022. “The wheels of change: Technology adoption, millwrights and the persistence in Britain’s industrialisation.” *The Economic Journal*, 132(645): 1894–1926.
- Moll, Benjamin.** 2021. “The Missing Intercept Problem When Going from Micro to Macro.” https://benjaminmoll.com/wp-content/uploads/2021/02/missing_intercept.pdf, Lecture slides, available at <https://benjaminmoll.com/lectures/>.
- Oberfield, Ezra, and Devesh Raval.** 2021. “Micro Data and Macro Technology.” *Econometrica*, 89(2): 703–732.
- Ohanian, Lee E., Musa Orak, and Shihan Shen.** 2023. “Revisiting Capital-Skill Complementarity, Inequality, and Labor Share.” *Review of Economic Dynamics*, 51: 479–505.
- Olley, G Steven, and Ariel Pakes.** 1996. “The Dynamics of Productivity in the Telecommunications Equipment Industry.” *Econometrica*, 64(6): 1263–1297.
- Persyn, Damiaan, and Wouter Torfs.** 2011. “Functional Labor Markets in Belgium: Evolution over Time and Intersectoral Comparison.” VIVES Research Centre for Regional Economics, KU Leuven Working Paper 17.
- RIA.** 2020. “Annual Report.” Robotic Industries Association.
- RoboJob.** 2026. “Customer Testimonials: More than 2,000 Installations.” RoboJob NV, Belgium. Available at <https://robojob.eu/en/2000-installations/customer-testimonials>.
- Roth, Jonathan, Pedro HC Sant’Anna, Alyssa Bilinski, and John Poe.** 2023. “What’s trending in difference-in-differences? A synthesis of the recent econometrics literature.” *Journal of Econometrics*, 235(2): 2218–2244.
- Solow, Robert M.** 1960. “Investment and Technical Progress.” In *Mathematical Methods in the Social Sciences, 1959: Proceedings of the First Stanford Symposium.*, ed. Kenneth J. Arrow, Samuel Karlin and Patrick Suppes, 89–104. Stanford, CA:Stanford University Press.
- Stuart, Elizabeth A, and Donald B Rubin.** 2008. “Best Practices in Quasi-Experimental Designs: Matching Methods for Causal Inference.” In *Best Practices in Quantitative Methods.*, ed. Jason W Osborne, 155–176. Thousand Oaks, CA:Sage Publications.
- Sun, Liyang, and Sarah Abraham.** 2021. “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects.” *Journal of Econometrics*, 225(2): 175–199.
- Trefler, Daniel.** 1995. “The Case of the Missing Trade and Other Mysteries.” *American Economic Review*, 85(5): 1029–1046.
- Winberry, Thomas.** 2021. “Lumpy Investment, Business Cycles, and Stimulus Policy.” *American Economic Review*, 111(1): 364–396.

Online Appendix

Capital-Skill Complementarity Beyond Firm Boundaries

Gert Bijnens	Anders Humlum	Stijn Vanormelingen
National Bank of Belgium	University of Chicago	KU Leuven

Contents

A	Conceptual Framework	3
A.1	Capital Investment and the Demand for Skills	3
A.2	Outsourcing of Labor Services	5
A.3	What Investment Spikes Identify	8
A.4	Skill Bias and Substitution Elasticities	12
B	Data and Measurement	14
B.1	Data	14
B.2	Measurement	19
C	New Machinery and the Demand for Skills	24
C.1	Empirical Strategy	24
C.2	Main Results	27
C.3	Robustness	33
D	Why Firms Outsource Capital-Complementary Skills	47
D.1	Firm Size	47
D.2	Application Sector	51
D.3	Capital Types	54
D.4	New versus Existing Machine Supply	57
D.5	Supply-Side Evidence	59
D.6	Wage-Savings Motive	63
E	Aggregate Implications	66
E.1	The Quantitative Model	66
E.2	Solving the Model	71
E.3	Estimation	72
E.4	Deriving the Sufficient Statistics	78
E.5	Measuring the Sufficient Statistics	85
E.6	Sensitivity	89
E.7	The Make-or-Buy Margin	91
	Online Appendix References	94

A Conceptual Framework

This appendix specifies firms' full dynamic program from Section 1 and provides formal proofs for the model predictions (Sections A.1 and A.2). Section A.3 states the conditions under which the investment-spikes design identifies the skill bias of new capital, and Section A.4 maps the estimated responses into substitution elasticities. Appendix E solves a quantitative general-equilibrium version of this model.

Firms face in-house labor costs $C_{jt}^H(\mathbf{L}_{jt}^H|\mathbf{L}_{jt-1}^H)$ (which may include fixed and adjustment components), external labor costs $C_{jt}^O(\mathbf{L}_{jt}^O)$, and investment costs $c_{jt}^I(X_{jt})$ that may include non-convex components. The firm's dynamic problem is to choose labor demands (hired and outsourced) and capital investments to maximize the present-discounted sum of profits:

$$\begin{aligned}
 V_{jt}(K_{jt}, \mathbf{L}_{jt-1}^H) = & \max_{\{\mathbf{L}_{jt}^H, \mathbf{L}_{jt}^O, X_{jt}\}} \underbrace{D_{jt}F(K_{jt}, \mathbf{L}_{jt}^P)}_{\text{Revenues}} - \underbrace{C_{jt}^H(\mathbf{L}_{jt}^H|\mathbf{L}_{jt-1}^H)}_{\text{In-house labor costs}} - \underbrace{C_{jt}^O(\mathbf{L}_{jt}^O)}_{\text{External labor costs}} \\
 & - \underbrace{c_{jt}^I(X_{jt})}_{\text{Investment costs}} + \underbrace{\beta \mathbb{E}_t V_{jt+1} \left((1-\delta)K_{jt} + G(X_{jt}, \mathbf{L}_{jt}^I), \mathbf{L}_{jt}^H \right)}_{\text{Continuation Value}} \quad (10) \\
 \text{s.t. } & \mathbf{L}_{jt} = \mathbf{L}_{jt}^H + \mathbf{L}_{jt}^O \text{ for } \mathbf{L} \in \{U^P, S^P, U^I, S^I\},
 \end{aligned}$$

where β is the discount factor, D_{jt} is the firm's output demand shifter, and $\mathbf{L}_{jt}^P \equiv (S_{jt}^P, U_{jt}^P)$ and $\mathbf{L}_{jt}^I \equiv (S_{jt}^I, U_{jt}^I)$ stack the production- and integration-stage labor inputs.

A.1 Capital Investment and the Demand for Skills

This subsection shows that, under the investment-skill and capital-skill complementarity hypotheses, profit-maximizing firms increase their relative demand for skilled workers in response to increases in investment I and capital K .

To distill our exposition, we make three simplifying assumptions. First, firms are price-takers with linear, static labor costs. Second, we abstract from the make-or-buy decision (Section A.2 considers it) by assuming identical labor costs across in-house and

external supply. Third, production and integration functions exhibit neutrality between skilled and unskilled labor within each activity, $F_{S_P U_P} = G_{S_I U_I} = 0$, allowing us to focus on complementarities with investment and capital.³³ Under these assumptions, the general program (equation 10) reduces to:

$$V(K) = \max_{\{X, L\}} DF(K, L_P) - w_S(S_P + S_I) - w_U(U_P + U_I) - c^I(X) + \beta \mathbb{E}V((1 - \delta)K + G(X, L_I)), \quad (11)$$

where $L_k = (S_k, U_k)$, $k \in \{P, I\}$, with $S_k \equiv S_k^H + S_k^O$ and $U_k \equiv U_k^H + U_k^O$ collapsing the make-or-buy distinction. Throughout this appendix, we write the activity (P or I) as a subscript, freeing the superscript for the in-house/external mode and keeping the derivative notation below legible. Appendix E.1 specifies functional forms for this program and solves it in general equilibrium.

Prediction 1 (Capital investment and the demand for skills). *Capital investment raises the relative demand for skills through two channels:*

Prediction 1(a) (Capital–skill complementarity in production). *Capital raises skilled production employment if $F_{S_P K} > 0$, and raises the relative demand for skilled production workers if the elasticity of skilled-labor demand with respect to capital exceeds that of unskilled labor: $\varepsilon_{S_P, K} > \varepsilon_{U_P, K}$ (Griliches, 1969).*

Proof. The first-order conditions for production labor are $DF_{S_P}(K, S_P, U_P) = w_S$ and $DF_{U_P}(K, S_P, U_P) = w_U$. Differentiating with respect to K (holding D, w_S, w_U fixed) and using $F_{S_P U_P} = 0$ yields

$$\frac{dS_P}{dK} = -\frac{F_{S_P K}}{F_{S_P S_P}}, \quad \frac{dU_P}{dK} = -\frac{F_{U_P K}}{F_{U_P U_P}},$$

³³The neutrality assumption is expositional: it removes cross-effects between skilled and unskilled labor from the comparative statics below. For example, the conclusion that capital raises skilled production employment when $F_{S_P K} > 0$ extends to $F_{S_P U_P} > 0$ provided $F_{U_P K} \geq 0$.

with $F_{S_P S_P}, F_{U_P U_P} < 0$ by strict concavity. Hence $\text{sign}(dS_P/dK) = \text{sign}(F_{S_P K})$ and symmetrically for U_P . The factor-demand elasticities are $\varepsilon_{S_P, K} = -F_{S_P K}K/(F_{S_P S_P}S_P)$ and $\varepsilon_{U_P, K} = -F_{U_P K}K/(F_{U_P U_P}U_P)$, and the skill-bias condition $d \log(S_P/U_P)/d \log K = \varepsilon_{S_P, K} - \varepsilon_{U_P, K} > 0$ reduces to the stated inequality. $\square$

Prediction 1(b) (Investment–skill complementarity in integration). *Investment in hardware raises skilled integration employment if $G_{S_I X} > 0$, and raises the relative demand for skilled integration workers if $\varepsilon_{S_I, X} > \varepsilon_{U_I, X}$ (Goldin and Katz, 1998).*

Proof. Differentiating the Bellman equation (11) with respect to integration labor yields the first-order conditions $q G_{S_I} = w_S$ and $q G_{U_I} = w_U$, where the shadow value of installed capital $q \equiv \beta \partial \mathbb{E}V / \partial K$ plays the role of the output price. Assuming G is strictly concave in labor and treating q as parametric (exact when the continuation value is locally linear in capital), the argument for Prediction 1(a) applies with G, X in place of F, K : differentiating the first-order conditions with respect to X (holding w_S and w_U fixed) under $G_{S_I U_I} = 0$ gives $\text{sign}(dS_I/dX) = \text{sign}(G_{S_I X})$ and the skill-bias condition follows. Both conditions scale with q , so their ratio, $G_{S_I}/G_{U_I} = w_S/w_U$, holds independently of the shadow price. $\square$

A.2 Outsourcing of Labor Services

We now study firms' make-or-buy decisions — the choice between $\mathbf{L}_{jt}^H$ and $\mathbf{L}_{jt}^O$ in the general program (equation 10) — and derive how specific features of the in-house cost function C^H shape outsourcing. We focus throughout on skilled integration labor; analogous arguments hold for production labor. Let total demand be $S_I = S_I^H + S_I^O$. Firms choose the in-house/external split to minimize total labor cost:

$$C_S^*(S_I) = \min_{\{S_I^H, S_I^O\}} \left\{ C_S^H(S_I^H) + C_S^O(S_I^O) : S_I^H + S_I^O = S_I \right\}. \quad (12)$$

We write w_S^H for the in-house marginal wage and w_S^O for the external supplier price, with $w_S^O > w_S^H$ reflecting a supplier markup.

Prediction 2 (Fixed costs and outsourcing). *If C^H includes a fixed cost of hiring capital-complementary skills, firms outsource these skills when their demand for them is too small to cover that cost. Smaller firms and firms in less capital-intensive sectors (e.g., services) should therefore rely more on external suppliers (Abraham and Taylor, 1996, § II.C).*

Proof. Let $C_S^H(S_I^H) = w_S^H S_I^H + f_S^H \cdot \mathbf{1}\{S_I^H > 0\}$ with fixed hiring cost $f_S^H > 0$ paid only when the firm hires any in-house labor, and $C_S^O(S_I^O) = w_S^O S_I^O$. Because $w_S^H < w_S^O$, total cost is decreasing in S_I^H whenever $S_I^H > 0$, so the firm chooses between corners: hire all in-house (cost $w_S^H S_I + f_S^H$) or outsource all (cost $w_S^O S_I$). Comparing:

$$w_S^H S_I + f_S^H \leq w_S^O S_I \iff S_I \geq S_I^* \equiv \frac{f_S^H}{w_S^O - w_S^H}. \quad (13)$$

Firms with demand above the threshold S_I^* amortize the fixed cost across enough labor to make in-house hiring worthwhile; firms with lower demand — the smaller ones, or those in the service sector — outsource. $\square$

Prediction 3 (Lumpy investments and outsourcing). *If C^H includes adjustment costs of hiring, firms outsource capital-complementary skills when their investment is lumpy enough that in-house staff would sit idle between investment episodes. Outsourcing should therefore be more prevalent for lumpier investment categories (e.g., machinery rather than intangibles) (Abraham and Taylor, 1996, § II.B).*

Proof. To highlight the comparative static, we make a series of simplifying assumptions. Suppose in-house costs include a prohibitive adjustment component, so that $C_S^H(S_{I,t}^H | S_{I,t-1}^H)$ penalizes deviations $S_{I,t}^H \neq S_{I,t-1}^H$ arbitrarily: in-house integration labor is effectively fixed over the investment cycle. Consider hardware investment that arrives in occasional spikes, $X_t = \bar{X}$ in spike years and zero otherwise, with average investment

$\tilde{X}$ over the cycle and lumpiness ratio $\lambda \equiv \bar{X}/\tilde{X} \geq 1$. Suppose further that hardware and skilled integrators are perfect complements ($G(X, S_I) = \min\{X, S_I\}$, as in Goldin and Katz, 1998).

To meet peak demand via in-house hiring, the firm hires $S_I^H = \bar{X}$ persistently, at annual wage cost $w_S^H \bar{X}$.³⁴ Full outsourcing scales with current demand at annual cost $w_S^O \tilde{X}$. The firm outsources iff

$$w_S^O \tilde{X} < w_S^H \bar{X} \iff \lambda > \frac{w_S^O}{w_S^H}. \quad (14)$$

For smooth profiles ($\lambda = 1$), outsourcing would be chosen only if external labor were cheaper per unit ($w_S^O < w_S^H$). For lumpy profiles ($\lambda > 1$), outsourcing can dominate despite being more expensive on a per-unit basis. $\square$

These simplifying assumptions are not essential to the comparative static. With a positive base load of investment, the firm covers the constant base with in-house labor (cheaper per unit and requiring no adjustment) and applies the same threshold to the transitory spike component, so lumpier profiles still tilt toward outsourcing.

Prediction 4 (Advanced capital and outsourcing). *If these fixed and adjustment costs are larger for more specialized skills, both outsourcing motives should be strongest for capital undergoing rapid technological change (Greenwood, Hercowitz and Krusell, 1997) (e.g., machinery rather than buildings, or machinery from new rather than existing suppliers).*

Proof. The outsourcing effects in Predictions 2 and 3 strengthen with the in-house cost parameters: the threshold for in-house hiring in Equation (13) rises with the fixed cost f_S^H , so more firms outsource, and larger adjustment costs in C_S^H make tracking lumpy investment demand in-house costlier, widening the range of profiles for which outsourcing dominates (Equation (14) is the limiting case of prohibitive adjustment costs). These cost

³⁴The firm would leave money on the table if $S_I^H \neq \bar{X}$, as it could then either decrease its peak investments (if $\bar{X} > S_I^H$) or lower its integration labor (if $\bar{X} < S_I^H$), without losing effective capital.

parameters are larger when the relevant skills are scarcer or more specialized, as typical for newer or more rapidly changing capital. $\square$

A.3 What Investment Spikes Identify

This subsection uses a canonical model of lumpy investment, the model that Appendix E takes to the data, to state the conditions under which investment spikes identify the skill bias of new capital (our parameter of interest). In the model, investments are subject to fixed costs, causing firms to lump them into spike events.

Skill Bias. As in Section A.1, we abstract from labor adjustment costs and the make-or-buy decision to home in on the implications of lumpy investment: the labor inputs S_P and U_P sum in-house and external labor, as in our outcome variables. Conditional on its capital stock K_{jt} , the firm faces a series of static labor demand choices:

$$\Pi_{jt}(K) = \max_{S_P, U_P} D_{jt} F(K, \theta_{jt}^S S_P, \theta_{jt}^U U_P) - w_{St} S_P - w_{Ut} U_P, \quad (15)$$

where $\theta_{jt} = (\theta_{jt}^S, \theta_{jt}^U)$ are labor-augmenting productivities, D_{jt} is the firm's output demand shifter, and w_t are market wages. We assume production F is increasing and homothetic (decreasing returns to scale are allowed),³⁵ and the integration technology G is homothetic with constant returns to scale.

Taking first-order conditions, the demand shifter D_{jt} cancels, and the relative demand for skills satisfies

$$\frac{\theta_{jt}^S F_S(K, \theta_{jt}^S S_P, \theta_{jt}^U U_P)}{\theta_{jt}^U F_U(K, \theta_{jt}^S S_P, \theta_{jt}^U U_P)} = \frac{w_{St}}{w_{Ut}}, \quad (16)$$

with F_S and F_U the derivatives of F with respect to effective skilled and unskilled labor. Because F is homothetic, these marginal products compare only the *proportions* of the

³⁵Homotheticity also accommodates a downward-sloping firm demand curve: with iso-elastic demand of elasticity $\varepsilon^Y > 1$, revenue is $D_{jt} F^{1-1/\varepsilon^Y}$, again an increasing homothetic function of inputs.

input bundle, never its size. Inverting Equation (16), the skill mix is determined by:

$$\ln \frac{S_{Pjt}}{U_{Pjt}} = \psi \left(\underbrace{K_{jt}/N_{jt}}_{\substack{\text{capital deepening:} \\ \text{variable of interest}}}, \underbrace{\theta_{jt}^S/\theta_{jt}^U}_{\substack{\text{factor-biased technology:} \\ \text{potential confounder}}}, \underbrace{w_{St}/w_{Ut}}_{\substack{\text{relative wages:} \\ \text{common across firms}}} \right), \quad (17)$$

where $N_{jt} \equiv \theta_{jt}^S S_{Pjt} + \theta_{jt}^U U_{Pjt}$ and the function ψ is determined by the curvature of F .³⁶

Equation (17) distills the identification problem: across firms, differences in relative skill demand can arise from capital deepening K_{jt}/N_{jt} , our variable of interest, and the factor-biased shifters θ_{jt} , the potential confounders. Relative wages are common to all firms and absorbed by time effects. Output demand is absent: D_{jt} moves the scale of the firm, not its skill mix.

