What is automation? Evidence on the role of capital heterogeneity from Swedish data

Chek Yin Choi * Isabella Maassen†

Abstract

Automation is often studied by relating investments in capital to labor market outcomes, yet production capital is highly heterogeneous, and different technologies need not interact with labor in the same way. This makes the empirical measurement of automation inherently sensitive to how automation capital is defined. Using matched Swedish administrative data linking firms, workers, and highly disaggregated imports of capital goods, we isolate the role of measurement by replicating the leading empirical approaches in the literature under a common institutional setting while systematically varying the definition of automation capital. We compare definitions ranging from industrial robots to broad automation measures and detailed capital classes. We find that capital heterogeneity matters in three ways. First, evidence based on industrial robots should be benchmarked against other capital goods before being interpreted as evidence on automation more generally, as many conventional technologies exhibit similar or stronger labor market effects while broader automation measures yield different conclusions. Second, aggregating heterogeneous technologies into broad automation measures masks economically meaningful variation. Third, under a common data environment, changing the definition of automation has a larger effect on empirical conclusions than changing the estimation strategy. Our results suggest that capital heterogeneity is a fundamental measurement issue and should be taken into account when drawing conclusions about the labor market effects of automation.

Keywords: automation, capital heterogeneity, labor demand, technology adoption, measurement.

I. Introduction

Automation has become one of the central topics in the debate on the future of work. Whether new technologies ultimately reduce or increase labor demand depends on the

*IIES, Stockholm University, chek-yin.choi@iies.su.se

†IIES, Stockholm University, isabella.maassen@iies.su.se

balance between displacement effects, which replace workers, and productivity effects, which allow firms to expand production and employment. Empirical evidence on the aggregate effects remains divided. While some studies conclude that automation reduces employment, others find a positive effect. Perhaps most strikingly, Acemoglu et al. (2020) (ALR) and Aghion et al. (2025) (AABJ) reach different conclusions using French firm-level data, a contrast highlighted in Aghion et al. (2022). This raises the question of whether part of the disagreement reflects differences in how automation itself is measured rather than differences in the underlying economic mechanisms.

Different capital goods embody different production technologies and therefore do not necessarily have the same relationship with labor demand. Treating them as interchangeable therefore affects what empirical measures of automation capture. Using matched Swedish administrative data linking firms, workers, and detailed import records on capital goods, we replicate the main empirical strategies used in the recent automation literature under a common setting. We then systematically vary the definition of automation capital, from industrial robots to broad automation measures and highly disaggregated capital classes. This allows us to answer the question: What role does capital heterogeneity play in the measurement of automation, and how does it affect empirical conclusions?

This paper shows that capital heterogeneity is central to the empirical measurement of automation. First, heterogeneity implies that results obtained from a single technology, such as industrial robots, should be benchmarked against other types of capital before being interpreted as evidence on automation more generally, and to avoid overemphasizing the importance of salient technologies. Second, aggregating different capital classes into a broad measure of automation can conceal heterogeneous effects at the granular level that partially offset one another. Third, because different types of capital interact differently with labor, the definition of automation itself becomes an important determinant of aggregate conclusions and helps explain why the literature reaches seemingly conflicting results.

Recent empirical work has approached this question using different measures of automation and different empirical strategies. At the firm level, ALR identify automation through industrial robot imports and estimate their effects using matched employer-employee data from France, concluding that while robot-adopting firms expand, negative spillovers on competitors imply a decline in aggregate labor demand. In contrast, Aghion et al. (2023) measure automation more broadly using investments in modern manufacturing capital and, employing an industry-level event-study design on the same country, find positive aggregate employment effects. More generally, the empirical literature spans a range of countries and settings, including firm-level evidence from Canada, Spain, China, Austria, and the United States, with conclusions differing across technologies, data, and empirical methods. This raises the possibility that part of the apparent disagreement reflects differences in how automation is measured, rather than differences in the underlying economic relationship between technology and labor.

Our contribution is to isolate the role of measurement where capital is heterogeneous, while holding the institutional setting and data constant. Replicating influential empirical strategies on Swedish data allows us to compare alternative definitions of automation directly. Applying the firm-level approach of ALR to detailed capital classes shows that the estimated effects of robots resemble those of several more conventional technologies, suggesting that focusing on a salient technology without benchmarking can overstate its general importance. Applying the broader automation measure used by Aghion et al. (2023) instead yields much weaker results because it combines capital classes with heterogeneous effects. Finally, comparing aggregate estimates across the alternative measures of automation in ALR and AABJ shows that the definition of automation capital is a more important source of disagreement than the estimation strategy in this debate.

The remainder of the paper proceeds as follows. Section 2 reviews the conceptual background and related literature. Section 3 describes the Swedish administrative data and the construction of the capital measures. Section 4 presents firm-level evidence and illustrates why capital heterogeneity matters at the disaggregated level. Section 5 compares aggregate estimates across alternative definitions of automation and discusses their implications for the broader debate. Section 6 concludes.

II. Conceptual background and related literature

As pointed out in a recent survey of the literature by Aghion et al. (2022), the literature has not converged on a single answer to the question: what are the effects of automation on labor? They emphasize that there are two views, depending on whether displacement effects or productivity gains prevail, and show that both conclusions have been supported in the literature. For example, ALR find that while robot adoption makes firms expand and hire more workers, the aggregate effect on labor demand is negative due to the effect on competitors. Using the same French data, however, Aghion et al. (2023) and AABJ find that automation has a positive employment effect at the aggregate level.

The literature has emphasized the role of various dimensions in which estimation methods diverge in explaining diverging findings. For example, Aghion et al. (2022) underlines that the level of analysis matters: automation can increase employment at the firm level but decrease it at the aggregate level due to business stealing. Relatedly, Zator, 2019 shows that the effects of automation differ by industry, where some service industries observe an increase in employment, while in manufacturing and retail the substitution effect dominates. This paper aims to isolate the role of alternative definitions of automation capital while holding data and estimation strategy fixed.

A natural starting point for thinking about this question is the literature on capital heterogeneity. If different types of capital embody different production technologies and interact differently with labor, then empirical measures of automation based on capital goods may capture different economic mechanisms. In that case, empirical measures of automation are not merely proxies for the same underlying concept. Instead, different definitions may capture different economic mechanisms, making measurement itself an important determinant of estimated labor-market effects. The importance of capital heterogeneity has long been recognized, all the way back to Robinson (1953), who criticized the literature for treating capital as a homogeneous input good. Krusell et al. (2000) further emphasize this using U.S. data to show that equipment capital is relatively more complementary with skilled labor than with unskilled labor. They attribute most of the rise in the skill premium in the second half of the 20th century to falling equipment prices. More recently, O'Mahony et al. (2021) show in a theoretical model that capital heterogeneity is central to understanding the decline in the labor share, and using industry-level OECD data, they demonstrate that different types of capital assets affect the labor share in opposite directions. Chen (2020) also highlight the importance of capital-skill complementarity and rising capital intensity in the goods sector for understanding structural transformation in the United States since the 1950s.

While there is little doubt that capital heterogeneity exists and is substantial, most of the empirical literature on automation has focused on robots as a benchmark for studying technology adoption and its labor market consequences. Many of the earlier papers use industry-level data from the 90s and early 2000s from the International Federation of Robotics (IFR), such as Graetz and Michaels (2018) across countries, Acemoglu and Restrepo (2020) across US commuting zones, and Dauth et al. (2021) across German local labor markets. Results are mixed: while Graetz and Michaels (2018) and Dauth et al. (2021) find no significant effect on employment, Acemoglu and Restrepo (2020) conclude that one additional robot per 1,000 workers reduced the employment-to-population ratio by .2 percentage points.

More recently, a growing literature has examined the effects of robot adoption at the firm and plant level. Three broad findings emerge. First, a recurring finding is that robot adoption increases capital intensity and reduces the labor share (Acemoglu et al. (2020), Dinlersoz and Wolf (2024), and Koch et al. (2021)). Evidence on employment is more nuanced: adopting firms often expand employment through scale effects (Acemoglu et al. (2020) and Dixon et al. (2020)), while incumbent workers may nevertheless experience greater separations and earnings losses (Bessen et al., 2023). Second, robot adoption improves firm performance along multiple dimensions. Adopting firms exhibit higher productivity and output, while also achieving gains in product quality, export performance, innovation, and broader measures of competitiveness (Acemoglu et al. (2020), Alguacil et al. (2022), Dixon et al. (2020), Koch et al. (2021), Luo and Qiao (2024), K. Wang et al. (2025), and Zhang et al. (2025)). Third, robot adoption often shifts composition toward higher-skilled and more technical occupations. Although the magnitude of these adjustments varies across countries, adopting firms generally employ relatively fewer production workers and relatively more engineers, technicians, and other skilled workers (Bonfiglioli et al. (2024), Dixon et al. (2020), Humlum (2019), Tverdostup et al. (2025), Voshaar et al. (2022), and T. Wang et al. (2024)).

While a great proportion of empirical work focuses on robots, there is work on other automation technologies as well. As mentioned above, Aghion et al. (2023) use a broad measure of automation capital, also used in Acemoglu and Restrepo (2022). They find an increase in aggregate employment, which contrasts with the patterns observed in the robot literature. Much of this literature fixes a definition of automation from the start; however, some have explicitly considered other types of automation capital. Caselli and Wilson (2004) use cross-country data on the composition of imported equipment and find that it can explain TFP differences. Bustos (2011) uses data on Argentinian firms' spending on various technologies and shows that they respond differently to a trade shock. Webb (2019) builds an algorithm for classifying patents into categories such as AI, robot, and software, and occupational task descriptions and shows how they exhibit different patterns of interaction with labor.

This paper brings these two literatures together. The capital heterogeneity literature suggests that different types of capital should not be expected to have identical relationships with labor, while the automation literature typically adopts a fixed definition of automation. We ask whether alternative definitions of automation therefore lead to systematically different empirical conclusions, holding the data setting and estimation strategy constant. By extending the methodology used in much of the robotization literature, namely using firm-level import data, to a wide array of capital goods, we can separate measure from method. Using matched Swedish administrative data linking firms, workers, and detailed capital imports, this paper contributes by providing novel evidence on automation that jointly captures firm performance and heterogeneous capital effects at a highly granular level, thereby complementing and extending existing firm-level studies.

III. Data and variable construction

The underlying data set is provided by Statistics Sweden (SCB) and includes matched data on firms, workers, and imports. Our baseline specification uses a four-year difference (2012–2016) rather than a five-year window. This choice is driven by the availability of consistent product and industry classifications, which change at five-year intervals. Using a four-year window ensures that firms are observed under a consistent classification system throughout the estimation period. To facilitate comparability with Acemoglu et al., 2020, who use a five-year window (2010–2015), we report corresponding robot estimates using a five-year difference in Table A.2 in the Appendix. For the years 2012 to 2016 which we use, consistent definitions of occupations, industry, and import goods codes are available. In the following, we will describe the data sources, processing, and variable construction.

A. Firm data

FEK The registry FEK (företagens ekonomi) contains firms' financial data. It comes from tax records and includes information on capital, labor, the wage bill, profits, revenue, and value added. All active Swedish firms except those in the financial sector are included. The data is filtered to only include manufacturing firms, which corresponds to SNI codes 10 – 33.[1] We further restrict the sample to firms that show up in both years 2012 and 2016 with positive revenue, paying wages, and a positive number of employees.

The register also includes information at the establishment level (LVE). This is used to add the commuting zone of the biggest establishment as a variable, relying on SCB's classification of commuting zones. This measure is based on information on workers' home location and location of their main employment.[2]

A corporate dummy is added from the corporation register KCR. We classify a firm as part of a corporation if the firm id shows up in that register. We measure the change in log TFPR, which is one of the outcome variables in the regressions, following Acemoglu et al., 2020. It is defined as:

$$\Delta TFPR = \Delta \log y - \lambda_l \Delta \log l - \lambda_m \Delta \log m - (1 - \lambda_l - \lambda_m)\Delta \log k,$$

where λ_l and λ_m are defined as the 2012 expenditure shares of labor and intermediates in revenue, respectively, y is total revenue, and k is the capital stock.

Value added is constructed as revenue minus intermediate inputs. There is an alternative measure for value added available from Statistics Sweden, which adjusts for changes in inventories, capitalized production, and some production expenses not included in intermediates. We choose to use the 'unadjusted' measure in accordance with how Acemoglu et al., 2020 construct value added. Similarly, we follow Acemoglu et al., 2020 and define labor productivity as value added divided by the number of workers, and the labor share as the share of the wage bill in value added.

[1] From 2008 onwards, the Swedish Statistical Office uses SNI2007 to classify industries. Documentation is available via https://www.scb.se/en/documentation/classifications-and-standards/swedish-standard-industrial-classification-sni/

[2] For documentation, see https://www.scb.se/hitta-statistik/statistik-efter-amne/arbetsmarknad/utbud-av-arbetskraft/befolkningens-arbetsmarknadsstatus/produktrelaterat/fordjupad-information/lokala-arbetsmarknader-la/ and https://www.scb.se/hitta-statistik/statistik-efter-amne/arbetsmarknad/utbud-av-arbetskraft/registerbaserad-arbetsmarknadsstatistik-rams/produktrelaterat/Fordjupad-information/lokala-arbetsmarknader-la/

B. *Worker Data*

LISA The SCB registry LISA, or *longitudinell integrationsdatabas*, contains employment-related administrative data for individuals aged 16 and older[3]. This data includes information on individuals and matches them to their main employer. We use this individual level information about gender, age, and occupation to construct variables capturing the workforce structure at the firm level. More specifically, we calculate the number of workers, female workers, skilled workers, and the average age of employees. Skilled workers in this context are defined as those with an occupation that requires at least a 2-year education after graduating from a *gymnasium*, the Swedish equivalent of high school. This corresponds to occupation codes starting with one, two, or three according to the Swedish classification system for occupations (SSYK)[4]. Note that the classification system was revised in 2012, effective from 2014, which did not affect the logic by which we classify 'skilled' and 'unskilled' labor.

C. *Import Data*

UHV The UHV, which stands for *utrikeshandel med varor*, contains information on both imports and exports at the country-firm-good level. Goods are classified according to the Combined Nomenclature (CN) of the EU with 8-digit identifiers[5]. The data is complete for imports from countries outside the EU, and is reported above a certain threshold for within EU imports, which has increased over time.

For identifying capital imports, we utilize the Broad Economic Categories (BEC) introduced by the UN. The BEC classification labels different products as "Consumption", "Intermediate" or "Capital". Within the group "Capital", products are further divided into "Generic" and "Specific" since 2012.

From the import data, which we merge to the firm dataset, we construct dummies for whether a company imported any capital between 2012 and 2016, whether it imported robots or automation capital, as well as each 8-digit product code within the class of capital goods. In the main regressions, the sample is restricted to capital importers. [6]

D. *Automation measures*

We introduce three main ways to measure automation. Where *robots* are referred to, we use the HS code "84795000 - Industrial robots, n.e.s.". *Automation capital* follows the automation technology list in Aghion et al. (2023), based on the definition of Acemoglu

[3]From 2016, the minimum age to be included in LISA is 15.

[4]Documentation can be found here: https://www.scb.se/dokumentation/klassifikationer-och-standarder/standard-for-svensk-yrkesklassificering-ssyk/

[5]CN codes can be looked up here https://cnwebb.scb.se/?languageId=GB

[6]HS 2012 product codes were obtained from the United Nations Statistics Division Classifications Portal: https://unstats.un.org/unsd/classifications/econ

and Restrepo (2022).[7] Finally, we define a *large machinery investment* as a change in the log value of the stock of machinery at the industry or firm level above the median or 50th percentile across all years and units of observation, respectively.

E. Descriptives

This matched dataset allows us to trace how the acquisition of specific capital goods, identified at the 8-digit CN level, relates to changes in firm performance and workforce composition over time.

There are 16,259 manufacturing firms in our sample, of which 4,665 imported any type of capital goods between 2012 and 2016. About 1.5% of these imported robots, and 11.8% imported automation capital. Note that there is some overlap between firms that automate and firms that imported robots.

Comparing the characteristics of the firms in these groups at baseline, it is clear that firms that import any type of capital are larger, whether measured in terms of revenue, value added, or the number of workers. They also generate more value added per worker and pay higher wages. Similarly, firms that automate are on average approx. 3.5 times larger than capital importers in terms of revenue, while robot adopters are more than 17 times the size of the average capital importer. Value added per worker and wages are higher among firms that automate, and even higher for robot adopters.

Turning to the workforce composition, capital importers have a slightly higher share of female workers than the average among all manufacturing firms. Robot adopters employ 22.8% women on average, while capital importers in general only have 18.4%. More sizable differences can be observed when looking at the skill share of the labor force, where the average among manufacturing firms is 15.2%, which increases to 23.7% among capital importers and is 31.6% for those who purchased robots. Firms that automate are more similar in terms of these workforce characteristics to the overall group of capital importers. Finally, it is worth noting that robot adopters have a much lower labor share, at 51.5%, than the average manufacturing firm, which has a labor share of 64.3%.

While the overall adoption rate of robots among capital importers is less than one percent, it is much higher within some industries. Motor vehicle manufacturing, production of iron and steel and ferroalloys, manufacture of optical instruments and photographic equipment, manufacture of consumer electronics, and manufacture of pipes, tubes, hollow profiles and accessories of steel all have robot adoption rates above 4%.

Examining instead the value of robots imported per firm, the highest variation is found in the same industries, with the exception of pipes, tubes, hollow profiles and accessories.

[7] A list of Automation Technologies https://www.openicpsr.org/openicpsr/project/184641/version/V1/view?path=/openicpsr/184641/fcr:versions/V1/replication_AEAPP-2023-1039/data/other

Instead, production of communication equipment is more varied in the value of robots imported.

While robot adopters make up for a small share of firms, they make up for approximately 17% of employment, 28% of value added, and 23% of revenue in the manufacturing sector. This concentration of adoption in large firms holds for automation capital as well: the probability of importing both automation capital and robots increases in firm size (see Figure A.1 in the Appendix).

In the following sections we will turn to how firms evolve over time, depending on the type of production capital that they import. [8]

IV. Firm-level estimates

A. Replication

In this section, we present results replicating the main specification in Acemoglu et al., 2020 in Sweden. We run these regressions on three definitions of robots and automation: the standard robot indicator; automation capital, used by Aghion et al., 2023; and a large change in the stock of machinery as in AABJ. The results are summarized in Table 1.

We regress the change in various firm characteristics between 2012 and 2016 on a dummy indicating whether the firm has adopted either robots or automation capital, or experienced a large machinery investment, and a number of controls:

$$\Delta Y_f = \alpha + \beta \text{Adopt}_f + X_f' \gamma + \mu_i + \lambda_z + \varepsilon_f, \tag{1}$$

where ΔY_f denotes the change in a given firm outcome between 2012 and 2016, and Adopt_f is an indicator equal to one if a firm adopted automation capital according to the relevant definition. X_f includes controls for initial firm size (log employment) and log value added per worker, as well as a dummy for whether the firm is part of a larger corporate group. μ_i and λ_z represent fixed effects for industry (at the 3-digit SNI level) and commuting zone of the largest establishment, respectively. Standard errors are clustered at the industry level. We also present employment-weighted specifications to account for firm size heterogeneity. The coefficient of interest, β, captures the average difference in growth of the outcome variable between adopters and non-adopters, conditional on baseline firm characteristics.

[8] Table A.1 in the Appendix contains additional summary statistics not discussed in the main text.