Investment Spikes. We now turn to what sets off an investment spike. Let the investment cost be $c_{jt}^I(X) = \nu_{jt} + c(X)$ for $X \neq 0$, with $\nu_{jt} > 0$ a fixed-cost draw paid upon any adjustment and $c(X)$ a convex component; the program in Equation (11) then becomes

$$V_{jt}(K) = \Pi_{jt}(K) + \max \left\{ \begin{aligned} &\beta \mathbb{E}_t V_{jt+1}((1 - \delta)K), \\ &\max_{X, S_I, U_I} \beta \mathbb{E}_t V_{jt+1}((1 - \delta)K + G(X, \theta_{jt}^S S_I, \theta_{jt}^U U_I)) \\ &\quad - \nu_{jt} - c(X) - w_{St} S_I - w_{Ut} U_I \end{aligned} \right\}, \quad (18)$$

where the first branch waits and the second adjusts. Comparing the two branches, the firm invests if and only if the value gain from optimally adjusting its capital covers the

³⁶Formally, let $\tilde{S} = \theta_{jt}^S S_P$, $\tilde{U} = \theta_{jt}^U U_P$, and $s = \tilde{S}/N_{jt}$. Because F is homothetic, the ratio F_S/F_U is homogeneous of degree zero in $(K, \tilde{S}, \tilde{U})$, so dividing its arguments by N_{jt} turns Equation (16) into $(F_S/F_U)(K_{jt}/N_{jt}, s, 1 - s) = (\theta_{jt}^U/\theta_{jt}^S)(w_{St}/w_{Ut})$: one equation in the skill split s given capital per effective worker and the two ratios, which pins down s under standard curvature conditions. The skill mix in Equation (17) follows, since $S_{Pjt}/U_{Pjt} = (\theta_{jt}^U/\theta_{jt}^S) s/(1 - s)$ is strictly increasing in s .

fixed-cost draw:

$$\max_{X, S_I, U_I} \left\{ \beta \mathbb{E}_t [V_{jt+1}((1 - \delta)K + G(\cdot)) - V_{jt+1}((1 - \delta)K)] - c(X) - w_{St}S_I - w_{Ut}U_I \right\} \geq \nu_{jt}. \quad (19)$$

The gain on the left grows as the gap between the firm's target and actual capital widens, through depreciation and the accumulation of shocks, while the fixed-cost draw on the right does not. Firms therefore wait until the gap is large enough for the gain to cover the draw and then close it in one step: adjustment follows an (S, s) -type policy, and investment arrives in the occasional, discrete episodes that our spikes pick up (Bloom, 2009; Cooper and Haltiwanger, 2006; Khan and Thomas, 2008; Winberry, 2021). The shadow value of installed capital rises as the gap widens and falls discretely once the firm closes it. Whatever provides the final push, a spike is therefore a discrete decline in the firm's effective price of capital: our ideal experiment.

The identification problem in Equation (17) extends to integration labor. Differentiating the adjust branch of Equation (18) with respect to S_I and U_I , the shadow value of installed capital cancels, and the relative demand for integration labor satisfies $\theta_{jt}^S G_S / (\theta_{jt}^U G_U) = w_{St} / w_{Ut}$ (see the proof of Prediction 1(b)). Because G is homothetic with constant returns to scale, inverting this condition gives the integration skill mix as a function of the same three arguments as in Equation (17), with hardware per effective integration worker in place of capital deepening.

Identification. What pushes a firm across the trigger? A low investment-cost draw ν_{jt} delivers this decline exogenously: it makes adjustment cheap yet appears nowhere in the labor demands of Equation (15). Accumulated growth in output demand D_{jt} also raises the gain from adjusting, but it is absent from the relative demand (17), so demand-triggered spikes still identify the skill bias. A neutral trigger moves all of a firm's labor

demands and capital proportionally, so it differences out of the skill differentials and leaves the mapping in Equation (8) unchanged. Factor-biased shocks θ_{jt} , by contrast, both set off spikes and move the skill mix through (17). This is the key potential confounder, and ruling it out is our identifying assumption:

Assumption 1 (Shocks and spike timing). *Conditional on the pre-spike state $\mathcal{I}_{j,e-1}$, factor-biased innovations are orthogonal to spike timing: $\mathbb{E}[\theta_{j,e+k} - \theta_{j,e-1} \mid \mathcal{I}_{j,e-1}, \text{spike at } e] = \mathbb{E}[\theta_{j,e+k} - \theta_{j,e-1} \mid \mathcal{I}_{j,e-1}, \text{no spike at } e]$ for all $k \geq 0$.*

Assumption 1 is the structural counterpart of the identification discussion in Section 3.1.1: the threat to identification is not that spikes are chosen, but that their timing coincides with factor-*biased* changes in fundamentals. This also gives a structural motivation for our pre-trend checks: diverging trajectories in factor-biased fundamentals would appear as pre-trends in skill bias through Equation (17).

Notably, Assumption 1 is implied by the assumptions of standard dynamic production-function estimators (Doraszelski and Jaumandreu, 2013; Olley and Pakes, 1996) that combine time-to-build (as in Equation (18)) with a Markov process for productivities θ_{jt} , whose innovations are mean-independent of the cost and demand draws. A spike at e is then decided in year $e - 1$, and conditioning on it adds nothing to what $\theta_{j,e-1}$ already predicts:

$$\mathbb{E}[\theta_{j,e+k} \mid \theta_{j,e-1}, \text{Spike}_{je}] = \mathbb{E}[\theta_{j,e+k} \mid \theta_{j,e-1}] \quad \text{for all } k \geq 0. \quad (20)$$

Because pre-spike outcomes reveal the pre-spike state through the factor demands in Equation (15), firms with similar pre-spike outcomes have similar states $\theta_{j,e-1}$, and Assumption 1 holds conditional on pre-spike outcomes.³⁷ Comparing such firms also requires

³⁷This inversion requires the static factor demand system implied by Equation (15) to be invertible in firm productivities. Berry, Gandhi and Haile (2013) show that a “connected substitutes” condition ensures invertibility; the set of production functions satisfying it includes nested CES (as in Appendix A.4), non-homothetic CES, mixed CES, and translog.

overlap: firms in the same pre-spike state must differ in whether they spike. This holds for states in which the value gain on the left of Equation (19) lies within the support of the fixed-cost draw ν_{jt} , so that some of these firms draw a cost low enough to invest and others do not.

In our setting, spiking and comparison firms have similar employment levels and skill mixes two years before the spike (Table 2), and the event studies show parallel trends before it; the only movement in year -1 , when firms decide to invest, is in external employment, as they begin engaging integrators for the upcoming installation. Hence, under time-to-build and Markov productivity dynamics, the two groups also have similar expected productivities in the post-event period, consistent with Assumption 1. Moreover, because comparison firms are matched only on industry and year, the balance on employment outcomes also serves as an overidentifying check in the spirit of Altonji, Elder and Taber (2005) and Oster (2019).

A.4 Skill Bias and Substitution Elasticities

This subsection shows how our estimated responses map into the substitution elasticities at the heart of the capital–skill complementarity literature, deriving Equation (8) of the main text.

Following Krusell et al. (2000), we specialize the homothetic production function of Appendix A.3 to a nested CES that combines machinery K with high-STEM production labor S_P in the inner nest, and that composite H with low non-STEM production labor U_P in the outer nest:

$$H = \left[c^K K^\rho + c^S S_P^\rho \right]^{1/\rho}, \quad \sigma_{KS} = \frac{1}{1-\rho}, \quad (21)$$

$$F = \left[c^H H^\lambda + c^U U_P^\lambda \right]^{1/\lambda}, \quad \sigma_{HU} = \frac{1}{1-\lambda}. \quad (22)$$

Here σ_{KS} is the elasticity of substitution between machinery and high-STEM labor, and

σ_{HU} the elasticity between the machinery–high-STEM composite and less-skilled labor. Capital–skill complementarity in the sense of Griliches (1969) and Krusell et al. (2000) corresponds to $\sigma_{HU} > \sigma_{KS}$: capital substitutes less readily for skilled than for less-skilled labor.

Derivation. We model the investment spike as a firm-specific decline in the capital user cost r and hold wages fixed within the event window, as the comparison of investing with non-investing firms does, because wage changes common to both difference out (Appendix E.4). The user cost r is the implicit rental of installed capital, the marginal value of capital services in the static problem of Appendix A.3 given the firm’s stock, which a cost-minimizing labor bundle reveals; the derivation therefore requires no interior first-order condition for capital. In the static approximation of Appendix E.4.2, $r = (i + \delta) p^K$. Log-differentiating the cost-minimizing factor demands implied by (21)–(22) yields four first-order relations, in which all changes are treated-minus-control responses and s_K^H is capital’s cost share in the inner nest:

$$(i) \text{ inner first-order condition: } \Delta \ln K - \Delta \ln S_P = -\sigma_{KS} \Delta \ln r, \quad (23)$$

$$(ii) \text{ inner price index: } \Delta \ln P_H = s_K^H \Delta \ln r, \quad (24)$$

$$(iii) \text{ outer first-order condition: } \Delta \ln U_P - \Delta \ln H = \sigma_{HU} s_K^H \Delta \ln r, \quad (25)$$

$$(iv) \text{ composite quantity index: } \Delta \ln H = s_K^H \Delta \ln K + (1 - s_K^H) \Delta \ln S_P. \quad (26)$$

Relation (23) is the relative factor demand within the inner nest, and (25) is its counterpart for the outer nest. Relation (24) follows from Shephard’s lemma: the elasticity of the inner-nest unit cost P_H with respect to r equals capital’s cost share s_K^H in that nest. Relation (26) is the CES quantity index. Substituting (23) and (26) into (25) eliminates

the latent shock $\Delta \ln r$, which cancels, leaving

$$\frac{\sigma_{HU}}{\sigma_{KS}} = 1 + \frac{1}{s_K^H} \cdot \frac{\Delta \ln S_P - \Delta \ln U_P}{\Delta \ln K - \Delta \ln S_P}. \quad (27)$$

The ratio exceeds one (capital–skill complementarity) if and only if high-STEM production employment rises more than low non-STEM production employment ($\Delta \ln S_P > \Delta \ln U_P$), given that capital deepens faster than skilled employment ($\Delta \ln K > \Delta \ln S_P$). The relations are exact for small changes; for the finite changes around a spike they hold with s_K^H averaged along the adjustment path, and since capital’s share rises as the firm deepens when $\sigma_{KS} > 1$, evaluating Equation (27) at the pre-spike share overstates the ratio (Appendix E.3.2). Equation (8) of the main text evaluates it with total high-STEM and low non-STEM employment, L_S and L_U , which also include integration labor. As in the main text’s identification argument, the mapping requires only that no factor-*biased* productivity shock coincide with the spike (Assumption 1): Hicks-neutral and pure scale shocks wash out of the *relative* responses in (27).

Estimation. Evaluated with total employment, the capital stock, and capital’s cost share at the firm’s optimum, Equation (27) overstates the ratio: on the model’s simulated event study it gives 1.37 against a true ratio of 1.15, largely because total employment still contains integration labor, so the skill differential combines both complementarities (Appendix E.3).

B Data and Measurement

B.1 Data

We link the following firm-level datasets for the period 2003–2021:

Annual Accounts. We obtain annual accounts for all incorporated non-financial firms in Belgium from the National Bank of Belgium. The dataset includes standard balance-sheet and profit-and-loss items, along with five-digit NACE industry codes. Crucially for our analysis, the book value of tangible fixed assets is disaggregated into (i) land and buildings, (ii) machinery and equipment, and (iii) rolling stock and office machinery. We retain firms that report positive employment in at least one year.

VAT Declarations. All VAT-liable firms in Belgium are required to report their annual totals for sales, investment, and intermediate inputs to the tax authorities. We use these VAT declarations to measure sales, since the annual accounts report sales figures only for large firms. These sales include both sales to consumers and businesses as well as exports. The data on investment refers to total investment and is not disaggregated by type (machinery, buildings, etc.).

Business-to-Business (B2B) Transactions. The VAT customer listings record domestic buyer–supplier links. Every VAT-liable firm must file an annual report listing the aggregate value of invoices issued to each customer. Imports and exports are not covered. We restrict the sample to pairs in which both firms file an annual account and the annual transaction value exceeds €1,000; this keeps the dataset manageable and drops economically negligible transactions. The resulting dataset contains 115.6 million transactions for 2003–2021. Our data construction follows Dhyne, Duprez and Komatsu (2023) and updates the dataset used by, among others, Bernard et al. (2022) and Baqaee et al. (2025).

Employee-Level Social-Security Data. We obtain worker-level information from the Crossroads Bank for Social Security (KSZ). Employees are linked to their employers’ social-security declarations, and each worker is assigned the level and field of their highest

diploma.³⁸ The classification of education fields changed in 2021, affecting the classification of low STEM and low non-STEM workers. Our employment results therefore focus on the period 2003–2020.

Tables B.1 and B.2 list the education categories by level and field.³⁹ For each firm-year, we observe both the headcount and average full-time-equivalent (FTE) hours for every education level-by-field cell, allowing us to construct FTE employment by education category.

We classify degrees into four categories based on the cross of these two dimensions:

- **High STEM:** tertiary education or higher in STEM fields.
- **High Non-STEM:** tertiary education or higher in non-STEM fields (including “not specified”).
- **Low STEM:** less than tertiary education (including undetermined or missing levels) in STEM fields.
- **Low Non-STEM:** less than tertiary education (including undetermined or missing levels) in non-STEM fields (including “not specified”).

³⁸Education information is sourced from the Belgian population census; hence, only employees residing in Belgium are included.

³⁹A sizable share of low-education records is coded as “not specified,” whereas the corresponding share among tertiary academic degrees is only 5.9%.

Table B.1: Levels of Education

Code	Description	Category	Share of Employment
0	Pre-primary education / less than primary education	Low	1.1%
02	Limited pre-primary education	Low	0.0%
10	Primary education	Low	5.9%
24	Lower secondary education – General	Low	9.0%
25	Lower secondary education – Vocational	Low	9.7%
34	Upper secondary education – General	Low	22.4%
35	Upper secondary education – Vocational	Low	13.3%
40	Post-secondary non-tertiary education	Low	2.6%
44	Post-secondary non-tertiary education – General	Low	0.4%
45	Post-secondary non-tertiary education – Vocational	Low	2.6%
55	Higher education – Short cycle – Vocational	High	0.2%
64	Higher education – Bachelor’s or equivalent – Academic	High	0.1%
65	Higher education – Bachelor’s or equivalent – Vocational	High	14.8%
74	Higher education – Master’s or equivalent – Academic	High	9.1%
84	Doctorate or equivalent – Academic	High	0.3%
98	Highest degree not determined for persons younger than 15	Low	0.0%
99	Highest education level could not be determined	Low	5.1%
.	Missing	Low	3.3%

Notes: This table shows how we classify education levels as “low” or “high.” The last column reports each level’s share of total full-time-equivalent employment in 2012, the midpoint of our sample period.

Table B.2: Fields of Education

Code	Description	Category	Share of Employment
00	Generic programmes and qualifications	Non-STEM	11.3%
01	Education	Non-STEM	1.0%
02	Arts and humanities	Non-STEM	3.9%
03	Social sciences, journalism and information	Non-STEM	2.1%
04	Business, administration and law	Non-STEM	11.3%
05	Natural sciences, mathematics and statistics	STEM	1.6%
06	Information and Communication Technologies	STEM	1.6%
07	Engineering, manufacturing and construction	STEM	26.7%
08	Agriculture, forestry, fisheries and veterinary	Non-STEM	1.1%
09	Health and welfare	Non-STEM	2.3%
10	Services	Non-STEM	3.6%
99	Not specified	Non-STEM	30.2%
NA	Not available	Non-STEM	0.0%

Notes: This table shows how we classify education field as “STEM” or “Non-STEM.” The last column reports each field’s share of total full-time-equivalent employment in 2012, the midpoint of our sample period.

B.1.1 Linking Datasets

Linking KSZ and Annual Accounts. We link the datasets using unique firm identifiers. Of the 3.9 million annual account observations, 2.7 million match to KSZ social-security records. The remaining 1.2 million observations almost entirely reflect firms with zero employees and therefore no KSZ record.⁴⁰ Restricting attention to firms that report positive employment, we obtain 2,717,110 observations on 326,554 firms for the period 2003–2021. Within manufacturing, the sample comprises 269,556 observations on 30,030 firms.

⁴⁰Only 83,000 unmatched observations report positive employment. These correspond to firms that file an annual report but lack a social-security number (for example, certain intermunicipal cooperatives).

Linking Annual Accounts and B2B Transactions. We link the annual accounts (complemented with KSZ employment data) to the B2B transactions using the same unique firm identifiers. Among the 115.6 million transactions in the B2B records, 90.6% of suppliers are matched to firms in our annual accounts dataset, corresponding to 91.8% of total transaction value. The unmatched suppliers are firms with zero employment.⁴¹ Dropping transactions with suppliers without employment has no bearing on first-order external employment and only drops higher-order chains that pass through firms without employees.