$\Delta\log(va)$	Δlabshare	Δskillshare	$\Delta\log(va/l)$	$\Delta\log(TFPR)$	$\Delta\log(l)$	$\Delta\log(w)$	n
A. Robot adoption							
Baseline (unweighted)							
0.321***	-0.216***	0.006	0.200***	0.025	0.121**	0.078***	68
Employment weighted							
0.232***	-0.128**	-0.004	0.124**	-0.026	0.108***	0.046**	68
B. Automation capital adoption							
Baseline (unweighted)							
0.041	-0.020	0.004	0.010	-0.015	0.032	0.013	549
Employment weighted							
0.011	-0.016	0.001	-0.026	-0.021	0.037	0.011	549
C. Large machinery investment							
Baseline (unweighted)							
0.337***	-0.074	0.005	0.078**	-0.214***	0.259***	0.043*	333
Employment weighted							
0.268***	-0.049	-0.007	0.038	-0.216***	0.230***	0.029**	333

Table 1: Robot and Automation Effects

Notes: The table reports coefficients from regressions of changes in firm outcomes between 2012 and 2016 on an indicator for whether the firm adopted either robots (Panel A) or automation capital (Panel B), and on an indicator for a large machinery investment (Panel C). The dependent variables are: $\Delta\log(va)$ change in log value added; Δlabshare change in the labor share; Δskillshare change in the share of skilled workers; $\Delta\log(va/l)$ change in log value added per worker (labor productivity); $\Delta\log(TFPR)$ change in log revenue productivity; $\Delta\log(l)$ change in log employment; and $\Delta\log(w)$ change in log average wages. All regressions include controls for the logs of labor and value added (in 2012), an indicator for group affiliation (whether the firm is part of a larger corporation), and fixed effects for industry (at the 3-digit SNI level) and commuting zone of the largest establishment. Employment-weighted regressions weight each firm by its initial employment. Standard errors are clustered at the industry level. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Comparing Panel A of Table 1 to the coefficients found in Acemoglu et al., 2020, we find overall similar effects using the Swedish data. The effects on value added, the labor share, value added per worker, employment, and wages go in the same direction and are significant. Quantitatively however, the effects are more pronounced in the Swedish data. For comparison, the coefficient for the change in value added is .204 in the French data, while we arrive at a value of .321. Two notable departures are worth pointing out: the skill share and the change in TFPR. For both, our results have opposite signs compared to the French results, and they are not significant at the 90% confidence level. We see a similar pattern after using employment weights.

In Panel B, we apply the same method to the adoption of automation capital. The results for value added, the labor share, and value added per worker have the same sign but are more muted and not significant at the 90% level. We interpret this as the bundling of multiple types of heterogeneous production capital partially netting out the aggregate effects. We will come back to this point in Section C, where we report results for granular categories of automation capital.

Finally, we use the change in the log stock of machinery and define a large investment as a log change that is above the 90th percentile across all years and firms in Panel C. Firms expanding their machinery stock show higher value added, employment, and wages, an unchanged labor share and skill share, and lower revenue productivity (TFPR). This mirrors the scale pattern of robots and automation capital: more machinery investment coincides with firm growth and hiring.

One key takeaway from this replication is that most of the main results are similar to those in the French data, but stronger. This also serves as a helpful benchmark for the following sections of this paper. At the same time, bundling multiple capital goods into one adoption indicator seems to mute effects, possibly because of heterogeneity among asset types within the bundle canceling each other out. This emphasizes the importance of taking capital heterogeneity into account in this type of analysis.

B. Robots in a mass of disaggregated capital

In this section, we present results from estimating the same specification as in the main replication for all 8-digit capital goods, with automation capital and industrial robots marked out. Together, there are 1,017 disaggregated capital classes. We can then observe, using standard methods in the literature, how robots compare with other capital goods.

In the main text, we focus on the results for value added and the labor share of value added for two reasons. First, they summarize neatly the effects on scale and the input mix between capital and labor. Second, these benchmark results are most clearly comparable across datasets. Additional results for employment and wages are in Appendix E.

Figure 1 puts robots in perspective by comparing their estimated effects with those of other 8-digit capital goods. Each point represents one product. The horizontal axis shows the change in value added following adoption, while the vertical axis shows the change in the labor share. For legibility the plot is trimmed at the 1st and 99th percentiles of both axes; this affects the display only, and all reported coefficients and counts use the untrimmed sample (see Appendix C).

A product is *natively* significant in a dimension when the unadjusted p-value of that coefficient is below .05, and FDR-significant when it survives a Benjamini–Hochberg correction at a 5% false discovery rate. The correction matters because testing over 1,000 products generates significant coefficients by chance. We apply it separately to each outcome, which controls the expected share of false discoveries among the coefficients classified as significant. The colors then distinguish six mutually exclusive categories, assigned from the most to the least demanding: FDR-significant in both dimensions; natively significant in both dimensions and FDR-significant in one; FDR-significant in one dimension; natively significant in both dimensions; natively significant in one dimension; and insignificant. Because the most demanding category a product qualifies for takes precedence, each product appears exactly once. Appendix C sets out the ranking procedure, the q-values, and the paired "and"/"or" definitions in more detail.

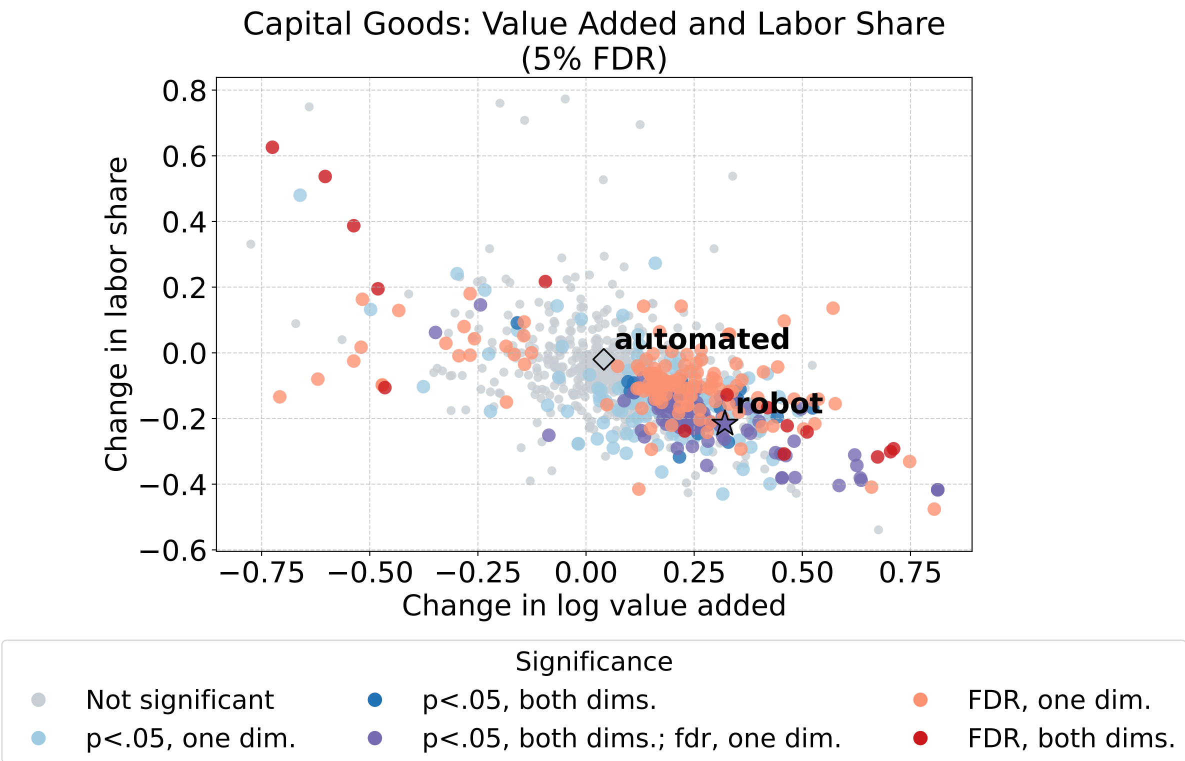

Figure 1: Estimated effects of capital adoption on value added and the labor share across 8-digit capital goods. Colors distinguish significance under the unadjusted tests and FDR correction.

Most products lie in the lower-right quadrant: adoption is associated with higher value added and a lower labor share. This suggests that production is scaled up more than labor inputs, as value added rises while labor compensation accounts for a smaller share of production. Whether this relative decline in labor's share translates into worse outcomes for workers is not immediately clear.

In fact, most significant estimated effects on the number of workers and wages per worker are positive (323 of 365 significant employment estimates and 167 of 215 significant wage estimates; see Figures A.4 and A.6). This means that a lower labor share that goes hand in hand with capital adoption does not necessarily imply negative effects for workers. This is in line with findings such as those in Hötte et al., 2023, who conclude that the displacement effect is often offset by scale effects. While labor plays a less important role in the production process, the increased productivity and associated scaling up of the production capacity lead to an increase in labor demand at the firm level.

Both industrial robots and automation capital fall in the lower-right quadrant. Robots fit a broader pattern across capital goods, alongside 718 other capital goods in this quadrant, of which 137 are significant, rather than representing a qualitatively distinct form of adoption. Automation capital, by contrast, has a small, insignificant effect. Moreover, across all products, 20 are significant under the 5% FDR threshold in both dimensions and 260 in at least one dimension, while under the native (unadjusted) definition 152 are significant in both dimensions and 426 in at least one. Robots are natively significant in both dimensions and survive the 5% FDR correction for value added but not for the labor share. Overall, robots are not that special, in terms of both economic and statistical significance.

We now dig deeper into which kinds of capital have effects similar to those of robots.

Figure 2 plots all products that are significant on both value added and the labor share, colored by HS chapter, with the robot benchmark marked.

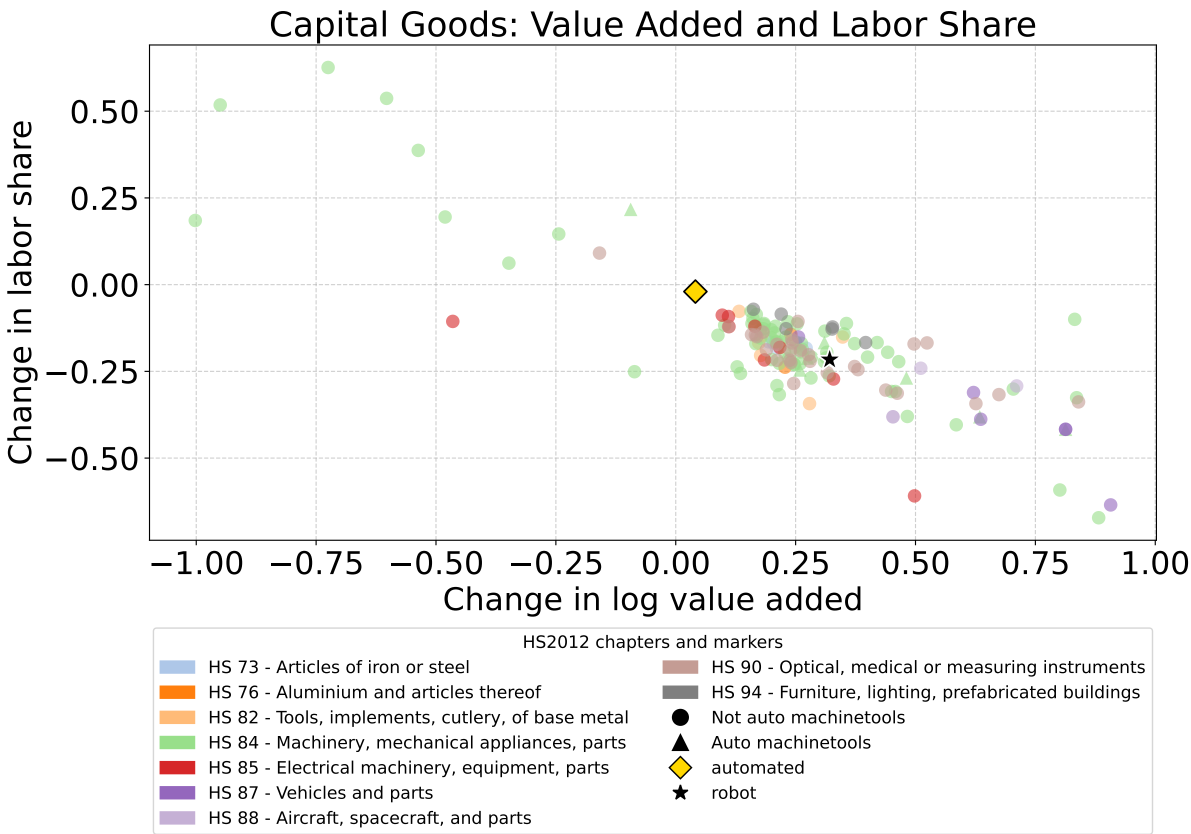

Figure 2: Significant results for value added and the labor share by HS Chapters.

11 capital goods that are significant on both dimensions lie in robots' immediate neighborhood, with effects on value added and the labor share of similar magnitude; Figure A.3 provides a zoomed-in view. These include precision and numerically controlled machine tools, measuring and analytical instruments, and pumps. Notably, several are more conventional technologies not directly related to automation, such as diesel engines, hand-held power tools, and air heaters, yet they show effects comparable to those of robots. This underscores that robots are part of a broader set of capital inputs that jointly shape firm performance and workforce composition.

This interpretation is further strengthened by the 34 capital goods with significant effects on both value added and the labor share that are stronger than those estimated for robots. Among these are ultrasound scanners, harvesting machinery, conventional grinding machines, electric generating sets, and road rollers.

Overall, we find that robots do not seem to be linked to stronger effects on firm scale and labor shares than other, more traditional types of capital and equipment. This suggests that the pattern may reflect a general feature of capital deepening rather than automation per se: adopting new capital that reshapes the production process tends to raise output and shift the input mix toward capital, a pattern visible even for older, conventional technologies.

C. Automation capital unpacked

From the results observed above, one extra striking fact stands out, that automation capital has a small and insignificant effect, which is not in line with previous literature. To understand how this result is generated, Figure 3 depicts the estimates for value added and the labor share of all the disaggregated capital classes included in automation capital. The aggregate estimate is marked with a diamond and is both economically and statistically insignificant as already shown, even though many of its constituent members are significant in both aspects.

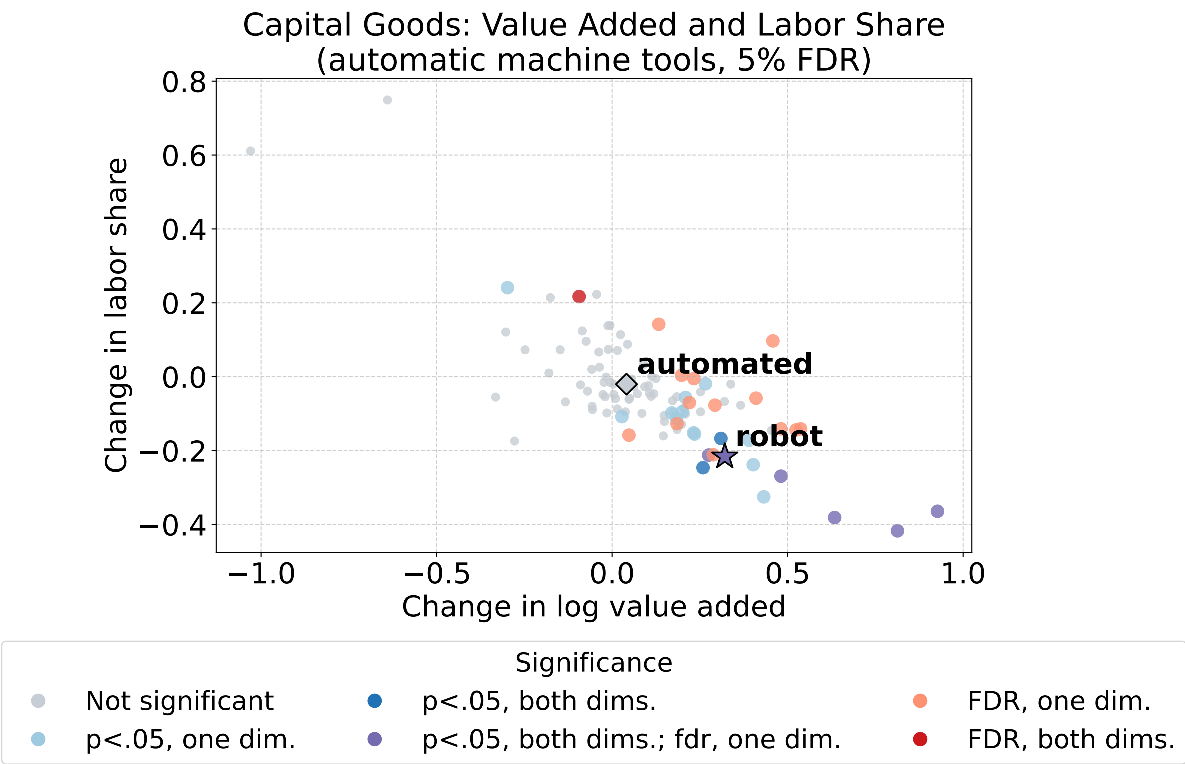

Figure 3: Capital classes included in *automation capital*, color coded by significance.

Within automation capital, 8 capital classes are significant for both outcomes (Figure 3); 7 increase value added while reducing the labor share. The strongest estimate is for gear-cutting, grinding, and finishing machine tools, whose estimated effect on value added is nearly three times that of robots.[9]

This illustrates how aggregating a number of capital classes into one broad group of *automation capital* affects the interpretation of results in two ways. First, if capital goods within the group have heterogeneous effects, with potentially different signs, then the aggregate effect can be centered around a near-zero effect size, even if the underlying capital goods have potentially significant and interesting effects. Second, the aggregation can lead to insignificant results even if (some of) the underlying estimates would be significant. Summarizing, the level of aggregation matters for both the effect size and its significance, which in turn plays a role in interpreting the interaction between capital and labor.

D. Too narrow or too broad?

In this section, we extend the previous analyses by taking a closer look at the estimated effects of disaggregated capital classes, with two main takeaways emphasizing the importance of taking into account heterogeneity in capital not just in theory, but empirically.

[9] KN2012 No. 84614071: "Machine-tools for gear cutting, gear grinding or gear finishing, working by removing metal, sintered metal carbides or cermets." Its estimated effect on value added is 0.927, compared with 0.321 for robots.

On the one hand, zooming in to one particular capital class, like *robots*, can lead to a lack of context for the interpretation of effect sizes, and a disproportionate weight being placed on that particular technology. On the other hand, aggregating a number of capital classes into one indicator as is the case with *automation capital*, can mask significant and sizable effects going on within the aggregated class. This point relates to Krusell et al. (2000) who demonstrate that lumping equipment and structures together is problematic as they have opposite labor-market implications.

V. Aggregate employment effects

While firm-level estimates deliver an insight into dynamics at adopting firms, such as effects on the structure of the workforce and the labor share, the debate rightfully centers on aggregate effects. At the core of this debate is the question whether automation replaces labor in the aggregate, or whether productivity gains and expansion dominate.

At this aggregate level, ALR and AABJ come to somewhat different conclusions using French firm-level data. While the former find that robot adoption decreases labor demand, employment grows as a consequence of automation in the latter work. As automation continues to be a much-debated topic in politics, with real-world consequences for policy, it is crucial to understand why this difference in conclusions based on data from the same country emerges.

To that end, this section first summarizes how ALR and AABJ differ in methodology, proceeds to replicate both in Swedish data, and finally explains how method and measure affect conclusions.