B.2 Measurement

B.2.1 A Robot-Integration Example

To fix ideas about the external skills required to integrate and operate new machinery, and how they enter our data, Box B.1 traces a typical robot-integration project through the buyer–supplier transactions we observe. The phases of the project, its split between hardware and integration labor, its milestone billing, and its recurring service relationship all follow standard industry practice and contract terms.⁴² Robot integration is a leading example of the machine-complementary services we measure, but only one instance of the broader skill bias of new machinery.

⁴¹Our annual accounts dataset includes only firms that report positive employment in at least one period, a restriction not applied to the B2B dataset.

⁴²We thank Søren Peters (HowToRobot), Maarten Daemen (KUKA), Peter Kiekens (Industrial Cobotics), Benjamin Parent (Stäubli), Alain Wayenberg (Agoria), Søren Peter Johansen (Danish Technological Institute, Robot Technology), Bo Hanfarn Eriksen (Region Syddanmark), Per Rasmussen (BILA Robotics), and Martin Jespersen (Odense Robotics) for conversations that greatly helped us understand the structure of the robot-integrator industry.

Box B.1: A Robot-Integration Project in the Transaction Data

A manufacturer decides to automate a production step and orders a robot cell from a specialized *integrator*. The installed system costs roughly €250,000 (Anandan, 2015). Only about one-third of this is the robot hardware; the remaining two-thirds pays for integration, covering the engineering and design of the cell, safety guarding, programming, installation and commissioning, and operator training (IFR, 2022; Anandan, 2015). A typical project runs over a few months, from design and engineering through build and factory testing to on-site installation, commissioning, and training (Elkins, 2025). The integrator’s engineers and technicians, many of them high-STEM, do this work, not the manufacturer’s own staff.

The project (spike year). Payments track the progress of the work. The European engineering industries’ standard contract conditions (Orgalime S 2012) stipulate one third of the price at the order, one third when the system is ready for delivery, and the remainder at delivery (Orgalime, 2012, Clause 19); bespoke contracts often tie further milestones to the project’s sign-off gates, such as design approval and acceptance testing (Elkins, 2025). The project is short, so almost all of the €250,000 shows up as the manufacturer’s purchases from the integrator within a single year, the investment-spike year, and the payments arrive in the same year as the integration labor they pay for. Under our proportionality assumption (Section 2.2.3), if the manufacturer accounts for a share s of the integrator’s sales that year, we attribute the share s of the integrator’s workforce to the manufacturer. In this way, the transaction data pick up the integration labor behind the investment even though none of these workers ever appear on the manufacturer’s payroll.

After commissioning (later years). The relationship does not have to end at handover: the manufacturer may keep buying preventive maintenance and spare parts, technical support, and, if it changes its product mix, reprogramming and re-tooling (Anandan, 2015). Industry estimates put such running costs on the order of 10–12% of the initial system cost per year (HowToRobot, 2024), with maintenance requirements rising as components wear. These smaller, recurring transactions sustain the buyer–supplier link, attributing a now-smaller share of the integrator’s engineers to the manufacturer for years after the installation.

The example clarifies the time profile of our estimates: external high-STEM demand at machine suppliers (Figure 3, Panel (a)) rises sharply in the spike year (the integration burst) and then declines without returning to baseline, reflecting the persistent servicing relationship.⁴³

⁴³Outsourcing integration is not legally mandated. Robot warranties typically require installation according to the manufacturer’s instructions, not by a certified external firm, and under the EU Machinery

B.2.2 Machine Supplier Industries

The value chain of machine supply includes producers, distributors, installers, and advisors of new machinery. While production is concentrated among a few large, international suppliers, the downstream stages of the value chain are primarily domestic and sometimes vertically integrated. For example, robot integrator firms often combine distribution, installation, and advisory services. Table B.3 lists the industry codes used to classify each of the machine supplier categories.

Table B.3: Machine Supplier Industries


Type	NACE
Builders	2562, 28
Installers	33, 4321
Distributors	4614, 466
Advisors	71, 7022

Notes: This table lists our classification of machine supplier industry codes.

Confirming the relevance of these industries, Figure B.1 shows that firms sharply increase their purchases from them around machinery investment spikes, especially for firms in manufacturing, construction, agriculture, and utilities.

Directive whoever assembles the final robot cell, an external integrator or the user itself, becomes its legal manufacturer, responsible for its risk assessment and CE marking (European Commission, 2024; International Organization for Standardization, 2025). Self-integration thus internalizes a fixed cost of integration capability that specialized integrators spread across clients (Prediction 2).

Figure B.1: B2B Transactions: Purchases from Machine Supplier Industries



Notes: This figure shows the share of purchases from machine suppliers around machinery investment spikes for all private-sector firms (Panel a), Manufacturing (Panel b), Agriculture/Utilities/Construction (Panel c), and Services firms (Panel d). The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.


We use two external data sources on robot suppliers to validate our industry classification. First, we use a market report on robot suppliers and integrators in Belgium (compiled by www.HowToRobot.com, a digital platform that connects manufacturing firms to machine suppliers).⁴⁴ Second, we gained access to the Customer Relationship Management (CRM) systems of three of Belgium’s largest robot producers and distributors. These lists identify 606 robot supplier firms in 2021.

Figure B.2 decomposes purchases from firms on the external lists by whether they fall in our machine-supplier industry codes. As the figure shows, our industry classification

⁴⁴Békés et al. (2025) also use the HowToRobot database to identify robot integrators.

captures most of the machine supply from the external lists. Compared to the external lists, the industry code classification has the advantage of being time consistent (the external lists are snapshots from 2021) and of covering machine supplies beyond robots (the external lists focus on *robot* suppliers).

Figure B.2: Share of Purchases from External List of Robot Supplier Firms




Notes: This figure splits the purchases of spiking firms from the robot supplier firms (in percent of total purchases) according to whether the suppliers are in our classified machine supplier industry codes. The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Manufacturing firms around their first machinery investment spikes. Balanced sample for five years before and after the spike.

To further validate our machine supplier classification and to underscore the importance of domestic integrators, Figure B.3 splits the accumulated imports of industrial robots (HS code 847950) by the industry of the importer. As Belgium does not produce industrial robots, all robots have to be imported from abroad. Indeed, customs records have been a common source of data on robot investments at the firm level; see e.g., Acemoglu, Lelarge and Restrepo (2020); Bonfiglioli et al. (2024) in France, Humlum (2022); Mann and Pozzoli (2022) in Denmark, Dixon, Hong and Wu (2021) in Canada, and Acemoglu, Koster and Ozgen (2023) in the Netherlands. However, as the figure shows, end-user industries (that is, manufacturing firms that are not machine suppliers) represent less than 10% of the robot imports in Belgium. By contrast, our machine supplier industries import more than 75% of all industrial robots in Belgium before these are sold

to end-users (usually bundled with integration services). Békés et al. (2025) document a similar pattern in Hungary.

Figure B.3: Accumulated Robot Imports by Sector of Importing Firm



Notes: This figure splits the accumulated aggregate imports of industrial robots (HS code 847950) to Belgium by the industry of the importer.

C New Machinery and the Demand for Skills

C.1 Empirical Strategy

Table C.1 shows how our analysis sample shrinks with the sample restrictions described in Section 3.1.

Table C.1: Machinery Investment Spikes Sample Reduction

Sample Step	Firms	Investment Spikes
	(1)	(2)
1. Private-sector firms with a machinery spike	37,379	57,783
2. At least five workers before the spike	9,463	15,542
3. First investment spike	9,463	9,463
4. Balanced sample (four years before and after)	3,863	3,863


Notes: This table shows how our analysis sample shrinks with the sample criteria listed in Section 3.1. Column (1) shows the number of eligible firms, and Column (2) shows their number of machinery investment spikes.

C.1.1 Proxy Controls

This subsection examines whether confounding output demand or productivity shocks drive our results, first by tracing sales around the spikes and then by controlling for their growth.

Figure C.1 shows the difference-in-differences estimate for log sales around machinery investment spikes, using the same specification as in Figure 2. Sales evolve smoothly around the spike, with no discernible pre-trend, and grow to about 20 percent four years after the spike, a relative increase similar to that for total employment. This smooth, gradual path is hard to reconcile with a discrete confounding demand or productivity shock coinciding with the investment spike. Even such a shock would not generate the high-STEM skill bias we document: output demand shocks raise all skill groups proportionally (Lindner et al., 2024), as Appendix A.3 shows.

Figure C.1: Sales around Machinery Investment Spikes




Notes: This figure shows the difference-in-differences estimate for log sales (in %) around machinery investment spikes, using the same specification as in Figure 2 (Equation (6)). The estimates are differences-in-differences in log sales (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2, with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

Figure C.2 adds controls for this sales growth directly to our employment specification.⁴⁵ Panels (a) and (b) reproduce the in-house and external employment effects from Figure 2; Panels (c) and (d) re-estimate them controlling for the firm’s growth in sales from year -2 to the event year. The controls lower absolute employment growth somewhat (most visibly for in-house employment in the spike year, where the year-0 effect falls by 1–2 percentage points off a 3–5% base), but the external high-STEM skill bias, our central finding, remains intact.

⁴⁵This control is consistent with the conceptual framework in Section 1: because capital takes a year to install (Equation (2)) and integration labor does not enter production (Equation (1)), year-0 sales reflect output produced with the firm’s existing, pre-spike capital stock, so controlling for them is not a “bad control”. We thank John Van Reenen for this suggestion.

Figure C.2: Employment around Machinery Spikes: Controlling for Sales Growth
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)




Notes: This figure reproduces the in-house and external employment effects from Figure 2 with and without controls for the investing firm’s sales growth. Panels (a) and (b) show the main specification (no controls); Panels (c) and (d) add controls for the firm’s growth in log sales from year -2 to the event year (year 0). The figures show employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2, based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

C.2 Main Results

Figure C.3 separates the external high-STEM bias of Figure 2 into two drivers, normalizing each component by its own pre-spike level rather than by total employment. Panel (a) shows that external employment grows faster than in-house employment around machinery investment spikes, while Panel (b) shows that the skill composition of external employment shifts toward high-STEM workers. Both channels reinforce each other in

generating the patterns observed in Figure 2: more outsourcing of an increasingly high-STEM-intensive supplier mix.

Figure C.3: Drivers of External Skill Bias around Machinery Investment Spikes




Notes: Panel (a) shows in-house employment in percent of pre-spike in-house employment and external employment in percent of pre-spike external employment. Panel (b) shows the external employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of pre-spike external employment of each skill group. Pre-spike employment is the average over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

C.2.1 Treated and Comparison Group Means

Figure C.4 plots total employment separately for spiking firms and their comparison firms, each in percent of its pre-spike level. The two groups start from the same level two years before the spike, the base year of our event studies, so the gap that opens between the lines is our difference-in-differences estimate (the sum of the in-house and external estimates in Figure 2): for high-STEM workers, 50% in the spike year and 33% four years later.

The sharp break at the spike comes entirely from the spiking firms. Comparison firms have not yet spiked, so their employment follows the smooth trend of their industry. A different comparison group can change this trend but not the break, which is why our estimates change little across the alternative comparison groups in Appendices C.3.4 and C.3.5. The comparison firms are also upgrading skills: between year -2 and year 4, their high-STEM employment grows by 15%, against 4% for low non-STEM workers. Our design differences out this common trend.

Figure C.4: Total Employment around Machinery Investment Spikes, by Group
(In Percent of Pre-Spike Total Employment by Skill Group)



Notes: This figure shows total employment (in-house plus external) around machinery investment spikes, separately for spiking firms and their matched comparison firms (the “spiking” and “non-spiking” firms of Table 2). Each panel covers one skill group, in percent of the firm’s total pre-spike employment in that group (Equation 7), so both lines are close to 100 before the spike. Pre-spike refers to the average over the years -3 to -1 . The two groups differ by at most 0.3 percentage points in year -2 , the base year of our event studies. At every horizon, the change in the vertical distance between the two lines relative to year -2 equals the sum of the in-house and external difference-in-differences estimates in Figure 2, which also reports their 95% confidence intervals. All four panels share the same vertical scale. *Sample:* Private sector firms around their first machinery investment spike and their matched comparison firms. Balanced sample for four years before and after the spike.

C.2.2 Decomposing Other Industries

Table C.2 decomposes the impact of machinery investment spikes on external employment (as in Figure 3) across subindustries. The high-STEM bias (last column) is positive in most supplier industries. Relative to their pre-spike share, it is most pronounced for computer programming, followed by utilities and construction; in absolute terms, wholesale and retail trade and construction contribute the most. Employment agencies,

a large supplier of external labor, contribute broadly across skill groups but add no skill bias.

Table C.2: Decomposing the Impact on External Employment through Other Industries

	Pre-Spike	Effect on External Employment (Year 4)				Skill
	Share	High STEM	High Non-STEM	Low STEM	Low Non-STEM	Bias
	(1)	(2)	(3)	(4)	(5)	$\frac{(2)-(5)}{(1)}$
Machine Suppliers	10.51	4.76	1.27	1.18	0.45	0.41
Other industries						
1-9 Agriculture/Mining	2.05	0.03	0.01	0.07	0.03	0.00
10-12 Food/Beverages	4.35	0.24	0.38	0.19	0.20	0.01
13-15 Textiles/Apparel	1.06	0.10	0.03	0.03	0.02	0.08
16-18 Wood, Paper, Print	3.25	0.30	0.24	0.23	0.20	0.03
19 Coke	0.04	0.00	0.00	0.00	0.00	0.00
20 Chemical and Pharma	1.31	0.34	0.12	0.05	0.03	0.24
22 Rubber-Plastic	1.21	0.22	0.13	0.08	0.05	0.14
23 Non-Metallic	2.54	0.41	0.34	0.23	0.17	0.09
24 Basic Metals	0.48	-0.01	0.00	0.00	0.00	-0.02
25 Fabricated Metals	3.25	0.60	0.32	0.19	0.27	0.10
26 Computer, Electronics	0.18	0.00	0.00	0.00	0.00	0.00
27 Electrical Equip.	0.35	0.04	-0.01	-0.03	-0.01	0.14
29-30 Transport Equip.	0.46	0.05	0.03	0.05	0.03	0.04
31-32 Other Manuf.	1.85	0.42	0.18	0.19	0.11	0.17
35-39 Utilities	1.79	0.68	0.36	0.21	0.17	0.28
41-43 Construction	9.09	3.08	0.55	1.82	0.57	0.28
45-47 Wholesale/Retail	24.11	6.33	5.35	2.40	1.76	0.19
49-53 Transport/Storage	6.46	0.21	0.24	0.68	0.63	-0.07
55-56 Hotels & Restaur.	1.13	0.02	0.04	0.05	0.04	-0.02
58-60 Publishing	0.60	0.07	0.17	0.01	0.02	0.08
61 Telecom	0.16	0.04	0.00	0.01	0.00	0.25
62 Computer Program	1.35	1.30	0.19	0.04	0.04	0.93
63 Information Serv.	0.10	0.01	0.01	0.00	0.00	0.10
64-66 Financial Serv.	0.29	0.01	-0.02	0.00	0.00	0.03
68 Real Estate	0.42	0.03	0.04	0.01	0.00	0.07
69 Legal Serv.	2.55	0.02	0.27	-0.01	0.01	0.00
70 Manag. Consult	0.72	0.14	0.28	0.02	0.06	0.11
72 R&D	0.03	0.00	0.00	0.00	0.00	0.00
73 Advertising	0.59	0.04	0.16	0.01	0.03	0.02
74-75 Other Serv	0.20	0.04	0.02	0.00	0.00	0.20
77 Rental/Leasing	0.63	0.09	0.07	0.04	0.05	0.06
78 Employment Act.	11.99	2.24	3.19	1.71	2.52	-0.02
79 Travel Agency	0.10	0.00	0.00	0.00	0.00	0.00
80 Security	0.38	0.03	0.02	0.04	0.03	0.00
81 Building Serv.	2.65	0.05	0.19	0.24	0.39	-0.13
82 Administration	0.69	0.00	0.06	0.01	0.01	-0.01
84+ Public	1.06	0.01	0.11	0.05	0.08	-0.07


Notes: This table decomposes the impact of machinery investment spikes on external employment (as in Figure 3) across two-digit industries and groups of them. Column (1) shows the percent share of the industry in external employment in the year before the spike. Columns (2)–(5) show the contribution of this industry to the effect on total employment of each skill group in year 4 after the spike (measured in % of pre-spike total employment). The final column shows the normalized contribution to skill bias, measured as the difference in effects on high-STEM (Column (2)) and Low Non-STEM (Column (5)), relative to the share of the industry in total external employment before the spike. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.

C.3 Robustness

C.3.1 Investment Spikes

Figure C.5 shows that the investment flow around a machinery spike is a one-off event: the machinery flow per worker in Panel (a) peaks in the spike year and falls back sharply the year after, even as the machinery stock it produces remains elevated (Figure 1). The machine investment flow is computed from the stock assuming a depreciation rate of 12%. Panel (b) plots the total fixed investment flow as a share of total sales taken from the VAT records, which shows a similar pattern. The VAT records used to produce the investment data do not split investment by asset type, so Panel (b) captures all fixed investment, not just machinery. Because Panel (b) uses a different data source, it also shows that the spikes are not an artifact of the annual accounts.

Figure C.5: Investment around Machinery Spikes




Notes: This figure plots the machine investment flow per worker (panel (a)) and total fixed investment as a share of total sales around machinery investment spikes (panel (b)). The machine investment flow per worker is $\frac{I_{jt}}{0.5(L_{jt-1} + L_{jt})}$, where $I_{jt} = K_{jt} - (1 - \delta)K_{jt-1}$ with a depreciation rate δ of 12%. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

Figures C.6 and C.7 show that our main findings are robust to varying the percentile cutoff for defining investment spikes. The 85th-percentile cutoff is €10,611 per worker and yields 5,539 spikes in the balanced panel; the 95th-percentile cutoff is €35,079 per worker and yields 1,826 spikes.


At both cutoffs, the qualitative pattern from Figures 2 and 3 carries through: machinery investment spikes raise the demand for high-STEM skills primarily through external suppliers, with the initial response driven by machine suppliers and the longer-run response by other industries.