A. On method and measure

The analyses in ALR and AABJ differ in two dimensions: estimation method and the measure of automation. ALR estimate firm-level effects of industrial robot adoption for the importing firm and for competitors in France between 2010 and 2015. This estimation happens jointly, where the exposure to robot adoption by competitors is measured as an employment-weighted exposure measure:

$$exposure = \sum_i m_{fi} * \sum_{f' \neq f} s_{if'} Robot_{f'}, \tag{2}$$

where $Robot_{f'}$ is an indicator for robot imports by firm f'. The share in total firm f sales that happens in industry i is defined as m_{fi}, while the share in total industry i sales generated by firm f' is defined as $s_{if'}$. Finally, the estimates of own (β_0) and competitor (β_1) effects are aggregated at the market level according to:

$$\Delta \log L = \beta_0 \sum_f \left(\frac{l_f}{L} \cdot \text{robot}_f \right) + \beta_1 \sum_f \left(\frac{l_f}{L} \cdot \text{robot}_f \cdot \sum_i m_{fi}(1 - s_{if}) \right), \tag{3}$$

where l_f is the number of workers at firm f in the first year, and $L = \sum_f l_f$ is total employment.

They conclude that the observed robot adoption by firms accounting for 20% of the workforce was associated with a 3.2% decrease in industry employment. As in the firm-level estimates, this is done by comparing outcomes in 2015 and 2010, with robot adoption being recorded for the whole time span.

Rather than aggregating up using estimates at the firm level, AABJ run an event study at the 5-digit industry level. They calculate the year-to-year log change in machinery stock at the industry level and define an automation event as an increase that is above the median for all year-industry combinations. They find that industry employment increases by about 10% within one or two years, and stays at that level afterwards until the end of the event horizon five years later, when looking at French data from 1995 to 2017.

Summarizing, there is a difference in what is being estimated, and a difference in measurement of automation. While ALR focus on robots as the automation capital, AABJ use a broader measure of machine capital. In addition, they rely on different time periods. Ex-ante it is unclear where the qualitatively different results originate. In what follows, we replicate both methods for both measures of automation for Sweden to tease out the role of method and measure.

B. Aggregate results

The results of this exercise are summarized in Table 2, which covers aggregation from firm-level estimates, and Figure 4, containing event graphs for both machinery stock increases and robot imports. We define the machinery adoption in Table 2 as an increase in machine stock at the firm level between 2012 and 2016. A substantial robot import event at the industry level in Figure 4 is defined as an industry importing robots relative to sales above the median across all years and industries between the years 2002 and 2022.[10]

We then estimate the following equation for 5-digit industries:

$$\log L_{it} = \sum_{k=-5}^{5} \delta_k E_{i,t-k} + \mu_i + \lambda_t + \varepsilon_{it}, \tag{4}$$

where L_{it} is industry i employment in year t, $E_{i,t-k}$ is an indicator that takes the value one if there is an event in industry i in year $t-k$ and zero otherwise. We include fixed effects for the industry (μ_i) and the year (λ_t). Then, all δ_k are normalized by the coefficient of the omitted period $k^{omit} = -1$.

[10]This refers to the median normalized input among positive year-firm import events.

A first take-away is that the exact replications yield results very similar to those in the original papers. Aggregating up the firm-level robot adoption estimates implies a decrease in employment of approx. 2.4% as a response to firms importing robots accounting for 16.6% of initial employment. The corresponding numbers in ALR are a 3.2% decline in employment as a response to 20% of workers being initially employed by robot adopters.

Similarly, the left panel of Figure 4 reveals a similar pattern for a large industry-level increase in the stock of machinery as the one in AABJ: an increase in employment in the event year and subsequent years, leveling off at around 0.125, compared to 0.1 in the French data. This means that on average, a large machinery investment in an industry is associated with an increase in employment by about 10% in France, and about 12.5% in Sweden.

	Own contribution	Competitor contribution	Employment share adopting (percent)	Aggregate effect (percent)
Robot adoption	0.006	-0.030	16.6	-2.44
Machinery adoption	0.127	-0.032	53.5	10.02

Table 2: Aggregate Employment Effects of Robot and Machinery Adoption
Notes: The table decomposes the aggregate employment effect into the employment-weighted contribution of a firm's own adoption and the employment-weighted contribution of competitors' adoption. The aggregate log effect is the sum of these contributions. The third column reports the initial-employment-weighted mean of the firm's adoption indicator: the share of employment accounted for by adopting firms. The final column converts the aggregate log effect to an exact percentage, $100[\exp(\widehat{\Delta \log l}) - 1]$. The regressions contain 16,251 observations.

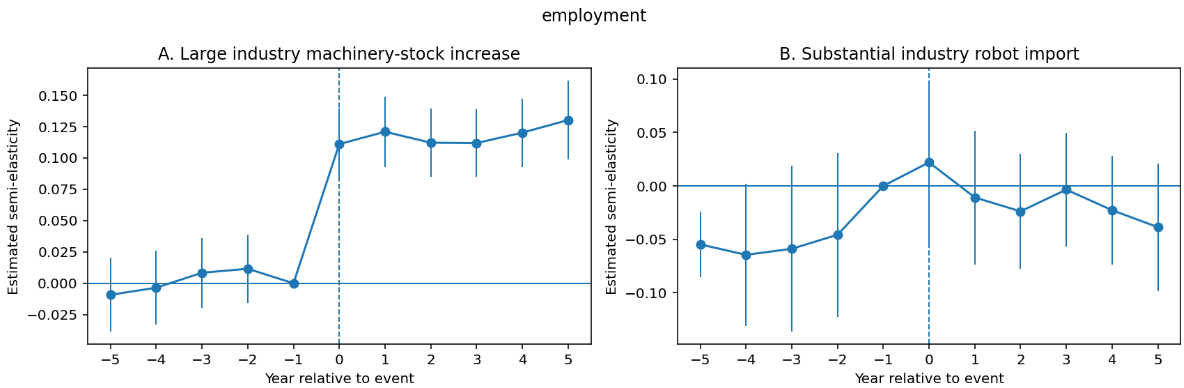

Figure 4: Industry-level event studies
Notes: The event in A., a large industry machinery stock increase, is defined as $d\log(k)_t = \log(k)_t - \log(k)_{t-1}$ above the median across all years 2002-2022 and 5-digit industries. A substantial industry robot import is defined as $\tilde{x}_t = x_t/y_t$ above the median, where x_t is the value of robot imports and y_t industry sales. The reference period is $t-1$.

While the replications using the same method and measure deliver the same qualitative results, crossing over the two delivers new insights. Applying the aggregation method from ALR to machinery adoption yields a positive effect on employment, as machinery increases by firms covering 53.5% of initial employment is associated with employment growing by 10%. The industry level event studies find no clear effects after a robot adoption event.[11]

This exercise shows that in the debate on the automation effects for workers, the definition of automation matters a lot. Using the same estimation strategy allows for opposing qualitative conclusions, while using a different estimation strategy does not seem to be decisive. We must therefore carefully assess what we mean by automation. Picking out a very concrete technology, such as robots, and taking the results as applying to automation in general does not seem to work. At the same time, simply measuring the effects of machinery investment more broadly might fail to observe the potential of modern, growing technologies to disrupt labor markets.

VI. Conclusion

This paper examines the role of capital heterogeneity in the empirical measurement of automation. Using matched Swedish administrative data linking firms, workers, and highly disaggregated imports of capital goods, we replicate the main empirical approaches in the recent automation literature under a common institutional setting while systematically varying the definition of automation capital. This allows us to separate differences arising from measurement from differences arising from estimation.

Our results show that capital heterogeneity is an important determinant of empirical conclusions. First, industrial robots are not exceptional when viewed alongside the broader distribution of capital goods. Although robot adoption is associated with higher value added, lower labor shares, and firm expansion, many other types of production capital exhibit effects of similar or even greater magnitude. This suggests that evidence based on a single salient technology should be interpreted relative to the broader distribution of capital investments before drawing conclusions about automation more generally.

Second, aggregating heterogeneous technologies into a broad measure of automation can conceal economically meaningful variation. Within the commonly used automation classification, individual capital classes display substantially different relationships with firm outcomes, with some exhibiting strong positive effects and others showing weak or offsetting patterns. As a result, broad automation measures may produce muted or insignificant estimates even when many of their constituent technologies have important effects individually.

Finally, comparing the aggregate approaches of Acemoglu et al. (2020) and Aghion et al. (2025) under the same data environment suggests that differences in the definition of automation capital matter at least as much as differences in estimation strategy. Applying alternative empirical methods to the same measure yields broadly similar qualitative

[11] Note that AABJ includes an exercise using robots as the measure of automation in the Appendix, but only for the firm-level results, where they find that their qualitative results prevail.

conclusions, whereas changing the measure of automation while holding the empirical approach fixed can reverse the sign of the estimated aggregate employment effect. This helps explain why the recent literature reaches seemingly conflicting conclusions despite studying closely related settings.

Beyond the automation literature, our findings illustrate two more general lessons for empirical work using increasingly granular data. The first is that focusing on a single, highly visible technology without benchmarking it against comparable observations risks overstating its broader importance. In the context of automation, the prominence of industrial robots has naturally attracted considerable attention, but our evidence suggests that they are one example of a wider process of capital deepening and technological adoption rather than a fundamentally distinct phenomenon. This distinction is important for policy discussions. For example, proposals such as taxes targeted specifically at robots implicitly assume that robots have effects that differ systematically from other forms of capital investment. Our results provide little evidence in support of that view.

The second lesson concerns aggregation. Researchers routinely combine detailed observations into broader categories in order to simplify empirical analysis, yet the appropriate level of aggregation is rarely obvious ex ante. When the underlying components are heterogeneous, aggregation may wash out meaningful variation and obscure the economic mechanisms of interest. The increasing availability of highly disaggregated administrative and transaction data makes it possible to examine these issues directly rather than imposing broad classifications from the outset.

These considerations are likely to become increasingly relevant as research shifts toward newer digital technologies, including artificial intelligence. Like production capital, AI encompasses a wide range of heterogeneous technologies that differ in how they interact with workers and production processes. Our results suggest that careful attention to measurement and technological heterogeneity will be essential for interpreting empirical evidence and for understanding the labor market consequences of future technological change.