Figure C.6: Employment around Machinery Spikes: 85th-Percentile Cutoff



Notes: This figure reproduces the employment effects of machinery investment spikes using the 85th-percentile cutoff (€10,611 per worker). Panels (a) and (b) consider in-house and external employment, respectively; Panels (c) and (d) decompose external employment into machine suppliers and other industries. The figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

Figure C.7: Employment around Machinery Spikes: 95th-Percentile Cutoff




Notes: This figure reproduces the employment effects of machinery investment spikes using the 95th-percentile cutoff (€35,079 per worker). Panels (a) and (b) consider in-house and external employment, respectively; Panels (c) and (d) decompose external employment into machine suppliers and other industries. The figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

C.3.2 Single-Spiking Firms

Figure C.8 reproduces the main employment effects of machinery investment spikes (Figures 2 and 3) restricting the sample to the firms that experience only one investment spike within the post-event window. The qualitative pattern is unchanged: in-house employment expands evenly across skill groups (Panel (a)), while external employment exhibits a sharp high-STEM bias (Panel (b)), driven by machine suppliers initially (Panel (c)) and by other industries over the longer run (Panel (d)). The estimates are smaller in the later


post-event years (Figure C.9), because follow-on investments add to later employment in the full sample, whereas here only control firms may still spike, and because conditioning on no follow-on spike selects firms with less favorable shocks (Section 3.1.1).

Figure C.8: Single-Spiking Firms: Employment around Machinery Spikes



Notes: This figure reproduces the employment effects of machinery investment spikes, restricting the sample to the 61% of firms that experience only one investment spike within the post-event window. Panels (a) and (b) consider in-house and external employment, respectively; Panels (c) and (d) decompose external employment into machine suppliers and other industries. The figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Private sector firms around their only machinery investment spike. Balanced sample for four years before and after the spike.

Figure C.9: In-House and External Employment around Machinery Spikes: All Spikes versus Single Spikes




Notes: This figure shows in-house and external employment growth of private sector firms around machinery investment spikes. Panel (a) shows results for all spiking firms; Panel (b) restricts the sample to single-spiking firms. In-house employment is expressed in percent of pre-spike in-house employment; external employment is expressed in percent of pre-spike external employment. Pre-spike values are averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

C.3.3 Unbalanced Panel

Figure C.10 reproduces the main employment effects of machinery investment spikes (Figures 2 and 3), including firm exits as zero in-house and external employment. The qualitative pattern is unchanged: in-house employment expands evenly across skill groups (Panel (a)), while external employment exhibits a sharp high-STEM bias (Panel (b)), driven by machine suppliers initially (Panel (c)) and by other industries over the longer run (Panel (d)). Conditioning on post-spike survival is therefore not driving the result.

Figure C.10: Including Exiting Firms: Employment around Machinery Spikes



Notes: This figure reproduces the employment effects of machinery investment spikes, assigning zero employment to firm exits. Panels (a) and (b) consider in-house and external employment, respectively; Panels (c) and (d) decompose external employment into machine suppliers and other industries. The figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Private sector firms around their first machinery investment spike, active for four years before the spike.

C.3.4 Cell-Average Comparison Group


Figure C.11 reproduces the main employment effects of machinery investment spikes (Figure 2) using as the comparison group the average of firms in the same industry-year that meet the same sample criteria but have not yet experienced a spike as of the event year, in place of a random non-spiking firm in the same industry-year cell. The results are highly similar to the baseline, with slightly lower standard errors.⁴⁶

⁴⁶This industry-year-average estimator shares the imputation logic of Borusyak, Jaravel and Spiess (2024): each spiking firm’s counterfactual is constructed from its own pre-spike outcome and the con-


temporaneous changes of all firms in the same industry-year that had not yet spiked as of the event year. Unlike their estimator, which gains further efficiency by leveraging all pre-treatment periods, it differences from a single pre-spike base year (as firms may engage integrators already in the year before the spike), uses industry-year cell averages rather than estimated fixed effects, and places no restriction on comparison firms' investment after the event year (Appendix C.3.5 shows robustness to imposing such a "clean control" restriction).

Figure C.11: Employment of Skill Groups around Machinery Investment Spikes
(Industry-Year Average as Comparison Group)

(a) In-House



(b) External




Notes: This figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Panel (a) considers in-house employment; Panel (b) considers external employment at supplier firms. Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. The comparison firms are the average of firms in the same industry-year cell that meet the sample criteria and have not yet spiked as of the event year. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

C.3.5 Clean Controls and Later Spikers

Neither the control group nor the treated group is restricted after the event: treated firms may undertake follow-on spikes, and control firms may undertake their first spike. As Figure C.12 shows, about 28% of treated firms undertake a follow-on spike within the post-event window, while about 8% of control firms undertake their first spike. Appendix C.3.2 shows that the main results are robust when limiting the sample to treated firms that do not undertake follow-on spikes within the post-event window. Figure C.13 reproduces the main results using a “clean controls” comparison group including only firms that do not spike within the post-event window, as in Cengiz et al. (2019). The results are highly similar, although the impact is somewhat larger, consistent with the fact that some of the comparison firms in the baseline spike in the post-event window.


Figure C.12: Follow-On and First Spikes after the Event




Notes: This figure shows the percentage of treated firms that undertake a spike (Panel (a)) and the percentage of control firms that undertake their first spike (Panel (b)) within the post-event window. The post-event window is defined as the four years after the first spike. *Sample:* Private sector firms with their matched controls around their first machinery investment spike. Balanced sample for four years before and after the spike.

Figure C.13: Employment of Skill Groups around Machinery Investment Spikes
(Clean Controls: firms that do not spike within the post-event window)

(a) In-House



(b) External




Notes: This figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Panel (a) considers in-house employment; Panel (b) considers external employment at supplier firms. Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. The “non-spiking firms” are random one-to-one matches within the same industry-year cell and exclude firms that spike within the post-event window. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

As an additional robustness check, we also estimate the effects using as comparison firms only firms whose first spike comes after the post-event window (Bessen et al., 2025). As Figure C.14 shows, the external high-STEM effect is, if anything, slightly larger at the spike, while by year four it returns to its baseline magnitude rather than the larger clean-controls one. This is consistent with later spikers beginning to engage integrators before their own investment, which makes them imperfect comparisons toward the end of the window.⁴⁷


⁴⁷Because the pool of comparison firms is much smaller, we match within the same NACE 2-digit industry instead of the same NACE 3-digit industry.

Figure C.14: Employment of Skill Groups around Machinery Investment Spikes
(Later Spikers as Comparison Firms: firms that spike after the post-event window)

(a) In-House



(b) External



Notes: This figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Panel (a) considers in-house employment; Panel (b) considers external employment at supplier firms. Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. The “non-spiking firms” are random one-to-one matches within the same NACE 2-digit industry-year cell and include only firms that spike after the post-event window. The estimates are based on the specification in Equation (6), with whiskers representing 95% confidence intervals. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

C.3.6 Higher-Order Links

Our main analysis measures external employment through firms' first-order suppliers. This subsection extends the analysis to third-order links (the suppliers of suppliers of suppliers) to assess whether the broader network amplifies or attenuates our results. We compute third-order external employment as


$$\mathbf{L}_{jt}^{O,3} = \sum_k \frac{T_{kjt}}{T_{kt}} \left(\mathbf{L}_{kt}^H + \sum_r \frac{T_{rkt}}{T_{rt}} \left(\mathbf{L}_{rt}^H + \sum_s \frac{T_{srt}}{T_{st}} \mathbf{L}_{st}^H \right) \right), \quad (28)$$

which iterates the first-order pass-through of in-house employment along the supply chain. Extending the iteration to all orders would apply the full Leontief inverse of the firm-to-firm input-output matrix, with each additional round damped by another layer of purchase shares.


Figure C.15 shows that our main findings are robust to including these higher-order links. Panel (a) reproduces the main external employment effects (Figure 2) and Panel (b) reproduces the drivers figure (Figure C.3), each measured using the third-order definition. Incorporating higher-order links amplifies the overall quantitative importance of external employment relative to in-house employment.

Figure C.15: Including Higher-Order Supplier Links (Third-Order Definition)

(a) External Employment in Percent of Pre-Spike *Total* Employment by Skill Group (Higher-Order Version of Figure 2)



(b) External Employment in Percent of Pre-Spike *External* Employment by Skill Group (Higher-Order Version of Figure C.3)



Notes: This figure shows third-order external employment growth around machinery investment spikes. Panel (a) shows external employment in percent of pre-spike total employment by skill group. Panel (b) shows external employment in percent of pre-spike external employment by skill group. Pre-spike values are averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

D Why Firms Outsource Capital-Complementary Skills

This appendix collects supplementary results for the analysis of why firms outsource capital-complementary skills in Section 4. Appendix D.1 reports balance and machinery-stock dynamics for the smaller- and larger-firm subsamples, Appendix D.2 does so for the manufacturing and services subsamples, Appendix D.3 reports additional results for building and intangible spikes, Appendix D.4 provides supplementary results on new versus existing machine supply, Appendix D.5 presents the exhibits behind the supply-side evidence, and Appendix D.6 examines the wage-savings motive.

D.1 Firm Size

Tables D.1 and D.2 report pre-spike employment for the smaller- and larger-firm subsamples, where smaller firms have pre-spike total employment (year -1) below the sample median (13.1 FTE). Spiking and non-spiking firms are well balanced on total employment and skill shares within each subsample, supporting the difference-in-differences design for both groups.

Comparing across subsamples, larger firms employ substantially more high-STEM workers in-house already before the spike: their in-house high-STEM share is roughly twice that of smaller firms (6.5% versus 3.2%). This pre-existing in-house capacity helps explain why larger firms rely less on external skills to integrate and operate more capital-intensive production, consistent with larger firms having already absorbed the fixed cost of in-house high-STEM hires.

Figures D.1 and D.2 show that the underlying investment shock is similar for smaller and larger firms: both subsamples experience an increase of roughly 40,000 euros in machinery stock per worker, and the difference-in-differences estimate of machinery per worker is not statistically different across firm size at the 5 percent level, although the point estimates diverge after year 1. Hence, the differences in in-house and external

employment responses across smaller and larger firms in Figure 4 reflect heterogeneous responses to a similar shock, not differences in the magnitude of the shock itself.

Figure D.3 shows that smaller firms rely more on purchases from external machine suppliers when upgrading their machinery: their share of machine suppliers in total purchases rises by six percentage points in the year of the spike, versus four percentage points for larger firms.

Table D.1: Smaller Firms: Firm Employment before Machinery Investment Spikes

	Spiking		Non-Spiking		Std. Difference	
	In-House (1)	External (2)	In-House (3)	External (4)	In-House (5)	External (6)
Total employment	5.7 (5.2)	2.2 (2.2)	6.2 (2.4)	1.9 (2.1)	-0.12	0.14
Share of employment (%)						
— High STEM	3.2 (10.1)	5.3 (5.8)	3.1 (9.4)	4.9 (5.6)	0.01	0.07
— High non-STEM	13.4 (21.3)	17.0 (11.3)	14.4 (21.6)	17.6 (11.9)	-0.05	-0.05
— Low STEM	26.2 (26.0)	25.2 (11.5)	23.9 (23.6)	24.3 (11.7)	0.09	0.08
— Low Non-STEM	57.1 (28.2)	52.5 (12.5)	58.5 (27.0)	53.2 (12.7)	-0.05	-0.06
Nr. Firms	1,926		1,926			


Notes: This table reports in-house and external employment in the treatment (Columns 1–2) and control (Columns 3–4) groups in year $t - 2$. The table reports total employment in full-time equivalents and the employment share of different skill groups. Columns 5–6 show the standardized difference between the treatment and control group means, i.e. the difference in means divided by $\sqrt{(s_T^2 + s_C^2)/2}$, where absolute values above 25% are a standard threshold for assessing imbalance (Stuart and Rubin, 2008). *Sample:* Smaller firms (pre-spike total employment below the sample median of 13.1 FTE) around their first machinery investment spike. Balanced sample for four years before and after the spike.

Table D.2: Larger Firms: Firm Employment before Machinery Investment Spikes

	Spiking		Non-Spiking		Std. Difference	
	In-House (1)	External (2)	In-House (3)	External (4)	In-House (5)	External (6)
Total employment	71.6 (456.0)	36.3 (183.4)	72.8 (339.8)	32.2 (166.7)	0.00	0.02
Share of employment (%)						
— High STEM	6.5 (10.5)	6.3 (5.7)	5.7 (9.3)	6.0 (6.0)	0.08	0.05
— High non-STEM	14.4 (16.2)	15.7 (9.3)	14.8 (16.3)	16.3 (9.7)	-0.02	-0.06
— Low STEM	25.9 (18.1)	24.4 (9.5)	24.6 (17.3)	23.2 (9.5)	0.07	0.13
— Low Non-STEM	53.2 (21.4)	53.6 (11.4)	54.9 (20.4)	54.5 (11.4)	-0.08	-0.08
Nr. Firms	1,926		1,926			


Notes: This table reports in-house and external employment in the treatment (Columns 1–2) and control (Columns 3–4) groups in year $t - 2$. The table reports total employment in full-time equivalents and the employment share of different skill groups. Columns 5–6 show the standardized difference between the treatment and control group means, i.e. the difference in means divided by $\sqrt{(s_T^2 + s_C^2)}/2$, where absolute values above 25% are a standard threshold for assessing imbalance (Stuart and Rubin, 2008). *Sample:* Larger firms (pre-spike total employment at or above the sample median of 13.1 FTE) around their first machinery investment spike. Balanced sample for four years before and after the spike.

Figure D.1: Machinery Stock per Worker around Spike




Notes: This figure shows the machinery capital stock per worker around years of machinery investment spikes. Panel (a) focuses on smaller firms, and Panel (b) on larger firms. Smaller firms have pre-spike total employment (year -1) below the median (13.1 FTE). *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

Figure D.2: Machinery per Worker: Larger versus Smaller Firms



Notes: This figure shows the machinery per worker around investment spikes. Smaller firms have pre-spike total employment (year -1) below the median (13.1 FTE). The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.

Figure D.3: Share of Machine Suppliers in Total Purchases: Smaller versus Larger Firms



Notes: This figure shows the share of machine suppliers in total purchases around the first machinery investment spike, separately for smaller and larger firms. Smaller firms have pre-spike total employment (year -1) below the median (13.1 FTE). *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

D.2 Application Sector

Tables D.3 and D.4 report pre-spike employment for the manufacturing and services subsamples. Spiking and non-spiking firms are well balanced on total employment and skill shares within each subsample, supporting the difference-in-differences design for both groups.

Comparing across subsamples, manufacturing firms are both larger (63.5 FTE in-house, versus 26.1 FTE in services) and employ a higher share of high-STEM workers in-house already before the spike (6.1% versus 4.3% in services). While both sectors display external high-STEM bias (Figure 5), the relative magnitudes are larger for services firms, consistent with manufacturing firms already having absorbed part of the fixed cost of in-house high-STEM hires before the investment spike. This pattern parallels the firm-size comparison (Appendix D.1).

Figure D.4 extends the analysis to the remaining private-sector industries (agriculture, utilities, and construction) and shows that the skill-biased external pattern from the main text holds in these sectors as well.

Table D.3: Manufacturing: Firm Employment before Machinery Investment Spikes

	Spiking		Non-Spiking		Std. Difference	
	In-House	External	In-House	External	In-House	External
	(1)	(2)	(3)	(4)	(5)	(6)
Total employment	63.5	27.3	59.4	23.3	0.02	0.05
	(215.6)	(95.8)	(148.6)	(81.0)		
Share of employment (%)						
— High STEM	6.1	6.5	5.0	5.9	0.12	0.12
	(9.5)	(5.5)	(8.4)	(4.4)		
— High non-STEM	10.6	14.9	11.0	15.1	-0.03	-0.02
	(13.3)	(8.1)	(13.1)	(8.4)		
— Low STEM	32.8	26.3	31.3	26.0	0.07	0.03
	(21.0)	(9.2)	(19.6)	(9.8)		
— Low Non-STEM	50.4	52.3	52.4	53.0	-0.09	-0.07
	(22.1)	(10.7)	(21.0)	(10.7)		
Nr. Firms	1,226		1,226			


Notes: This table reports in-house and external employment in the treatment (Columns 1–2) and control (Columns 3–4) groups in year $t - 2$. The table reports total employment in full-time equivalents and the employment share of different skill groups. Columns 5–6 show the standardized difference between the treatment and control group means, i.e. the difference in means divided by $\sqrt{(s_T^2 + s_C^2)}/2$, where absolute values above 25% are a standard threshold for assessing imbalance (Stuart and Rubin, 2008). *Sample:* Manufacturing sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

Table D.4: Services: Firm Employment before Machinery Investment Spikes

	Spiking		Non-Spiking		Std. Difference	
	In-House	External	In-House	External	In-House	External
	(1)	(2)	(3)	(4)	(5)	(6)
Total employment	26.1 (388.7)	13.6 (82.2)	37.8 (425.1)	17.8 (212.5)	-0.03	-0.03
Share of employment (%)						
— High STEM	4.3 (10.9)	5.6 (6.3)	4.9 (11.9)	5.1 (6.1)	-0.05	0.08
— High non-STEM	18.0 (21.9)	18.4 (11.6)	18.9 (21.2)	19.7 (12.4)	-0.04	-0.11
— Low STEM	19.7 (19.6)	22.1 (10.4)	18.2 (18.2)	21.1 (13.5)	0.08	0.08
— Low Non-STEM	58.0 (25.4)	54.0 (12.7)	58.0 (24.5)	54.1 (12.8)	0.00	-0.01
Nr. Firms	2,011		2,011			

Notes: This table reports in-house and external employment in the treatment (Columns 1–2) and control (Columns 3–4) groups in year $t - 2$. The table reports total employment in full-time equivalents and the employment share of different skill groups. Columns 5–6 show the standardized difference between the treatment and control group means, i.e. the difference in means divided by $\sqrt{(s_T^2 + s_C^2)}/2$, where absolute values above 25% are a standard threshold for assessing imbalance (Stuart and Rubin, 2008). *Sample:* Service sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

Figure D.4: Employment around Investment Spikes: Agriculture/Utilities/Construction



Notes: This figure shows the in-house and external employment around machinery investment spikes for firms in agriculture, utilities, and construction. The figures show the employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. *Sample:* Firms in agriculture, utilities, and construction around their first machinery investment spike. Balanced sample for four years before and after the spike.