References

Acemoglu, D., Lelarge, C., & Restrepo, P. (2020). Competing with Robots: Firm-Level Evidence from France. *AEA Papers and Proceedings*, *110*, 383–388. https://doi.org/10.1257/pandp.20201003

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

Acemoglu, D., & Restrepo, P. (2022). Demographics and automation. *The Review of Economic Studies*, *89*(1), 1–44.

Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2023). The local labor market effects of modern manufacturing capital: Evidence from france. *AEA Papers and Proceedings*, *113*, 219–223.

Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2022). The effects of automation on labor demand: A survey of the recent literature. In *Robots and ai* (pp. 15–39). Routledge.

Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2025). Modern manufacturing capital, labor demand, and product market dynamics: Evidence from france. https://www.xavierjaravel.com/_files/ugd/bacd2d_0453767d96e14aeb8c4c06818e118b23.pdf#page=19.23

Alguacil, M., Turco, A. L., & Martínez-Zarzoso, I. (2022). Robot adoption and export performance: Firm-level evidence from spain. *Economic Modelling*, *114*, 105912.

Bessen, J., Goos, M., Salomons, A., & Van den Berge, W. (2023). Automatic reaction-what happens to workers at firms that automate? *The Review of Economics and Statistics*, (Feb. 6, 2023).

Bonfiglioli, A., Crino, R., Fadinger, H., & Gancia, G. (2024). Robot imports and firm-level outcomes. *The Economic Journal*, *134*(664), 3428–3444.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *American economic review*, *101*(1), 304–340.

Caselli, F., & Wilson, D. J. (2004). Importing technology. *Journal of monetary Economics*, *51*(1), 1–32.

Chen, C. (2020). Capital-skill complementarity, sectoral labor productivity, and structural transformation. *Journal of Economic Dynamics and Control*, *116*, 103902. https://doi.org/https://doi.org/10.1016/j.jedc.2020.103902

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, *19*(6), 3104–3153.

Dinlersoz, E., & Wolf, Z. (2024). Automation, labor share, and productivity: Plant-level evidence from us manufacturing. *Economics of Innovation and New Technology*, *33*(4), 604–626.

Dixon, J., Hong, B., & Wu, L. (2020). *The employment consequences of robots: Firm-level evidence*. Statistics Canada Ontario.

Graetz, G., & Michaels, G. (2018). Robots at work. *Review of economics and statistics*, *100*(5), 753–768.

Hötte, K., Somers, M., & Theodorakopoulos, A. (2023). Technology and jobs: A systematic literature review. *Technological Forecasting and Social Change*, *194*, 122750.

Humlum, A. (2019). Robot adoption and labor market dynamics. *Princeton University*.

Koch, M., Manuylov, I., & Smolka, M. (2021). Robots and firms. *The Economic Journal*, *131*(638), 2553–2584.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., & Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, *68*(5), 1029–1053.

Luo, H., & Qiao, H. (2024). Exploring the impact of industrial robots on firm innovation under circular economy umbrella: A human capital perspective. *Management Decision*, *62*(9), 2763–2790.

O'Mahony, M., Vecchi, M., & Venturini, F. (2021). Capital heterogeneity and the decline of the labour share. *Economica*, *88*(350), 271–296.

Robinson, J. (1953). The production function and the theory of capital. *The review of economic studies*, *21*(2), 81–106.

Tverdostup, M., Ghodsi, M., & Leitner, S. M. (2025). *Migration vs. automation as an answer to labour shortages: Firm-level analysis for austria* (tech. rep.). wiiw Working Paper.

Voshaar, J., Loy, T. R., & Koch, M. (2022). Firm-level robot use and labor cost behavior. *Available at SSRN 4315287*.

Wang, K., Zhou, J., Li, G., Hu, Y., & Hu, F. (2025). Industrial automation and product quality: The role of robotic production transformation. *Applied Economics*, *57*(34), 5164–5179.

Wang, T., Zhang, Y., & Liu, C. (2024). Robot adoption and employment adjustment: Firm-level evidence from china. *China Economic Review*, *84*, 102137.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Zator, M. (2019). Digitization and automation: Firm investment and labor outcomes. *Available at SSRN 3444966*.

Zhang, H., Chen, Z., & Wei, Y. (2025). Robot adoption and export sophistication: Firm-level evidence from china. *Journal of Asian Economics*, *98*, 101891.

Appendix

A. *Summary Statistics*

This is the summary statistics of the dataset by subsample.

	All manufacturing	Capital importers	Automation	Robots
Revenue (*)	103.706	326.730	1140.692	5743.977
No. of workers	29.1	82.2	284.3	1156.6
% Robot adoption	0.4	1.5	4.2	100.0
% Automation	3.4	11.8	100.0	33.8
Labor productivity (*)	0.690	0.920	0.951	1.248
Value added (*)	39.613	125.465	469.099	2662.568
Labor share	0.643	0.561	0.549	0.515
Wage (*)	0.307	0.355	0.373	0.392
% Skilled labor	15.2	23.7	25.0	31.6
% Female labor	13.4	18.4	17.9	22.8
No. of firms	16,259	4,665	549	68

All variables measured at the 2012 baseline. (*) In 1 million SEK.

Table A.1: Summary statistics by group at baseline

Figure A.1 shows the average adoption rate of automation capital and robots, respectively, by revenue size quintile within an industry, illustrating that adoption of both types of capital is concentrated among larger firms.

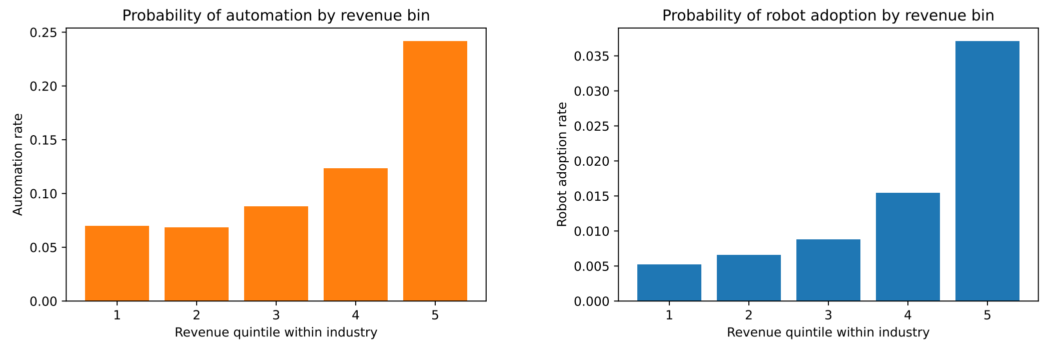

Figure A.1: Adoption rates of automation capital and robots by size decile within an industry.

$\Delta\log(va)$	Δlabshare	Δskillshare	$\Delta\log(va/l)$	$\Delta\log(TFPR)$	$\Delta\log(l)$	$\Delta\log(w)$	n
A. Robot adoption							
Baseline (unweighted)							
0.260	-0.412	-0.014	0.223	0.009	0.037	0.071	80.0
**	**		**				
Employment weighted							
0.191	-0.385	-0.015	0.168	0.044	0.023	0.013	80.0
*	*		**				
B. Automation capital adoption							
Baseline (unweighted)							
0.017	-0.044	0.004	-0.013	-0.007	0.030	0.004	580.0
Employment weighted							
0.006	-0.078	0.003	-0.030	-0.011	0.036	0.001	580.0
C. Machinery stock change in logs							
Baseline (unweighted)							
0.054	-0.001	0.000	0.006	-0.093	0.049	0.003	-
***			*	***	***		
Employment weighted							
0.066	-0.026	-0.002	0.005	-0.082	0.061	0.000	-
***				***	***		

Table A.2: Robot and Automation Effects (2010–2015)

Notes: The table reports coefficients from regressions of changes in firm outcomes between 2010 and 2015.
Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

These results use a five-year difference to match the specification in Acemoglu et al. (2020). The magnitude and direction of effects are broadly comparable to our baseline results, suggesting that differences in time aggregation do not drive the main findings.

B. Additional regression results

Table A.3 reports the specification of Table 1 for additional outcomes.

$\Delta\log(y)$	Δlabshare(y)	$\Delta\log(y/l)$	$\Delta\log(l)$	Δskillshare	$\Delta\log(l^f)$	$\Delta\log(age)$	n
A. Robots							
Baseline (unweighted)							
0.195	0.000	0.064	0.131	-0.000	0.088	-0.011	68.0
**			*				
Employment weighted							
0.149	0.004	0.026	0.122	-0.012	0.102	0.001	68.0
*			***		*		
B. Automation							
Baseline (unweighted)							
0.044	-0.002	0.013	0.031	0.065	0.002	0.003	549.0
				**			
Employment weighted							
0.023	0.002	-0.014	0.037	0.051	0.026	-0.002	549.0
				*			

Table A.3: Robot and Automation Effects (additional outcomes)

Notes: The table reports coefficients from regressions of changes in firm outcomes between 2012 and 2016 on an indicator for whether the firm adopted either robots (Panel A) or automation capital (Panel B). The dependent variables are: $\Delta\log(y)$ — change in log sales; $\Delta labshare(y)$ — change in the labor share of sales; $\Delta\log(y/l)$ — change in log sales per worker (labor productivity); $\Delta\log(l)$ — change in log employment; Δskillshare — change in the share of skilled workers; $\Delta\log(l^f)$ — change in log number of female employees; and $\Delta\log(age)$ — change in log firm age. All regressions include controls for the logs of labor and value added (in 2012), an indicator for group affiliation (whether the firm is part of a larger corporation), and fixed effects for industry (at the 3-digit NACE level) and commuting zone of the largest establishment. Employment-weighted regressions weight each firm by its initial employment. Standard errors are clustered at the industry level.
Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

C. Technical Details Behind the Scatter Plots

For each outcome, we estimate one coefficient and its p-value for every 8-digit capital product. Under the unadjusted, or "native," definition, a coefficient is significant when its p-value is below 0.05, which is the same as carrying at least two stars in the disaggregated results tables. This rule is appropriate for reading a single regression, but applying it across many products generates false positives by chance.