D.3 Capital Types


Figure D.5 plots the capital stock per worker around each type of spike, mirroring Figure 1 for machinery. Each spike captures a discrete jump in the corresponding capital stock. The figure uses the 85th-percentile cutoff for defining spikes.

Figures D.6 and D.7 address the concern that building spikes coincide with machinery installation. Figure D.6 confirms that machinery stock per worker does rise around building spikes, so part of the external skill bias from buildings in Figure 6 likely reflects simultaneous machinery investment. Figure D.7, which excludes building spikes that coincide with the top 25% of machinery-per-worker increases, shows that the external skill bias narrows but does not disappear when buildings do not coincide with large machinery investment, plausibly reflecting the involvement of architects and building engineers.

Figure D.8 extends the building analysis to the full private sector at the 85th-percentile cutoff (€137,000), which yields enough spikes to cover the full private sector. The qual-


itative pattern from Figure 6 carries through: building spikes generate a less persistent external skill bias and a limited in-house response.

Figure D.5: Capital Stocks per Worker around Investment Spikes




Notes: Panels (a)–(c) plot the capital stock per worker around, respectively, machinery, building, and intangible investment spikes, mirroring Figure 1 for machinery. Machinery spikes are defined as years in which the firm’s change in machinery stock per worker exceeds the 85th percentile of positive changes (€11,000). Building spikes are defined analogously using the change in building stock per worker (85th percentile: €137,000). Intangible spikes are defined analogously using the change in intangible capital per worker (85th percentile: €14,000). *Sample:* Private-sector firms around their first investment spike, using a balanced panel spanning four years before and after the event.

Figure D.6: Machinery Stock per Worker around Building Investment Spikes




Notes: The figure plots the machinery capital stock per worker around years of building investment spikes. Building spikes are defined as years when firms' change in building stock per worker exceeds the 90th percentile (€220,195). *Sample:* Private sector firms around their first building investment spike. Balanced sample for four years before and after the spike.

Figure D.7: External Employment around Building Spikes (Excluding Machine Spikes)



Notes: This figure shows the external employment of manufacturing sector firms around investment spikes in buildings, excluding building spikes that coincide with the top 25% of machinery-per-worker increases. Building spikes are defined at the 90th percentile of positive changes in the building stock per worker. The figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. *Sample:* Manufacturing sector firms around their first building investment spike. Balanced sample for four years before and after the spike.

Figure D.8: Building Spikes: 85th-Percentile Cutoff




Notes: This figure shows the in-house and external employment of private-sector firms around investment spikes in buildings, defined at the 85th percentile of positive changes in the building stock per worker (€137,000). The figure shows employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). Total pre-spike employment is the sum of in-house and external employment, averaged over years -3 to -1 . The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms”. *Sample:* Private sector firms around their first building investment spike. Balanced sample for four years before and after the spike.

D.4 New versus Existing Machine Supply

Figure D.9 shows the in-house employment estimates that correspond to Figure 7 from the main text. The in-house response is similar across spikes driven by new versus existing suppliers, in line with the limited in-house effects of machinery spikes overall (Section 3.2).

Figure D.10 decomposes the external employment response by supplier industry. The differential high-STEM bias for new-supplier spikes shows up both at machine suppliers (Panel (i)) and at firms in other industries (Panel (ii)), consistent with new machinery embodying new technologies that demand specialized skills throughout the supply chain.


Figure D.9: In-House Employment around Machinery Spikes: New versus Existing Machine Supply




Notes: This figure shows in-house employment around machinery investment spikes, in percent of total pre-spike employment by skill group (Equation 7). Panel (a) focuses on spikes driven by new suppliers; Panel (b) focuses on spikes driven by existing suppliers. The classification is based on the ratio of the change in purchases from new machine suppliers to the sum of the absolute changes in purchases from new and existing machine suppliers at the time of the spike, $S_{\text{new}} = \Delta_{\text{new}} / (|\Delta_{\text{new}}| + |\Delta_{\text{exist}}|)$. Spikes with S_{new} above the sample median (excluding firms with non-positive changes in total purchases from machine suppliers) are classified as *new-supplier* spikes; the remainder are classified as *existing-supplier* spikes. The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.

Figure D.10: External Employment around Machinery Spikes: By Supplier Industry

(i) Machine Suppliers



(ii) Other Industries



Notes: This figure decomposes the external employment response from Figure 7 (main text) by supplier industry: machine suppliers (Panel (i)) and other industries (Panel (ii)). Subpanels (a) focus on spikes driven by new suppliers; Subpanels (b) focus on spikes driven by existing suppliers. The estimates are differences-in-differences in outcomes (measured relative to year -2) between the “spiking” and “non-spiking firms” from Table 2. *Sample:* Private sector firms around their first machinery investment spikes. Balanced sample for four years before and after the spike.

D.5 Supply-Side Evidence

Figure D.11 documents that each investing firm’s demand for capital-complementary high-STEM labor is small in absolute terms, supporting the fixed-cost rationale for outsourcing. Even summing across all of its suppliers, the median spiking firm sources only 0.05 high-STEM full-time equivalents (FTEs) in the spike year, and even the 90th percentile reaches only 0.77, both below one full-time worker, a natural minimum scale for an in-house


integration role.

Figure D.12 documents that machine suppliers absorb the lumpy demand of individual investing firms by spreading their workforce across many imperfectly correlated clients, supporting the adjustment-cost rationale. The figure plots machine suppliers' absolute year-over-year changes in own high-STEM employment against the sum of their client-level absolute changes in the same quantity. Both quantities are scaled by the midpoint of the supplier's own high-STEM employment in the two adjacent years,⁴⁸ such that the x-axis runs between 0 and 2 by construction. If machine suppliers' workforce volatility were a one-for-one reflection of their clients' demand changes, the binned means would lie on the 45-degree line. Instead, gross client-level changes are 15.4 times the supplier-side changes, because one client's expansion absorbs another's contraction. The figure also shows non-machine suppliers serving spiking firms as a benchmark.

Table D.5 reports estimates of the relationship between machine suppliers' skill mix and the share of their revenue accruing from spiking-firm clients. Across all four skill groups, the estimated coefficients are economically and statistically close to zero: machine suppliers' skill mix does not vary with their customer composition, consistent with a structurally specialized workforce. Machine suppliers also rely more on spiking-firm clients: they derive 1.76% of their revenue from these firms, more than double the 0.69% at non-machine suppliers.


⁴⁸Davis and Haltiwanger (1992) use a similar midpoint normalization in measuring job creation and destruction.

Figure D.11: External High-STEM Labor per Spiking Firm



Notes: This figure shows the distribution of external high-STEM employment at spiking firms in the spike year, summed across all of the firm's suppliers (not only its machine suppliers). The vertical reference line marks one full-time equivalent (FTE), an arguable minimum efficient scale for an in-house integration role. The median is 0.05 FTEs and the 90th percentile is 0.77 FTEs, both well below one FTE. *Sample:* Spiking firms; year of investment spike.

Figure D.12: Workforce Smoothing at Machine Suppliers



Notes: This binned scatter plots, at the supplier-year level, the supplier's absolute change in high-STEM employment from year $t - 1$ to year t (x-axis, normalized by the midpoint of the two years) against the sum across the supplier's clients of absolute year-on-year changes in their high-STEM labor allocations (y-axis, with the same supplier-midpoint denominator). Bins have fixed width, and marker size is proportional to total high-STEM machine-supplier employment ($\bar{S}$) in the bin. The 45-degree reference line corresponds to perfectly correlated client demand changes; points above 45 degrees indicate smoothing. The weighted-average ratio of the y- to x-axis quantities (weighting bins by $\bar{S}$) is 15.4 for machine suppliers. Non-machine suppliers serving spiking firms are shown as a benchmark. *Sample:* Belgian machine and non-machine suppliers, 2003–2020.

Table D.5: Skill Mix Stability across Customer Mixes (Machine Suppliers)


	High STEM (1)	High Non-STEM (2)	Low STEM (3)	Low Non-STEM (4)
Spiking-Client Revenue Share (%)	.0018 (.0042)	−.0012 (.0067)	−.0012 (.0048)	.0006 (.0063)
Mean of dep. var. (%)	13.3	22.8	23.0	40.9
Supplier FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	257,332	257,332	257,332	257,332

Notes: This table reports estimates from supplier-year-level regressions of machine suppliers' skill-group employment shares on the share of their revenue accruing from spiking-firm clients. The spiking sample includes all spikes for which total employment was larger than 5 FTEs before the first spike. Each column corresponds to a separate regression with a different dependent variable (employment share in the indicated skill group, in percent). The regressor of interest is the share of the supplier's revenue coming from spiking-firm clients in that year. All specifications include supplier and year fixed effects. Standard errors clustered at the supplier level are shown in parentheses. On average, machine suppliers derive 1.76% of their revenue from spiking-firm clients, compared to 0.69% at non-machine suppliers. *Sample:* Belgian machine supplier firms, 2003–2021.

D.6 Wage-Savings Motive

Figure D.13 shows in-house and external employment by skill group around machinery investment spikes, split by whether the investing firm’s residualized pre-spike wages are above or below the sample median. A direct wage-savings motive would predict that higher-wage firms outsource more, as they have more to save by sourcing capital-complementary labor externally. We find the opposite: lower-wage firms display, if anything, a *larger* external high-STEM skill bias. This pattern partly reflects that lower-wage firms are also smaller, which Section 4.1 documents are more likely to outsource their capital-complementary skills. Either way, the evidence weighs against direct wage savings as the primary driver of the outsourcing we document.

Figure D.13: Employment by Pre-Spike Investing-Firm Wage
(Diff-in-Diffs, In Percent of Pre-Spike Total Employment by Skill Group)



Notes: This figure shows in-house and external employment by skill group around machinery investment spikes, split by whether the investing firm has residualized pre-spike wages above or below the sample median. Firms are split by residualized pre-spike wages: firm-level log wages are residualized on skill shares and year fixed effects (estimated on the unbalanced sample of spiking firms); the residual in the year before the spike is then compared to the median residualized wage in the balanced sample. The figures show employment of high-STEM, high non-STEM, low STEM, and low non-STEM workers in percent of total pre-spike employment by skill group (Equation 7). *Sample:* Private sector firms around their first machinery investment spike. Balanced sample for four years before and after the spike.

Table D.6 turns to the second part of the wage-savings analysis: a direct comparison of the wages paid by machine suppliers and by investing firms in the year before the spike. We regress log average firm wages on a machine-supplier indicator. Column (1) controls only for year fixed effects: machine suppliers pay 13.8% higher wages than investing firms. Column (2) adds controls for the firm's skill shares, isolating the wage gap for a comparable workforce composition: machine suppliers still pay about 3.6% higher wages. This premium is the opposite of what direct wage savings would predict; it is consistent

with efficiency gains from outsourcing to specialized suppliers and with the internal-wage-structure motive discussed in Section 4.6.

Table D.6: Wage Premia of Machine Suppliers vs. Investing Firms

	<i>Dependent variable: log average wage</i>	
	(1)	(2)
Machine supplier	.138*** (.008)	.036*** (.007)
Year fixed effects	✓	✓
Skill share controls		✓
Observations	86,196	86,196

Notes: Sample consists of machine suppliers and firms that experience an investment spike, observed in the year preceding the spike. The sample is restricted to firms with more than five FTE employees. Standard errors clustered at the firm level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

E Aggregate Implications

This appendix lays out the model behind Section 6. Section E.1 states the model, Section E.2 solves it, and Section E.3 estimates it. Section E.4 derives the sufficient statistics behind Equation (9) and decomposes them into firm-level elasticities, Section E.5 measures them, from the event study and in the estimated model, and Section E.6 shows how the results change with the elasticities we set outside the model. Finally, Section E.7 extends the model to the make-or-buy margin.

E.1 The Quantitative Model

The model is the firm's dynamic program of Appendix A (Equation (10)) with functional forms, firm heterogeneity, and a general-equilibrium closure. As in Appendix A.1, its baseline abstracts from the make-or-buy decision, so S_P , U_P , S_I and U_I count in-house and external workers alike; Appendix E.7 endogenizes that decision, letting firms hire these workers in-house or buy them from suppliers.

Technology. Production and integration are both nested CES with constant returns. Firm j produces output with machinery K , high-STEM production labor S_P and low non-STEM production labor U_P :

$$H_j = \left[c_j^K K_j^\rho + c_j^S \left(\theta_j^S S_{Pj} \right)^\rho \right]^{1/\rho}, \quad F_j = \Omega \left[c^H H_j^\lambda + c^U \left(\theta_j^U U_{Pj} \right)^\lambda \right]^{1/\lambda}. \quad (29)$$

This specializes the production function F of Equation (10) to the nested CES of Krusell et al. (2000), as in Appendix A.4, with the factor-augmenting productivities θ of Appendix A.3. Production abstracts from other inputs, such as other capital, materials, high non-STEM labor and low STEM labor. Our conclusions about skill bias carry over if F_j is a homothetic aggregate that is weakly separable from these inputs: they then scale the demands for K , S_P and U_P proportionally, as output demand does in Appendix A.3,

and leave relative skill demand unchanged. Krusell et al. (2000) treat structures in the same way: they enter as a Cobb–Douglas factor multiplying the nested CES, so the skill premium does not depend on them.

Installed capital, $I_j = G_j$, is produced from hardware X and the integration labor of both skill groups, S_I and U_I :

$$J_j = \left[c^X X_j^\phi + c^{S_I} S_{Ij}^\phi \right]^{1/\phi}, \quad G_j = \left[c^J J_j^\xi + c^{U_I} U_{Ij}^\xi \right]^{1/\xi}. \quad (30)$$

This specializes the integration technology G of Equation (10). Appendix A.3 lets the productivities θ augment integration labor as well; the quantitative model sets $\theta = 1$ in integration, the technology of specialized suppliers, so the price of installed capital is common to firms that face the same hardware price. The elasticities of substitution are $\sigma_{KS} = 1/(1 - \rho)$, $\sigma_{HU} = 1/(1 - \lambda)$, $\sigma_{XS} = 1/(1 - \phi)$ and $\sigma_{JU} = 1/(1 - \xi)$. Capital–skill complementarity is $\sigma_{HU} > \sigma_{KS}$; investment–skill complementarity is $\sigma_{JU} > \sigma_{XS}$. The integration technology generalizes Goldin and Katz (1998), in which hardware and skilled labor are perfect complements ($\sigma_{XS} = 0$) and low-skilled labor does not enter integration ($c^{U_I} = 0$). The coefficients c are normalizations that give each input its calibrated cost share at the base prices (Table E.1), and Ω scales output so that the base input bundle produces one unit.

Demand and Shocks. Each firm sells a differentiated variety with demand elasticity ε^Y . Revenue is $R_j = \tilde{R} A_j F_j^{1-1/\varepsilon^Y}$, where A_j combines the demand shifter D_j of Appendix A with Hicks-neutral productivity, and the industry shifter $\tilde{R}$ clears industry demand with elasticity η : $\tilde{R} = \mathcal{A}^{1/\eta} Y^{1/\varepsilon^Y - 1/\eta}$, with $Y = \left[\int A_j F_j^{1-1/\varepsilon^Y} d\mu \right]^{\varepsilon^Y/(\varepsilon^Y - 1)}$. The neutral index combines a permanent neutral productivity type and a persistent shock, $\ln A_j = \ln \bar{A}_j + a_j$, where $\ln \bar{A}_j \sim N(0, \sigma_A^2)$ and $a' = \rho_A a + u$ with $u \sim N(0, \sigma_A^2)$. Firms also differ permanently in the machinery share of H , whose logit is normal around the

logit of s_K^H with standard deviation σ_{s_K} . Each firm pays its own hardware price $p^X \zeta_j$, where $\ln \zeta_j \sim N(-\sigma_z^2/2, \sigma_z^2)$, and the firm redraws ζ_j from this distribution with probability p_z each year. The model also allows for firm-level skill-biased productivity shocks, an AR(1) in $\ln(\theta_j^S/\theta_j^U)$ with persistence ρ_b , innovation standard deviation σ_b , and correlation ρ_{bA} with u . All draws are independent across firms. The skill-bias shock is Markov, so Assumption 1 holds conditional on a firm's state in the year it decides to invest (Appendix A.3). We set $\sigma_b = 0$ in the baseline.

The Firm's Problem. Firms take as given the wages w_S and w_U , the industry shifter $\tilde{R}$, and the hardware price p^X , all constant in a stationary equilibrium. Because Equation (30) has constant returns, installed capital has a unit cost $p^K = p^K(p^X \zeta_j, w_S, w_U)$, and integration labor per unit of installed capital depends only on prices. The frictionless user cost of capital, $r = (i + \delta) p^K$, is the price at which the static approximation of Appendix E.4.2 lets firms trade capital services against labor; the firm's problem below has no rental price, since the firm weighs the cost of installing capital today against the profits it will earn later. Capital is fixed within a year and production labor is chosen each year, so profit given capital is

$$\Pi(K, e) = \max_{S_P, U_P} \tilde{R} A_j F_j^{1-1/\varepsilon^Y} - w_S S_P - w_U U_P, \quad (31)$$

where e collects the firm's exogenous states. Capital takes a year to build and investment is irreversible. The convex adjustment cost is denominated in installed capital, so every cost scales with prices; like the fixed cost, it is a pure cost that adds neither to the stock nor to the demand for hardware and integration labor, which scale with gross investment I alone. A firm that invests more than a maintenance allowance mK also pays that fixed

cost, a random share ν of its revenue R . The value of a firm is

$$V(K, e) = \Pi(K, e) + \mathbb{E}_\nu \max \left\{ \max_{0 \leq I \leq mK} [\beta \mathbb{E}[V(K', e') | e] - p^K C(I, K)], \right. \\ \left. \max_{I \geq mK} [\beta \mathbb{E}[V(K', e') | e] - p^K C(I, K)] - \nu R \right\}, \quad (32)$$

where $K' = (1 - \delta)K + I$, $C(I, K) = I + \frac{\gamma}{2}(I/K)^2 K$, $\nu \sim U[0, \bar{\nu}]$ drawn each year, and $\beta = 1/(1+i)$. The first branch invests within the maintenance allowance; the second pays the fixed cost and moves capital to its target. The fixed cost makes investment lumpy: firms let a gap open between their capital and its target, and close it in an episode. Firms invest only on low draws: the fixed cost they pay averages 11 percent of annual revenue per episode. Because G has constant returns, the cost of installed capital bundles the hardware and integration-labor costs that Appendix A.3 writes separately, and the fixed cost νR is the fixed-cost draw of Equation (19).

Definition 1 (Stationary equilibrium). *Given the hardware price p^X , a stationary equilibrium consists of wages w_S and w_U , an industry shifter $\tilde{R}$, a value function with investment and labor policies, and a distribution μ of firms over capital and exogenous states, such that:*

- (i) *the value function and the investment policy solve Equation (32), production labor solves Equation (31), and integration labor is gross investment I times the requirement per unit of installed capital implied by Equation (30);*
- (ii) *μ is invariant under the investment policy and the processes for the exogenous states;*
- (iii) *the industry shifter is consistent with aggregate output, $\tilde{R} = \mathcal{A}^{1/\eta} Y^{1/\varepsilon^Y - 1/\eta}$, with Y aggregated over μ ;*
- (iv) *both skill markets clear:*

$$\int [S_P + S_I] d\mu = B_S w_S^{\varepsilon_S}, \quad \int [U_P + U_I] d\mu = B_U w_U^{\varepsilon_S}. \quad (33)$$

Labor supply to each skill group is iso-elastic in its own wage, with the same elasticity ε^S for both groups. The supply shifters B_S and B_U are set so that the initial wages clear both markets, and they do not change afterwards. Market clearing counts production and integration labor alike, and in-house and external workers alike. Wages are thus endogenous, whereas the price of hardware is the exogenous driving force in the model.⁴⁹


The Counterfactual. We compare two stationary equilibria that differ only in the hardware price.⁵⁰ As in Krusell et al. (2000), a falling price of equipment drives the rise in the skill premium. The price decline is calibrated to deliver the 108.0 log point rise in machinery per worker in the Belgian private sector, from €19,498 in 2003 to €57,441 in 2021 (Figure E.1).⁵¹ This requires the quality-adjusted hardware price to fall to 0.3027 of its initial level, because integration labor, which does not become cheaper, accounts for two thirds of the cost of a machine (Appendix B.2.1).

⁴⁹A natural assumption for Belgium, a small open economy facing world machinery prices: industrial robots, the canonical case in our setting, are not produced domestically and are entirely imported (Appendix B.2.2).

⁵⁰Comparing stationary equilibria is a standard approach to evaluating the equilibrium impact of new technologies in firm-dynamics models (see, e.g., Hubmer and Restrepo, 2026; Rubinton, 2026).

⁵¹The aggregate series runs one year beyond our 2003–2020 estimation sample.

Figure E.1: Aggregate Machinery Per Worker



Notes: This figure plots the aggregate machinery stock per worker in the total private sector and in manufacturing. Machinery capital is deflated using the equipment price deflator from Federal Reserve Bank of St. Louis (2025).

E.2 Solving the Model

Solving the Firm's Problem. We discretize capital on a log-spaced grid of 240 points, from 0.05 times the smallest to 10 times the largest frictionless capital target. The exogenous states take 225 combinations: the persistent productivity shock takes five values by the method of Rouwenhorst (1995), and the neutral productivity type, the machinery-share type, and the hardware price take three, three, and five equal-probability nodes of their (log or logit) normal distributions. Given capital, labor solves the static first-order conditions. We solve Equation (32) by value function iteration, accelerated with Howard's improvement step (Howard, 1960). The continuation value is interpolated linearly in capital, and in each iteration the investment target maximizes that interpolated objective exactly; the expectation over the fixed-cost draw has a closed form. Because the contraction property of the exact Bellman operator (Stokey and Lucas, 1989) does not carry over automatically to this discretization, we verify the solution directly: the Bellman

residual, the optimality of the policy against a finer search, concavity, homogeneity in prices, stability of the aggregates when the grid is doubled, and that virtually no firms reach the top of the capital grid. We find the stationary distribution of firms by iterating the solved policies forward until the distribution no longer changes, assigning each off-grid capital choice to its two neighboring grid points (Young, 2010).

Solving for Equilibrium. Heterogeneous firms interact only through prices: each firm’s problem depends on the rest of the economy only through the wages w_S and w_U and the industry shifter $\tilde{R}$. A stationary equilibrium is therefore a fixed point in these three prices, which we find as follows:

1. Guess w_S , w_U and $\tilde{R}$.
2. Solve Equation (32) at these prices and find the stationary distribution, as described above.
3. Aggregate over the distribution and compute the log deviations from the industry-demand and skill-market conditions of Definition 1.
4. Update the three prices with a quasi-Newton step and return to step 2, until the largest log deviation is below 10^{-7} .

For the counterfactual we add the hardware price as a fourth unknown and the condition that machinery per worker rises by 108.0 log points.

E.3 Estimation

We estimate the model in three steps: we set the parameters that the event studies cannot inform, run our investment-spike design on simulated firms, and choose the remaining 12 parameters so that the simulated event study matches the estimated one. Section E.3.2 then assesses the fit, reads the estimates, and compares the implied elasticity ratio with aggregate estimates.

E.3.1 Method

Parameters Set Outside the Model. The two production cost shares come from the average spiking firm before the spike (Table 2). Machinery services, valued at a user cost of $i + \delta = 0.17$ on pre-spike machinery, and a high-STEM wage bill at a fully loaded cost of €75,000 per full-time equivalent (Eurostat, 2026*a,b*) give s_K^H ; the same two flows against total cost give s_H (Table E.1). The remaining parameters are listed in Table E.1: σ_{HU} takes its macro-labor value (Krusell et al., 2000; Katz and Murphy, 1992), ε^S is set equal to it as a symmetric baseline and varied in Table E.3, δ and i are conventional, and χ_X follows the rule of thumb of Appendix B.2.1. σ_{XS} and σ_{JU} are set rather than estimated, because the event study does not sharply distinguish them,⁵² and alternative values move the production ratio only between 1.07 and 1.20 (Table E.5).

The Simulated Event Study. We simulate 20,000 firms for 60 years from the stationary distribution and apply our investment-spikes design to them. We apply the spike cutoff in euros rather than as a percentile, because in the model only investment episodes raise the stock. Machinery per worker divides by all labor a firm uses, in-house and external, and we convert the data’s machinery moments and cutoff to that basis. As in the data, each firm’s first spike is matched one-to-one to a not-yet-spiking firm with at least five workers, and we estimate Equation (6) with the outcomes, base year, and winsorization of Section 3.1; firm size is scaled so that the median spiking firm employs 13.1 workers. Because the baseline model has no make-or-buy decision, we target each skill group’s total employment response, with the dynamics informing the split between integration and production labor.

⁵²The spike-year responses are informative about integration, but they cannot separate its substitution elasticities from the skill share of integration costs and the dispersion of hardware prices.

Moments and Objective. We estimate the model by indirect inference (Gourieroux, Monfort and Renault, 1993), targeting 26 moments: the responses of machinery and of total (in-house plus external) high-STEM and low non-STEM employment at event years -4 , -3 and -1 to 4 , the pre-spike level of machinery per worker at spiking firms, and the share of spiking firms that spike again within four years (Appendix C.3.5). The 12 estimated parameters minimize

$$Q(\vartheta) = \frac{\sum_k w_k \left[(\hat{m}_k(\vartheta) - m_k) / \text{se}_k \right]^2}{\sum_k w_k}, \quad (34)$$

where m_k is an estimated moment, se_k its standard error, and $\hat{m}_k(\vartheta)$ its simulated counterpart. The weights w_k are one for the event-study responses and pre-spike machinery, and 0.1 for the follow-on share, which describes the investment distribution rather than the response to an episode. We minimize Q with the Nelder–Mead simplex method (Nelder and Mead, 1965), searching on a smaller panel of 8,000 firms over 50 years with fixed random draws; all reported results use the 20,000-firm panel above, on which $Q = 4.8$ at the estimated parameters.

Standard Errors. Standard errors come from the usual formula for indirect inference,

$$\Sigma_{\vartheta} = (G'WG)^{-1}G'W(\Sigma_m + \Sigma_{\text{sim}})WG(G'WG)^{-1}, \quad (35)$$

where G is the Jacobian of the simulated moments with respect to the parameters, W the weight matrix of Equation (34), Σ_m the sampling variance of the event-study estimates, and Σ_{sim} the covariance of the simulated moments across 64 random draws of the panel used in the search. We take Σ_m as diagonal, since the estimates come with one standard error per coefficient and no covariances. Because simulated moments are not smooth in the parameters, we compute G from finite differences over a step of ± 4 standard errors rather than as a local derivative, verified not to bend over that range. We report standard

errors conditional on six parameters that our moments inform only weakly, treating them as known: five that end near their search bounds ($\bar{\nu}$, γ , ρ_A , σ_A , m) and $\sigma_{\bar{A}}$, which changes firm size but not skill mix. This also puts the remaining estimates at an interior optimum, as the formula requires. By the delta method, with the hardware-price decline re-solved to the observed rise in machinery at each perturbed parameter, the rise in the skill premium in Table 3 has a standard error of 0.57 log points on the same basis.

E.3.2 Results

Model Fit. Figure E.2 compares the simulated and estimated event studies for machinery per worker and for employment by skill group. The model fits the data well, with two exceptions that both relate to time to build. The model installs new machinery within a single year, whereas in the data an investment episode is more spread out. First, some firms engage integrators in the year before the machinery arrives, raising the demand for skilled workers already in year -1 , which the model cannot generate. Second, some firms install over several years, which registers as a follow-on spike, most often in the year after the event (Figure C.12); the model produces fewer of these, 0.01 against 0.28. The model’s spiking firms also start from less machinery per worker than in the data, €16,967 against €19,086.

Parameter Estimates and Related Moments. Table E.1 lists the parameters and the moment most closely associated with each. The 12 estimated parameters fall into four groups, each tied to a feature of the event study.

1. *The skill bias of machinery.* The model separates the two stages by their timing: integration labor rises in a burst around installation, production labor with the stock.⁵³ At fixed wages, the firm’s skill mix in production satisfies $\Delta \ln S_P -$

⁵³Applied to the simulated event study with the capital stock and the cost share at the firm’s optimum, Equation (27) gives 1.37 with total employment, which still contains integration labor four years after

$\Delta \ln U_P = (\sigma_{HU}/\sigma_{KS} - 1) s_K^H (\Delta \ln K - \Delta \ln S_P)$, which contains neither the neutral shock nor the hardware price. The high-STEM skill differential per unit of machinery deepening from year 1 on, once the new machinery is in use, therefore pins down σ_{KS} given the cost share and σ_{HU} : with σ_{HU} at its macro-labor value, σ_{KS} is 1.306 (s.e. 0.058), a production ratio of 1.15. More precisely, the differential identifies the product $(\sigma_{HU}/\sigma_{KS} - 1) s_K^H$, which is 0.054 at the estimate, so the ratio is conditional on the cost share: letting the other parameters re-adjust, the ratio at $s_K^H = 0.25$ or 0.45 is 1.19 or 1.12. With time to build, the spike-year responses are integration labor, and their skill differential pins down χ_S , the skilled integrators' share of the cost of installed capital.


2. *The lumpiness of investment.* The size of the machinery response at the spike informs the fixed-cost draw $\bar{\nu}$, its timing the convex cost γ , and its decay the maintenance allowance m .
3. *Selection into spiking.* The pre-spike machinery level of spiking firms, relative to the economy's, together with the size of the spike, pins down the dispersion of machinery intensity σ_{s_K} and of hardware prices σ_z : machinery-intensive firms spike from a higher stock, firms whose hardware becomes cheap from a lower one. The probability of a new hardware-price draw, p_z , is tied to the follow-on spike share and is the least precisely estimated.
4. *Firm scale and neutral shocks.* Employment growth from year 1 on pins down the demand elasticity ε^Y given σ_{HU} , and the flat employment pre-trends pin down the dispersion of the neutral shock σ_A : with time to build and Markov productivity, they imply that investment spikes are mostly triggered by hardware prices and depreciation rather than by demand shocks. Its persistence ρ_A and the dispersion

a spike, and 1.24 with production labor alone, the remaining gap reflecting dynamics the static formula ignores.

of the permanent size type $\sigma_{\bar{A}}$ are not informed by scale-free moments, and the standard errors below are conditional on them.

Aggregate Capital–Skill Complementarity. To compare with aggregate estimates, we fit the aggregate nested CES of Krusell et al. (2000) to the model economy. Lowering the hardware price at fixed wages and applying Equation (27) to aggregate machinery, employment, and cost shares gives 1.21 when integration workers count as production labor, still well below their 2.5.

Figure E.2: The Simulated Event Study against the Estimated One



Notes: The left panel plots the difference-in-differences in machinery per worker (in euros, per worker in-house and external) and the right panel in high-STEM and low non-STEM employment (in percent of pre-spike employment of each skill group), relative to year -2 . *Model* lines estimate Equation (6) on the simulated panel of the estimated model; *data* lines are the targeted event-study estimates.

Table E.1: Parameters

Description		Value	Source or moment
<i>Measured outside the model</i>			
δ	Depreciation rate of machinery	0.120	BEA
i	Interest rate ($\beta = 1/(1+i)$)	0.050	Conventional
s_K^H	Machinery's share of the high-STEM nest H	0.360	Table 2; Appendix E.3
s_H	Share of H in production cost	0.115	Table 2; Appendix E.3
χ_X	Hardware share of installed-capital cost	0.333	Installation rule of thumb
σ_{HU}	Elasticity, H vs. low non-STEM labor	1.500	Katz and Murphy (1992); Krusell et al. (2000)
<i>Chosen by indirect inference on the event study</i>			
σ_{KS}	Elasticity, machinery vs. high-STEM labor	1.306	Skill differential, years 1–4
χ_S	High-STEM integrators' share of capital cost	0.055	Skill differential in year 0
$\bar{\nu}^\dagger$	Maximum fixed-cost draw (share of revenue)	4.943	Machinery response at the spike
$\gamma^\dagger$	Convex adjustment cost	0.026	Timing of the machinery response
σ_z	S.d. of log hardware-price multiplier	1.801	Pre-spike machinery level; response at the spike
p_z	Annual probability of redrawing ζ_j	0.015	Follow-on spike share
$m^\dagger$	Fixed-cost-exempt maintenance (share of K)	0.104	Decay of the machinery response
$\rho_A^\dagger$	Persistence of the log neutral shock	0.984	Employment growth
$\sigma_A^\dagger$	Innovation s.d. of the log neutral shock	0.014	Employment pre-trends
$\sigma_{\bar{A}}$	S.d. of the permanent log neutral type	0.498	Employment growth
σ_{s_K}	S.d. of the logit of machinery's share of H	0.141	Pre-spike machinery level
ε^Y	Demand elasticity for a firm's variety	3.828	Employment growth, years 1–4
<i>Set a priori</i>			
ε^S	Elasticity of relative skill supply	1.500	Varied in Table E.3
η	Industry demand elasticity	1.000	No effect on relative demand
σ_{XS}	Elasticity, hardware vs. high-STEM integrators	0.900	Varied in Table E.5
σ_{JU}	Elasticity, J vs. low non-STEM integrators	1.500	Varied in Table E.5
ρ_b	Persistence of the skill-bias shock	0.900	Off in the baseline
σ_b	Innovation s.d. of the skill-bias shock	0.000	Off in the baseline
ρ_{bA}	Correlation of skill-bias and neutral innovations	0.000	Zero in the baseline

Notes: Parameters chosen by indirect inference are estimated jointly; the source listed is the moment most closely associated with each, and Appendix E.3 describes what each estimate is conditional on. $\dagger$ At a bound of the search. $\sigma_{\bar{A}}$ is not identified by scale-free moments.

E.4 Deriving the Sufficient Statistics

We compare two equilibria of the labor market, before and after the rise in machinery. Let

$L_S \equiv \int [S_P + S_I] d\mu$ and $L_U \equiv \int [U_P + U_I] d\mu$ denote aggregate high-STEM and low non-

STEM employment, $\omega \equiv \ln(w_S/w_U)$ the skill premium, and M machinery per worker. We

first derive Equation (9), then decompose its sufficient statistics into firm-level elasticities.

E.4.1 The Sufficient Statistics

Supply. Dividing the two supply schedules in Equation (33) and taking logs gives relative supply,

$$\ln \frac{L_S}{L_U} = \ln \frac{B_S}{B_U} + \varepsilon^S \omega, \quad (36)$$

which holds exactly, for a change of any size, because the supply shifters do not change between the two equilibria.