We therefore apply the Benjamini–Hochberg (BH) procedure separately to each outcome. Let $p_{(1)} \leq p_{(2)} \leq \cdots \leq p_{(m)}$ denote the ordered product-level p-values, where m is the number of products. For a target FDR level α, let

$$k = \max\left\{i : p_{(i)} \leq \tfrac{i}{m}\,\alpha\right\},$$

and classify all coefficients with rank $i \leq k$ as significant. Equivalently, the BH-adjusted

q-value of the coefficient ranked i is

$$q_{(i)} = \min\left\{1,\ \min_{j \geq i} \frac{m}{j}\, p_{(j)}\right\},$$

and a coefficient is FDR-significant when $q_{(i)} \leq \alpha$. Under the standard BH conditions, this controls the expected share of false discoveries among the coefficients classified as significant. The correction is applied to each dimension separately.

For an outcome pair (a, b), the "and" definition classifies a product as significant when both outcomes survive the correction ($q_a \leq \alpha$ and $q_b \leq \alpha$), while the "or" definition requires at least one. The native versions of these rules apply the unadjusted $p < 0.05$ rule in place of the q-value rule. We use a 5% FDR threshold ($\alpha = 0.05$) in the main results.

Outlier trimming. A small number of 8-digit products have extreme estimated coefficients that would otherwise compress the remaining observations into an unreadable cluster. For legibility, the scatter plots are therefore trimmed: within each plotted sample we compute the 1st and 99th percentiles of both axis variables and drop products that fall outside that range on either axis. Trimming affects only what is displayed. All coefficients, counts, and significance classifications reported in the text and in the summary tables are computed on the untrimmed sample.

Two features of this rule are worth noting. First, because the percentiles are computed on the sample actually being plotted, the cutoffs differ across figures: for the full cloud of 1,020 products the trim removes 36 products, whereas for the subsample that is significant on both value added and the labor share it removes 6. The corresponding cutoffs and the list of dropped products are stored alongside each figure. Second, the robot and automation-capital benchmarks are drawn separately from the cloud, so they are always displayed regardless of whether the trim would otherwise exclude them.

The one exception is Figure 3, which plots the members of *automation capital*. It is left untrimmed so that every member of the group appears, since the point of that figure is the dispersion within the aggregate.

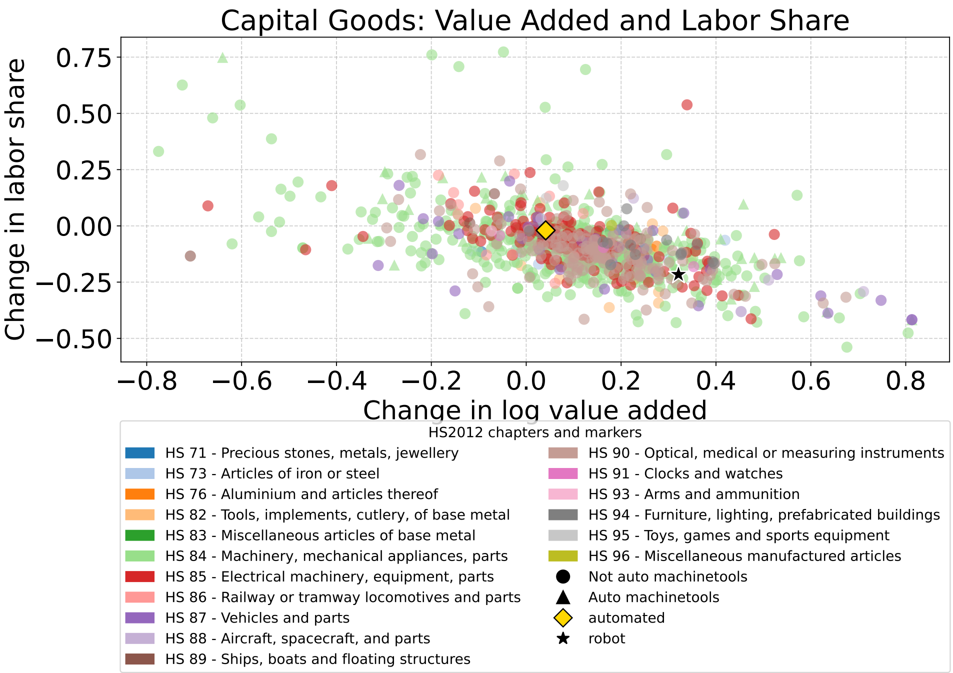

Figure A.2: Value added and the labor share for all 8-digit capital goods, regardless of significance. Points are colored by HS chapter.

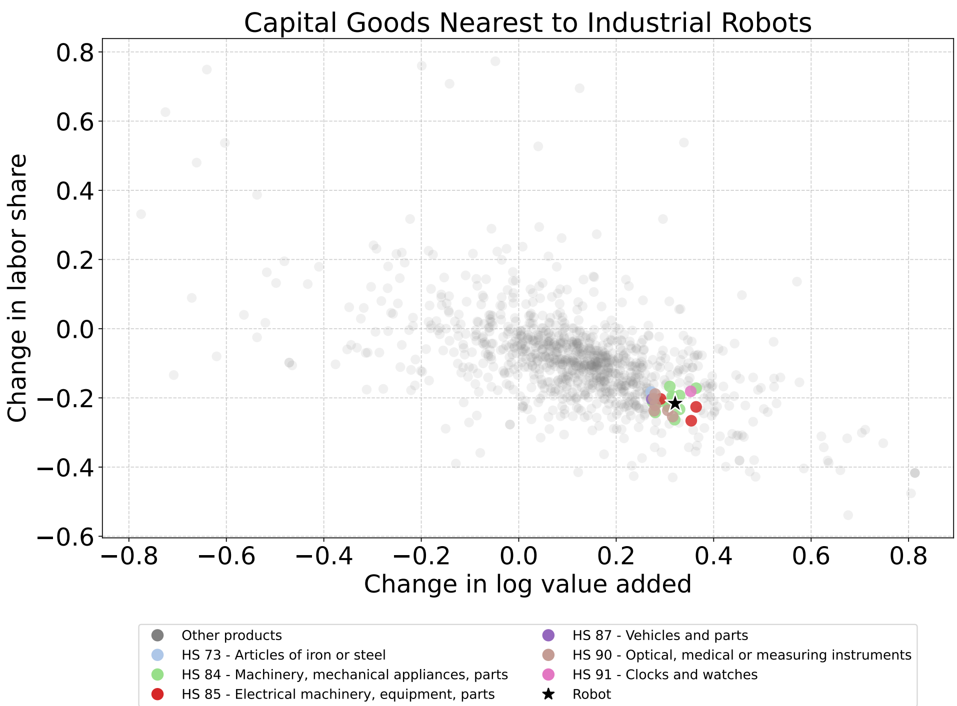

Figure A.3: Value added and the labor share for the capital goods whose effects are most similar to robots, regardless of significance. Points are colored by HS chapter.

E. Employment and Wages Across Capital Goods

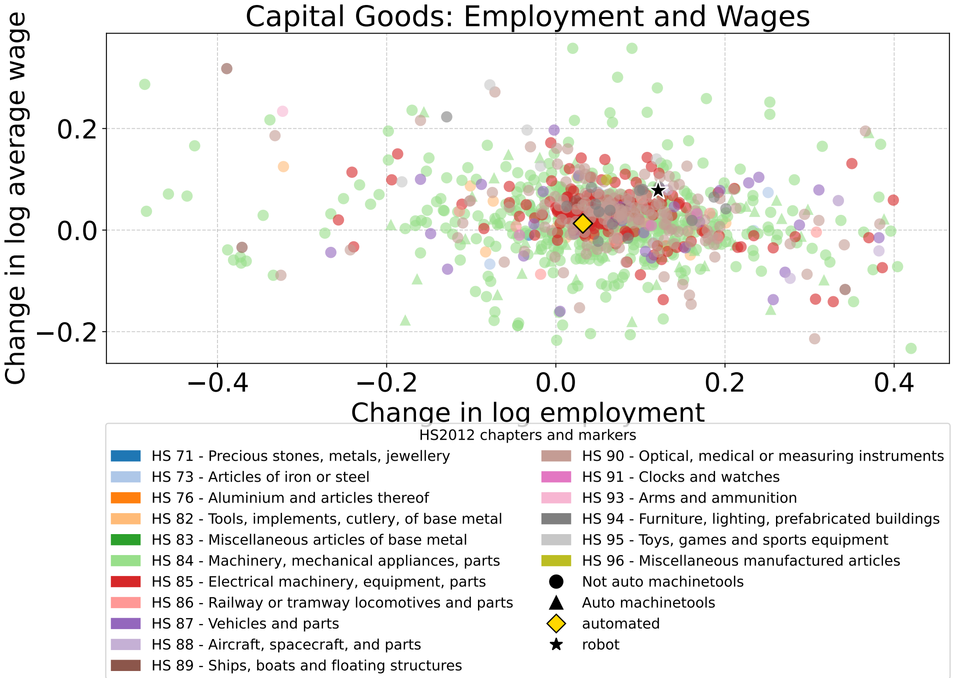

Figure A.4: Employment and wages for all 8-digit capital goods, regardless of significance. Points are colored by HS chapter.

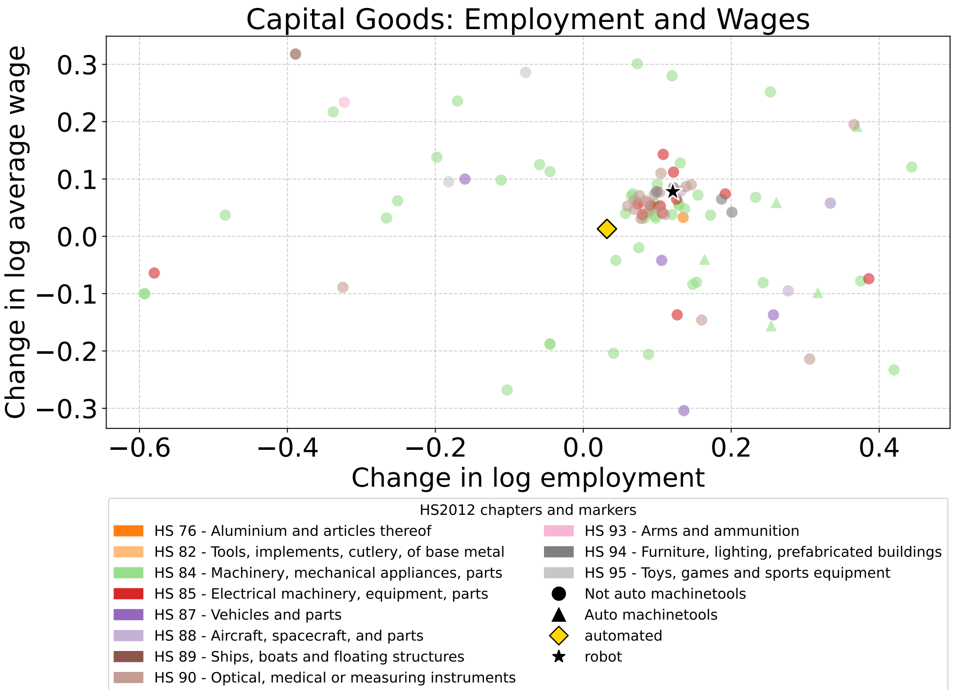

Figure A.5: Employment and wages for capital goods significant at the 5% level (unadjusted) for both outcomes. Points are colored by HS chapter.

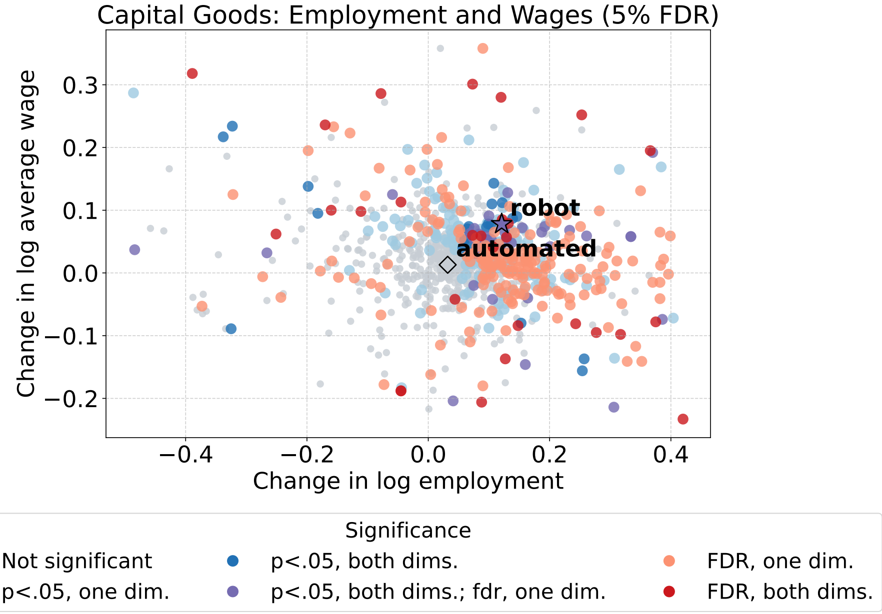

Figure A.6: Employment and wages for all capital goods, colored by significance bucket. Buckets combine unadjusted and 5% FDR significance; where a product qualifies under both, the unadjusted (two-dimension) label takes priority.

F. *Aggregation and the Automation Capital Group*

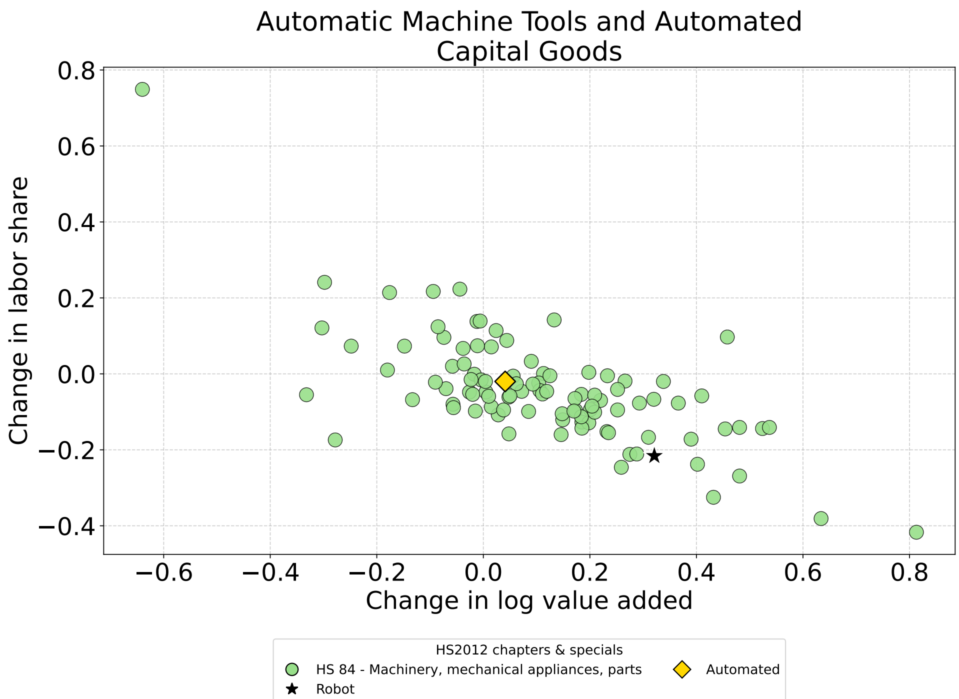

Figure A.7: Value added and the labor share for the capital goods within *automation capital*, regardless of significance. The aggregate automation-capital estimate is marked with a star.

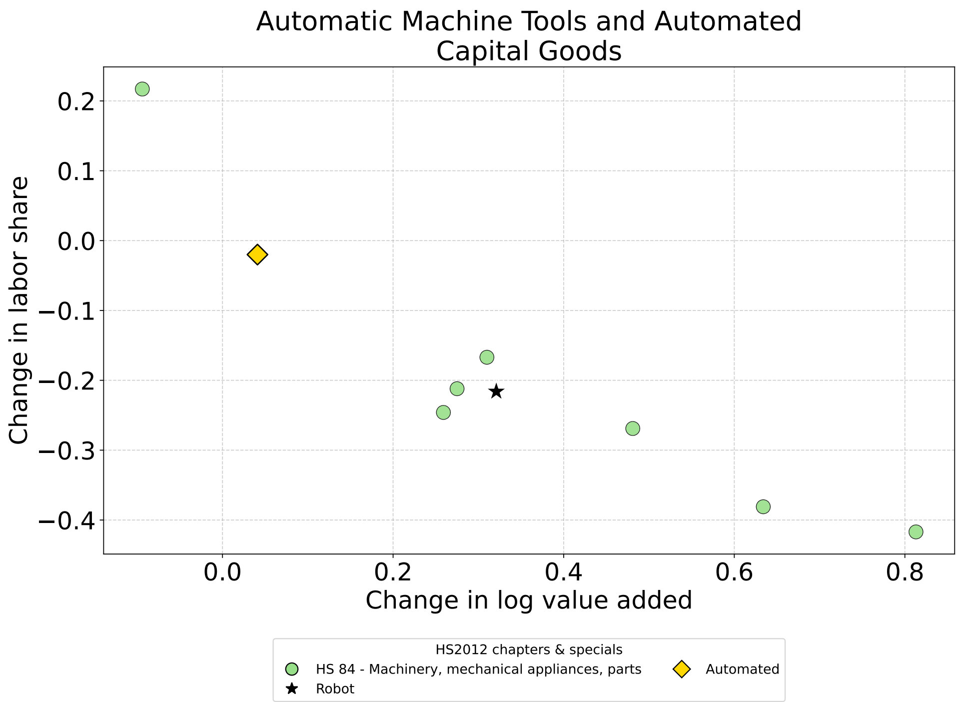

Figure A.8: Value added and the labor share for the automation-capital goods significant at the 5% level (unadjusted) for both outcomes. Points are colored by HS chapter.

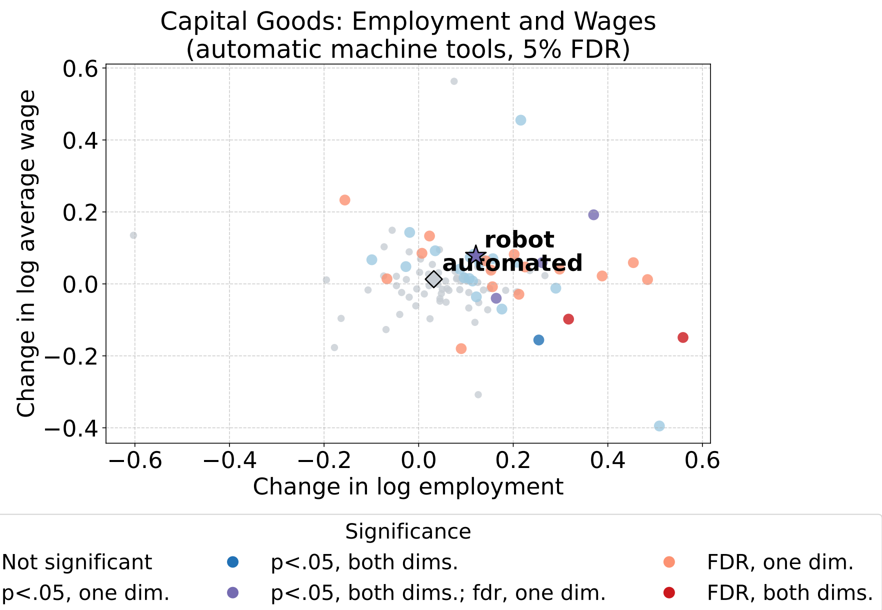

Figure A.9: Employment and wages for the automation-capital goods, colored by significance bucket. Buckets combine unadjusted and 5% FDR significance; where a product qualifies under both, the unadjusted (two-dimension) label takes priority.