Demand. Every cost in the firm's problem scales with the wages and the hardware price, and because production has constant returns and every cost is proportional to a firm's size, the industry revenue shifter scales firms up or down without changing their skill mix or machinery per worker. Relative labor demand, the aggregate counterpart of Equation (17), therefore depends only on the skill premium ω and x , the log hardware price relative to the geometric mean of the two wages. Given ω , a lower hardware price raises machinery per worker, so over the range of our counterfactual we can replace x by $\ln M$ and write $\ln(L_S/L_U) = L(\omega, \ln M)$. This lets us define the counterfactual by the observed rise in machinery rather than by an unobserved price. Go from $(\omega_0, \ln M_0)$ to $(\omega_1, \ln M_1)$ in two steps, first raising machinery at the initial skill premium and then raising the premium at the final machinery:

$$\Delta L = \underbrace{\frac{L(\omega_0, \ln M_1) - L(\omega_0, \ln M_0)}{\Delta \ln M}}_{\Psi_M} \cdot \Delta \ln M - \underbrace{\frac{L(\omega_1, \ln M_1) - L(\omega_0, \ln M_1)}{-\Delta \omega}}_{\varepsilon_M^D} \cdot \Delta \omega. \quad (37)$$

Ψ_M is the shift in relative demand at the initial wages, the aggregate counterpart of what our event study measures: comparing investing with non-investing firms holds wages fixed, because wage changes common to both difference out. Because $\Delta \ln M$ is the observed

rise in machinery, which already contains the response of machinery to wages, ε_M^D holds machinery per worker fixed rather than letting it adjust.⁵⁴

Equilibrium. In equilibrium, the change in relative demand in Equation (37) equals the change in relative supply from Equation (36), $\varepsilon^S \Delta\omega$:

$$\Psi_M \cdot \Delta \ln M - \varepsilon_M^D \cdot \Delta\omega = \varepsilon^S \cdot \Delta\omega.$$


Solving for the premium, and using the supply schedule for employment,

$$\Delta\omega = \frac{\Psi_M \cdot \Delta \ln M}{\varepsilon^S + \varepsilon_M^D}, \quad \Delta \ln \frac{L_S}{L_U} = \varepsilon^S \cdot \Delta\omega,$$

which is Equation (9); Figure E.3 illustrates the adjustment. In the model, $\Psi_M = 0.0817$ and $\varepsilon_M^D = 1.379$ when both are measured as averages along the path between the two equilibria, which makes Equation (9) exact; the derivatives at the initial equilibrium are 0.0740 and 1.399. Section 6.3 reports the path average for Ψ_M , which is the shift in Table 3, and the initial derivative for ε_M^D , which Section E.4.2 decomposes.

⁵⁴Describing relative demand by (ω, x) instead gives an elasticity that lets machinery adjust, $\varepsilon_A^D \equiv -\partial L / \partial \omega|_x$. The local derivatives are linked by $\varepsilon_M^D = \varepsilon_A^D + \Psi_M \cdot \partial \ln M / \partial \omega|_x$, so pairing ε_A^D with the observed rise in machinery would count the response of machinery to wages twice.

Figure E.3: General-Equilibrium Adjustment in the Relative Labor Market



Notes: Illustration of the general-equilibrium adjustment in the market for high-STEM labor relative to low non-STEM labor. The shift in relative demand at the initial wages, $\Psi_M \cdot \Delta \ln M$, shifts the relative demand curve rightward from D_0 to D_1 . The new equilibrium E_1 implies a relative employment response smaller than the partial-equilibrium shift, with the residual absorbed by growth in the skill premium. The figure is drawn with equal slopes, close to our calibration ($\varepsilon^S = 1.5$, $\varepsilon_M^D = 1.399$), in which relative employment absorbs about half of the shift.

E.4.2 Decomposition into Firm-Level Elasticities

Firm-Level Responses. We now express Ψ_M and ε_M^D in terms of the model's parameters. The expressions are a static approximation: each firm responds as if it sat at its capital target, facing the current user cost, with its investment proportional to its capital. The expressions are also first order: cost shares and employment weights are held at their initial values, so they describe local responses, which Section E.4.1 averages along the path between the two equilibria when it evaluates finite changes. Hold w_U fixed and let $\hat{w} \equiv \Delta \ln w_S$ and $\hat{p} \equiv \Delta \ln p^X$. A firm has five cost shares: machinery's share of H , s_K^H ; the share of H in production cost, s_H ; the shares of hardware and of skilled integration

labor in the price of installed capital, χ_X and χ_S ; and the share of hardware in J , s_X^J . The firm's price of capital and unit cost change by

$$\hat{r} = \chi_X \hat{p} + \chi_S \hat{w}, \quad \hat{c} = s_H \left[s_K^H \hat{r} + (1 - s_K^H) \hat{w} \right].$$

Three firm-level elasticities summarize the technology:

$$\begin{aligned} \sigma^P &= \sigma_{KS} s_K^H + \sigma_{HU} (1 - s_K^H), & \kappa^P &= \sigma_{KS} (1 - s_K^H) + \sigma_{HU} s_K^H, \\ \sigma^I &= \sigma_{XS} s_X^J + \sigma_{JU} (1 - s_X^J). \end{aligned} \tag{38}$$

σ^P is the elasticity of substitution between high-STEM labor and the other production inputs, κ^P the one between machinery and the other production inputs, and σ^I the one between high-STEM labor and the other integration inputs. Cost minimization gives the firm's relative responses:

$$\begin{aligned} g^P &\equiv \Delta \ln(S_P/U_P) = -\sigma^P \hat{w} - (\sigma_{HU} - \sigma_{KS}) s_K^H \hat{r}, \\ g^K &\equiv \Delta \ln(K/U_P) = -\kappa^P \hat{r} - (\sigma_{HU} - \sigma_{KS}) (1 - s_K^H) \hat{w}, \\ g^I &\equiv \Delta \ln(a_S/a_U) = -\sigma^I \hat{w} - (\sigma_{JU} - \sigma_{XS}) s_X^J \hat{p}, \\ \Delta \ln a_U &= \sigma_{JU} \hat{r}, \end{aligned} \tag{39}$$

where a_S and a_U are integration labor per unit of investment. At fixed wages ($\hat{w} = 0$), the production responses reduce to the nested-CES relations (23)–(26) of Appendix A.4, with $\hat{r}$ the change in the user cost. The firm's scale changes its low non-STEM production employment by $-(\varepsilon^Y - \sigma_{HU}) \hat{c}$, plus a term common to all firms. Production and integration have the same structure: when the input that complements skill becomes cheaper, demand tilts toward skill in proportion to that input's cost share and to the gap between the two elasticities.

Aggregation. Let $\mathbb{E}_Z$ denote the average over firms weighted by each firm's share of aggregate Z . Let ι_s be the share of integration labor in the employment of skill group s ,

and n_S the high-STEM share of employment. Adding up the firms' responses gives the changes in aggregate relative employment and in machinery per worker:

$$\begin{aligned}
\Delta \ln \frac{L_S}{L_U} &= \underbrace{(1 - \iota_S) \mathbb{E}_{S_P}[g^P]}_{\text{within production}} + \underbrace{\iota_S \mathbb{E}_{S_I}[g^I]}_{\text{within integration}} + \underbrace{(\iota_S \mathbb{E}_{S_I} - \iota_U \mathbb{E}_{U_I}) [\Delta \ln a_U + g^K]}_{\text{across stages}} \\
&\quad - \underbrace{(\varepsilon^Y - \sigma_{HU}) (\mathbb{E}_S - \mathbb{E}_U) [\hat{c}]}_{\text{across firms}}, \tag{40} \\
\Delta \ln M &= \underbrace{\mathbb{E}_K[g^K]}_{\text{within firms}} - \underbrace{n_S(1 - \iota_S) \mathbb{E}_{S_P}[g^P]}_{\text{production skill mix}} \\
&\quad - \underbrace{\left\{ n_S \iota_S \mathbb{E}_{S_I}[g^I + \Delta \ln a_U + g^K] + (1 - n_S) \iota_U \mathbb{E}_{U_I}[\Delta \ln a_U + g^K] \right\}}_{\text{integration employment}} \\
&\quad - \underbrace{(\varepsilon^Y - \sigma_{HU}) (\mathbb{E}_K - \mathbb{E}_N) [\hat{c}]}_{\text{across firms}}. \tag{41}
\end{aligned}$$

The scale term common to all firms drops out because every set of weights sums to one. The across-stages term is the change in integration employment relative to production employment, $\Delta \ln a_U + g^K$: cheaper hardware raises capital per production worker by $\kappa^P \chi_X$ and lowers integration labor per unit of investment by $\sigma_{JU} \chi_X$, so its sign is that of $\kappa^P - \sigma_{JU}$. With σ_{JU} set equal to σ_{HU} , that difference is $-(\sigma_{HU} - \sigma_{KS})(1 - s_K^H)$, negative under capital-skill complementarity. That step uses aggregate flow balance firm by firm. Aggregate investment equals δK exactly in a stationary equilibrium, but when a firm's hardware price, and with it its integration requirements a_s , changes over time, integration labor is $\mathbb{E}[a_s I]$ rather than $\delta \mathbb{E}[a_s K]$; the gap belongs to the dynamic residual. The across-firm terms are those of Oberfield and Raval (2021): firms whose costs fall most expand, and this changes aggregate relative employment to the extent that these firms use one input more intensively than the other.

The Demand Shift and the Demand Elasticity. Ψ_M is the ratio of Equation (40) to Equation (41) when the hardware price falls with wages unchanged, $(\hat{w}, \hat{p}) = (0, -1)$:

$$\Psi_M = \mathcal{G}/\mathcal{H},$$

$$\begin{aligned}\mathcal{G} &= (\sigma_{HU} - \sigma_{KS})(1 - \iota_S) \mathbb{E}_{S_P}[s_K^H \chi_X] + (\sigma_{JU} - \sigma_{XS}) \iota_S \mathbb{E}_{S_I}[s_X^J] \\ &\quad + (\iota_S \mathbb{E}_{S_I} - \iota_U \mathbb{E}_{U_I}) [(\kappa^P - \sigma_{JU}) \chi_X] + (\varepsilon^Y - \sigma_{HU}) (\mathbb{E}_S - \mathbb{E}_U) [s_H s_K^H \chi_X], \\ \mathcal{H} &= \mathbb{E}_K[\kappa^P \chi_X] - n_S(1 - \iota_S)(\sigma_{HU} - \sigma_{KS}) \mathbb{E}_{S_P}[s_K^H \chi_X] \\ &\quad - n_S \iota_S \mathbb{E}_{S_I} [(\sigma_{JU} - \sigma_{XS}) s_X^J + (\kappa^P - \sigma_{JU}) \chi_X] \\ &\quad - (1 - n_S) \iota_U \mathbb{E}_{U_I} [(\kappa^P - \sigma_{JU}) \chi_X] + (\varepsilon^Y - \sigma_{HU}) (\mathbb{E}_K - \mathbb{E}_N) [s_H s_K^H \chi_X].\end{aligned}\quad (42)$$

$\mathcal{G}$ is the skill bias of cheaper hardware and $\mathcal{H}$ the deepening it causes. For ε_M^D , set $\hat{w} = 1$ and let the hardware price change by $\hat{p}^*$, the value that sets Equation (41) to zero, so that machinery per worker does not change. With $\hat{r}^* = \chi_S + \chi_X \hat{p}^*$ and $\hat{c}^*$ the implied changes,

$$\begin{aligned}\varepsilon_M^D &= \underbrace{(1 - \iota_S) \mathbb{E}_{S_P}[\sigma^P]}_{\text{within production}} + \underbrace{(\sigma_{HU} - \sigma_{KS})(1 - \iota_S) \mathbb{E}_{S_P}[s_K^H \hat{r}^*]}_{\text{capital-price adjustment}} \\ &\quad + \underbrace{\iota_S \mathbb{E}_{S_I} [\sigma^I + (\sigma_{JU} - \sigma_{XS}) s_X^J \hat{p}^*]}_{\text{within integration}} \\ &\quad + \underbrace{(\iota_S \mathbb{E}_{S_I} - \iota_U \mathbb{E}_{U_I}) [(\kappa^P - \sigma_{JU}) \hat{r}^* + (\sigma_{HU} - \sigma_{KS})(1 - s_K^H)]}_{\text{across stages}} \\ &\quad + \underbrace{(\varepsilon^Y - \sigma_{HU}) (\mathbb{E}_S - \mathbb{E}_U) [\hat{c}^*]}_{\text{across firms}}.\end{aligned}\quad (43)$$

The first four terms are the average of firms' own responses. The last is reallocation across firms: it moves the aggregate elasticity away from that average by $\varepsilon^Y - \sigma_{HU}$ times the gap in cost exposure between the firms that employ high-STEM workers and those that employ low non-STEM workers, in either direction depending on that gap.

Identical Firms without Integration. With identical firms and no integration labor, Equation (42) becomes

$$\Psi_M = \frac{(\sigma_{HU}/\sigma_{KS} - 1) s_K^H}{1 + (1 - n_S)(\sigma_{HU}/\sigma_{KS} - 1) s_K^H}, \quad (44)$$

which depends on the two elasticities only through their ratio, the object that Equation (27) maps from the event-study responses.

E.5 Measuring the Sufficient Statistics

E.5.1 From Firm-Level Event Studies to the Aggregate Demand Shift

The Event-Study Ψ_M . Taking the event study at face value, Ψ_M is the skill differential four years after a spike per log point of machinery deepening, with machinery per worker on the same basis as our employment outcomes, all the labor a firm uses. High-STEM employment is 33 percent above its pre-spike level and low non-STEM employment 19 percent, a differential of 0.113 in logs, over a rise in machinery per worker of 1.051: $\Psi_M = 0.108$ (standard error 0.023). The log differential removes the investing firm's own growth, which all skill groups share, and a spike raises machinery per worker by about as much as the aggregate rise we study (105.1 against 108.0 log points), so the ratio is applied over the range it is measured on. Measured against the comparison firms' own growth instead of the pre-spike level, the ratio is about 16 percent lower. Applied to this Ψ_M , Equation (9) gives 4.0 log points in the skill premium and 6.0 in relative employment, against 3.1 and 4.6 in the model.

From Investing Firms to the Aggregate. The model's own simulated event study gives $\Psi_M = 0.1176$, within one standard error of our estimate, so the gap to the aggregate $\Psi_M = 0.0817$ reflects a difference between the investing firms that enter our event studies and the economy-wide response to the rise in machinery per worker. The difference

has two causes. First, the event study follows selected firms through one long jump, from below the economy’s machinery per worker to well above it, where production labor responds more per log point of deepening than for the average firm. Second, the year-four differential misses two pieces of integration labor that pull in opposite directions. A firm with more machinery must install more each year to replace depreciation, but four years after a spike investing firms replace only about a quarter of that flow; adding the missing installation labor, at the labor per euro installed that the spike year reveals, raises the shift from 11.7 to about 14.1 log points. The ratio also takes a machine to need as much integration labor in 2021 as in 2003, whereas in the model cheaper hardware lowers it, which takes back about 2.2 log points. In our event studies the two roughly offset, leaving about 12 log points; in the estimated model the second is larger, which, together with the selection of investing firms, accounts for the gap to 8.8.

In-House Employment. A researcher with payroll data would instead divide in-house employment by in-house pre-spike employment. We approximate that log change by $\ln(1 + \beta_s^H / h_s)$, where β_s^H is our in-house coefficient and h_s the in-house share of pre-spike employment (0.63 for high-STEM and 0.68 for low non-STEM workers, from Table 2). Four years after a spike, in-house employment has risen by 8.7 percent of pre-spike employment for high-STEM workers and by 12.5 percent for low non-STEM workers, which gives the negative payroll Ψ_M of Section 6.4. The conversion approximates aggregate in-house shares, the right ones for an aggregate of payrolls, by ratios of Table 2’s means; at the average in-house share across investing firms, 0.28, the payroll Ψ_M would instead be positive, 0.093.

E.5.2 The Decomposition in the Estimated Model

Table E.2 evaluates the terms of Equations (42) and (43) in the estimated model. Substitution within production accounts for most of both Ψ_M and ε_M^D . Integration contributes

35 percent of the skill bias in Ψ_M with only 4.3 percent of high-STEM employment, mainly because the gap between the two substitution elasticities is wider in integration than in production. That gap rests on elasticities we set rather than estimate, and integration's share ranges from 8 to 68 percent across the alternatives in Table E.5; what disciplines them is that, under them, integration remains a roughly stable share of what a machine costs as hardware gets cheaper (Appendix B.2.1). Reallocation across firms is small because firms that employ high-STEM workers and firms that employ low non-STEM workers see much the same cost decline: machinery intensity is the model's only source of variation in skill intensity. Skill shares are therefore less dispersed across firms in the model than in the data; at the dispersion we measure, ε_M^D would be 1.67 and the equilibrium effects about 9 percent smaller (Appendix E.6, Table E.6). The dynamic residual, the gap between the static formulas and the model's own derivatives, is -0.0007 for Ψ_M and 0.003 for ε_M^D .

Table E.2: The Decomposition of Ψ_M and ε_M^D in the Estimated Model

	Value	Share
Ψ_M : <i>skill differential, numerator</i>		
within production	0.0383	63.4%
within integration	0.0213	35.2%
across stages	-0.0017	-2.9%
across firms	0.0026	4.3%
total	0.0605	
Ψ_M : <i>machinery deepening, denominator</i>		
within firms	0.7817	96.6%
production skill mix	-0.0021	-0.3%
integration employment	-0.0004	0.0%
across firms	0.0302	3.7%
total	0.8094	
Ψ_M , formula	0.0747	
Ψ_M , the model's own derivative	0.0740	
dynamic residual	-0.0007	
ε_M^D		
within production	1.3568	97.2%
capital-price adjustment	-0.0021	-0.2%
within integration	0.0379	2.7%
across stages	0.0031	0.2%
across firms	0.0008	0.1%
ε_M^D , formula	1.3964	
ε_M^D , the model's own derivative	1.3993	
dynamic residual	0.0028	

Notes: The table evaluates the terms of Equations (42) and (43) on the stationary distribution of the estimated model. Ψ_M is the ratio of the first block to the second: both are responses to a fall in the hardware price at unchanged wages, so the price change cancels and the ratio is the skill differential per log point of machinery deepening. Shares are of each block's formula total. "Formula" is the closed form above, "the model's own derivative" is the same object computed numerically from the model, and the dynamic residual is their difference: uncertainty, time to build, irreversible and lumpy investment, and the timing of investment across hardware prices, which the closed forms do not carry. All derivatives are taken at the initial equilibrium, so Ψ_M here differs from the 0.0817 used in the counterfactual (Appendix E.4), which averages responses along the path between the two equilibria.

E.6 Sensitivity

Table E.3 solves the model under other supply elasticities: with perfectly inelastic supply the whole shift raises the skill premium, and as supply becomes more elastic more of it raises relative employment. The results also depend on σ_{HU} , which our event study does not pin down: in the special case of Equation (44) only the ratio σ_{HU}/σ_{KS} matters for Ψ_M , and although the full model depends on levels too, ε_M^D moves almost one for one with σ_{HU} . Holding Ψ_M and the ratio fixed, raising σ_{HU} from 1.25 to 1.5 raises ε_M^D from 1.18 to 1.40 and lowers the rise in the skill premium from 3.3 to 3.0 log points (Table E.4). The two elasticities inside integration are set a priori. Re-estimating all parameters for each pair in Table E.5, the model's rise in the skill premium ranges from 3.1 to 4.4 log points, against 3.1 at the baseline.

Table E.3: The Results under Other Supply Elasticities

ε^S	Skill premium	Relative employment
0.0	6.4	0.0
0.5	4.7	2.3
1.0	3.7	3.7
1.5	3.1	4.6
3.0	2.0	6.0

Notes: The model solved numerically, in log points, on the model's own demand shift (Table 3). Each row lowers the hardware price until machinery per worker has risen by the observed amount.

Table E.4: The Results under Other Values of σ_{HU}

σ_{HU}	σ_{KS}	ε_M^D	Skill premium	Relative employment
1.05	0.914	1.000	3.5	5.3
1.25	1.088	1.179	3.3	4.9
1.50	1.306	1.396	3.0	4.6

Notes: Equation (9) with the model's demand shift $\Psi_M = 0.0817$ (Table 3), $\varepsilon^S = 1.5$ and the observed rise in machinery per worker, in log points. σ_{KS} moves with σ_{HU} so that their ratio stays at its estimated value; ε_M^D is the closed form of Equation (43) evaluated in the model with the other parameters at their estimated values. Away from the baseline the first-order approximation is less accurate: in the model, the formula's error against the exact counterfactual reaches 9 percent at the ends of this range.

Table E.5: The Results under Other Integration Elasticities

σ_{XS}	σ_{JU}	σ_{KS}	σ_{HU}/σ_{KS}	Equation (27)	Integration share	Skill premium	Relative employment
0.90	1.50	1.306	1.15	1.57	35%	3.1	4.6
0.50	1.50	1.399	1.07	1.46	68%	4.4	6.6
0.90	1.25	1.249	1.20	1.68	21%	3.7	5.6
1.40	1.50	1.277	1.17	1.57	8%		

Notes: For each pair of integration elasticities, all 12 parameters are re-estimated by indirect inference on the event study; the first row is the baseline. The column headed Equation (27) applies that equation to each pair's simulated event study as a researcher would, with machinery per worker and the pre-spike cost share. The skill premium and relative employment are the model's general-equilibrium counterfactual, in log points; the integration share is its share of the skill differential in Ψ_M . The model requires $\sigma_{JU} > \sigma_{XS}$. For $(\sigma_{XS}, \sigma_{JU}) = (1.4, 1.5)$ the two general-equilibrium columns are blank because the model cannot be solved at the counterfactual hardware price; the base economy, which the other columns use, does solve. Omitted because the model cannot be solved for them at all: $(\sigma_{XS}, \sigma_{JU}) = (1.25, 1.5)$; $(\sigma_{XS}, \sigma_{JU}) = (0.9, 2)$.

Reallocation across Firms. The last term in Equation (43) is reallocation across firms, and it depends on how much more skill-intensive the firms that employ high-STEM workers are than those that employ low non-STEM workers. That difference is small in the model, where firms differ in productivity, which changes their scale but not their skill mix, and in machinery intensity, which is the only thing that does. Skill shares are therefore less dispersed than in the data: the high-STEM share of total employment has a standard deviation of 6.3 points around a mean of 5.0 across firms, against 0.5 in the model. Table E.6 shows what this is worth: with the model's skill-intensity gap replaced

by the one the measured dispersion implies within industries, and everything else at its estimated value, ε_M^D rises from 1.40 to 1.67 and the rise in both the skill premium and relative employment is about 9 percent smaller. The two shrink together because a more elastic relative demand moves the equilibrium less for the same shift; the split between wages and employment is set by the supply elasticity (Table E.3).

Table E.6: The Results under the Measured Dispersion of Skill Shares

	Skill-intensity gap	ε_M^D	Skill premium	Relative employment
the model	0.000	1.399	3.0	4.6
measured dispersion, within industry	0.118	1.674	2.8	4.2
measured dispersion, raw	0.179	1.816	2.7	4.0

Notes: Equation (9) with the model’s demand shift $\Psi_M = 0.0817$ (as in Table E.4), in log points. The skill-intensity gap is the difference in skilled cost shares between the firms that employ high-STEM workers and those that employ low non-STEM workers, the last term of Equation (43); ε_M^D adds that term to the model’s value. “Within industry” uses the standard deviation of the high-STEM share of total employment demeaned by NACE 2-digit industry and year, “raw” the undemeaned one. The implied gap assumes equal firm sizes and compares high-STEM workers with all other workers; it moves only ε_M^D , with Ψ_M held fixed.

E.7 The Make-or-Buy Margin

The baseline model sums in-house and external workers in each firm’s labor demand. We now bring back the make-or-buy choice of Equation (3) for high-STEM production labor, in a quantitative version of the fixed-cost mechanism behind Prediction 2: each firm chooses every year whether to hire its high-STEM production workers in-house, paying the wage w_S plus a fixed cost $f w_S$, or to buy them from suppliers at the price $(1 + \tau) w_S$ with no fixed cost.⁵⁵ Given capital and states, the two options yield profits

$$\pi^H(K, e) = \Pi(K, e; w_S) - f w_S, \quad \pi^O(K, e) = \Pi(K, e; (1 + \tau) w_S), \quad (45)$$

⁵⁵Both the fixed cost and the markup are pure costs, denominated in the skilled wage so that they scale with prices: neither is paid in labor, and bought workers enter market clearing in the same skill pool at w_S , so the markup acts like an iceberg cost on bought labor. It stands for the higher wage bill and the overhead behind a supplier’s price.

where $\Pi(K, e; w)$ is the profit of Equation (31) when high-STEM production labor costs w . In the notation of Appendix A.2, the external price is $w_S^O = (1 + \tau) w_S$ and the fixed cost is $f_S^H = f w_S$. A mean-zero logit shock with scale ς smooths the choice, so the firm buys with probability $\Pr(O)$ and earns in expectation $\tilde{\Pi}$:

$$\Pr(O) = \frac{\exp(\pi^O/\varsigma)}{\exp(\pi^H/\varsigma) + \exp(\pi^O/\varsigma)}, \quad \tilde{\Pi} = \varsigma \ln [\exp(\pi^H/\varsigma) + \exp(\pi^O/\varsigma)]. \quad (46)$$

$\tilde{\Pi}$ replaces Π in Equation (32), and ς is half a percent of the firm's in-house skilled wage bill, so the choice turns on the cost saving rather than on noise. Firms that need few high-STEM workers buy them, as in Prediction 2. Bought workers come from the same skill pool, so Equation (33) is unchanged, with each firm's labor demand the probability-weighted average of the two options; integration labor is supplied externally throughout. For each markup τ , we set f so that 0.63 of high-STEM employment is in-house, as in the data, hold the other parameters at their estimated values, and recompute Ψ_M and ε_M^D (Table E.7).

The margin makes high-STEM demand more responsive in both directions. Cheaper hardware pushes some firms to hire in-house, lowering their price of skill by the markup and raising Ψ_M ; a higher skilled wage pushes some to buy, raising ε_M^D . The two effects largely offset in the skill premium. With a markup of 3.6 percent, the machine-supplier wage premium we estimate (Table D.6), 77 percent of firms buy, and Ψ_M , ε_M^D and the premium rise by 0.9, 1.7 and 0.4 percent. Because suppliers' prices also cover overhead and margins, the wage premium is a lower bound on the markup; with a markup of 15 percent, the changes are 3.8, 9.3 and 1.5 percent.

Table E.7: The Results with an Endogenous Make-or-Buy Decision

	τ	f	External share	Firms buying	Ψ_M	ε_M^D	Skill premium
Baseline					0.0740	1.3993	3.06
Make-or-buy	0.036	0.0096	0.371	0.774	0.0747	1.4230	3.08
Make-or-buy	0.150	0.042	0.371	0.800	0.0768	1.5296	3.11

Notes: τ is the supplier markup and f the fixed cost of in-house high-STEM labor, in units of the high-STEM wage. The model without the margin describes each firm's total labor demand, in-house and external together, and has no split, so its external share and share of firms buying are empty. With the margin, the external share counts high-STEM production labor bought from suppliers plus integration labor, which is assigned to suppliers; f is set so that this share matches the data, where 0.63 of high-STEM employment is in-house. Ψ_M and ε_M^D are the model's local derivatives. The skill premium is the model's general-equilibrium counterfactual, in log points, as in Table 3. All other parameters are held at their estimated values.

Online Appendix References

- Abraham, Katharine G, and Susan K Taylor.** 1996. “Firms’ use of outside contractors: Theory and evidence.” *Journal of Labor Economics*, 14(3): 394–424.
- Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo.** 2020. “Competing With Robots: Firm-Level Evidence From France.” Vol. 110, 383–388, American Economic Association.
- Acemoglu, Daron, Hans RA Koster, and Ceren Ozgen.** 2023. “Robots and Workers: Evidence From the Netherlands.” National Bureau of Economic Research 31009.
- Altonji, Joseph G, Todd E Elder, and Christopher R Taber.** 2005. “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools.” *Journal of Political Economy*, 113(1): 151–184.
- Anandan, Tanya M.** 2015. “Calculating Your ROI for Robotic Automation: Cost vs. Cash Flow.” Robotic Industries Association (A3), Industry Insights, March 19, 2015. Available at <https://www.automate.org/robotics/industry-insights/calculating-your-roi-for-robotic-automation-cost-vs-cash-flow>.
- Baqae, David, Ariel Burstein, Cédric Duprez, and Emmanuel Farhi.** 2025. “Consumer Surplus from Suppliers: How Big is it and Does it Matter for Growth?” *Econometrica*, 93(6): 2043–2081.
- Békés, Gábor, Alessandra Bonfiglioli, Rosario Crinò, and Gino Gancia.** 2025. “Integrators and Robot Adoption: Facts from Hungary.” Centre for Economic Policy Research Discussion Paper 20026.
- Bernard, Andrew B., Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes.** 2022. “The Origins of Firm Heterogeneity: A Production Network Approach.” *Journal of Political Economy*, 130(7): 1765–1804.
- Berry, Steven, Amit Gandhi, and Philip Haile.** 2013. “Connected Substitutes and Invertibility of Demand.” *Econometrica*, 81(5): 2087–2111.
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge.** 2025. “What Happens to Workers at Firms that Automate?” *The Review of Economics and Statistics*, 107(1): 125–141.
- Bloom, Nicholas.** 2009. “The Impact of Uncertainty Shocks.” *Econometrica*, 77(3): 623–685.
- Bonfiglioli, Alessandra, Rosario Crinò, Harald Fadinger, and Gino Gancia.** 2024. “Robot Imports and Firm-Level Outcomes.” *The Economic Journal*, 134(664): 3428–3444.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess.** 2024. “Revisiting event-study designs: robust and efficient estimation.” *The Review of Economic Studies*, 91(6): 3253–3285.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer.** 2019. “The Effect of Minimum Wages on Low-Wage Jobs.” *The Quarterly Journal of Economics*, 134(3): 1405–1454.
- Cooper, Russell W, and John C Haltiwanger.** 2006. “On the Nature of Capital Adjustment Costs.” *The Review of Economic Studies*, 73(3): 611–633.
- Davis, Steven J, and John Haltiwanger.** 1992. “Gross Job Creation, Gross Job Destruction, and Employment Reallocation.” *The Quarterly Journal of Economics*, 107(3): 819–863.
- Dhyne, Emmanuel, Cédric Duprez, and Toshiaki Komatsu.** 2023. “The Belgian business-to-business transactions dataset 2002–2021.” National Bank of Belgium 444.
- Dixon, Jay, Bryan Hong, and Lynn Wu.** 2021. “The Robot Revolution: Managerial and Employment Consequences for Firms.” *Management Science*, 67(9): 5586–5605.
- Doraszelski, Ulrich, and Jordi Jaumandreu.** 2013. “R&D and Productivity: Estimating Endogenous Productivity.” *Review of Economic Studies*, 80(4): 1338–1383.
- Elkins, Dean.** 2025. “Manufacturing Automation Project Guide: 9 Essential Steps from

- Discovery to Success.” Control System Integrators Association (CSIA) / Automation World, August 4, 2025. Available at <https://www.automationworld.com/control/article/55307156/control-system-integrators-association-csia-manufacturing-automation-project-guide-9-essential-steps-from-discovery-to-success>.
- European Commission.** 2024. “Guide to Application of the Machinery Directive 2006/42/EC.” Edition 2.3, April 2024. Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs, Brussels.
- Eurostat.** 2026a. “Labour Cost, Wages and Salaries, Direct Remuneration by NACE Rev. 2 Activity [lc_ncost_r2].” *Labour Cost Survey, Eurostat*, Retrieved 2026-07-17 from https://ec.europa.eu/eurostat/databrowser/view/lc_ncost_r2/default/table.
- Eurostat.** 2026b. “Mean Annual Earnings by Sex, Economic Activity and Educational Attainment [earn_ses18_30].” *Structure of Earnings Survey, Eurostat*, Retrieved 2026-07-17 from https://ec.europa.eu/eurostat/databrowser/view/earn_ses18_30/default/table.
- Federal Reserve Bank of St. Louis.** 2025. “Equipment Deflator [EDEF].” *FRED, Federal Reserve Bank of St. Louis*, Quarterly index (2017 = 100), retrieved 2025-11-10 from <https://fred.stlouisfed.org/series/EDEF>.
- Goldin, Claudia, and Lawrence F Katz.** 1998. “The Origins of Technology-Skill Complementarity.” *The Quarterly Journal of Economics*, 113(3): 693–732.
- Gourieroux, Christian, Alain Monfort, and Eric Renault.** 1993. “Indirect Inference.” *Journal of Applied Econometrics*, 8(S1): S85–S118.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell.** 1997. “Long-run implications of investment-specific technological change.” *The American Economic Review*, 87(3): 342–362.
- Griliches, Zvi.** 1969. “Capital-Skill Complementarity.” *The Review of Economics and Statistics*, 51(4): 465–468.
- Howard, Ronald A.** 1960. *Dynamic Programming and Markov Processes*. Cambridge, MA:MIT Press.
- HowToRobot.** 2024. “Robot Price: How Much Does a Robot Solution Really Cost?” HowToRobot Expert Insight, December 12, 2024. Available at <https://howtorobot.com/expert-insight/robot-price-how-much-does-robot-solution-really-cost>.
- Hubmer, Joachim, and Pascual Restrepo.** 2026. “Not a Typical Firm: Capital–Labor Substitution and Firms’ Labor Shares.” *American Economic Journal: Macroeconomics*, 18(2): 34–71.
- Humlum, Anders.** 2022. “Robot Adoption and Labor Market Dynamics.” ROCKWOOL Foundation Research Unit Study Paper 175.
- IFR.** 2022. “IFR Annual Report.” International Federation of Robotics.
- International Organization for Standardization.** 2025. “Robotics — Safety Requirements — Part 2: Industrial Robot Applications and Robot Cells.” ISO 10218-2:2025, Edition 2, February 2025. ISO, Geneva.
- Katz, Lawrence F, and Kevin M Murphy.** 1992. “Changes in Relative Wages, 1963–1987: Supply and Demand Factors.” *Quarterly Journal of Economics*, 107(1): 35–78.
- Khan, Aubhik, and Julia K Thomas.** 2008. “Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics.” *Econometrica*, 76(2): 395–436.
- Krusell, Per, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante.** 2000. “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis.” *Econometrica*, 68(5): 1029–1053.
- Lindner, Attila, Balázs Muraközy, Balázs Reizer, and Ragnhild Schreiner.** 2024. “Firm-level Technological Change and Skill Demand.” *American Economic Review*. conditionally accepted.
- Mann, Katja, and Dario Pozzoli.** 2022. “Automation and Low-Skill Labor.” IZA Institute

- of Labor Economics Discussion Paper 15791.
- Nelder, John A., and Roger Mead.** 1965. “A Simplex Method for Function Minimization.” *The Computer Journal*, 7(4): 308–313.
- Oberfield, Ezra, and Devesh Raval.** 2021. “Micro Data and Macro Technology.” *Econometrica*, 89(2): 703–732.
- Olley, G Steven, and Ariel Pakes.** 1996. “The Dynamics of Productivity in the Telecommunications Equipment Industry.” *Econometrica*, 64(6): 1263–1297.
- Orgalime.** 2012. “General Conditions for the Supply of Mechanical, Electrical and Electronic Products (S 2012).” Brussels: Orgalime, the European Engineering Industries Association.
- Oster, Emily.** 2019. “Unobservable Selection and Coefficient Stability: Theory and Evidence.” *Journal of Business & Economic Statistics*, 37(2): 187–204.
- Rouwenhorst, K. Geert.** 1995. “Asset Pricing Implications of Equilibrium Business Cycle Models.” In *Frontiers of Business Cycle Research.*, ed. Thomas F. Cooley, 294–330. Princeton, NJ: Princeton University Press.
- Rubinton, Hannah.** 2026. “The Geography of Business Dynamism and Skill-Biased Technical Change.” *The Review of Economic Studies*, 93(3): 2099–2132.
- Stokey, Nancy L, and Robert Lucas.** 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Stuart, Elizabeth A, and Donald B Rubin.** 2008. “Best Practices in Quasi-Experimental Designs: Matching Methods for Causal Inference.” In *Best Practices in Quantitative Methods.*, ed. Jason W Osborne, 155–176. Thousand Oaks, CA: Sage Publications.
- Winberry, Thomas.** 2021. “Lumpy Investment, Business Cycles, and Stimulus Policy.” *American Economic Review*, 111(1): 364–396.
- Young, Eric R.** 2010. “Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell–Smith Algorithm and Non-Stochastic Simulations.” *Journal of Economic Dynamics and Control*, 34(1): 36–41.